Visual Analytics to Support Evidence-Based Decision Making

Vom Fachbereich Informatik
der Technischen Universität Darmstadt
genehmigte

DISSERTATION

zur Erlangung des akademischen Grades eines
Doktor-Ingenieurs (Dr.-Ing.)
von

Dipl.-Math. Tobias Ruppert
geboren in Flörsheim am Main, Deutschland

Referenten der Arbeit:	Prof. Dr. Dieter W. Fellner
	Technische Universität Darmstadt
	Prof. Dr. Jörn Kohlhammer
	Technische Universität Darmstadt
	Prof. Dr. Silvia Miksch
	Technische Universität Wien

Tag der Einreichung:	14.09.2017
Tag der mündlichen Prüfung:	23.11.2017

Darmstädter Dissertation
D 17
Darmstadt 2018

Abstract

The aim of this thesis is the design of visual analytics solutions to support evidence-based decision making. Due to the ever-growing complexity of the world, strategical decision making has become an increasingly challenging task. At the business level, decisions are not solely driven by economic factors anymore. Environmental and social aspects are also taken into account in modern business decisions. At the political level, sustainable decision making is additionally influenced by the public opinion, since politicians target the conservation of their power. Decision makers face the challenge of taking all these factors into consideration and, at the same time, of increasing their efficiency to immediately react on abrupt changes in their environment. Due to the digitization era, large amounts of data are digitally stored. The knowledge hidden in these datasets can be used to address the mentioned challenges in decision making. However, handling large datasets, extracting knowledge from them, and incorporating this knowledge into the decision making process poses significant challenges. Additional complexity is added by the varying expertises of stakeholders involved in the decision making process. Strategical decisions today are not solely made by individuals. In contrast, a consortium of advisers, domain experts, analysts, etc. support decision makers in their final choice. The amount of involved stakeholders bears the risk of hampering communication efficiency and effectiveness due to knowledge gaps coming from different expertise levels. Information systems research has reacted to these challenges by promoting research in computational decision support systems. However, recent research shows that most of the challenges remain unsolved. During the last decades, visual analytics has evolved as a research field for extracting knowledge from large datasets. Therefore, combining human perception capabilities and computers' processing power offers great analysis potential, also for decision making. However, despite obvious overlaps between decision making and visual analytics, theoretical foundations for applying visual analytics to decision making have been missing.

In this thesis, we promote the augmentation of decision support systems with visual analytics. Our concept comprises a methodology for the design of visual analytics systems that target decision making support. Therefore, we first introduce a general decision making domain characterization, comprising the analysis of potential users, relevant data categories, and decision making tasks to be supported with visual analytics technologies. Second, we introduce a specialized design process for the development of visual analytics decision support systems. Third, we present two models on how visual analytics facilitates the bridging of knowledge gaps between stakeholders involved in the decision making process: one for decision making at the business level and one for political decision making. To prove the applicability of our concepts, we apply our design methodology in several design studies targeting concrete decision making support scenarios. The presented design studies cover the full range of data, user, and task categories characterized as relevant for decision making. Within these design studies,

we first tailor our general decision making domain characterization to the specific domain problem at hand. We show that our concept supports a consistent characterization of user types, data categories and decision making tasks for specific scenarios. Second, each design study follows the design process presented in our concept. And third, the design studies demonstrate how to bridge knowledge gaps between stakeholders. The resulting visual analytics systems allow the incorporation of knowledge extracted from data into the decision making process and support the collaboration of stakeholders with varying levels of expertises.

Zusammenfassung

Ziel dieser Arbeit ist das Design von Visual Analytics-Lösungen für die Unterstützung evidenzbasierter Entscheidungsfindung. Aufgrund der stetig wachsenden Komplexität der Welt, wird die strategische Entscheidungsfindung zu einer immer größeren Herausforderung. Auf Unternehmensebene werden Entscheidungen nicht mehr auf Basis rein ökonomischer Faktoren getroffen. Umweltbezogene und soziale Aspekte werden ebenfalls berücksichtigt in modernen Unternehmen. Auf politischer Ebene wird zudem die öffentliche Meinung in die politische Entscheidungsfindung mit einbezogen, da Politiker ihren Machterhalt durch Wiederwahl anstreben. Entscheidungsträger stehen vor der Herausforderung diese unterschiedlichen Faktoren in ihrer Entscheidungsfindung zu berücksichtigen und gleichzeitig die Dauer des Entscheidungsprozesses zu beschleunigen, um auf sich immer schneller ändernde Anforderungen reagieren zu können. Das Wissen, das benötigt wird, um diese Herausforderung zu bewältigen, steckt in großen Datenmengen, die dank der Digitalisierungs-Ära auch digital zur Verfügung stehen. Allerdings rufen das Bearbeiten großer Datenmengen, das Extrahieren von Informationen aus diesen, sowie das Verwenden des erhaltenen Wissens im Entscheidungsprozess weitere Herausforderungen hervor. Der Entscheidungsprozess wird zudem durch unterschiedliche Wissensstände der beteiligten Personen erschwert. Strategische Entscheidungen werden nur noch selten von einzelnen Personen getroffen. Im Gegenteil, ganze Konsortien bestehend aus Beratern, Analysten, Domänenexperten und anderen Interessenvertretern werden in heutige Entscheidungen mit einbezogen. Unterschiedliche Expertisen sorgen für Wissenslücken, die eine effiziente und effektive Kommunikation zwischen den involvierten Personen erschweren. Die Forschung im Bereich der Informationssysteme hat auf diese Herausforderungen reagiert und theoretische Grundlagen für die computergestützte Entscheidungsfindung geschaffen. Nichtsdestotrotz bleiben viele Herausforderungen ungelöst. Während des letzten Jahrzehnts wurde das neue Forschungsfeld Visual Analytics geschaffen. Visual Analytics zielt auf das Extrahieren von Wissen aus großen Datensätzen ab. Dabei werden die menschlichen Stärken in der visuellen Mustererkennung mit den Stärken von Computern bei der Bearbeitung großer Datenmengen verknüpft. Das eröffnet großes Potenzial für die Datenanalyse und damit auch für die Entscheidungsfindung. Trotz vieler Anknüpfungspunkte zwischen computergestützter Entscheidungsfindung und Visual Analytics fehlen in der wissenschaftlichen Literatur theoretische Grundlagen für das Anwenden von Visual Analytics im Entscheidungsprozess.

Der Beitrag dieser Dissertation beinhaltet die Definition einer theoretischen Grundlage für das Erweitern von Entscheidungsunterstützungssystemen durch Visual Analytics-Technologie. Im ersten Schritt charakterisieren wir computergestützte Entscheidungsfindung im Allgemeinen. Dafür beschreiben wir potenzielle Nutzergruppen, relevante Datenkategorien und Aufgaben im Entscheidungsprozess, die mit Visual Analytics-Technologien unterstützt werden können. Im zweiten Schritt präsentieren wir

einen spezialisierten Designprozess für das Entwickeln von an den Entscheidungsprozess angepassten Visual Analytics-Systemen. Im dritten Schritt beschreiben wir, wie Visual Analytics dafür genutzt werden kann, Wissenslücken zwischen im Entscheidungsprozess involvierten Personen zu schließen. Um die Anwendbarkeit unseres Konzepts zu demonstrieren, präsentieren wir darauffolgend sechs Designstudien, die sich auf konkrete Entscheidungsunterstützungsszenarien beziehen. Die vorgestellten Designstudien decken das komplette Spektrum an für die Entscheidungsfindung als relevant charakterisierten Daten-, Nutzer und Aufgaben-Kategorien ab. In jeder Designstudie nutzen wir die im Konzept beschriebene allgemeine Domänencharakterisierung, um das vorliegende konkrete Entscheidungsproblem zu beschreiben. Dabei zeigen wir, dass unser Konzept eine konsistente Charakterisierung von Nutzer-, Daten- und Aufgabentypen unterstützt. Außerdem verwenden wir in jeder Designstudie, den im Konzept präsentierten Designprozess. Und schließlich zeigen wir mit den Designstudien, dass unser Konzept das Überbrücken von Wissenslücken unterschiedlicher Nutzergruppen unterstützt. Die resultierenden Visual Analytics-Systeme ermöglichen das Generieren und Einbeziehen von Wissen in den Entscheidungsprozess und unterstützen die Kollaboration zwischen Personen mit unterschiedlichen Erfahrungswerten.

Acknowledgements

The creation of this thesis would not have been possible without the support of mentors, colleagues, co-authors, family, friends, and my girlfriend, Tina.

First, I would like to thank my primary supervisor Prof. Dr. Dieter W. Fellner, director of Fraunhofer Institute for Computer Graphics Research (IGD), who always asked the difficult questions that made me strengthen the concept of my thesis. Second, I want to thank Prof. Dr. Jörn Kohlhammer, my secondary supervisor and head of the "Information Visualization and Visual Analytics" department at Fraunhofer IGD, for his trust in me and his support in pushing my work in the right direction. It was a pleasure working in his department. Third, I want to thank Prof. Dr. Silvia Miksch from the TU Wien, the tertiary supervisor of my thesis, for completing the advisory body as an external professor with a high reputation in the Visual Analytics community.

Next, I want to express my appreciation to Dr. Thorsten May, Dr. Jürgen Bernard, and Prof. Dr. Arjan Kuijper. My scientific expertise has also been shaped by many fruitful discussions with Dr. May and Dr. Bernard, while Prof. Kuijper supported me from a high-level perspective on science. I also want to thank my colleagues Andreas Bannach, Hendrik Lücke-Tieke, Alex Ulmer. Without their deep technical knowledge and support, my work at Fraunhofer IGD would not have been such a success. Additionally, I want to thank my colleagues Martin Knuth, Marco Hutter, Sebastian Maier, Martin Steiger, Kawa Nazemi, Dirk Burkhardt, and Michel Krämer. They were always open to review and discuss my scientific work. A strong thank you goes to Gabriele Knöß and Patricia Häg, the secretaries of our department. By creating a supporting and friendly work environment, they allowed me to focus on my scientific work.

Last but not least, I want to thank my family and friends for their help throughout my doctoral phase. While my parents and my two sisters always believed in me reaching my goals, they also taught me the humility to always question myself and my scientific work. The rest of my family including my grandparents, aunts and uncles, cousins, brothers-in-law, nieces and nephews, and also family Jakobus have taught me the importance of family in my life, creating a safe and supportive environment for me to thrive in. In addition, my friends have helped me to step out of the scientific world allowing me to get the head clear for fresh ideas. I especially want to mention Isabel Woelk, who proof-read the introduction of my thesis. And finally, I want to thank my girlfriend Tina. She left me the space I needed to write this thesis, and at the same time pushed me through the difficult phases. I am very grateful for her patience and positive charisma that continues to inspire me every day.

Tobias Ruppert

November 2017

Contents

1. Introduction

Contents

1.1. Motivation and Problem Description

The increasing complexity of societal, economic, and environmental problems in the last decades has brought new challenges to decision makers. At the business (or organizational) level, decisions are not solely driven by economic factors anymore. Today, the triple bottom line (TBL) calls for the consideration of economic, social, and environmental factors to achieve sustainable decisions [Elk97]. At the political level, this triplet is augmented with value systems of public societies articulated through the 'public opinion' [DMRI16], e.g., via social media channels. The consideration of multiple and often conflicting factors within the decision making process imposes great challenges to decision makers. In order to make sustainable decisions a profound analysis of the problems and possible solutions needs to be conducted.

Fortunately, the digitization era we live in today produces massive amounts of data that are available in digital format. These data sources bear a great potential to support decision making processes with knowledge as scenarios beyond the human memory load can be tackled. As an example, the German railway company 'Deutsche Bahn' has installed sensors at most of their locomotives and uses the data to schedule maintenance routines with the goal to reduce train delays or failures [Qua15]. Among others, underlying data formats include unstructured texts, empirical/statistical data, and data artificially generated by computational models. Currently, a great portion of the information relevant for decision making is hidden in textual reports. However, the sheer amount of textual information is difficult to grasp in an efficient and effective way. Empirical data derived, e.g., from scientific studies helps to understand the problem domain and design alternative solutions. However, the interpretation of this information and its integration in the decision process is a complex task. Computational models can be used, e.g., to simulate the potential impacts of decisions. However, the complexity of underlying models and the interpretation of generated data is challenging. Example data from the business level include sales figures, share values, CO^2-emissions, energy consumption, customer reviews, employee surveys, etc. Examples at the political level include employment figures, economic growth figures, public opinions communicated at social media channels, statistical records on human well-being, earth observation data, environmental pollution data, etc. Hence, one of the main research questions addressed in this thesis is: How can we make use of the digitally available data to improve evidence-based decision making?

1.1.1. Problems with Data

The inclusion of data into the decision making process bears several problems to be addressed. Without claiming completeness, we identified the following:

Access: Datasets need to be made accessible for the decision makers. Although large amounts of data are being generated and stored every day, this does not imply the accessibility of data for all stakeholders involved in the decision making process. Often datasets are spread over different data sources with various data formats. In some cases, new datasets have to be collected. [Bel09] [CCS12]

Complexity: Not every dataset can be treated in the same way. Several data types exist, from unstructured data like texts, video, audio, or images to structured data stored in tabular databases. Additionally, computational models are applied in the decision making process. This abundance of different data types makes the analysis all the more complex. [Cou01] [CCS12]

Quantity: The sheer amount of available data cannot be handled solely by the decision maker. Relevant data needs to be distinguished from irrelevant data. If users can formulate a query for the data they are searching for, classical search methodologies can be applied. However, in many cases the relevance of data for a given decision is not obvious and users need to explore the data space with the help of appropriate aggregation techniques to search for relevant information. [CCS12] [WR09]

Quality: To make reliable data-driven decisions, the data quality needs to be assessed. Among others, characteristics like missing values, small sample sizes, unreliable sources, outdated creation or collection dates are indicators for poor data quality. In case of model-driven data, computational models only attempt to approximate reality. Hence, decision makers need to be aware of uncertainty in the data to make reliable decisions. [CCS12] [Hov07]

Trust: Decision makers need to have trust in the data before they base their decisions on them. Decision makers are likely to question data-driven information or the results of a data analysis process, if they do not know their origin, or if they cannot reconstruct the data analysis process itself. Therefore, a comprehensible presentation of the knowledge extraction process is required. [Hov07] [SSK*16]

Usability: Another problem remains in the usability of data. In order to make use of large datasets, the core information hidden in the data needs to be extracted. This has to be processed and presented to the decision maker in a meaningful way. In most cases, this is a non-trivial task. [Cou01] [PS07]

These problems need to be addressed to successfully incorporate knowledge extracted from data into the decision making process. They build the foundation for the definition of research challenges being addressed in this thesis.

1.1.2. Problems with the Involvement of Multiple Stakeholders

Besides the described problems related to the incorporation of data-driven knowledge in the decision making process, additional problems lie in the collaboration of numerous stakeholders involved in decision making scenarios. Today, evidence-based decisions cannot be handled solely by the decision maker himself. First, most decisions are made by a consortium of decision makers. Second, analysts and advisers support the decision makers with information and alternative decision options to be considered. Third, depending on the complexity of the topic, external domain experts are consulted to contribute external knowledge to the decision process. Fourth, modeling experts design simulation or optimization models to estimate the impact of alternative decision options or mitigate trade-offs. Fifth, different external stakeholders like lobbyists or investors influence the decision makers towards their interests. Due to varying expertises of these stakeholders, the whole process may suffer from knowledge gaps [RBK13].

Competence gap: Possibly, all stakeholders involved in the decision making process differ with respect to their expertise in the targeted domain and the supporting methods. This competence gap can hinder an efficient communication flow, which may provoke time loss and misunderstandings.

Analysis gap: The data analysis process is often distributed over several stakeholders. For example, domain experts provide data sources, modeling experts design computational models, and the analysts derive alternative solutions for the targeted problem. Finally, the decision maker only gets a condensed perspective on the decision options. This bears the risk of the underlying data not being exploited in an optimal way.

Iteration gap: As described above, involving several stakeholders in a process can reduce the efficiency. The analysis of a problem and possible solutions is an iterative process. However, due to time constraints and inefficient communication a critical amount of process iterations can be undercut, which results in suboptimal solutions. Hence, the necessity to involve several stakeholders in the decision process provokes inefficient and ineffective analysis cycles. These three knowledge gaps and how to tackle them will be addressed in the concept of our thesis.

1.1.3. Problems with Visual Analytics Applied to Decision Making

During the last decades, visual analytics has evolved as a discipline to analyze large and complex datasets. Combining human perceptual abilities, exploited through information visualization, and the processing power of computers via data mining allows to address complex data analysis tasks. Therefore, visual analytics bears great potential to augment and improve existing decision support systems. However, until today, visual analytics research has rarely targeted decision making explicitly. Some related aspects have been tackled. However, there are still open problems to be considered:

Collaborative visual analytics: As already discussed, decision making is a process that involves several stakeholders with differing expertises. Visual analytics can support the synchronous / asynchronous and co-located / distributed collaboration between stakeholders [IES*11]. However, limited research has been conducted on collaborative visual analytics with respect to the users' expertise.

Specific tasks in the decision process: The decision making process requires specific tasks to be addressed, e.g., the creation of alternative solutions to a problem, etc. Although several task taxonomies related to visual analytics have been published (e.g., [BM13]), we could not identify a taxonomy that explicitly addresses the specific tasks related to decision making. For example, visual analytics research has been focusing mainly on the exploration and analysis of data. The presentation of results to non-expert has rarely been considered in this research field [KM13].

Incorporation of alternative models: Finally, visual analytics research is mainly exploiting data mining techniques as models to support the analysis process. However, decision making requires alternative models like simulation or optimization models to be included in the decision making process [Pow02]. These models are rarely applied in visual analytics research. Moreover, most of the approaches are designed for expert users.

In addition to the aspects above visualization literacy needs to be taken into consideration. Although visualization is already applied in several application domains, users still need to acquire expertise in interpreting the visually presented information. The visual encoding and the interaction design needs to be learned in order to fully exploit the power of data visualization [BRBF14].

In summary, visual analytics bears great potential to support evidence-based decision making. However, the visual analytics domain is lacking a theoretical foundation on how to design visual analytics systems to support the decision making process considering the described problems. Moreover, only few visual analytics approaches exist that explicitly target decision making. In the following sections, we derive research challenges from the discussed problems and summarize the contributions of this thesis that aim at addressing these challenges.

1.2. Challenges

In the previous section, problems related to (a) the incorporation of knowledge extracted from data in the decision process (Section 1.1.1), (b) the involvement of multiple stakeholders in the decision process (Section 1.1.2), and (c) the application of visual analytics to decision making (Section 1.1.3) have been presented. The main research challenges of this thesis are reflecting these general problems. In the following, we[1] briefly discuss the identified research challenges. More details are provided in Chapter 3. We differentiate between conceptual and technical challenges.

[1] I believe that in computer science, collaboration is essential. Therefore, to acknowledge the contributions of my colleagues, paper co-authors, partners, students, and domain experts to my work, I decided to use the 'we'- instead of the 'I'-form.

Conceptual Challenges

The conceptual challenges of this thesis tackle theoretical foundations on how to design visual analytics solutions targeting decision making.

C_{VDSS} Design methodology for visual analytics decision support
From our state-of-the art review and our experience in collaborating with decision makers, we have learned that visual analytics bears great potential for informing decision makers with knowledge extracted from data. However, as we have discussed in Section 1.1.1 (Problems with Data) and Section 1.1.3 (Problems with Visual Analytics Applied to Decision Making), today's decision making scenarios impose several problems to be considered during the design of visual analytics systems for decision making support. Despite this, no dedicated design methodology for specific visual analytics systems targeting decision making support exists. However, we claim that decision support research and visual analytics research would strongly benefit from such a methodology.

C_{BKG} Bridge knowledge gaps between involved stakeholders
Strategical decisions on the business and political levels are not solely made by an individual decision maker. In contrast, multiple stakeholders with varying expertises are involved in today's decision processes. As a consequence, knowledge gaps between these stakeholders impede an efficient and effective decision making process, which we discussed in Section 1.1.2 (Problems with the Involvement of Multiple Stakeholders). Hence, one major challenge for successful decision support remains in bridging these knowledge gaps.

Technical Challenges

The technical challenges describe specific combinations of tasks and data categories to be supported in the decision making process.

C_{Proc} Explore and monitor decision processes
In practice, most decision making processes are unstructured. Grasping the current status, identifying relevant stakeholders, and keeping up to date remain challenging tasks. Most processes can be structured along text documents that document intermediate results. However, often these documents are distributed among various sources. We claim that decision making would benefit from a system bringing together all relevant documents and stakeholders and thus providing an overview of the process.

C_{Doc} Explore and analyze text document collections
A rule of thumb says that most of the information relevant for decision making is stored in textual formats. However, in most cases, a decision maker does not have the time to read all the textual content gathered on a specific topic. Hence, automatic text analysis methods are frequently used to create content-based overviews of large amounts of text documents. However, these automatically generated overviews often do not match the specific users' needs and the target at hand. Therefore, supporting users in creating overviews of large document collections with text analysis methods is another challenge related to decision making support.

C_{Deb} Explore, analyze, and compare stakeholder opinions and arguments
As described in the introduction, the incorporation of social factors like the public opinion is critical for today's decision making. The public debate is being documented daily through various social media channels. Especially, estimating the relevance of alternative solutions, favoring or opposing arguments, and the sentiment towards specific topics would be beneficial for decision making. However, extracting relevant information out of these large data sources remains a complex and time-consuming task.

C_{Dat} Explore, analyze, and compare empirical performance indicators
Due to the digitization era, large structured numerical datasets exist that are beneficial for the decision making process. Several approaches like business intelligence, business analytics, or policy analytics already target the extraction of knowledge from these structured datasets. However, most approaches do not cover all tasks relevant for the decision making process. For example, the creation of decision options and their comparison is seldom supported.

C_{Imp} Explore, analyze, and compare the impacts of solutions
A core challenge in the decision making process remains in the estimation of a decision's impact. The comparison of alternative solutions' impacts supports decision makers in choosing the most appropriate solution to a given problem. Computational simulation or regression models are applied to estimate decision impacts. However, in most cases these models are complex and difficult to assess for the decision makers.

C_{Opt} Create, analyze, and compare optimal solutions
Finally, decision processes often involve the balancing of trade-offs. Optimization techniques support decision makers in finding optimal solutions with respect to a given target function and constraints. However, similar to the impact assessment, incorporating computational optimization models in the decision process is a challenging task. We claim that a transparent access to both impact assessment and optimization models would improve decision making.

These research challenges build the motivation for this thesis. Throughout the thesis, we present solutions to these research challenges in the form of scientific contributions.

1.3. Contributions

In the following, we briefly summarize the contributions of this thesis. We differentiate between two types of contributions: conceptual and technical. We present two conceptual contributions that build the theoretical foundation of this thesis and address challenges $\mathbf{C}_{VDSS}$ and $\mathbf{C}_{BKG}$. The first concept targets the definition of a methodology for designing visual analytics systems that support evidence-based decision making (see Challenge $\mathbf{C}_{VDSS}$). We call these systems visual analytics decision support systems. The second concept explains how visual analytics decision support systems enable the bridging of knowledge gaps between stakeholders involved in the decision making process (see Challenge $\mathbf{C}_{BKG}$). In addition, we contribute six technical contributions that prove the applicability of our concept. As proofs of concept, we apply the presented concepts on different decision making-related scenarios and address the remaining challenges of the thesis (Challenges $\mathbf{C}_{\text{Proc}}$, $\mathbf{C}_{\text{Doc}}$, $\mathbf{C}_{\text{Deb}}$, $\mathbf{C}_{\text{Dat}}$, $\mathbf{C}_{\text{Imp}}$, $\mathbf{C}_{\text{Opt}}$). The main contributions of this thesis to the state of the art are summarized in the following table:

Conceptual Contributions		
Concept for the design of visual analytics decision support systems	$\mathbf{C}_{\text{VDSS}}$	Chapter 3
Concept for bridging knowledge gaps between involved stakeholders	$\mathbf{C}_{\text{BKG}}$	Chapter 3
Technical Contributions		
Proof of Concept: Visual-interactive access to decision making processes	$\mathbf{C}_{\text{Proc}}$	Chapter 4
Proof of Concept: Visual-interactive access to text document collections	$\mathbf{C}_{\text{Doc}}$	Chapter 5
Proof of Concept: Visual-interactive access to online debates	$\mathbf{C}_{\text{Deb}}$	Chapter 6
Proof of Concept: Visual-interactive access to empirical datasets	$\mathbf{C}_{\text{Dat}}$	Chapter 7
Proof of Concept: Visual-interactive access to simulation models	$\mathbf{C}_{\text{Imp}}$	Chapter 8
Proof of Concept: Visual-interactive access to optimization models	$\mathbf{C}_{\text{Opt}}$	Chapter 9

Table 1.1.: Contributions of this thesis.

Conceptual Contributions

In Chapter 3, we present our first conceptual contribution, a concept for the design of visual analytics decision support systems, which addresses Challenge $\mathbf{C}_{\text{VDSS}}$. For this concept, we first characterize the decision making process as defined by Simon [Sim60] from the visual analytics perspective. Second, we provide an abstract characterization of the decision making domain. This includes a definition of data, users, and task categories specifically relevant in decision making. We distinguish between the data categories textual data, empirical data, and model-driven data. User types are categorized into decision makers, analysts, modeling experts, domain experts, and stakeholders. Finally, we define a task taxonomy comprising the abstract tasks exploration, creation, analysis, comparison, and presentation. Based on these taxonomies, we complete our concept with the introduction of a design process dedicated to visual analytics decision support system. The process is structured into four distinct stages and discusses goals and validation methods at every stage.

Our second conceptual contribution, presented in Chapter 3, targets the bridging of knowledge gaps between stakeholders involved in the decision making process, which addresses Challenge $\mathbf{C}_{\text{BKG}}$. We introduce two models for the bridging of knowledge gaps. The first is dedicated to organizational decision making. The second model targets political decision making. The models describe how visual analytics simplifies the communication of information and knowledge extracted from data. In addition, we describe how different decision making tasks are supported by different visualization disciplines to take account of the stakeholders' varying expertise levels. Finally, we recapitulate on synergy effects created by the incorporation of visual analytics into the decision making process.

Technical Contributions

In Chapter 4, we present a proof of concept for enabling users to access the decision making process visually and interactively. The designed visual analytics system provides a visual overview of all relevant text documents. Stakeholders are allowed to access and rate existing documents, or augment the process with additional documents. The design study serves as a proof of concept for applying our concept on the meta-data of text documents and addresses Challenge $\mathbf{C}_{\text{Proc}}$.

In Chapter 5, we introduce a visual analytics system providing visual-interactive access to document collections. The system enables analysts to create text document overviews via content-based clustering. It serves as a proof of concept for applying our design methodology on textual data. As a result, analysts are able to provide overviews on document collections relevant for the decision making process. This addresses Challenge $\mathbf{C}_{\text{Doc}}$.

In Chapter 6, we present a system for the visual-interactive exploration of online debates on decision-related topics. The system allows monitoring the relevance of policy domains, policies, and arguments from textual social media statements. It serves as a proof of concept on how visual analytics can support the inclusion of public opinions into the decision process and targets Challenge $\mathbf{C}_{\text{Deb}}$.

In Chapter 7, we present a visual analytics system that provides visual-interactive access to country-specific performance indicators in the mining sector. The underlying empirical dataset was collected to improve the transparency in the mining sector and attract investments in resource-rich countries. Our visual analytics system provides intuitive access to this data for investor, governmental, or public decision makers. It serves as a proof of concept on how to apply our design methodology on empirical datasets. Challenge $\mathbf{C}_{\text{Dat}}$ is addressed with this design study.

In Chapter 8, we present a visual analytics system providing visual-interactive access to a simulation model targeting the estimation of decision impacts. Users are enabled to explore different decision scenarios simulated by the model and to analyze their impacts. The design study demonstrates the applicability of our concept to model-driven data. It addresses Challenge $\mathbf{C}_{\text{Imp}}$ by incorporating impact assessment methods in the decision making process.

In Chapter 9, we introduce a visual analytics system that provides visual-interactive access to an optimization model that supports mitigating trade-offs between different decision targets. It allows the calculation of optimal solutions based on the definition of target function(s) and constraints. In this design study, we applied our design methodology on model-driven data and address Challenge $\mathbf{C}_{\text{Opt}}$.

1.4. Outline

This thesis is structured into three main parts: In Chapter 2, we summarize the theoretical foundations that build the baseline for this thesis. This includes a review of related work on decision making, policy making, and visual analytics. In Chapter 3, we derive challenges from the problem description and the reviewed related work. Moreover, we present the concept of this thesis addressing the two conceptual challenges of our approach. We introduce a novel concept for the design of visual analytics decision support systems and describe how the concept simplifies the bridging of knowledge gaps between stakeholders involved in the decision making process. Chapters 4 – 9, are dedicated to prove the applicability of our concept to different data categories and tasks in the decision process. In Chapter 4, we present a proof of concept for providing visual-interactive access to the decision making process. Chapter 5 targets visual-interactive text document clustering to assess the content of document collections. A visual analytics system for estimating the relevance of policy options, opinions, and arguments from public online debates is presented in Chapter 6. In Chapter 7, we explain how to provide visual-interactive access to country performance indicators in the mining sector, which addresses the incorporation of empirical data into the decision process. Chapter 8 targets the assessment of decision impacts by combining visualization techniques with an agent-based simulation model. A visual analytics system that provides visual-interactive access to an optimization model for the creation of optimal solutions to a given problem, and thereby, balance trade-offs, is presented in Chapter 9. Finally, in Chapter 10, we recapitulate the findings of this thesis and provide suggestions for future work.

2. Foundations in Decision Making, Policy Making, and Visual Analytics

In this chapter, we will present and discuss theoretical foundations in decision making, policy making, and visual analytics. First, we describe decision making in general discussing prominent decision making models. Then, we review decision support system (DSS) research as an attempt to apply information and communication technology on the decision making process. We also tackle business intelligence and business analytics technology as descendants from decision support systems. Second, we present policy making as political decision making. We describe the policy cycle as the underlying process, policy analysis as a specific discipline within the policy cycle, and policy analytics as a data-driven decision support discipline. In the third part of this chapter, we discuss foundations in visual analytics and information visualization related to this thesis. We review existing data, user, and task taxonomies in visualization research that build the basis of our concept. Finally, in the summary of the section on visual analytics we review existing visual analytics approaches addressing decision making support.

Contents

2.1. Decision Making

In this thesis, we aim at supporting evidence-based decision making via visual analytics. As a baseline of our concept, we review theoretical approaches on decision making and its computational support in particular. We begin with scientific approaches in decision support system theory. This includes a characterization of the decision making process, involved stakeholders, and computational methods being applied in the decision making process to support decision makers. We extend our review to business intelligence and business analytics as descendants of decision support systems. Our approach is restricted to unstructured strategical decision problems in contrast to naturalistic decision making, which we briefly discuss. Finally, we provide advice on further readings and summarize our findings.

2.1.1. Decision Support System Theory

Decision making is defined as the process of selecting a course of action among a set of alternatives to address a given problem. Computational support for decision making has been a subject of methodological research from the early 1960s. Herbert Simon's theoretical view on decision making builds the foundation for most approaches in this research field. Therefore, we also select his work as the baseline for this thesis. In his frequently cited work "The new science of management decision" he describes three principal phases in rational decision making processes: "finding occasions for making a decision; finding possible courses of action; and choosing among courses of action" [Sim60]. Simon calls these phases intelligence, design, and choice. Figure 2.1 shows this process, which is often augmented with an additional 'implementation' stage. Some approaches even add a fifth 'evaluation' stage. Simon clarifies that the model is only an abstraction of real world decision making. Complex decision processes often contain multiple sub-processes with multiple iterations until a final decision is made. However, he emphasizes that principally all complex decision making processes can be characterized with the three mentioned stages.

Moreover, in his work, Simon distinguishes between programmed and nonprogrammed decisions. Programmed decisions are repetitive and routine. These decisions can be easily automated. Nonprogrammed decisions are novel, unstructured and consequential. The terminology of nonprogrammed and programmed decisions was re-used but also renamed by several scientists. Commonly, the terms unstructured and structured are applied (e.g., [GSM71]). In an outlook on 'modern' decision-making techniques, Simon names the following computer-supported methods that will be applied to decision making: operations research comprising mathematical analysis, models, and computer simulation; electronic data processing; and heuristic problem-solving techniques that can be implemented in heuristic computer programs. This assessment is taken up and refined by scientists in consecutive research approaches. For example, the term 'electronic data processing' (EDP) that Simon refers to is replaced by management information systems (MIS) [LL10] during the 1960s. Following Alter, during the 1970s decision support systems research evolved from management information systems research [Alt80]. Although it has been defined years ago, Simon's model greatly influences the design of

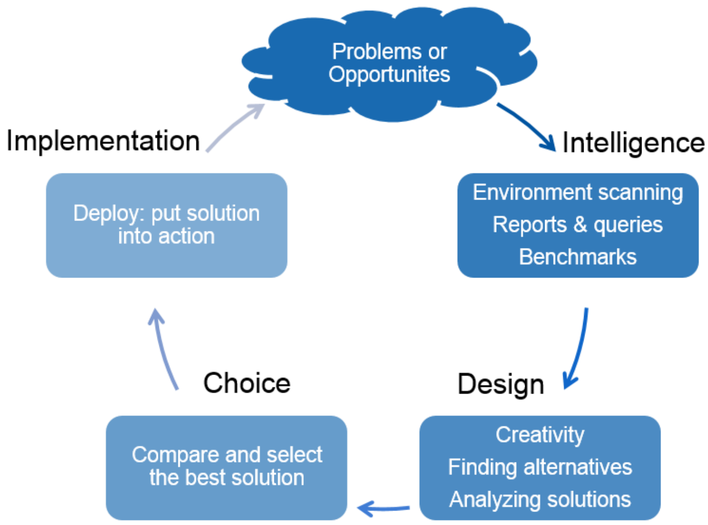

Figure 2.1.: The decision making process steps introduced by Simon [Sim60]. The diagram is adapted from Turban et al. who provide more details on the consecutive stages [TSD14].

computational decision support systems (DSS) [PA04]. In the concept chapter of this thesis, we re-use his decision making model for characterizing visual analytics support in decision making.

One of the first concepts on decision support systems presented by Gorry and Scott Morton evolved from management information system (MIS) research [GSM71]. The authors introduce a framework for information systems along two orthogonal axes. On the vertical axis, systems are ordered based on the decision type they support from structured over semi-structured to unstructured decisions. The decision types are based on Simon's nonprogrammed and programmed decision types [Sim60]. Gorry and Scott Morton add an additional intermediate decision type: semi-structured decisions. On the horizontal axis, systems are classified based on the management activity level, from 'operational control' at the bottom of the organizational hierarchy via 'management control' to 'strategic planning' on the top level. This classification was derived by the work of Anthony [Ant65]. Gorry and Scott Morton define the information systems that tackle semi- or unstructured decision problems as decision support systems. Information systems that address structured decisions are classified as management information systems (MIS). With their framework, they identified the need for a concentration of efforts on the development of information systems that support managers in unstructured strategic planning decisions. This motivated further research in the area of decision support systems.

In the foreword of Alter's book on decision support systems, Keen and Stabell characterize decision support science as the concept for the development of tools that "address nonstructured rather than structured tasks; support rather than replace judgment; focus on effectiveness rather than efficiency in decision processes" [Alt80]. Keen and Stabell, as most scientists in the field, emphasize that an important factor of a decision support system is that it only 'supports' decision makers in making decisions, it is not replacing their final choice. This notion fits very well to visual analytics approaches that attempt to keep humans in the analysis loop. Alter himself describes the heterogeneity of decision support systems, represented through three case studies, with the following opposing characteristics: use by decision makers $\leftrightarrow$ use by intermediaries; online $\leftrightarrow$ offline use; systems based on data $\leftrightarrow$ systems based on models; systems with one primary user $\leftrightarrow$ systems with many users [Alt80, p.71]. These characteristics are also relevant to our approach. We differentiate between user roles, we ensure the accessibility of our support systems via web technology, we differentiate between model-driven and data-driven (we name the latter empirical data) approaches, and we consider collaborative aspects that allow multiple users to participate in decision making.

In another survey on decision support systems by Shim et al. the past, present, and future of decision support technology is discussed [SWC*02]. The authors provide an abstract definition of decision support systems: "decision support systems (DSS) are computer technology solutions that can be used to support complex decision making and problem solving" [SWC*02]. In this definition the relation between decision making and problem solving is evident. They also introduce a refinement of the decision making process introducing the steps: problem recognition, problem definition, alternative generation, model development, alternative analysis, choice, and implementation. The refinement of the process model informed the definition of the individual process steps presented in our approach. Shim et al. distinguish between four decision support tools: data warehousing, online analytical processing (OLAP), data mining, and web-based decision support systems. Moreover, they expand this set of tools by optimization-based and collaborative decision making tools.

Having presented various decision support system categorizations in the previous paragraphs, we conclude with the categorization that has the highest influence on our approach. Power's framework structures decision support systems into five categories: data-driven, model-driven, knowledge-driven, document-driven, and communications-driven decision support systems [Pow02]. These can be described as follows:

- **Data-driven DSS** support the analysis of structured datasets. Examples include reporting systems, data warehouses, and business intelligence systems.
- **Model-driven DSS** focus on providing support for accessing and manipulating models. Examples include statistical, financial, optimization, or simulation models.
- **Knowledge-driven DSS** suggest or recommend actions to managers via specialized business rules or knowledge bases.
- **Document-driven DSS** support users in gathering, retrieving, classifying, and managing unstructured documents. Examples of unstructured documents include text documents, images, sounds, and video.
- **Communication-driven DSS** support the communication and collaboration of a team for decision making. Group Decision Support Systems (GDSS) are an example of communication-driven DSS. Further examples include online communication, scheduling, and document sharing tools.

Power derives some additional technologies from these categories. For example, data mining is defined as a support technology for building hybrid data-driven and knowledge-driven DSS. The data-driven

DSS provides data that can be analyzed with data mining techniques in order to extract rules that build the basis for a knowledge-driven DSS. Moreover, Power names the combination of a document-driven DSS and a search engine as an example of a knowledge management system. In our work, we re-use the DSS categorization by Power to characterize the data types of our concept. We distinguish between empirical data, textual data, and model-driven data. Empirical data builds the basis for data-driven DSS. We use the term empirical data, since in most cases structured data originates from empirical measurements. In our definition, the foundation of document-driven DSS is restricted to textual data. We do not consider video and audio data in our approach. Finally, we name the data originating from models model-driven data, referring to the model-driven DSS category.

A general definition of computational model-driven approaches is given by Hill et al. [HCSG01]: "a set of computational codes, executable in some software/hardware environment, that transform a set of input data into a set of output data, with the input, output, and transformation typically having some interpretation in terms of real-world phenomena." Model-driven DSS were often associated to the domain of operations research [Pow03]. Following Power and Sharda, model-driven DSSs may also include algebraic, decision analytic, financial, simulation, and optimization models [PS07]. Although Turban et al. state that "no universally accepted definition" for decision support systems exist [TSD14], we rely on the frequently cited definition by Keen and Scott Morton: "(Decision support systems) are computer-based support for management decision makers who are dealing with semi-structured problems." [KSM78, p.97].

In a recent book, Power characterizes modern decision support systems with the following attributes [Pow13]: (1) access capabilities from any location at anytime; (2) access very large historical datasets almost instantaneously; (3) collaborate with multiple, remote users in real-time using rich media; (4) receive real-time structured and unstructured data when needed; (5) view data and results visually with excellent graphs and charts. These attributes of modern DSS also served as requirements to the solutions, we present in this thesis. All of our approaches are implemented as web applications, which allows an easy access. The client-server applications allow the processing of large datasets. Several users can access the data simultaneously and collaborate in finding solutions. We support both access to structured (e.g., numerical) and unstructured (e.g., textual) data. Finally, all of our approaches provide an intuitive visual and interactive access to the data or models provided.

Power further expands these ideas and introduces basic characteristics for computerized DSS:

1. "Facilitation. DSS facilitate and support specific decision-making activities or decision processes, or both.
2. Interaction. DSS are computer-based systems designed for interactive use by decision makers or staff users who control the sequence of interaction and the operations performed.
3. Ancillary. DSS can support decision makers at any level in an organization. They are not intended to replace decision makers.
4. Repeated Use. DSS are intended for repeated use. A specific DSS may be used routinely or used as needed for ad hoc decision support tasks.

5. Task-Oriented. DSS provide specific capabilities that support one or more tasks related to decision making, including intelligence and data analysis, identification and design of alternatives, choice among alternatives, and decision implementation.
6. Identifiable. DSS may be independent systems that collect or replicate data from other information systems or subsystems of a larger, more integrated information system.
7. Decision Impact. DSS are intended to improve the accuracy, timeliness, quality, and overall effectiveness of a specific decision or a set of related decisions." [Pow13, p.39]

Finally, from an online article in which Power provides a comprehensive historical overview on DSS research, we also want to share a decision support system definition that summarizes the lessons learned during our literature review on decision support systems in a compact way: "A DSS is an interactive computer-based system or subsystem intended to help decision makers use communications technologies, data, documents, knowledge and/or models to identify and solve problems, complete decision process tasks, and make decisions. Decision Support System is a general term for any computer application that enhances a person or group's ability to make decisions. Also, Decision Support Systems refers to an academic field of research that involves designing and studying Decision Support Systems in their context of use. In general, Decision Support Systems are a class of computerized information system that support decision-making activities. Five more specific Decision Support System types include: Communications-driven DSS, Data-driven DSS, Document-driven DSS, Knowledge-driven DSS, Model-driven DSS." [Pow03]

2.1.2. From DSS to Business Intelligence and Business Analytics

During the 1990s, DSS terminology was adapted. In their frequently cited book, Turban et al. synonymously use the terms business intelligence and business analytics for online analytical processing (OLAP), and model-base management systems and models, respectively. The term business intelligence was promoted by Howard Dresner within Business and IT communities in 1989 [Pow03] [CCS12]. As denoted by Power, business intelligence systems can be interpreted as data-driven DSS [Pow03]. Chen et al. define business intelligence and analytics "as the techniques, technologies, systems, practices, methodologies, and applications that analyze critical business data to help an enterprise better understand its business and market and make timely business decisions." [CCS12] In their frequently cited survey article on business intelligence and analytics, the past, present, future trends in these lines of research are discussed. They also discuss application areas for business intelligence and analytics including e-commerce and market intelligence, e-government and politics 2.0, science and technology, smart health and wellbeing, and security and public safety. The authors also discuss enabling technologies like big data analytics, text analytics, web analytics, and network analytics. Moreover, they provided an intensive review of 3602 research papers published between 2000-2011 that contain the keywords business intelligence and analytics to better understand the current trends in the respective scientific fields. Interestingly, none of the reviewed papers was published at one of the major visualization conferences.

Having reviewed theoretical foundations on DSS, we also want to shed light on visual support techniques as discussed in decision support research. From the very beginning of decision support research, the importance of user interfaces was highlighted. In their work, Sprague and Carlson structure DSS into three components, the data component, the model component, and the user interface component [SC82, p.313]. These components are also considered in the visual analytics model by Keim et al., comprising data, visualization, models, and knowledge [KAF*08]. Courtney promotes a new decision-making paradigm for DSS [Cou01]. He proposes the usage of "diagramming tools" to communicate the complexity of a given problem to all relevant stakeholders. Technically, he proposes the incorporation of mathematical models into online DSS systems to allow several stakeholders to test their hypothesis with the models by themselves. Power and Sharda promote advancing the state of the art of visual-interactive DSS [PS07]. In the domains of visual analytics and information visualization decision making is an important application area [TC05] [KKEM10]. A more specific research agenda in the field of geovisual analytics is provided by Andrienko et al. [AAJ*07]. Liu et al. review existing decision support systems from the integration perspective [LDWB10]. They differentiate between five integration perspectives: data and information integration, model integration, process integration, service integration, and presentation integration. By the latter, he explicitly promotes the presentation layer of a DSS.

Despite the fact that user interfaces and result visualization are described as fundamental components in decision support systems, the respective visual analytics and information visualization research branches are seldom consequently embedded in the decision support theory. As a counter example in the area of business intelligence and analytics we propose the work by Kohlhammer et al. who also reflect decision support concepts from the visualization research perspective [KPW13]. Still, this work is lacking a clear orientation on Simon's or any other decision making process model.

2.1.3. An Alternative Decision Making Theory

Finally, for the sake of completeness, we also want to emphasize that besides the classical normative branch of decision making research a second branch was promoted by scientists like Gary Klein. Naturalistic decision making theories evolved as an alternative decision making theory in the late 1980s [KOCZ93]. While classical decision making research targets the identification of a problem, the definition of alternatives, and the rational choice between the alternatives, naturalistic decision making mainly relies on the experience of the decision maker. Orasanu and Connolly describe the differences between classical and naturalistic decision making. They claim that in naturalistic decision making "much effort is devoted to situation assessment, or figuring out the nature of the problem; single options are evaluated sequentially through mental simulation of outcomes; and options are accepted if they are satisfactory (rather than optimal)" [OC93]. Moreover, they describe a naturalistic decision making setting with the following attributes: "ill-structured problems; uncertain dynamic environments; shifting, ill-defined, or competing goals; action/feedback loops; time stress; high stakes; multiple players; organizational goals and norms" [KOCZ93]. Numerous visual analytics approaches focus on supporting naturalistic decision making. Examples include applications in time-critical situations (e.g., [KMH09] [MJR*11] [AME11]). This line of research is highly relevant for visual analytics. However, in this thesis, we target the visual-interactive support for classical decision making problems.

2.1.4. Summary of Decision Making Theory

In this section, we reviewed the related work in computerized decision support theory. The theory grounds on the model by Herbert Simon, who separates the decision making process into the three stages design, intelligence, and choice. Following his model decision making comprises the steps "finding occasions for making a decision; finding possible courses of action; and choosing among courses of action". We re-use Simon's model to structure visual analytics tasks along the decision making process. Additionally, Simon's work motivated the later differentiation between structured, semi-structured, and unstructured decision types. Structured decisions are repetitive and routine, which allows their automation. Unstructured decisions are novel and consequential. They cannot be automated. However, human decision makers can be supported by computerized decision support systems in addressing unstructured decision types. This characterization allows us to describe the focus of our concept targeting visual analytics support for unstructured decisions on a strategic planning level. We also reviewed multiple decision support tool categorizations. Eventually, we chose the well-known categorization of Power to motivate our concept. He distinguishes between data-driven, model-driven, knowledge-driven, document-driven, and communications-driven and group decision support systems. Our concept primarily covers visual analytics support for data-driven, model-driven, and document-driven decision support systems. In the remainder of the section, we reviewed further characteristics of decision support systems, including the descendants of DSS, business intelligence (data-driven DSS) and business analytics (model-driven DSS). Moreover, we identified the need for visualization and visual analytics to support DSS research. Finally, we further restricted our focus on classical normative decision making in contrast to alternatives like naturalistic decision making.

2.2. Policy Making

In the previous section, we discussed decision making support applied to the business level. We reviewed scientific approaches that describe how information systems can support decision making in organizations. In this section, we review theoretical foundations in political decision making. In most cases, political decision making, or policy making, results in new or adapted policies or directives. We discuss the specifities of policy making by reviewing the definition of public policy research and the creation of policies via the public policy cycle. We draw similarities between decision making and the policy analysis task within the policy cycle and discuss challenges in applying rational decision making in the policy process. In addition, we review approaches on how to incorporate big data analytics and policy analytics (as analogy to business analytics) in the political decision making process. We continue by reviewing approaches that promote the application of visualization techniques in policy analysis. The section is concluded with a summary of our findings.

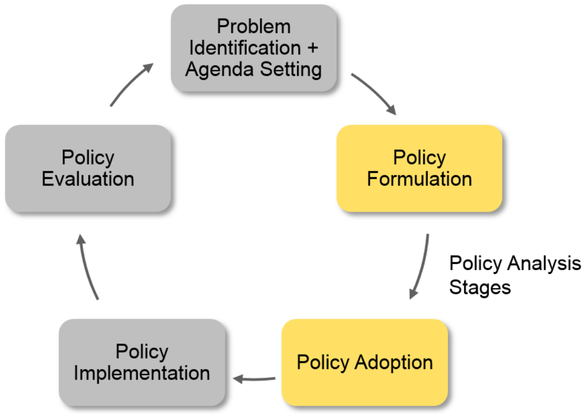

Figure 2.2.: Policy Cycle adapted from Anderson [And75]. Policy analysis is mainly conducted in the policy formulation and the policy adoption stage. Alternative solutions to a given problem are defined in the formulation stage. One of these alternatives is selected for implementation in the adoption stage.

2.2.1. Public Policy and the Policy Cycle

We begin this section by characterizing political decision making processes in order to identify relations to classical decision support theory that mainly targets the business sector, as discussed in the previous section. In most cases, political decision making processes are resulting in a policy that is implemented and applied by the administration of a government. Following Anderson's definition, "a policy is defined as a relatively stable, purposive course of action followed by an actor or set of actors in dealing with a problem or matter of concern" [And75]. Reviewing the terminology that Simon used, the choice of a course of action is also the result of a decision making process. An alternative definition of public policy in general is provided by Chocran and Malone: "public policy can be described as the overall framework within which government actions are undertaken to achieve public goals" [CLC14]. In the scientific literature, political decision making processes are often structured into so called policy cycles. Although several policy cycle definitions exist, most of them share similar steps in the process. Some refine aspects on different levels of detail. Most cycles presented in the literature are based on the concept of Lasswell who compared policy making to problem solving. In 1956, the political scientist introduced a decision process model divided into the seven phases intelligence, promotion, prescription, invocation, application, termination, and appraisal [Las56]). In the intelligence phase, relevant information is gathered. In the promotion phase, policy alternatives are defined and arguments for and against these alternatives are formulated. In the prescription phase, a specific policy is selected. In the invocation phase, this policy is implemented. In the application phase, the implemented policy is applied by the administration. In the termination phase, the policy process terminates. And finally, in the appraisal phase, the policy is evaluated towards the achievement of the specified objectives. Anderson simplified this policy process to five distinctive stages in a policy cycle (see Figure 2.2): problem identification and agenda setting, formulation, adoption, implementation, and evaluation [And75]. In the first stage public problems are identified, and the political agenda is set by prioritizing societal problems. In the second stage alternative solutions to these problems are explored and evaluated. In the third stage, these policy options are compared, and it is decided which option to choose. In the fourth stage the selected policy option is implemented through legal process. In the last stage of the cycle the implemented policy is evaluated with respect to the objectives defined in the first stage of the cycle. This simplified policy cycle is often referred to as the standard cycle. We take up this process definition and re-use it in the concept of this thesis. Nevertheless, alternative policy cycles appear in the literature that we briefly discuss in the following. For example, Brewer's policy cycle comprises six stages: invention/initiation, estimation, selection, implementation, evaluation, and termination [Bre74]. Patton and Sawicki's 6-step policy analysis cycle mainly focuses on the policy analysis step (see Section 2.2.2) and disregards the formal implementation step. The six stages are: verify, define, and detail the problem; establish evaluation criteria; identify alternative policies; evaluate alternative policies; display and distinguish among policy alternatives; monitor the implemented policy [PS83]. In the Australian Policy Handbook an 8-step cycle is promoted with the following stages: identify issues, policy analysis, policy instruments, consultation, coordination, decision, implementation, evaluation [ABD08]. This cycle also reflects legislative and political sub-processes. The reviewed cycles introduce specific perspectives on policy making and informed the characterization of policy making in our concept.

Numerous alternative policy cycle definitions exist. However, we rely on Jann and Wegrich, who claim that "today, the differentiation between agenda-setting, policy formulation, decision making, implementation, and evaluation (eventually leading to termination) has become the conventional way to describe the chronology of a policy process." [JW07] Stakeholders involved in the process steps were added by Howlett et al., which results in the following cycle [HRP09]: (1) agenda-setting by policy universe, (2) policy formulation by policy subsystem, (3) decision-making by government decision-makers, (4) policy implementation by policy subsystem, (5) policy evaluation by policy universe.

Agenda Setting and Problem Definition: At the initial stage of the policy cycle public problems that shape the agenda for policy making are identified. The entire policy universe participates in this stage. This also includes the civil society.

Policy Formulation: In the policy formulation stage, policy options to address a given problem are formulated and discussed by the policy subsystem. Only actors with a profound knowledge of the given problem are involved, e.g. policy analysts, knowledge workers, etc. The stage includes policy analysis.

Decision Making: In the decision-making stage, one of the alternative options defined in the policy formulation stage is adopted or no action is taken. The underlying decision is made by governmental decision makers based on an analysis of the alternative decision impacts.

Policy Implementation: The policy implementation stage describes the administrative act in the policy cycle. The selected policy is put into practice by the policy subsystem.

Policy Evaluation: Finally, the resulting policy is evaluated in the policy evaluation stage. At this stage, again the entire policy universe is involved. After the evaluation step the problems and the solutions may be reviewed, which results in returning to the first stage of the policy cycle.

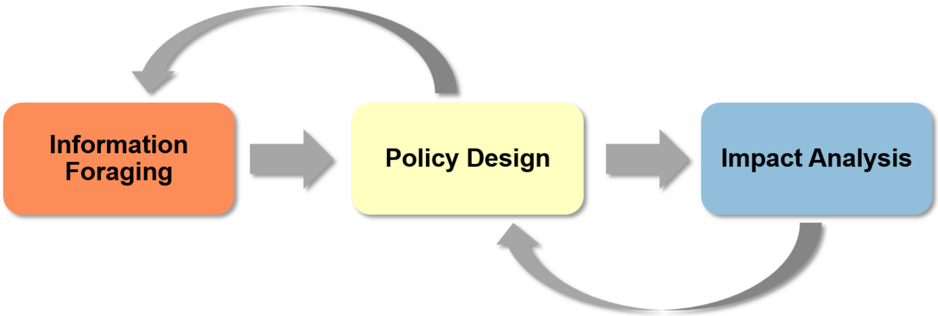

Figure 2.3.: Simplified policy process as introduced by Kohlhammer et al. [KNRB12]

Howlett et al. further specify different domestic policy actors that influence policy processes and outcomes. These are: elected politicians, the public, bureaucracy, political parties, interest or pressure groups, think tanks and research organizations, mass media, academic policy experts and consultants. [HRP09] An adaptation of the presented policy cycle to the field of eParticipation with a focus on the usage of information and communication technology (ICT) was presented by Macintosh [Mac04a]. She names the stages: agenda setting, policy analysis, policy creation, implementation, and monitoring. Höchtl et al. adapted an alternative policy cycle in order to explain how big data analytics may support public policy making [HPS15]. Our first attempt to introduce interactive visualization to the policy cycle is described in Figure 2.3 (extracted from Kohlhammer et al. [KNRB12]). Here, we simplify the policy cycle to the steps information foraging, policy design, and impact analysis, and promote the inclusion of different visualization disciplines into the policy cycle.

2.2.2. Policy Analysis and the Policy Analysis Paradox

In the previous section, we briefly introduced the policy cycle as it is discussed in the political sciences. In this section, we will provide further details on policy analysis as a fundamental method during the policy cycle. Within the policy cycle at the policy formulation stage different policy options to address a public problem are created and compared (see also Figure 2.2). "Policy formulation clearly is a critical phase of the policy process. Certainly designing the alternatives that decision makers will consider directly influences the ultimate policy choice" [Sid07]. Jenkins-Smith provides a similar definition replacing the term policy formulation with policy analysis: "policy analysis is a set of techniques and criteria with which to evaluate public policy options and select among them" [JS90]. Similarly, Howlett et al. name policy analysis as a critical component of the policy formulation stage [How09]. Policy analysis as a discipline of the policy sciences was introduced by Lasswell and Lernen in their work "The Policy Sciences" in 1951 [LL51]. Various interpretations of policy analysis exist in the literature.

A historical overview of policy making starting from the concept of Lasswell, with an outlook to the future of policy analysis is given by Hoppe [Hop99]. Two theoretical perspectives on public policy analysis exist: positivism and post-positivism. Positivist approaches before 1990 originate from welfare economists and interpret policy analysis from a rational perspective. The theory is based on the assumption that individuals make the most social decisions through market mechanisms. The role of the government is interpreted as a regulator for so-called market failures. Following Howlett et al., “they (positivists) argue that political institutions can act to supplement or replace markets to produce better outcomes in terms of enhancing overall social welfare” [How09, p.22]. Positivists analyze policy making with empirical and quantifiable facts grounded in economic theory. For example, Weimer and Vining state that policy analysis has to be enriched by craft skills for gathering information, structuring analysis, and cost-benefit analysis among others [WV05]. Positivism is interpreted as societal problem-solving discipline with the higher goal to support rational decisions in policy making.

During the 1990s critics on this technocratic perspective emerged. The main breach of positivism approaches results from the experience that decisions solely based on rational perspectives are not sufficiently considering external factors within real-world scenarios. “Even if one could identify the most efficient and effective policy, which is difficult given the limitations innate to the social sciences, the actual policy choice is a political, not a technical, one, bound by political institutions and made by political actors in response to political pressures, ideologies, and self-interests, among other factors.” [How09, p.26] Following the critics’ argumentation, there is no single objective solution to a problem. From this experience post-positivism approaches evolved. Post-positivists aim at letting citizens participate in the policy process and deriving policies through policy debates and combinations of normative and empirical arguments. “In politics, politicians and policy decision-makers put forth proposals about what to do based on normative arguments. Empirical analysis comes into play but only when there are reasons to question or explore the factual aspects of the argument.” [Fis07, p.227] In Shulock’s view, “policy analysis is more a tool of the democratic process than the problem-solving process. Its value lies in its contribution to the understandings that citizens have of issues and the political process... Analysis can lead to better policies if by ‘better’ we mean more responsive to, and supported by, the public” [Shu99]. Moreover, she states that “policy analysis has changed, right along with the policy process, to become the provider of ideas and frames, to help sustain the discourse that shapes citizen preferences, and to provide the appearance of rationality in an increasingly complex political environment”.

Complementary to the positivists’ and the post-positivists’ perspectives on policy analysis, further theories have been introduced that describe public policy making from different perspectives. Mayer et al. introduce a complex framework (a hexagon model) that attempts to combine all existing policy analysis theories. They differentiate between six policy advice activities and six policy analyst styles which they organize in a linked hexagon. The hexagon model is further augmented and provides a comprehensive characterization of policy analysis models. The framework unifies, beyond others, positivism and post-positivism approaches. The six policy analysis activities are: research and analyze; design and recommend; clarify values and arguments; advise strategically; democratize; and mediate. The six policy analyst styles are: rational style (application of scientific methods to generate knowledge; what is good knowledge?), argumentative style (illustrate arguments and justifications thereof; what is good for the debate?), client advice style (study environment and provide advice on stakeholders and positions; what is good for the client/problem owner?), participatory style (representation of stakeholder perspectives not involved in the process; what is good for society?), process style (focusing on the procedural aspects of the decision making process; what is good for the process?), and interactive style (facilitating consultations between different actors; what is good for mutual understanding?). By grouping subsets of theses activity-style combinations, existing policy analysis disciplines can be characterized. Finally, Mayer et al. strongly recommend to combine different policy analysis styles within one policy process. [MvDB13]

Besides scientific work on the characterization of policy analysis in general, especially the inclusion of scientific research results in political decision making has been heavily discussed. Engels provides examples that describe where scientific expertise is relevant for decision making: scientific warning and awareness creation, problem definition, ex ante impact assessment for policy options, ex post evaluation of policy choices, monitoring of implementation [Eng05]. However, despite the expected importance of considering scientific knowledge in the policy formulation phase, Shulock identifies a main deficit of policy making: the 'policy analysis paradox'. It describes the asymmetry between the amount of knowledge generated by scientific experts, and the actual amount of knowledge effectively used in the decision making process [Shu99]. Van den Hove names a number of "theoretical problems" that emerge at the intersection between science and policy. Among others, she names the complexity, uncertainty, and indeterminacy of scientific outputs. Scientific models attempting to simulate the complex reality are seldom precise or accurate. Furthermore, in most cases this uncertainty is not communicated to the user. This fails to raise the awareness of the model's uncertainty, and as a consequence, reduces the credibility of scientific outputs.

Concepts have been introduced to mitigate these problems, which evolve from bringing two contrary systems together, politics and science. Most of these concepts can be summarized under the term science-policy interface. Van den Hove defines science-policy interfaces as "social processes which encompass relations between scientists and other actors in the policy process, and which allow for exchanges, co-evolution, and joint construction of knowledge with the aim of enriching decision-making." [Hov07]. The positive aspects out of this are: a) rationality and legitimation through knowledge in politics, b) exploration of policy alternatives with focus on cause and effect, c) communication between two fields – e.g., research assignment, and scientific advice. In order to realize these aspects, the concept of "knowledge brokers" is propagated. Their goal is to mediate between the two systems [HW09]. Still, these concepts contain the risk of subjectivity. As a consequence, the propagation of a merely technocratic model has to be replaced by a concept with high interaction possibilities between knowledge brokers and decision makers. We identified the policy analysis paradox as another motivation for including visual analytics in the policy cycle. We assume that visual analytics systems that provide access to information relevant to policy making could serve as an online science-policy interface, and bridge knowledge gaps between stakeholders from science and politics.

2.2.3. Big Data Analytics and Policy Analytics

In the 'big data' era, policy analysis is reviewed from a data-driven perspective. However, Chen et al. state: "Despite the significant transformational potential for BI&A (business intelligence & analytics) in e-government research, there has been less academic research than, for example, e-commerce-related BI&A research" [CCS12]. Höchtl et al. noted that there is only little discussion on how data analytics may support policy making. Therefore, they discuss how big data analytics might be applied to a slightly adapted policy cycle with the stages agenda-setting, policy discussion, policy formation, policy acceptance, provision of means, implementation, and evaluation. The authors postulate that "many of the claims with respect to goals, benefits, and perils must be adopted from a business-related domain and imposed on government action and policy making, which is justified, as good governance means putting the citizen into the focus of consideration" [HPS15]. They also refer to the work by Chen et al. on business intelligence and analytics that we discussed in the previous section [CCS12]. In our approach, we rely on the ideas of Höchtl et al. on how big data analytics (BDA) may be applied to the decision making process: In the agenda setting stage, BDA can be used to identify emerging topics via social media and online news analysis, although the authors note a "problematic practice of simply equating online tweets, blogs, and comments with 'public opinion' in general". In the policy discussion stage, policy options to tackle identified problems are discussed. BDA can support citizens' participation by monitoring and aggregating the information provided by the public via social media. In the policy formation and policy acceptance stages, BDA can be used to predict the acceptance of a formulated policy by the public. In the provision of means stage, the authors claim, e.g., that "there is already some empirical evidence that the use of big data in budgeting can increase efficiency and effectiveness while reducing costs". In the implementation stage, Höchtl et al. identify two ways how BDA can support policy making. First, problematic areas can be identified prior to implementation that allows for an adaption of the policy's intensity in that area. Second, the new data generated in the implementation process can be used to evaluate the current policy in order to improve future policy processes. The continuous evaluation of the described stages allows the authors to eliminate the last evaluation step: "BDA enables evaluation, instead of being a well-defined process step at the very end of the policy cycle, to happen at any stage and to happen opaque to the affected stakeholders." Höchtl et al. also provide an example of the UK governance performance program that utilize visualization to continuously monitor and evaluate government policy making. [HPS15]

De Marchi et al. [DLT16] and Daniell et al. [DMRI16] go even one step further describing the emerging field 'policy analytics'. The latter define policy analytics as "new analytic methods that can be used to support public policy problem-solving and decision processes... (while balancing) the need for robust and convincing analysis with the need for satisfying legitimate public expectations about transparency and opportunities for participation" [DMRI16]. This definition nicely separates organizational (or business) decision making from public policy making by adding to the latter the requirement to satisfy public expectations on governmental decisions. In addition, the authors extend the model by Mayer et al. [MvDB13] with exemplary policy analytics techniques applied to each activity level.Applied methods include text mining, exploratory data analysis, game theoretic models, large-scale mathematical optimization, clustering, support vector machines, spreadsheet models, and argumentation theory.

A research and practice agenda for policy analytics is presented by Tsoukias et al. [TMLB13]. The authors discuss facets that distinguish policy making processes from other decision making processes. The concept of 'policy analytics' is comprehensively defined by the following statement: "To support policy makers in a way that is meaningful (in a sense of being relevant and adding value to the process), operational (in a sense of being practically feasible) and legitimating (in the sense of ensuring transparency and accountability), decision analysts need to draw on a wide range of existing data and knowledge (including factual information, scientific knowledge, and expert knowledge in its many forms) and to combine this with a constructive approach to surfacing, modelling and understanding the opinions, values and judgements of the range of relevant stakeholders" [TMLB13]. In addition, five major complexities in policy making are identified: Use of public resources, multiple stakeholders, long-time horizon, legitimation and accountability, deliberation.

2.2.4. Visualization for Policy Analysis

Despite the fact that data- and model-driven approaches find their way into policy analysis, only little research focuses on the inclusion of visual analytics and information visualization into political decision making. In the Data4Policy state-of-the-art report by Poel et al. data-driven approaches applied to policy making in the European Union are examined [PST*15]. Besides the literature review, the team of authors present a detailed stakeholder analysis, interviews with selected experts, and a review of innovative data-driven initiatives in the policy making field. Their study reveals "substantial opportunities in moving to advanced data analytics and visualization". Lindquist also sheds light on how visualization might be applied in the policy making era. In his research Lindquist provides a survey on visualization [Lin11b] and introduces a discussion on how visualization may support policy making [Lin11a]. Lindquist requests more investments by public institutions in visualization research targeting policy making. He describes this from the perspective of a political scientist. McInerny et al. criticize a lack of applications of information visualization to science and policy [MCF*14]. They promote the application of information visualization in science and science policy to avoid "an increased potential for missed discoveries, miscommunications, and, at worst, creating a bias towards the research that is easiest to display... Visualisation should be supporting the whole information pipeline; from acquiring and exploring data and analysing models, to the visual analytics used to reason across research and assessment activities, all the way to storytelling for communicating background information, results, and conclusions" [MCF*14].

2.2.5. Summary of Policy Making Theory

In this section, we reviewed public policy making theory and the advent of computer-based support technologies for political decision making. A policy is defined as governmental purposive course of action for dealing with a societal problem. To understand how policies evolve, we reviewed various policy (making) cycles. The most representative one by Anderson comprises the five steps problem identification and agenda setting, policy formulation, policy adoption, policy implementation, and policy evaluation. In addition, some approaches discuss the types of stakeholders involved in the process. Among others, decision (or policy) maker, analysts (or advisers), knowledge brokers, scientists, and public stakeholders are named. During the policy cycle the policy formulation stage is defined as the most important stage. Similar to the decision making model by Simon, at this stage alternative policy options to address a specific problem are designed, analyzed, compared. In the policy adoption stage one of these options is chosen for implementation. This process is often referred to as policy analysis. However, the definition of policy analysis has changed over time. Two main lines of research exist: positivism and post-positivism. While positivism approaches mainly focus on the economic aspects of policy making, post-positivism approaches see policy analysis as a discipline to support the political discourse. For post-positivists the inclusion of the public into policy making is of utmost importance. As a consequence of controversial theories in policy analysis hybrid approaches evolved that attempt to unify both theories. Despite these efforts, researchers identify a policy analysis paradox: the refusal of including scientific knowledge into policy making despite its accepted importance. As a consequence science-policy interfaces are promoted that should support the collaboration between scientific and political stakeholders. Nevertheless, the policy analysis paradox remains a challenge to address in policy making. Through the advent of big data analytics, the potential of including data-driven decision support methods into policy making is emphasized. Therefore, the emerging field of policy analytics is introduced. However, scientists report that more efforts need be spent on applying data analytics methods efficiently and effectively in the policy making domain. Visual analytics techniques are scarcely discussed in this domain, although visualization and data analysis is deemed to be of utmost importance within the coming years.

2.3. Visual Analytics to Support Decision Making

In the previous section, we reviewed scientific approaches in decision support and policy making theory. We discussed decision processes, involved stakeholders, and computerized technologies. In the following, we describe the fields of information visualization and visual analytics.

2.3.1. Visualization Disciplines

First, we review scientific research in visualization-related disciplines that are relevant for data-driven decision making. We review theoretic approaches on the two disciplines information visualization and visual analytics in more detail, since they are most important for our approach. Moreover, we briefly discuss the disciplines information design and knowledge discovery in data bases (short: KDD), since they play a role in some related approaches. Finally, we compare the four disciplines and their relevance for decision making.

An overview on data visualization in general is presented by Few [Few09]. In his view, the purpose of information visualization is to support the exploration, sense-making, and communication of data. Figure 2.4 shows his view on the broader data visualization field. The activities addressed by data visualization are exploration and sense-making as analysis tasks, and communication as knowledge transfer task. While exploration and sense-making have the goal to extract knowledge from data, the purpose of communication is the transfer and presentation of the extracted knowledge. Few differentiates between the data visualization disciplines information visualization (for abstract data), scientific visualization (for physically-based data), and graphical presentation (which we use as a synonym for information design). As an intermediate goal data visualization attempts to support the understanding of information hidden in the massive amount of data. The ultimate goal is to support good decision making based on the knowledge extracted from data.

Data Visualization		
Activities	Exploration Sense-Making	⇨ Communication
Technologies	Information Visualization Scientific Visualization	⇨ Graphical Presentation
Immediate Goal	Understanding	
End Goal	Good Decisions	

Figure 2.4.: A characterization of data visualization by Stephen Few [Few09]. Distinction between activities addressed by data visualization. Exploration and sense-making as analysis tasks. Communication as knowledge transfer task. Technologies differ with respect to presented data (information visualization for abstract data and scientific visualization for physically-based data) and interaction capabilities of the technologies, e.g., graphical presentation alone does not imply user interaction. Understanding of provided information as intermediate goal. Good decisions based on derived knowledge as end goal.

Information Visualization

Information visualization is defined as "the use of computer-supported interactive, visual representations of abstract data to amplify cognition" [CMS99]. The definition emphasizes several aspects of information visualization. First, as a research discipline from computer science its solutions are provided as software. Second, these software solutions address the visual representation of data through visual artifacts or diagrams. The artifacts consist of basic visual elements that have been presented in the theory of graphics by Bertin [Ber83]. Third, in contrast to static visual representation, information visualization deals with the interactive visual representation of data. Hence, an important aspect lies in the possibility of the user to interact with graphics generated by the software. Zoom and filter operations on the data are an example for user interaction. Fourth, information visualization deals with abstract data in contrast to scientific data. While scientific data is typically physically-based, reflecting at least some geometric information, abstract data such as economic data or text document collections is not. Finally, the goal of information visualization is to amplify cognition. Cognition is defined as the acquisition of knowledge and insights about the world. With information visualization the user is enabled to gain knowledge about the internal structure of the data and causal relationships within. Thereby, vision as the human sense with the highest bandwidth is exploited to support the comprehension of information. "Visual representations and interaction techniques take advantage of the human eye's broad bandwidth pathway into the mind to allow users to see, explore, and understand large amounts of information at once" [TC05]. Card et al. provide further details on how visualization can amplify cognition: "1) by increasing the memory and processing resources available to the users. 2) by reducing the search for information, 3) by using visual representations to enhance the detection of patterns, 4) by enabling perceptual inference operations, 5) by using perceptual attention mechanisms for monitoring, 6) by encoding information in a manipulable medium." [CMS99, p.16]

Following Shneiderman and Bederson's definition, information visualization emerged from interdisciplinary research in human-computer interaction, computer science, graphics, visual design, psychology, and business methods [SB03]. It allows to intuitively access results of complex models, even for non-experts, while not being limited to intrinsic application fields. As discussed in the previous sections, the extraction of information from data is one of the key challenges to be addressed by data-driven decision support systems which makes information visualization a valuable component for decision support systems.

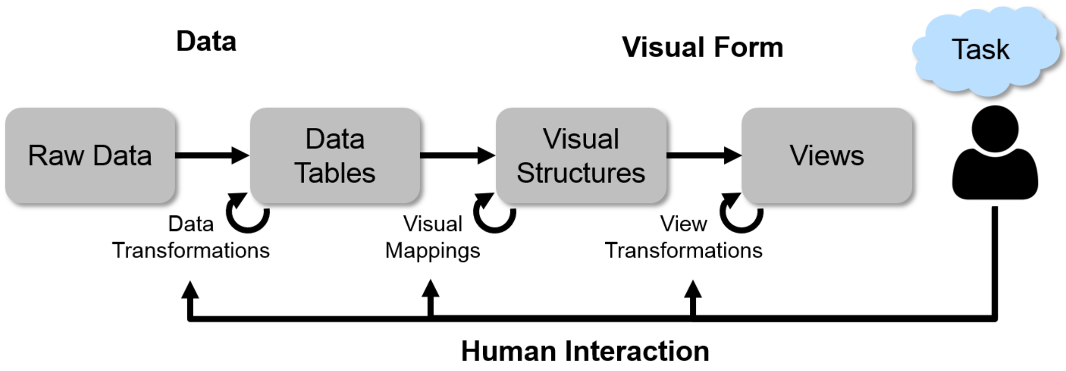

Figure 2.5.: Information Visualization Pipeline as introduced by Card et al. [CMS99]

Card et al. introduce a representative process model for information visualization: the InfoVis Pipeline [CMS99]. As shown in Figure 2.5, the InfoVis pipeline illustrates how raw datasets are processed through several transformation steps until they are visually presented to the user. An important feature of the pipeline it the interaction loop which allows users to interactively control each data transformation step in the process. An extension of interaction capabilities by incorporating interactive access to model-based parameters is realized by visual analytics approaches which we discuss in the following.

Visual Analytics

The growing amount of data collected and produced in modern society contains hidden knowledge that needs to be considered in decision making. Due to the data's volume and complexity information visualization can no longer be applied alone. A new research discipline within information visualization was introduced. Visual analytics is defined as "the science of analytical reasoning facilitated by interactive visual interfaces" [TC05]. An alternative definition is presented in the European roadmap for visual analytics: "Visual analytics combines automated analysis techniques with interactive visualisations for an effective understanding, reasoning and decision making on the basis of very large and complex datasets" [KKEM10].

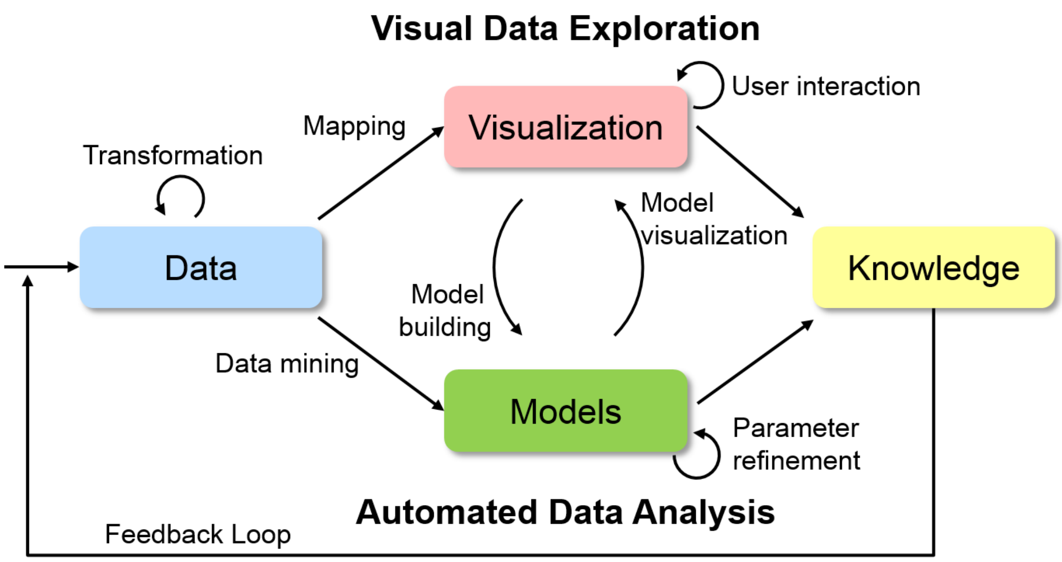

Figure 2.6.: Visual analytics process adapted from Keim et al. [KAF*08] [KKEM10, p.10]. Connecting the information visualization pipeline (top) and the KDD process (bottom).

The goal of visual analytics is the creation of tools and techniques to enable the user to "(1) synthesize information and derive insight from massive, dynamic, ambiguous, and often conflicting data; (2) detect the expected and discover the unexpected; (3) provide timely, defensible, and understandable assessments; (4) communicate assessment effectively for action" [KAF*08]. In contrast to pure information visualization, visual analytics combines interactive visualization with automated data analysis methods to provide scalable interactive decision support.

Figure 2.6 shows an adaptation of Keim's widely accepted visual analytics process model [KAF*08] [KKEM10]. The visual data exploration process from information visualization research (InfoVis pipeline, upper part) [CMS99] and automated data analysis methods for the extraction of knowledge from data (KDD pipeline, lower part) [FPSS96] are combined to one visual and interactive analysis process model. The user is directly included in the model by interactive access to the process steps. This generic process model makes visual analytics applicable to a variety of data-oriented research fields such as engineering, financial analysis, public safety and security, environment and climate change, as well as socio-economic applications and policy making, respectively. The scope of visual analytics can also be described in terms of the incorporated information and communication technologies (ICT) like information visualization, data mining, knowledge discovery or modeling and simulation [KAF*08]. In its framework program seven, the European commission emphasized visualization as a key technology in the objective "ICT for governance and policy modeling" [Eur10]. Recent approaches in visual analytics focus on the questions how to simplify the access to the analysis functionality of visual analytics techniques, and how to present analysis results. This includes the analysis process with its intermediate steps, and the findings derived with the visual analytics techniques [KM13].

Comparison of Visualization Disciplines

We already discussed the two visualization disciplines information visualization and visual analytics. To provide a full picture we add the two disciplines information design and knowledge discovery in databases (KDD). We now briefly discuss the differences between information design, information visualization, visual analytics, and knowledge discovery in databases (KDD). In Figure 2.7, we characterized the visualization disciplines based on the involvement of humans and computers in the reasoning process [KNRB12]. Table 2.1 gives a more detailed differentiation between the disciplines.

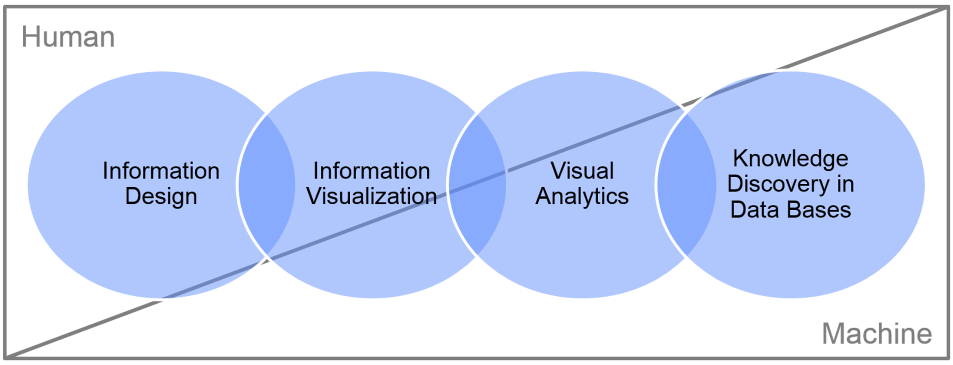

Figure 2.7.: Visualization disciplines categorized by involvement of human and computer. Adapted from Kohlhammer et al. [KNRB12].

Information design targets presenting and communicating the key insights of an already concluded analysis. It is designed after the analysis was already conducted. In most cases the static visualizations are printed on paper. Little to no user interaction is involved. The reasoning, or the consumption of the information is only done by the user without the need of a computer.

Vis Discipline	Goal	User Type	Data Abstraction	Interaction	Pattern Detection	DM equivalent
Information Design	present insights	decision maker	aggregated reporting data	low	human (guided)	Reporting
Information Visualization	explore insights	analyst	structured data	medium	human	Business Intelligence
Visual Analytics	explore and extract insights	expert user	unstructured data	high	combined human+machine	Visual Business Analytics
KDD	extract insights	data mining expert	unstructured data	low	machine	Business Analytics

Table 2.1.: Visualization Disciplines

Information visualization targets the interactive exploration and analysis of structured data. It provides more interactivity than information design and allows the exploration of larger datasets via overviews, zooming, filtering, and drill-downs. The user is involved in the analysis via interaction. Automated data analysis techniques (e.g., data mining and machine learning) are only loosely integrated in the workflow. Mainly, the results of an automated analysis are shown.

Knowledge discovery in databases (KDD) targets extracting knowledge from data via automatic data analysis methods like data mining and machine learning [FPSS96]. It is an important component of visual analytics. Scientists in this field access the underlying algorithms via command lines or basic graphical user interfaces. Visualization techniques are only applied for the presentation of results.

Visual analytics combines information visualization and data mining to support the analysis of large complex datasets. This allows the simultaneous exploitation of the humans' perceptual power and the machines' processing power. The tight integration between visualization and automated data analysis supports the highest level of interaction.

In the remainder of this thesis, which primarily focuses on information visualization and visual analytics approaches, we will use visual analytics as an umbrella term for interactive visualization disciplines. Visual analytics is defined as a discipline that combines information visualization and data mining (or: KDD). Therefore, we interpret information visualization as a special case of visual analytics with limited integration of data mining, while we interpret data mining as a special case of visual analytics with limited integration of information visualization.

2.3.2. The Visual Analytics Design Process

Our approach aims at integrating visual analytics approaches into the decision making process to improve data-driven decision making. The design of visual analytics systems is a non-trivial task. As van Wijk states "successful solutions are often found through trial and error, with solid guidelines and findings still lagging" [vW13]. As a core idea to improve the likelihood of developing successful solutions, the user-centered design approach introduced by Norman and Draper is applied throughout this thesis [ND86]. User-centered design promotes the participation of users throughout the entire design process, and the validation of each design step via user-centered evaluation. Isenberg et al. introduce evaluation to the development life cycle by repetitively iterating through the stages design, implementation, and evaluation [IZCC08]. In the following, we review existing design approaches that support the successful design of visual analytics systems. The review influences the definition of our specific design process targeting visual analytics decision support systems.

Our concept is mainly influenced by the design process definitions of Andrews [And08], Munzner [Mun09], Lam et al. [LBI*11], and Sedlmair et al. [SMM12]. They presented methodologies on how to design, implement, and evaluate visual analytics solutions for data-driven challenges of domain experts. Due to their reflections upon practical experiences of hundreds of information visualization and visual analytics research papers, the value of the introduced methodologies is widely recognized. In these design methodologies visualization researchers are guided through the analysis of a specific real world problem faced by domain experts, the design of visualization systems that support solving this problem, and the validation of the design.

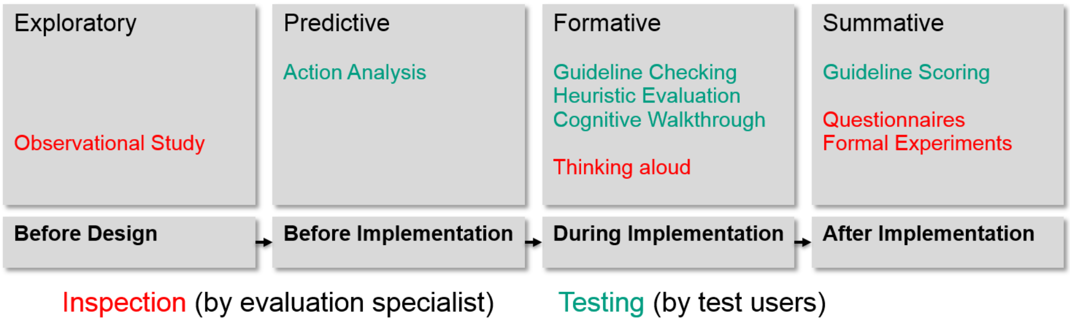

Figure 2.8.: Andrews' evaluation goals and methods in the design process [And08].

Andrews structures the design process into the stage before design, the stage before implementation, the stage during implementation, and the stage after implementation (see Figure 2.8) [And08]. Similar process stages are introduced by Lam et al.: pre-design, design, prototype, deployment, and re-design [LBI*11]. Sedlmair et al. reduce the design and evaluation process to three stages – pre-design, during-design, post-design [SIBB11]. From this analysis, we learned that the defined design stages vary with respect to their terminology. However, the underlying design stages are similar, and can be mapped on each other in most cases. In addition to the definition of design stages, Andrews maps the following evaluation goals to the stages: exploratory (how will the system be used and which tasks is it supposed to address), predictive (how effective and efficient might the user be with the system), formative (how could the current implementation of the system be improved), and summative (how well is the system finally performing). Finally, Andrews provides suggestions for evaluation methods to be applied at the distinct stages. [And08] In our concept, we re-used Andrews' structure of the design process and the underlying evaluation goals.

Munzner provides an alternative perspective on visualization design. She presents a nested model that differentiates between abstraction levels of a visualization design. The nested levels are: domain

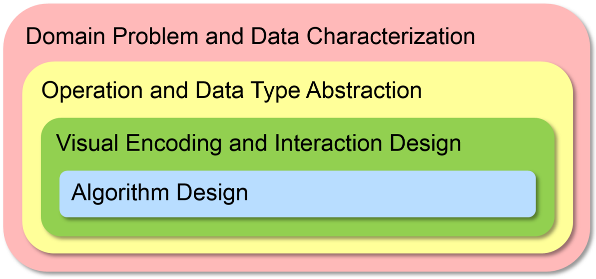

Figure 2.9.: Munzner's nested model for the design and validation of visualization techniques [Mun09].

problem and data characterization; operation and data type abstraction; visual encoding and interaction design; and algorithm design (see Figure 2.9). In the following, we present Munzner's nested layers and draw parallels to the seven evaluation scenarios by Lam et al. [LBI*11]. An initial step in the visual analytics design process is the domain problem and data characterization stage. "At this first level, a visualization designer must learn about the tasks and data of target users in some particular target domain" [Mun09]. This stage is also presented as one of seven evaluation scenarios by Lam et al.: "understanding environments and work practices" [LBI*11]. The first level implicitly describes the components to be characterized before the visualization design: data, users, and tasks. These 'ingredients' are also referred to by several other theoretical approaches. We highlight the approaches by van Wijk [vW13] and Miksch et al. [MA14]. Van Wijk describes the three components from an evaluation perspective. Miksch et al. organize them into a design triangle and add quality criteria to the triangle edges: appropriateness (between users and tasks), expressiveness (between tasks and data), and effectiveness (between users and data). Although they apply the design triangle only on time-oriented visual analytics approaches, it can be used in a broader context. The concept of this thesis is based on a general decision making domain characterization. Our domain characterization heavily relies on the data-user-task-triangle. As a scientific foundation of this characterization, we review existing data, user, and task taxonomies in visual analytics in the following sections. The second layer of Munzner's nested model is called "operation and data type abstraction". This layer is represented via the "evaluating user experience" scenario by Lam et al. [LBI*11]. It aims at deriving concrete data types from the relevant domain data identified in the first layer. In addition, the tasks and the underlying data operations to be supported in order to address the domain problem are formulated. At the third layer of the nested model the visual encodings and the interaction design is developed. This involves the design on how the data is represented at the display and how the user can interact with the resulting visualization. This layer can be evaluated by the Lam's scenario "evaluating user performance". Finally, at the fourth and inner layer of Munzner's model, the algorithm design takes place. Lam et al. also present a scenario called "evaluating visualization algorithms" [LBI*11]. Munzner promotes the development of the visual analytics design from the outer to the inner layer, validating each layer at each stage. However, some the outer levels require an additional validation after the inner layers are addressed. Munzner's model is very influential for our approach. We used several of the proposed validation techniques and focused on addressing the layers separately during the design of the system. We augment the nested model with a domain characterization including user, data, and task characterizations dedicated to the specifities of decision making.

We could map four out of seven evaluation scenarios introduced by Lam et al. to the nested model of Munzner. The three additional evaluation scenarios are orthogonal to the nested layer but still relevant for designing visual analytics decision support systems. These are “evaluating visual data analysis and reasoning”, “evaluating communication through visualization” and “evaluating collaborative data analysis”. We tackle data analysis and reasoning via the exploration and analysis tasks, and the communication via our presentation task described in our task taxonomy in Section 3.2. The collaborative data analysis scenario is tackled in Section 3.3 via the bridging of knowledge gaps between stakeholders.

Finally, we highlight the relevance of the design study methodology by Sedlmair et al. The authors present a list of 32 pitfalls to be avoided during a design process. Moreover, they expand the design process by including a last stage: the publication of the findings via a design study [SMM12]. The authors refer to the specific paper type design study that is also explained in a publication on how to write information visualization papers by Munzner [Mun08]. We emphasize that most of the technical contributions in this thesis have been published as design studies, since this paper type is suited best to prove the applicability of our concept to decision making. For further reading in design and validation of visualization approaches, we refer to the book by Munzner ‘Visualization Analysis and Design’ that aggregates previous work by the author and provides a comprehensive overview on visualization design and validation for researchers and practitioners [Mun14].

In summary, existing design study methodologies present concrete guidelines for the design and implementation of visual analytics expert systems that we could re-use in this thesis. However, in contrast to these methodologies in the field of decision making several stakeholders with different levels of expertise need to be included into the process. Moreover, decision making is a very time-critical process, which provokes fast development cycles with short requirement analysis and evaluation stages. Hence, we identify a need for an adapted design study methodology for implementing visual analytics systems in order to support decision making.

2.3.3. Data, User, and Task Taxonomies

In the following, we review data, user, and task taxonomies that serve as a theoretical foundation for the decision making domain characterization introduced in our concept chapter (Chapter 3).

Data Taxonomies

Abstract data can be described from different perspectives. Stevens categorizes data based on the underlying scales [Ste46]. He distinguishes between ratio, interval, ordinal, and nominal data scales that are characterized based on the absence of the formal properties category, magnitude, equal interval, and absolute zero. Nominal scales only have a category that differentiates a data entity from others (e.g., file name). Ordinal scales additionally a magnitude that allows to bring entities in an order (e.g., the seasons in a year). Interval scales additionally allow the concrete calculation of a difference between two entities (e.g., points in time). Ratio scales additionally have an absolute point of zero, which allows the calculation of a ratio between two items (e.g., weight). Although this categorization was criticized by statisticians [VW93], adopted version are still used in information visualization.

An adapted version of this categorization was presented in the 'Semiology of Graphics' by Bertin [Ber83]. He differentiates between three levels of organization: qualitative (or nominal), ordered, and quantitative (or interval-ratio). The latter unifies Stevens' interval and ratio scales. Card, Mackinlay, and Shneiderman reuse this definition. They describe three basic variable types: nominal, ordinal, and quantitative [CMS99]. In addition, they define four variable subtypes: quantitative spatial (as used in scientific visualization), quantitative geographical (as used in geographic maps), quantitative time (as used in temporal visualizations), and ordinal time. These reflect important conditions that need to be considered in the visualization design. Ware provides a comprehensive overview on different aspects of data categorizations [War13]. He differentiates between entities (objects of interest) and relations (that describe the relationships between these objects). Both entities and relations may have attributes that provide details on the properties of these entities or relations. An attribute can have 1 to n dimensions. The quality of each attribute value can be describes with Stevens' data scale definition (see above). Finally, Ware states that in visualization research mostly the three basic data classes category data (nominal), integer data (ordinal), and real-number data (the combined interval and ratio scales) are applied.

Shneiderman also introduced an alternative categorization of data types relevant for information visualization [Shn96]. He distinguishes between seven data types: 1-dimensional (e.g., text documents), 2-dimensional (e.g., geographic maps, or page layouts), 3-dimensional (e.g., real world objects, or volume data), multi-dimensional (as stored in relational or statistical databases), temporal (time series data), tree (acyclic graphs like hierarchies), and network data (graphs). This data type taxonomy was adapted by Keim [Kei02] who distinguishes between 1-dimensional (e.g., temporal data), 2-dimensional (e.g., geographical maps), multi-dimensional (e.g., relational tables), text/web (e.g., news articles and web documents), hierarchies/graphs (e.g., telephone calls or linked web documents), and algorithm/software data (e.g., software code). He uses a similar terminology with minor variations. For example, time series are assigned to 1-dimensional data. No explicit temporal data category is named. Moreover, Keim adds text and algorithm/software as additional categories, while he merges trees and graphs and excludes 3-dimensional data as an explicit category. In the US roadmap of visual analytics Thomas and Cook distinguish between textual data, databases, image data, sensor data and video data without explicitly defining a data taxonomy [TC05].

More recently, Meyer et al. differentiate between dataset types and attribute types. As data types they enumerate tables, networks, and text. As attribute types, they name categorical, ordered, and quantitative (similar to Stevens' definition) [MSQM15]. This characterization is refined in Munzner's book describing data types and dataset types. She defines five basic data types: items, attributes, links, positions, and grids [Mun14]. In her definition a dataset is a collection of information. Four basic dataset types are described: tables, networks, fields, and geometry. As an additional way to group data items, Munzner names the categories clusters, sets, and lists.

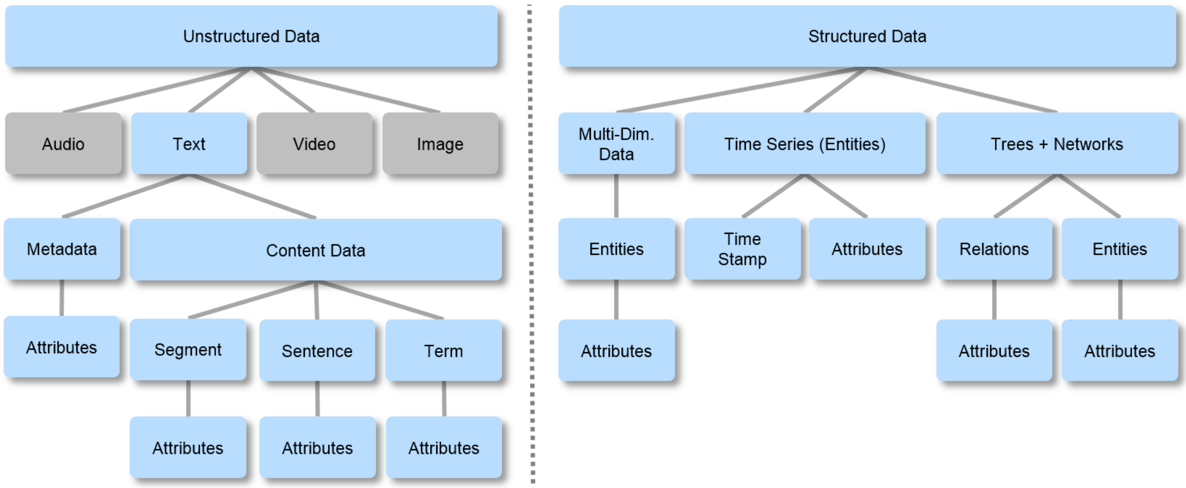

Figure 2.10.: Unified data taxonomy. Top level: structured and unstructured data. Second level: data categories audio, video, image (in gray, since not considered in this thesis), text, multidimensional data, time series, graphs and networks. Mid-level: entities and relations. Leaf notes: attributes.

In Figure 2.10, we show a unified data taxonomy informed by the reviewed data taxonomies from related work. At the first level of the hierarchy, we distinguish between structured and unstructured data categories. This differentiation can be found, e.g., in the US roadmap for visual analytics [TC05]. At the second level, the data categories as discussed by Shneiderman [Shn96] or Keim et al. [Kei02] are presented. These include audio, video, image (which are not considered in this thesis), text, multidimensional data, time series, tree and network data. Below, the specifities of the respective data category is shown, considering Ware's entities and relations [War13]. Finally, the leaves of our hierarchical data taxonomy are attributes of the respective entity or relation. Geographic data is not included as an own data category, since any data entity in our taxonomy can have an additional attribute geo-location.

The data type taxonomy described in this thesis is also inspired by the decision support system characterization by Power [Pow02]. As discussed in Section 2.1.1, he differentiates between the following DSS: communication-driven DSS, data-driven DSS, document-driven DSS, knowledge-driven DSS, and model-driven DSS. Disregarding the communication-driven DSS category, these categories relate to specific data types. Data-driven DSS handle data that is stored in a tabular, and therefore, structured format. Document-driven DSS address textual, video, or audio data comprised in documents. Knowledge-driven DSS contain set of rules or similar structures that represent extracted knowledge. And finally, model-driven DSS use data or parameters as input and produce data as output. Examples include simulation or optimization models. Based on this categorization, we distinguish between model-driven, textual, and empirical data. With empirical data, we mean structured data that is, in contrast to model-driven data, empirically collected and not artificially generated by a computer.

User Taxonomies

User characterization can be found in different domains. In our approach, we focus on user taxonomies from of decision support theory, policy making, and visual analytics research. Sprague et al. discuss different stakeholders in decision support theory [Spr80]. He states that "knowledge workers are the clientele (of DSS). This group includes managers, professionals, staff analysts, and clerical workers whose primary job responsibility is the handling of information in some form." More concretely, he differentiates between the roles manager (or user), intermediary (or staff assistant), DSS builder, technical supporter, and toolsmiths. Kandel et al. differentiate between three types of data analysts depending on their technical skills: hackers, scripters, and application users [KPHH12]. Hackers have advanced programming skills, which enables them to integrate different data sources easily. They mainly work on data analysis tasks prior to the modeling stage. Scripters are less experienced in programming but work with the advanced models by combining existing analytical packages via scripts. Application users rely on existing software with smaller datasets. They are dependent on IT staff that support them in data acquisition and preparation tasks.

In the policy making area, as discussed in the previous chapter, Howlett et al. describe three main stakeholder groups involved in the policy making process: the policy universe, the policy subsystem, and the decision makers [HRP09]. The policy universe comprises diverse stakeholder interested or affected by a policy process, including, e.g., NGOs, influencers, and the public. The policy subsystem only comprises policy actors that have sufficient knowledge in the addressed problem area. The decision makers are elected government officials. Howlett et al. further distinguish between more concrete stakeholder groups. These are elected politicians, the public, bureaucracy, political parties, interest or pressure groups, think tanks and research organizations, mass media, academic policy experts and consultants. Another common distinction of stakeholders in policy analysis are knowledge suppliers (or producers), knowledge brokers, and knowledge users (or decision makers) [HW09]. Knowledge suppliers originate from academia and research institutes. Knowledge brokers work as intermediaries between suppliers and users. They work as specialized advisers within governments, temporary in commissions and task forces, or in non-governmental organizations like think-tanks or interest groups. A categorization of stakeholders from the policy analytics perspective is presented by de Marchi et al. [DLT16]. The authors differentiate between policy makers, experts, citizens, and stakeholders. Here, the term stakeholder represents humans that are influenced by a policy.

In the context of information visualization, van Wijk identifies a gap between visualization researcher and domain expert [Wij06]. He discusses the different objectives of both stakeholder roles - publication of new results vs. solving a real-world problem - that not necessarily contribute to a successful collaboration. In their design study methodology, Sedlmair et al. discuss several stakeholders that need to be identified and involved in the design of domain-specific visualization systems. These include front-line analysts (similar to van Wijk's definition, the targeted domain experts and real end users of the system), gatekeepers (the managers that make final decisions but not necessarily use the system by themselves), connectors (people that connect visualization experts with relevant people), translators (who help in the abstraction of the problem), and others (co-authors and fellow tool builders). [SMM12] An analysis of collaborator roles relevant for the design of visual analytics approaches applied to the digital library domain is introduced by Bernard. He distinguishes between domain experts, data curators, digital librarian, gatekeeper, political & financial stakeholder, and data scientists as the designer of the visual analytics system. [Ber15] Andrienko et al. provide the characterization of user roles which is closest to the one presented in our concept. The authors review spatial decision making support in the context of geovisual analytics. They identify four user roles, the decision makers who select from a small number of options that are presented to them in a condensed way. Analysts who create the options by analyzing the problem space. Consultants who support the analyst in defining alternative options with specific domain knowledge. And stakeholders that are affected by a made decision, and therefore, try to influence the decision process. [AAJ*07]

In our concept, we slightly modify this taxonomy. We also name decision makers, analysts, and stakeholders. In addition, we differentiate between two types of consultants: domain experts and modeling experts. In our definition, the domain expert adds domain-specific knowledge encoded in data to the decision making process, while the modeling expert contributes computational models like simulation or optimization models. We make this differentiation to take the different data categories relevant for decision making into account. From our review of the related work we identified three characteristics that support the categorization of stakeholders: their technical expertise in advanced analysis tasks, their domain expertise in the specific target domain, and finally, their discretionary competence.

Task Taxonomies

Finally, as a third ingredient of our general decision making domain characterization, we review existing task taxonomies from the visual analytics field. The identified taxonomies serve as a baseline for the characterization of decision-specific tasks to be supported by visual analytics approaches applied to decision making.

The most prominent guidelines for designing a visual analytics system are introduced by Shneiderman [Shn96] and Keim et al. [KAF*08]. Shneiderman's visual information seeking mantra contains the guidelines "overview first, zoom and filter, then details-on-demand" [Shn96]. These guidelines can be translated to the visualization tasks provide an overview, allow zoom and filter, and support drill-down to raw data. Shneiderman augment the mantra with three more tasks: relate (find relations between data objects), history (store the user interactions to retrieve and refine analysis steps), and extract (support the extraction of data subsets and underlying parameters). Keim et al. adapted the visual information seeking mantra to the visual analytics mantra "analyze first, show the important, zoom, filter and analyze further, details on demand" [KAF*08]. These guidelines – or respectively tasks – build a solid foundation for the tasks that need to be addressed by any visual analytics system, which is why we also considered them during the design of our visual analytics approaches.

The knowledge crystallization model by Card et al. explains how "a person gathers information for some purpose, makes sense of it by constructing a representational framework [or schema], and then packages it to some form of communication or action" [CMS99]. Figure 2.11 (left) illustrates the circular concept containing the following meta-steps: task; forage for data; search for schema;

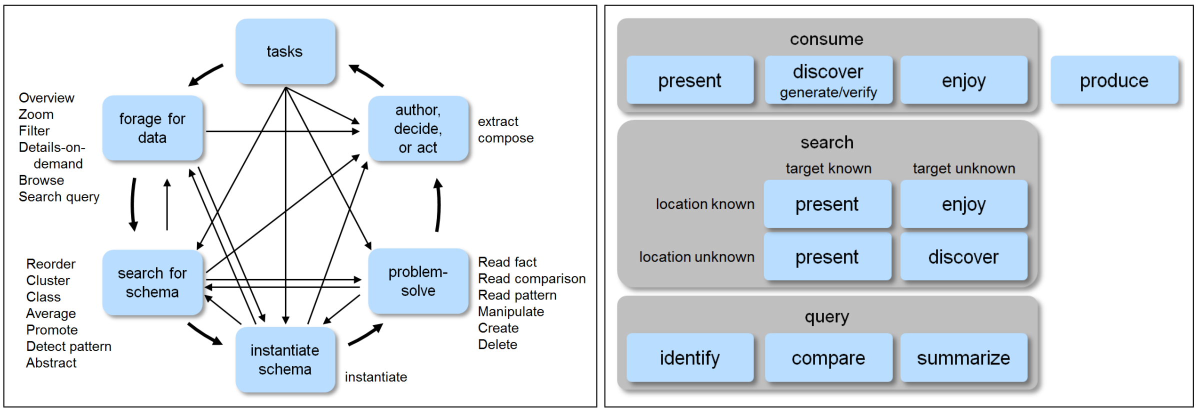

Figure 2.11.: Left: Knowledge crystallization model adapted from Card et al. [CMS99]. Right: Task typology adapted from Brehmer et al. addressing the question "why is a task performed" [BM13].

instantiate schema; problem-solve; author, decide, or act. Furthermore, the authors concretize these abstract steps by additional sub-tasks. For example, the forage-for-data-step contains Shneiderman's tasks plus browse and search query. The problem-solve step contains the sub-tasks read fact, read comparison, read pattern, manipulate, create, and delete.

A comprehensive review of the related work on visualization tasks is presented by Brehmer et al. [BM13]. The authors compare 30 task taxonomies in order to extract a multi-level tasks typology. Their typology is structured along three questions: why is a task performed, how is the task performed, and what is the input and output of a task. Figure 2.11 (right) shows the 'why'-level. In our domain characterization, we only tackle this abstract task level. However, the related work provides numerous task taxonomies on how to address the abstract task level by sequences of low-level tasks like filter, navigate, delete, etc. [WL90] [AES05] [YKSJ07] [BM13].

Schulz et al. introduce an encompassing approach on characterizing the design space of visualization tasks [SNHS13]. They characterize the visualization task space from several perspectives. Similarly to Brehmer et al., they ask questions on specific task characteristics and provide different specifications as answers to each of the questions. Table 2.2 summarizes their approach. The main findings relevant for our approach are the definition of goals to achieve with a given task. Schulz et al. distinguish between exploratory analysis, confirmatory analysis, and presentation. These high-level tasks are also represented in our approach via the tasks exploration, analysis, and presentation.

An alternative perspective on visualization tasks is provided by Amar and Stasko [AS04]. They identify analytic gaps "which represent obstacles faced by visualizations in facilitating higher-level analytic tasks, such as decision-making and learning." The authors argue that research has too much focused on low-level visualization tasks that focus on the correct and comprehensive representation of data instead of high-level analysis tasks such as decision making and learning. They present three

Question	Explanation	Options
why is a task pursued?	goal	exploratory analysis, confirmatory analysis, presentation
how is a task carried out?	means	navigation, (re-)organization, relation
what does a task seek?	data characteristics	low-level, high-level
where in the data does	target	attribute relations, structural relations
a task operate?	cardinality of data entities	single, multiple, or all instances

Table 2.2.: Design space of visualization tasks summarized from Schulz et al. [SNHS13]

rationale tasks that focus on "expressing confidence in the correctness and utility of [identified] relationships" in the data. And they introduce three worldview tasks that focus on "what needs to be shown to the user to draw a representational conclusion for making a decision". The six tasks are expose uncertainty, concretize relationships, formulate cause and effect (as rationale tasks), determine domain parameters multivariate explanation, and confirm hypotheses (as worldview tasks). Amar and Stasko's approach supports our notion in formulating high-level tasks to support decision making. Moreover, the worldview and rationale tasks are implicitly covered by our concept.

In the context of the business domain, Kandel et al. describe high-level tasks to be supported by visualization and data analysis [KPHH12]. The authors distinguish between the five data analysis tasks: discovery (identifying data necessary to address analysis task), wrangling (pre-processing and integrating data into desired format), profiling (verifying quality and suitability of data), modeling (applying summarization or prediction models on data), and reporting (presenting findings to decision makers). We considered these data analysis tasks in our concept. The following mapping explains how: discovery ↔ exploration, wrangling ↔ pre-processing (only implicitly tackled in our approach), profiling ↔ analysis, modeling ↔ model-driven data (implicitly covered through the incorporation of computational models), reporting ↔ presentation.

In summary, we have reviewed existing task taxonomies from visual analytics and related fields that have a relevance to decision support. These taxonomies inspired our own task taxonomy relevant for the design of visual analytics decision support systems.

2.3.4. Summary of Visual Analytics for Decision Support

In this section, we reviewed literature from visual analytics and related fields that are relevant for the design of visual analytics decision support systems. We started with an introduction to the visualization disciplines information visualization and visual analytics. Information visualization is defined as "the use of computer-supported interactive, visual representations of abstract data to amplify cognition" [CMS99]. "Visual analytics combines automated analysis techniques with interactive visualisations for an effective understanding, reasoning and decision making on the basis of very large and complex datasets" [KKEM10]. The difference between information visualization and visual analytics lies in the deeper integration of automated data analysis techniques, which makes visual analytics a technology that is, until today, mainly applied by expert users. Furthermore, we compared these two disciplines with information design and knowledge discovery in databases to provide a comprehensive overview of visualization-related techniques relevant for decision support. For the remainder of the thesis, we use visual analytics as an umbrella term for the visualization disciplines discussed in the related work.

After having reviewed the visualization disciplines relevant for supporting decision making, we reviewed existing work on the design of visual analytics systems. We characterized different stages in the design process: before design, before implementation, during implementation, after implementation. Furthermore, the related work highlights the necessity of considering user-centered design principles for visual analytics research. Moreover, we discussed Munzner's nested model with the four nested layers that need to be compiled and validated in any visual analytics design process. These layers include: domain problem and data characterization, operation and data type abstraction, visual encoding and interaction design, algorithm design. In addition, we reviewed data, user, and task taxonomies from the visual analytics field. These characterizations are a key task during the domain and problem characterization step. The review of existing taxonomies built the bases for our concept that aims at defining a decision making-specific domain and problem characterization required for the design of visual analytics decision support systems.

Finally, we summarize related work in visual analytics explicitly targeting decision making support. These approaches are still surprisingly scarce compared to the number of approaches presented in other analysis-driven fields [KNRB12]. From the conceptual perspective, Andrienko et al. provide the work most related to our concept [AAJ*07]. The authors present a research agenda for supporting spatial decisions with geovisual analytics. Besides the characterization of involved stakeholders, they highlight future challenges to be addressed by geovisual analytics for spatial decision support. Among others they promote (a) to support spatial decision making as a process, (b) to support the exploration of the problem and the potential solutions, (c) to support the consideration of heterogeneous information, (d) to support rational choice by multi-criteria decision analysis, (e) to support reasoning, deliberation, and communication, (f) to support time-critical decision making, and (g) to support different actors. These future research directions also hold for decision making in general and motivated this thesis, in which all of the presented challenges are at least partially addressed.

Moreover, we identified several visual analytics approaches that name decision making as an abstract task to be supported. Some of these approaches target the definition of an optimal solution for a given problem (e.g., [SME08] [OSS*17]). Others incorporate simulation models to predict the impacts of potential decision strategies (e.g., [BMPM12] [MMT*14] [KWS*14]). In addition, visual analytics approaches are applied to support users in weighting the relevance of empirical attributes and rank alternative solutions (e.g., [CL04] [GLG*13]). However, most of these approaches lack a detailed explanation on how the decision making is conducted and supported by visual-interactive solutions. In fact, most of them focus on the externalization of knowledge without providing details on how this knowledge is transferred into a decision. The only approach that maps visual analytics tasks to the decision making process is introduced by Andrienko and Andrienko [AA03]. They contribute a geo-spatial decision support system, combining visual analytics and methods from multi-criteria decision making. Despite this work, we could not find visual analytics approaches focusing explicitly on decision making processes.

3. Concept for Visual Analytics Decision Support

This thesis aims at supporting complex decision making with visual analytics. By 'complex' we mean decisions that (a) involve many stakeholders, (b) are unstructured with various dependencies and influencing factors, and (c) have an impact on many internal and external factors. Targeted decisions include unstructured strategic planning decisions at the organizational (or business) level (decision making, cf. Section 2.1) and policy decisions at the political level (policy making, cf. Section 2.2). In this chapter, we first summarize challenges that emerged from our review of the state of the art in decision making, policy making, and visual analytics, and our experience gathered in several projects with domain experts from the policy making domain. Second, we provide a concept and guidelines for designing visual analytics systems for the support of decision making. Third, we explain how the realization of the concept supports the bridging of knowledge gaps between stakeholders involved in decision making. Finally, we give an outlook on remaining challenges that are addressed in the remaining chapters of this thesis. This chapter is partially based on our previous work published in [KNRB12], [RBK13], [RBLTK14], [RDK*15].

Contents

3.1. Challenges for Visual Analytics Decision Support Systems

From our study of the related work and the experience we gathered in collaborating with domain experts in the field, we extracted the following challenges that need to be addressed to realize visual analytics support for decision making. We distinguish between conceptual and technical challenges. The conceptual challenges describe (1) a missing methodology for the design of visual analytics solutions that support decision making, and (2) the existence of knowledge gaps between stakeholders involved in the decision making process that need to be bridged. The technical challenges describe specific decision making scenarios that need to be supported with data-driven decision support. In the following table both conceptual and technical challenges are presented. More details are given below.

Conceptual Challenges	
$\mathbf{C}_{\text{VDSS}}$	Design methodology for visual analytics decision support
$\mathbf{C}_{\text{BKG}}$	Bridge knowledge gaps between involved stakeholders
Technical Challenges	
$\mathbf{C}_{\text{Proc}}$	Explore and monitor decision processes
$\mathbf{C}_{\text{Doc}}$	Explore and analyze text document collections
$\mathbf{C}_{\text{Deb}}$	Explore, analyze, and compare stakeholder opinions and arguments
$\mathbf{C}_{\text{Dat}}$	Explore, analyze, and compare empirical performance indicators
$\mathbf{C}_{\text{Imp}}$	Explore, analyze, and compare the impacts of solutions
$\mathbf{C}_{\text{Opt}}$	Create, analyze, and compare optimal solutions

Table 3.1.: Research challenges to be addressed in this thesis.

Conceptual Research Challenges

$\mathbf{C}_{\text{VDSS}}$ Design Methodology for Visual Analytics Decision Support
From the literature, we learned that there is a strong need for visual analytics techniques to support decision making. However, despite that fact no model for the design of visual analytics systems with the target to support decision making exists. Existing visual analytics models focus on the data processing pipeline [CMS99], knowledge generation [CMS99] [SSS*14], sensemaking [PC05], the integration of information visualization and data mining techniques [KAF*08] [KKEM10], or the value of visualization in general [vW05]. These approaches finish after the extraction of knowledge from data. To the best of our knowledge, no visual analytics model explains how the derived knowledge is integrated in the decision making process, or how the user is guided from knowledge to decisions. Various design methodologies exist that describe how visual analytics systems need to be built and validated to support the users in addressing real world problems. Examples include the design study methodology by Seldmair et al. [SMM12], Munzner's Nested Model [Mun09] and its derivatives (e.g., [MSQM15]), or the seven evaluation scenarios by Lam et al. [LBI*11] to name a few. However, none of them adapt to the specific needs in policy making or strategic decision making. First, the involvement of stakeholders with varying expertises is neglected. Second, the creation, analysis, and comparison of alternative solutions to a given problem are not considered. And third, specific data types to be used in the deci-

sion making process are not described. A visual analytics methodology targeting decision making is required to support the development of visual analytics decision support systems.

C_{BKG} Bridge knowledge gaps between involved stakeholders
In order to make the right decisions a profound analysis of the problems and possible solutions has to be performed. Therefore, decision makers rely on the advice of analysts. These analysts need to collaborate with external experts consulted as advisers. Due to different stakeholder expertises the entire decision process may suffer from knowledge gaps. From interviews and discussions with domain experts, we identified three main challenges, described as gaps in the decision making process. (1) The competence gap: In the decision making process, competence gaps appear between stakeholders with different knowledge and expertise. For a collaborative decision making process a transparent and unmitigated information transfer is required. Depending on the type of the communication medium and the intensity of its usage, there is a latent risk of a suboptimal information flow. The problem increases if knowledge has to be communicated via mediators. Then, additional information loss effects may appear. (2) The analysis gap: As an intermediate result of a decision making process, alternative solutions to a given problem are created. These complex solutions need to be summarized and communicated to the decision maker, who has to decide which alternative to be implemented. This bears the risk that either the applied model is not exploited in an optimal way, some results are not communicated, or, even worse, that the model is ill-defined without a sufficient validation by the involved stakeholders. The exploration of analysis results by the decision maker would be a valuable feature in an efficient decision making process, since it would support rational decision making. (3) The iteration gap: Subprocesses within the decision process like the simulation of potential impacts may not be sufficiently repeated and improved by feedback loops. Only in a few cases, a first draft is already fully developed and workable. Moreover, an iterative improvement of the analysis model and thereby, its results, raises the chance to conclude in feasible solutions. Time-consuming communication efforts contribute to this gap. A fundamental challenge in decision making is to bridge these and other knowledge gaps between stakeholders involved in the decision making process.

Technical Research Challenges

C_{Proc} Explore and monitor decision process from intelligence to implementation stage
Decision processes stretch from the identification of a problem to the implementation and evaluation of a solution to the problem. Different factors contribute to the complexity of a decision process. First, a solution needs to pass different steps within the process. Second, the processes vary in their durations, some processes may last for years. Third, an alternative solution within the process may run through several iterations. Finally, a large number of stakeholders with differing expertise are involved in the process. The decision process can be structured along text documents that record intermediate results (reporting) and discussions within the process. Furthermore, social media sources (like Twitter, LinkedIn, etc.) are playing an increasing role in strategic decision making. All stakeholders involved in the decision process need to get a comprehensive overview of these documents in an efficient way

due to temporal pressure. The following sub-challenges impede the creation of overviews. Since no standardized process steps exist, the identification of the current status is complicated. Due to long process durations stakeholders need to constantly update their knowledge about the addressed problem. The sheer amount and variety of relevant documents hampers stakeholders to condense the information and get the full picture. This also has an impact on the required amount of iterations until the implementation of a solution. A further challenge remains in achieving a critical mass of expertise to be included in decision making. Relevant experts in the field need to be identified and integrated into the process. Visual analytics can help solving these problems. However, to the best of our knowledge the decision process monitoring has hardly been subject to visual analytics research.

C_{Doc} Explore and analyze text document collections

The volume of digitally available textual data relevant for decision making is continuously increasing. Following a rule of thumb, about 80 % of the enterprise information originates from unstructured (textual) data. Examples for document collections include newspaper articles, scientific papers, technical reports, patents, legislative documents or social media entries like tweets, blog posts or customer reviews. These documents are highly relevant for many types of decision makers in political and business environments. Methods from information retrieval are the means of choice, if stakeholders can specify their information need precisely, e.g., by formulating a *search* query. However, these fact retrieval or known-item search techniques often become ineffective, if document collections are large, complex, or unknown. In such scenarios, the goal to gain an overview of the document collection can be achieved via the *exploration* of structural information within the collection. The mechanisms needed to enable the exploration of document collections strongly differ from classical search methods [WR09]. Among others, data aggregation methods support the generation of content-based overviews, by condensing large numbers of documents into a small set of representatives. Examples include topic modeling, document clustering, or document classification with a plethora of techniques. However, analysis approaches based on these methods are confronted with a variety of challenges. A key challenge lies in the transformation of unstructured textual content into machine-readable numerical data. Content-based overviews of text document collections strongly depend on the chosen methodology for representing text documents, and the chosen aggregation algorithms. The underlying text analysis workflow must be adapted to the given data, user, and task at hand. However, text analysis requires a high level of technical expertise. Satisfying the specific demand of decision makers is a challenge that experts in text analysis need to face.

C_{Deb} Explore, analyze, and compare stakeholder opinions and arguments

Today, decision makers are requested to integrate large amounts of external knowledge and public opinions in their decision making processes. Examples include customer reviews, or public discussions on policy topics in social media. Large parts of the information to be considered are available on the web. However, the manual monitoring and analysis of this data is time-consuming and therefore not applicable in real-world scenarios. Automatic methods for the mining and analysis of textual content exist. These include classical text analysis methods like clustering, classification, topic modeling, sentiment analysis, argument extraction and summarization, and document summarization. To make use of these powerful tools, some challenges need to be tackled. First, the methods need to be combined to a workflow and adapted to the addressed domain. Second, the results of the workflow need to be presented to decision makers in an intuitive way. Third, since the accuracy of text analysis methods is not necessarily satisfying the users' expectations, concepts for incorporating user feedback into the text analysis workflow need to be considered.

C_{Dat} Explore, analyze, and compare empirical performance indicators

Besides the challenges related to textual data, empirical (or: structured) data also needs to be considered in the decision making process. Examples include sensor data, financial data, sales figures, earth observation measurements, demographics, etc. This data needs to be incorporated in the decision making process. Challenges lie in the collection, integration, processing, analysis, and visualization of these datasets. Due to the increasing amount of collected data, the big data problematic [CCS12] becomes relevant for decision making. In the business context methods from business intelligence, business analytics, and visual business analytics come into play. In the policy context policy analytics is evolving as a new discipline. However, despite the fact that numerous visual analytics approaches target the analysis and exploration of structured tabular data, a specific focus on the decision making process reflecting specific decision making tasks is missing.

C_{Imp} Explore, analyze, and compare the impacts of solutions

During the decision making process the impact of alternative solutions to a given problem needs to be explored, analyzed and compared before a course of action is selected and implemented into practice. Impacts can be described by economic, environmental and social determinants. Among others, simulation and regression models are common choices for assessing the impact of a decision. Analysts use modeling results to compare the impacts of alternative solutions and communicate them to the decision maker. Additional alternative solutions are created by experimenting with alternative parameterizations of models and the comparison of the model results. So called what-if scenarios help to exploit the design space of alternative solutions. However, competences and responsibilities are distributed over different stakeholders, e.g., modeling experts, analysts, decision makers, etc. Moreover, the entire process from the model type selection, over the model creation and parameterization, to the result analysis, contains design decisions to be made and discussed among the stakeholders. Exploring the multitude of possible alternative solutions, understanding the decision space, and identifying the most appropriate solution are non-trivial tasks and therefore, challenges related to computer-based decision support.

C_{Opt} Create, explore, analyze, and compare optimal solutions

Complex decisions are often depending on multiple criteria which include trade-offs between important factors. One possible way to address multi-criteria decision problems is mathematical optimization. Established models and algorithms that support multidimensional problems exist, but they have to be transferred to the decision making process. Moreover, trade-offs evolving through contradicting target functions cannot always be resolved with optimization. In these scenarios, multiple optimal solutions exist, e.g., along a Pareto frontier. Handling the complexity of these models and the underlying degrees of freedom poses a challenge to decision makers. On the organizational level methods from operations research are incorporated in the decision making process. On the political level strategic environmental assessment (SEA) is applied. Nevertheless, both lack visual interfaces that provide an intuitive access to decision makers.

With the identification of the described challenges, we have set the stage for our concept for visual analytics decision support. In the remaining part of this chapter, we address the first two challenges ($\mathbf{C}_{\text{VDSS}}$ and $\mathbf{C}_{\text{BKG}}$). First, we present a design methodology for visual analytics decision support ($\mathbf{C}_{\text{VDSS}}$). Second, we explain how visual analytics systems that follow our guidelines help to bridge knowledge gaps between stakeholders in the decision process ($\mathbf{C}_{\text{BKG}}$). Finally, we provide an outlook on the following chapters that address the presented technical challenges ($\mathbf{C}_{\text{Proc}}$, $\mathbf{C}_{\text{Doc}}$, $\mathbf{C}_{\text{Deb}}$, $\mathbf{C}_{\text{Dat}}$, $\mathbf{C}_{\text{Imp}}$, $\mathbf{C}_{\text{Opt}}$).

3.2. Design Methodology for Visual Analytics Decision Support

In this section, we present a methodology for the design of visual analytics decision support systems. This addresses the first challenge $\mathbf{C}_{\text{VDSS}}$ identified in the previous section. Visual analytics systems targeting real-world problems need to be carefully designed. As an initial step in the design process, the targeted domain and problem need to be characterized. Since we are aiming at a general domain characterization for decision making, we first describe the representative decision making process as introduced by Simon [Sim60] (Section 3.2.1). Second, as emphasized in several design methodologies from the visual analytics research community (e.g., [vW13], [MA14], [Mun09]), the main ingredients for a visual analytics design need to be characterized: users, data, and tasks. Hence, we introduce a general characterization of the data to be considered, the user roles to be involved, and the tasks to be supported in the decision making process (Section 3.2.2). Finally, we present a visual analytics design process that guide visual analytics experts in how to combine visualization and automatic data analysis functionality to support the decision making process (Section 3.2.3).

3.2.1. The Decision Making Process

Simon's representative decision making process model consists of three major steps: intelligence, design, and choice [Sim60]. This model was later extended by two further steps - implementation and review. As described in the related work section, several alternative extensions, refinements, and adaptations of Simon's model exist. However, at their core most decision making models are derived from this initial model, which allows us to generalize our concept to strategical decision making problems. In the following, we describe the consecutive steps in the decision making process in more detail.

Intelligence

The intelligence step is the first step in the decision making process. It comprises the identification of a problem and the foraging of information relevant for the decision. As an initial task, the need for making a decision is realized. Therefore, conditions in the environment that call for action are identified. This implies the characterization of the underlying problem. Information relevant to the problem is gathered. Since in many cases the scope of the problem is not clear, the information gathering might involve exploratory search processes. As a first result of the Intelligence step, the problem characterization based on the extracted information is specified. This includes the identification of relevant parameters and dependent variables. As a second result, the objectives of the targeted decision are defined. In addition to the objectives of the decision, constraints on possible solutions are specified. In summary, the results of the Intelligence step include (a) a concrete problem statement, (b) the objectives of the decision, (c) constraints on possible solutions, and (d) a collection of information that support the consecutive steps of the decision making process.

Design

The design step is the second step of the decision making process. It targets the creation and analysis of alternative solutions to the problem defined in the Intelligence step. Based on the data and information gathered in the design step, alternative solutions to address the problems are created. These solutions consider the specified objectives and constraints. Moreover, the impacts of potential solutions are evaluated. Additional factors might be identified, that have not been considered in the Intelligence step, but are affected by a created solution.

Choice

In the choice step, the alternative solutions created in the Design step are compared to each other and the an optimal (if available) solution is chosen for action. The individual performances of the alternative solutions towards the previously defined objectives are compared. Moreover, the impacts of individual solutions on external variables are considered and compared to the impacts of alternatives. In decision making processes, opposing objectives can invoke trade-offs. Therefore, objectives are weighted based on their relevance. In these cases, optimization models can help to balance trade-offs. Finally, a decision is made. In most cases, it is finally a human who decides which alternative solution is chosen. This step concludes the decision making process.

Simon's decision making process model is often extended by two additional steps: the Implementation and the Review step. For the sake of completeness, we briefly present these two steps, although we only consider the three main steps in our concept.

Implementation

The implementation step contains the realization of the chosen alternative solution. This might also include the planning and coordination of actions to be taken to approach the defined solution. This planning can be supported by several visualization techniques (e.g., GANTT-charts, UML-diagrams). Research and practical guidance on this phase of the decision making process exist. Therefore, in this thesis, we will not focus on this step.

Review

The review step is focused on the a posteriori evaluation of the implemented alternative. Since, in most cases the impacts of a decision can only be predicted, or estimated, it is necessary to measure the real effect of a decision in real-world scenarios. This monitoring activity is covered by business intelligence (BI) applications in practice. The comparison of target and effective performance can lead to the refinement of a decision in a consecutive decision making process. Therefore, the decision making process can also be seen as a 5-step process cycle.

3.2.2. Decision Making Domain Characterization

The described decision making process builds the foundation for our decision making domain characterization. In the following, we characterize data, users, and tasks along this process. First, we introduce a data taxonomy adapted to the specifities of decision making processes. Second, we characterize different stakeholders involved in the decision making process. These stakeholders are described as potential users of visual analytics decision support systems. Third, we summarize the main tasks in the decision making process that need to be supported with visual analytics. The resulting task taxonomy serves as a guideline for the identification of visualization tasks to be considered during the design of visual analytics decision support systems.

3.2.2.1. Data Characterization for Decision Support

In the context of decision making, we differentiate between three major data categories: textual data, empirical data, and model-driven data (see Figure 3.1). This differentiation is motivated by our study of related work in decision support system theory (e.g., [Pow02]) and the data taxonomies found in visual analytics research (e.g., [Mun14]). In contrast to existing data categories from the visual analytics field, we distinguish data based on two characteristics: the structure of the data, and the origin of the data. When we write about data, we premise its availability in digital form, which means that it is stored in some digital way. A large range of research approaches on the digitization of non-digital data exist. However, this line of research is not subject of this thesis.

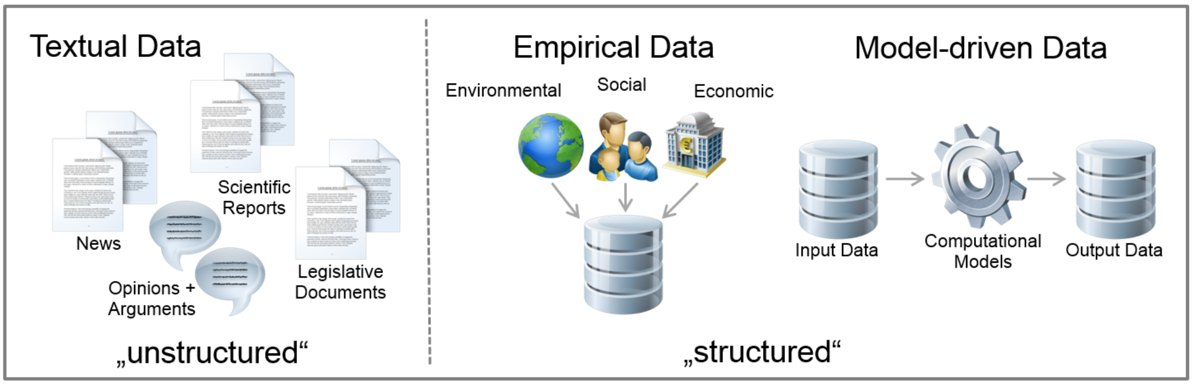

Figure 3.1.: Data categories separated into unstructured (textual) and structured (empirical and model-driven) data. Empirical and model-driven data differ with respect to origination and structure. Empirical data is collected from the real-world and contains only entities and attributes. Model-driven data is artificially created by a computational model and characterized by the relation between input and output data.

Textual Data

By textual data, we refer to a sequence of words ordered in a meaningful way. The order of words follows grammar rules (syntax), and the meaning of the text is interpretable by a human reader (semantic). Textual data originates from intentionally written text or text transcribed from speech.

In this thesis, we consider the basic entity of textual data, the text document, as a consecutive sequence of one or more words. A text document can strongly vary in its size ranging from short statements (e.g., social media comments, or SMS), over medium sized articles (e.g., research papers, or news articles) up to full-sized books. Textual data is an unstructured data type, which makes it difficult to process by computers. For example, text documents cannot be automatically compared without pre-processing the unstructured text. However, in many cases text documents are complemented by metadata. Metadata provide information about the text document and is classified into three different types: administrative, structural, and descriptive metadata [Org04]. Administrative metadata helps to manage the data resource and conveys mainly technical information like file type, or rights management information. Structural metadata provides information about the internal structure of the data, e.g., how a book is separated into chapters. Finally, descriptive metadata contains information about the content data, e.g., author, or keywords, which simplifies the categorization and search of documents. Exemplary metadata fields include ‘title’, ‘author’, ‘creation date’, ‘topic’, ‘document type’, or ‘keywords’. These meta-tags simplify the indexing of text documents. Figure 3.2 shows the structure of textual data.

Textual data is highly relevant for decision making, since many decisions are based on information and knowledge encoded in textual data. Moreover, most intermediate results of the decision making process are documented in written protocols or reports. In fact, textual data is the most prominent data category in decision making. Following a highly cited but not scientifically validated rule of thumb, about 80 % of the enterprise information originates from unstructured (textual) data. Examples include books, newspaper articles, scientific reports, research articles, legislative texts, blogposts, and statements posted through social media channels like Twitter, LinkedIn, Facebook etc. However, the inclusion of textual data in the decision making process induces several sub-challenges.

SC$_{\text{Text-1}}$ Identification of relevant documents: The initial task for a decision maker, mainly executed in the intelligence step, is to forage documents that are relevant to a given problem or topic space. Relevant documents need to be distinguished from non-relevant documents. However, the sheer amount of available textual data imposes a challenge. Methods from computational text analysis address this challenge. From the perspective of text analysis research, the identification of documents relevant for a given topic can be described as a categorization task. Three different approaches on text categorization exist: information retrieval, supervised learning, and unsupervised learning. If the target of the search is already known, hence, if a concrete search query can be specified, techniques from *information retrieval* are the means of choice. Examples include classical search engines like Google etc. These techniques take a search query as input, and provide an ordered list of documents fitting to the search query as output. However, in some cases, the search query cannot be defined, e.g., if a decision maker does not know the scope of the underlying problem yet. Moreover, the user might want to define a threshold that separates the ranking of documents into relevant and irrelevant documents. If the user already knows some exemplary documents deemed to be relevant, these documents can be used as training data in a *supervised learning* approach. Supervised learning aims at generating a model based on labeled training data that automatically assigns labels to unseen data. Text classification is a prominent example of supervised learning. Examples are spam filtering, genre classification, etc. In decision making, documents can be labeled as relevant or non-relevant. Moreover, documents can be classified into thematic categories. This supports the structuring of document collections. Finally, if neither the target of the search nor the categories including example documents are known, *unsupervised learning* is applied. These techniques group similar documents within a document collection into clusters. That way, a document collection is organized into groups without prior knowledge about their content.

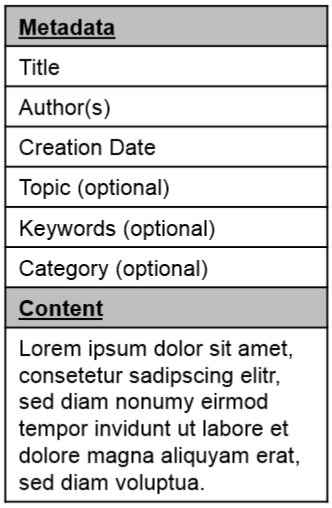

Figure 3.2.: Textual Data: the structure of a text document is separated into metadata and content data.

SC$_{\text{Text-2}}$ Extraction of key information: After the categorization of documents, the key information needs to be extracted from the documents. This includes document summaries, identified topics, stated facts, arguments, and opinions that may contribute to the decision making process. In most cases, decision makers cannot read the entire document collection considered relevant. The automatic extraction of the key information about a given problem is required. Text mining research provides various methods that support this task. Techniques from the field of *topic modeling* are able to detect word clusters as thematic groups of terms with high co-occurrences within document collections. Topic modeling approaches are often used to provide an overview on the topics discussed within a document collection. These results can be augmented with result from *information extraction* techniques, which yield at extracting named-entities like dates, persons, countries, companies, etc. from documents. More recently, this also includes document *summarization* methods that summarize the content of single or multiple documents *Opinion mining* (or sentiment analysis) techniques aim at identifying subjective information within textual data and calculate the underlying polarity (neutral, positive, or negative). Some approaches add a more detailed scale to the three classes. Moreover, aspect level opinion mining approaches even provide more detail by extracting the polarity towards entities and aspects described in the text. Finally, more recently some research has been focused on argument extraction. These methods aim at extracting arguments from textual data.

SC$_{\text{Text-3}}$ Assessment of quality: Third, text documents normally strongly vary in their quality. For example, a peer-reviewed scientific paper is more reliable, and therefore, should have a higher impact on a decision than a Tweet of a user profile with twenty followers. Several aspects provide hints on the quality of a text document. The metadata of a document can already help to judge on its quality: the author, the source, the creation date, the document type. In addition, automatic text analysis approaches aim at estimating the syntactical quality of documents. However, the final judgment on the quality and credibility of a written text has to be made by humans. Nonetheless, research in this direction is ongoing.

We described challenges related to incorporating textual data in the decision making process. In addition, we reviewed several text analysis methods that support solving these challenges. However, most of these analysis methods have a high level of complexity, which hinders decision makers to apply them in the decision making process. Therefore, we promote to facilitate the application of text analysis methods in the decision making process by providing an intuitive visual-interactive access to the underlying algorithms for non-experts in the field.

Empirical Data

We define empirical data as a second data category to be considered in the decision making process. By empirical data, we mean structured data that was empirically collected. In our definition, ‘structured’ means that the data can be stored in some tabular way. Figure 3.3 shows the internal structure of empirical data. In most cases, the tabular structure is derived from organizing data entities in rows, and their attributes in columns (cf. Section 2.3.3). This also holds for graph-based data. For this special data type, nodes and edges can be organized as rows and their attributes in columns. This makes it distinguishable from textual data. Facing the origin of the data, we differentiate between three types of empirical data: social, environmental, and economic. In political decision making, the relevance of these categories is obvious. Nevertheless, in modern business theory, the three categories are also considered [Elk97]. The triplet is derived from sustainable development theory and is also known as TBL or 3BL (triple bottom line) as defined by John Elkington [Elk97]. We provide an example to clarify the distinction between these three data origins. In the example, a company collects empirical data to inform the decision whether to continue the production of a specific car. Relevant social data include the number of workers employed in the production process, or the safety of the car with respect to crashes. Environmental aspects include the carbon emission of the car. Finally, the profit generated with the car is an example for economic data. In this example different data origins are interlinked and correlate to each other.

Empirical data is highly relevant for the decision making process, since it brings factual knowledge into the analysis of a given problem. While textual data adds the receiver's human interpretation of written, empirical data contains the raw original data measured in the environment, the economy, or the society. The usage of empirical data in the business environment is already established via the fields of Business Intelligence (BI) and Business Analytics (BA). Recent approaches are also focusing on providing an intuitive visual access to this information (Visual Business Analytics) [KPW13]. In summary, of the three data categories described in this thesis, empirical data is probably the most established in the context of decision making. Still, several sub-challenges concerning this data category remain.

Object/Attribute	Attr. 1	Attr. 2	Attr. 3	…
Data Object 1				
Data Object 2				
Data Object 3				

Figure 3.3.: Empirical Data: structured in tabular format, containing data entities (organized in rows) and data attributes (organized in columns).

$SC_{Data\text{-}1}$ Identification of relevant data: The available amount of empirical data is constantly increasing. Foraging relevant data to support the decision making process is still a challenging task. Open data initiatives are supporting the collection and distribution of freely available data. However, the decentralized storage further impedes the identification of relevant data.

$SC_{Data\text{-}2}$ Complexity of data: Although empirical data in our definition is available in a tabular, and therefore, structured way, it may contain hundreds of dimensions and many more entities. Extracting information from large datasets is subject to extensive research in the fields of KDD (knowledge discovery in databases) and visual analytics. However, the available techniques are mainly focusing on the exploration and analysis of data. The presentation of analysis results in a comprehensible way by reducing the complexity of information is often disregarded, although this is an important task required in the decision making process.

$SC_{Data\text{-}3}$ Availability and Uncertainty of data Although vast amounts of empirical data exist, in many cases the data might not exactly fit the given problem. As a consequence, some data needed for the decision making process might not be available. Therefore, strategies for collecting new social, economic, or environmental datasets need to be established. Moreover, data might be available, but erroneous or incomplete. Raising the awareness of uncertainty in the data is a challenge related to empirical datasets.

Model-Driven Data

The third data category relevant for the decision making process is model-driven data. By model-driven data, we mean data that is used as an input or created as an output of a computational model (e.g., simulation, optimization, etc.). Figure 3.4 shows the internal structure of model-driven data. Hill et al. provide a definition for computational models: "a set of computational codes, executable in some software/hardware environment, that transform a set of input data into a set of output data, with the input, output, and transformation typically having some interpretation in terms of real-world phenomena" [HCSG01].

The definition emphasizes the transformation of input into output data and the relation to real-world phenomena. In the decision making process, computational models can be applied to simulate or predict social, environmental, or economic behavior as response to a strategic decision. The decision is encoded into specific input parameters of the model. The output of the model the resulting behavior after the decision has been made. The difference to empirical data lies in the artificial origin of the output data and in the dependencies between input parameters and output data. Each model run is stored as a scenario (see Figure 3.4). We distinguish between deterministic and probabilistic (stochastic) models. While in the deterministic case, for given input data the model always produces the identical output data, this does not hold for probabilistic models. In the probabilistic case, the model is run several times with the same input variables. The output is aggregated and represented, e.g., as a probability distribution. In addition to simulation and prediction, we count optimization models to the class of computational models relevant for the decision making process. Optimization models can support the choice step by finding the optimal among a set of alternative solutions. Here, the input is given by an objective function and constraints, while the output is the calculated optimum. With these models, what-if scenarios can be executed that can inform the decision and support the search of an optimal solution, mitigating trade-offs between opposing variables.

	Input Variables		Internal Variables		Output Variables	
	In 1	In 2	Int 1	Int 2	Out 2	Out 3
Scenario 1						
Scenario 2						
Scenario 3						

Figure 3.4.: Model-Driven Data: can be divided into input data, computational code and output data. The computational code may contain additional real-world variables that are hidden from the user. In many cases the model is fed with empirical data.

$\text{SC}_{\text{Model-1}}$ Identification and development of appropriate models: Although a variety of computational models exist, the choice of an appropriate model for a given problem is challenging. The selection must be based on the data and task at hand. In most cases, this choice can only be made by an expert with profound knowledge in modeling techniques. Furthermore, since no one-fits-all solution exists, the models must be adapted to the decision domain, and fed with data from the domain. In most cases, this implies immense development efforts.

$\text{SC}_{\text{Model-2}}$ Complexity of models: A second sub-challenge lies in the complexity of computational models. Computational models are often implemented as black box systems. For a given input an output is calculated. For the end user the interpretation of the output is hampered by the complexity of the model. This fact reduces the trust of the end users in such systems. Moreover, the complexity of models might prevent decision makers to include computational models in the decision making process.

SC$_{\text{Model-3}}$ Uncertainty of models: Another sub-challenge relates to the uncertainty present in computational model outputs. Computational models are designed to simulate the real world under laboratory conditions. They are not able to reproduce it entirely. Hence, this uncertainty has to be quantified and communicated to the decision maker. In some cases, uncertainty is quantifiable, and can be calculated as accuracy, precision, recall, etc. However, this only holds for the available data that the model is optimized for. Especially, human behavior is still very challenging to predict.

3.2.2.2. User Roles

As a second ingredient of our decision making domain characterization, relevant stakeholders involved in the decision making process are characterized. These stakeholders are also potential users of visual analytics decision support systems. From our experience gathered in several design studies in the field of decision making, we identified five main user roles: decision makers, analysts, domain experts, modeling experts, and stakeholders (or influencers) (see Figure 3.5). In general, these user roles can be characterized based on their technical expertise, domain expertise, and decision competence.

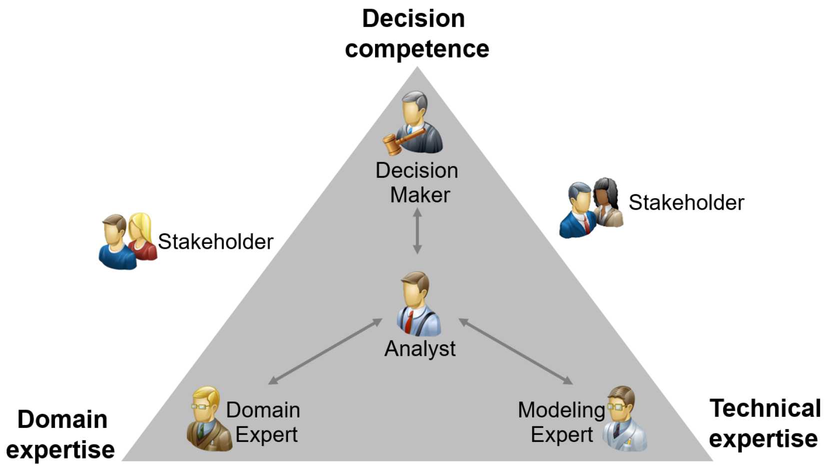

Figure 3.5.: User roles in the decision making process

Decision Makers play the key role in the decision making process. They initiate the decision process by identifying a critical situation and grasping the conditions that call for action. And finally, they decide which alternative solution is chosen to address the given problem. In most cases, decision makers do not have the time and the technical background to execute the entire decision making process by themselves. For making profound decisions they have advisers, namely analysts, that conduct the analysis, and provide summaries of the analysis in the form of reports, or presentations. Still, decision makers have the key responsibility in the decision making process. In the intelligence step, they call for action by initiating the process. Moreover, they provide input for the concrete definition of the problem. In the design step, they provide guidance in defining alternative solutions to the problem. And finally, in the choice step, they decide which solution to be implemented. However, during the entire process the detail level of the information viewed by the decision maker has to be compact and with reduced complexity.

Analysts are the coordinators of decision making process. Their task is to conceptualize the problem based on the objectives and constraints specified by the decision makers. In the intelligence step, they search for information sources and consult external advisers, e.g., domain experts and/or modeling experts, that provide assistance in the decision making process. Analysts are also responsible for the creation of alternative solutions in the design step. In the choice step, they provide comparisons of alternative solutions. Finally, in the choice step, they present the results of the analysis to the decision makers including a proposition of the alternative to be chosen. In most cases this is done via written reports or presentations. The final decision is made by the decision maker.

Domain Experts are in most cases external advisers recruited by the analyst. They contribute additional domain knowledge relevant for the specific problem to the decision making process. This can be provided in the form of textual data (e.g., scientific papers, technical reports, etc.), empirical data (e.g., research data, sensor data), or model-driven data resulting from existing computational models. In addition, domain experts might contribute solutions to the given problem, or provide insights on possible impacts of these solutions. In many cases, computational models applied in the decision making process have to be fed with domain knowledge and data, that the domain experts can provide. Domain experts do not necessarily have expertise in computational modeling techniques. They mainly serve as information and data provider.

Modeling Experts are, like the domain experts, external advisers recruited by the analyst. They have profound knowledge in computational modeling techniques. Expertise in the addressed domain is not necessarily required from modeling experts. However, they might have to feed their models with empirical data from the domain that can be provided by domain experts. The computational models are used, e.g., for simulating the impacts of the alternative solutions created by the analyst. In addition, modeling experts can contribute optimization models that support the choice step by identifying optimal solutions in trade-off scenarios.

Stakeholders are not explicitly involved in the decision making process, but try to influence decision makers since they are affected by the results of the decision. They play an increasingly important role in today's decision making processes. The transparency requested in many decision making processes allows external stakeholders to monitor the progress of the process. Moreover, by providing an intuitive access to intermediate results of the process, the underlying data, and models, even non-experts can understand the impacts of a decision. This further increases the transparency of the whole decision making process, supports democratic ideas, and raises the trust in decision makers. Still, stakeholders will also try to influence the process, e.g., by sharing their opinions and arguments.

In our concept, we distinguish between these user roles. However, in a real-world scenario, a person might hold several roles at once. Decision making processes can principally be executed by a single person. However, in a worst case scenario, each role is represented by an individual stakeholders which requires collaboration efforts to guarantee a successful decision making process. Due to the stakeholders' different backgrounds and knowledge collaboration is a challenge to be addressed.

3.2.2.3. Tasks

In the related work Section 2.3.3 we discussed several task taxonomies from visual analytics research. In the following, we define an abstract five-level task taxonomy that matches the specific needs of decision making support. Figure 3.6 shows our concept that associates the three decision making steps intelligence, design, and choice with the general decision making tasks: exploration, creation, analysis, comparison, and presentation. In general, each task should be considered during the design of a visual analytics decision support system. This concept can be applied to all data categories presented in our data taxonomy – textual data, empirical data, and model-driven data. The technical contributions of this thesis, presented in Chapter 4 – 9, are based on this concept.

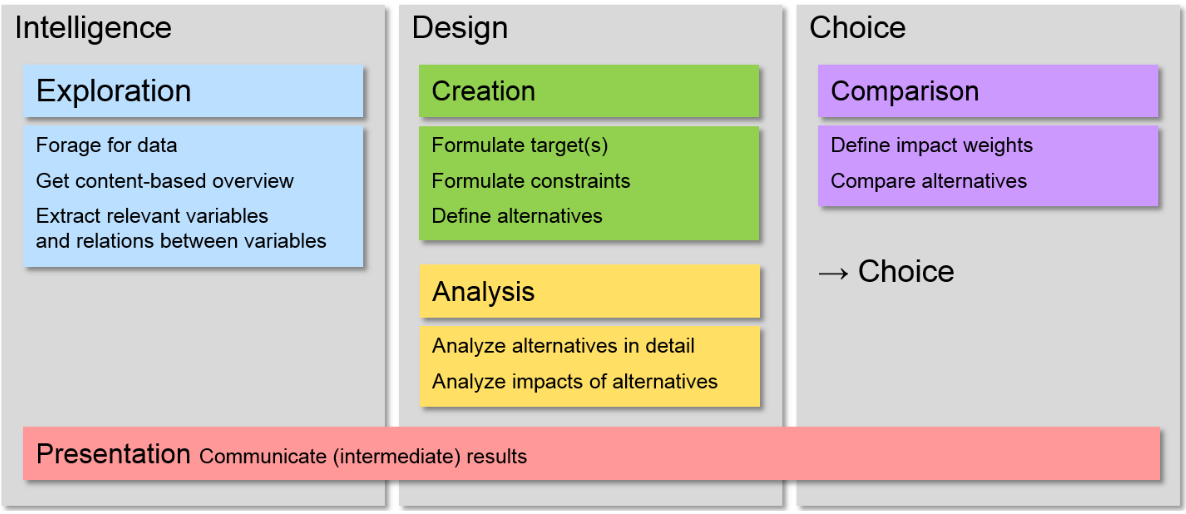

Figure 3.6.: Decision making tasks to be supported with visual analytics. Throughout this thesis, we use this concept for the definition and characterization of decision-related user tasks.

Exploration

We define the exploration task as the undirected discovery of information from data. In contrast to search, the target to be specified by a search query is unknown. Tominski provides a comprehensive definition for visual exploration that also fits the definition of this thesis: "The aim pursued with visual exploration is to give an overview of the data and to allow users to interactively browse through different portions of the data. In this scenario users have no or only vague hypotheses about the data; their aim is to find some. In this sense, visual exploration can be understood as an undirected search for relevant information within the data. To support users in the search process, a high degree of interactivity must be a key feature of visual exploration techniques" [Tom06].

From the perspective of decision making, the exploration task is mainly executed in the intelligence phase. In this phase data relevant for decision making is foraged and knowledge is extracted from data. In the case of textual data, this implies the foraging of relevant text documents. Available text document collections are explored with the goal to get an overview of the targeted domain. Then, relevant documents are separated from irrelevant documents. The relevant documents are aggregated and the underlying content is summarized. In addition, knowledge to support the definition making process is extracted from the text documents. Examples include discussed topics, key facts, arguments, and opinions to support the definition making process. In the case of empirical data, the exploration task supports users to get an overview of available empirical datasets. These datasets are explored to discover relevant variables and dependencies between variables to be considered in the decision making process. The knowledge extracted from empirical datasets allows decision makers to define reasonable decision targets and constraints on potential solutions. Finally, the exploration of model-driven data within the decision making process is two-fold. First, in the different model types are explored in order to judge on their appropriateness for the targeted problem. Second, for a selected model, the exploration of insights related to input-output dependencies is of utmost importance for the creation of alternative solutions.

Creation

The creation task aims at defining alternative solutions to a given problem targeted by the decision making process. This also involves the augmentation of the available data collection with additional data that characterize the solutions. In addition, prior to the definition of a solution the decision targets and constraints are specified, which we also define as creation tasks. The creation task is rarely presented in task taxonomies from visual analytics research. However, we consider it relevant for the decision making process, especially in the design step.

For textual data, the creation task includes the augmentation of the decision process with additional text documents (e.g., scientific reports, newspaper articles, etc.) that inform decision making. Alternative solutions are often documented in written reports that are shared among the involved stakeholders. Rating these documents via online platforms is also defined as creation task. In addition, the sharing of facts, arguments, opinions and judgments on documents, as increasingly executed within social networks and eParticipation platforms, is considered a creation task. In the context of empirical data, the collection of new or the addition of existing datasets to the decision making process are defined as creation tasks. Moreover, we interpret the definition of target variables and constraint variables in empirical datasets as creation tasks. Finally, the extraction of data subsets that describe a potential decision options also fall into the creation task category. In visual analytics research, this task is mostly defined as filter and/or lookup task. For model-driven data, we provide two interpretations of the creation task: First, the specification of a computational model to be included in the decision making process. And second, the creation of model data by running the model with a specified input that generates a depending output.

Analysis

In this thesis, we define the analysis task as the detailed analysis of a created alternative solution to the underlying problem. This includes the drill-down to data relevant for the specific alternative and the evaluation of possible impacts of a given alternative. The analysis task is executed in the design step.

For textual data the analysis task describes the extraction of facts, arguments and opinions towards the defined alternatives and their relevance with respect to quantitative (how many?) and qualitative (by whom?) occurrences. That way, e.g., the social impact of a given solution can be analyzed. For empirical data, this means the drill-down to details and dependencies of a given alternative. In many cases the analysis task can be addressed via zooming and detail-on-demand interaction. In addition, data mining and machine learning approaches support the extraction of patterns from data that might provide more details and allow the assessment of impacts. Finally, computational models are often applied to estimate the impacts of potential solutions. The analysis task supports the detailed analysis of the model output to judge on the appropriateness of the solutions. During this task, the robustness of the model is also considered, e.g., via sensitivity analysis. Allowing users to assess the uncertainty in computational models is of utmost importance for the decision making process.

Comparison

The comparison task is executed in the choice step before the final decision is made. Alternative solutions are compared based on the impacts derived from the preceding analysis task. As an optional step, decision makers might define weightings on the individual impacts that specify their relevance for the decision.

Concerning textual data, the ratings, extracted facts, arguments, and opinions towards the alternative solutions are compared. In the case of empirical data, data subsets defining alternative solutions and their value ranges are compared. Finally, the data input and output of the computational models are compared to choose the optimal solution. In this phase, e.g., cost-benefit calculations help to make the final decision.

Presentation

The presentation of data is a horizontal task in this taxonomy required throughout the entire decision making process. The presentation task is needed for the communication of intermediate results between different stakeholders. As an example, the analyst presents intermediate analysis results to the decision maker to share the current state of the process. The final decision in the choice step is made by the decision maker based on the condensed presentation of the decision options and their impacts. In most cases, the complexity needs to be reduced in order to bridge knowledge gaps between stakeholders. To increase the transparency of decision making, they promote the presentation task being executed throughout the entire decision making process.

With the definition of tasks, we conclude our decision making domain characterization. The general domain characterization serves as a baseline for defining a process for the design of visual analytics decision support systems. Moreover, we apply the domain characterization during the design of visual analytics decision support systems that we present as proof of concepts in chapters 4 – 9.

3.2.3. Visual Analytics Design Process

In the following, we introduce a methodology for the user-centered design and evaluation of visual analytics decision support systems (see Figure 3.7). Our methodology is based on Andrews' concept who uses four distinct design stages to define evaluation goals along the process [And08] and Munzner's 'nested model' [Mun09] about the design and validation of visualization systems. As described in the related work, Munzner attempts to unify the treatment of visualization design and its corresponding evaluation. The nested model is organized into four nested layers (see Figure 2.9). The results of a higher layer form the input for the layers beneath. Errors on the upper layers are, thus, propagated to the layers beneath. Every layer has different requirements on the evaluation. The four layers are: domain problem and data characterization; operation and data type abstraction; visual encoding and interaction design; and algorithm design. Andrews assigns evaluation goals to the visualization design process. He distinguishes between the stages before design, before implementation, during implementation, and after implementation. We re-use Andrews' stages to structure our design process. In the following we propose a design process based on the concept of Andrews. Before the visualization design process, the decision maker and the analyst have to formulate the problem and identify relevant external experts for supporting the decision making process. We propose to include visual analytics experts in the process from the very beginning, since they need to participate in the requirement analysis to understand the domain problem. The following design process stages are characterized with respect to the specific goals to be achieved within each step and exemplary evaluation methods to validate the outcome. We emphasize that the validation methods are examples that, from our point of view, match best to the respective stage. However, as Munzner states, alternative mappings of validation techniques to design stages may, if correctly applied, provide meaningful results [Mun09].

Stage before design: The stage before design involves the gathering of information that sets the stage for the design of the system. The design process starts with the identification and characterization

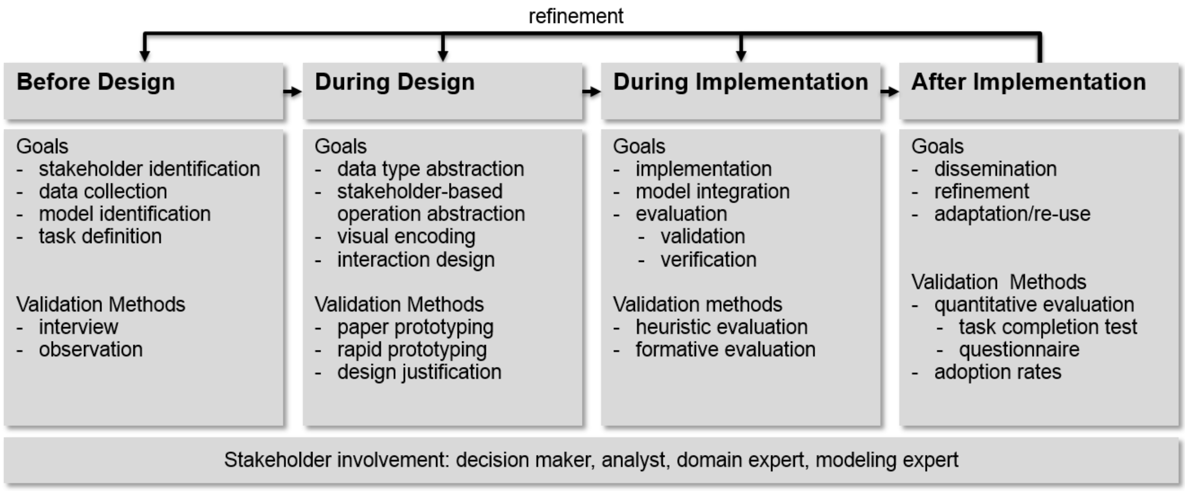

Figure 3.7.: Process for the design of visual analytics decision support systems.

of relevant datasets (including computational models), stakeholders, and tasks. The general decision making domain characterization introduced in the previous section serves as a guideline for the concrete problem characterization to be prepared in the first stage of our design process. As a first step, relevant stakeholders are identified. We recommend to cover the user roles defined in the previous section with real-world persons. This includes the identification of decision makers, analysts, domain experts, and modeling experts. Second, available datasets that support the decision process are identified. Textual document collections and empirical datasets are collected. If required, computational models that support decision making are identified together with modeling experts. Finally, the general tasks of our domain characterization are reviewed and mapped to the concrete problem at hand. Validation methods of choice for the initial stage in the design process include those from qualitative analysis, as described by [IZCC08], which aim at capturing the environment in which the visual analytics decision support system will be used. Munzner states that users should be observed and interviewed [Mun09], which matches well with Andrews stipulation that users should be surveyed in an exploratory evaluation to establish their requirements for the system [And08]. Munzner's domain problem characterization layer is relevant at this stage [Mun09]. In addition, the first four steps of the design study methodology by Sedlmair et al. – learn, winnow, cast, and discover – should be considered [SMM12].

Stage before implementation: At the stage before implementation the findings from the previous stage are used to derive an initial visual analytics design targeting the support of the decision scenario characterized in the previous stage. The initial visual analytics design comprises the visual encodings of the underlying data (how to display the data) and the interaction design (how to interact with the data to support the tasks defined in the previous stage). For the preparation of the design the following steps are required. As a first step, the collected datasets are characterized with respect to the underlying data types. Figure 2.10, presented in the foundations chapter, serves as a baseline for the characterization of data. The identification of data types facilitates the selection of appropriate visual encodings. Moreover, a stakeholder-based operation abstraction is derived. The data type and operation abstractions are also represented in the second layer of Munzner's nested model [Mun09]. By operations, we mean data transformation steps that are (a) automatically executed in the back-end, or (b) executed on-demand via user interaction. At this step, a mapping of operations on user roles is required. The visual analytics expert has to decide which operations are made accessible via interaction to which user role. A complex interaction design improves the analysis functionality while decreasing the usability of the system. The resulting systems are more likely to be used by analysts or modeling experts. On the other hand, a simple interaction design supports an intuitive access to the data while reducing the analysis capabilities of the system. Systems that realize these designs are more appropriate for decision makers and stakeholders. If the visual analytics system is supposed to be used by all user roles, a trade-off solution between the two opposites needs to be defined. Hence, we argue that the user-based operation abstraction is the most critical step in the design process, especially, if the visual analytics solution targets decision support with multiple user roles involved. Based on the data abstraction and operation abstraction the initial visual encodings and the interaction design are developed. These tasks are also represented at the third layer of Munzner's nested model [Mun09]. We promote the application of paper prototyping or rapid prototyping techniques to allow multiple iterations with the involved stakeholders. We suggest to discuss the resulting designs with all stakeholders involved applying focus group techniques [IH01].

Stage During Implementation: At this stage, the designs created in the previous stage are implemented and integrated into a functional system. The system is implemented in an iterative process together with the involved stakeholders. Intermediate versions of the system are shown to the users and improved based on the users' feedback regarding initial and upcoming requirements. The system implementation task includes the integration of data sources, data transformation pipelines, computational models, and visualization techniques. Interfaces between components created from different stakeholders are specified and implemented. For example, the computational model provided by the modeling expert needs to be connected to the visual interface. This allows analysts to access model parameters via the visual interface. Sedlmair et al. recommend the assessment of the system parallel to its implementation; as components are implemented, they should be provided to users to collect direct verbal feedback [SMM12]. We emphasize this suggestion by proposing the execution of several consecutive implementation and evaluation cycles. For example, qualitative heuristics can be used, which enable users to describe their impressions and experiences. In addition, various methods that fall into the category of formative evaluation can be applied. Examples include cognitive walkthrough and thinking aloud techniques. Besides testing the interface with the involved stakeholder, external stakeholders not being involved in the design of the system should be invited for the evaluation. According to Munzner's model, this assessment uncovers, whether or not the wrong abstractions are chosen in the previous stages [Mun09]. In case of inappropriate designs, the initial design is refined by returning to the previous stage. The stage during implementation is a highly collaborative stage. Modeling experts are enabled to use initial visual analytics implementations to validate their computational models. Additional requirements might evolve during their application of the system. Decision makers and analysts may identify aspects not considered in the initial design stage. Objectives or constraints on the underlying problem to be tackled with the visual analytics system may provoke an update. In summary, the design of the visual analytics system already constitutes an initial step in the decision making process. Collaboration is enforced and the formulation of requirements on the system supports the clarification of the underlying decision problem.

Stage After Implementation: At the stage after the implementation the system is deployed, and tested regarding usability and utility within the target environment. It is measured whether the users can solve their tasks with the system, and how easily and intuitively it is adopted by the user. The goal of this stage is to examine to what extent the system is accepted by users and whether it meets the specified requirements. According to [SMM12] this would be the reflection phase, in which the insights gained in the previous stages are summarized. In Munzner's model, this stage is situated back on the upper layer – the domain and problem characterization [Mun09]. Following Andrews, questionnaires and formal experiments with real-world users should be applied to validate the overall approach. Additionally, guideline scoring techniques are proposed [And08]. From our experience, we suggest to conduct a user testing session including task completion tests and a usability questionnaire to be filled-in afterwards. We learned that users are able to provide more detailed and constructive feedback in a questionnaire after they have executed specific tasks with the system. Depending on the openness of the decision making scenario, the user evaluation can support the dissemination of the system to external users. We identify three core benefits in sharing the system with the public. First, a broad range of users is reached which increases the amount of feedback required for improving the system. Second, sharing the decision support system allows public users to create alternative scenarios that might even provide better results than those created by a smaller set of experts. Third, involving the public in the decision making process increases (a) the acceptance of decisions through transparency, and (b) the trust in the decision makers in general.

With the finalization of the design process the decision making process starts. However, since all stakeholders have been involved in the design process, some aspects of the decision making process are already covered. Examples include the identification of data relevant for the decision, the definition of abstract targets and constraints to be defined within the design phase, and the analysis of potential solutions to the addressed problem that might have been identified during the evaluation of the system.

3.3. Bridging Knowledge Gaps in Decision Making with Visual Analytics

The research challenges targeted in this thesis are (1) the extraction of knowledge from data with visual analytics to support the decision making process, and (2) the bridging of gaps between involved stakeholders. The first challenge has been addressed in the previous section. In this section, we present a concept on how visual analytics can support the knowledge transfer between stakeholders in the decision making process. The concept is strongly related to the ideas of collaborative visualization as presented by Isenberg et al. among others [IES*11]. However, so far collaborative visualization has been addressed from the temporal (synchronous vs. asynchronous work) and the spatial (co-located vs. distributed work) perspective. Varying user expertises and roles, as present in the decision making process, have rarely been considered. In the following, we first describe an adapted decision making process (see Figure 3.8) and an adapted policy cycle (see Figure 3.9) that support the bridging of knowledge gaps between involved stakeholders by the extensive use of visual analytics technology. Second, we describe how the complexity of decision support systems can be reduced by applying different visualization disciplines. And finally, we discuss synergy effects by incorporating visual analytics into the decision making process.

3.3.1. Bridging Knowledge Gaps in Organizational Decision Making

Our first concept extends the ‘classical’ decision making process with visual analytics technology (see Figure 3.8). The concept addresses the bridging of knowledge gaps in model-driven data analysis scenarios. We select a model-driven scenario, since it involves a large variety of user roles: decision makers, analysts, and modeling experts. However, we claim that the concept can be adapted to textual data (with text analysis methods as models) and empirical data (with data mining techniques as models). In model-driven scenarios, a computational model is applied in the decision making process to create and evaluate alternative solutions to a given problem. However, the models’ complexity impede the consideration of the created knowledge in the decision making process. Hence, we suggest connecting visualization techniques to computational models developed by modeling experts to improve their usability for non-experts. That way, the complexity of the models is hidden in the computational back-end, while only the information necessary for providing user input (e.g., control parameters, etc.)

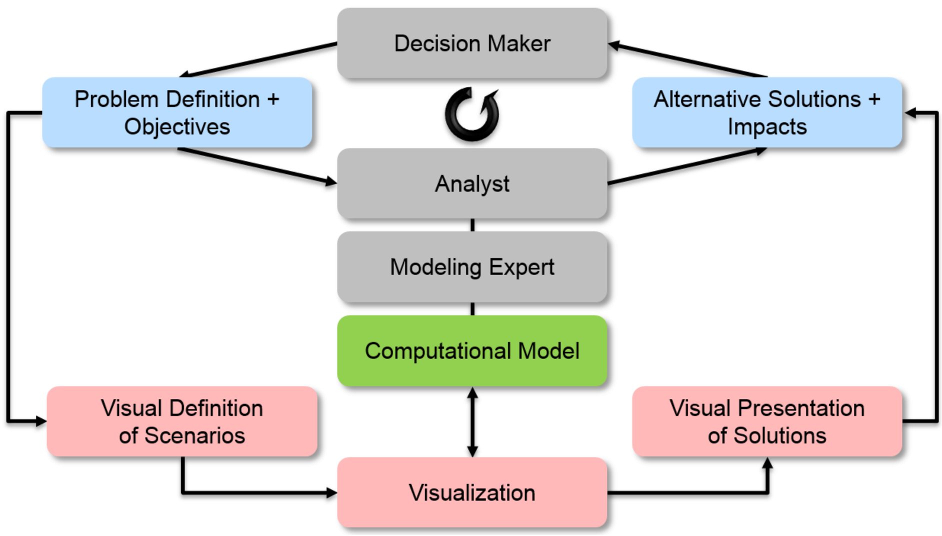

Figure 3.8.: A concept to bridge knowledge gaps between stakeholders with visual analytics. The classical communication workflow from the decision maker over the analyst to the modeling expert and back is bridged with visualization techniques.

and analyzing the model output (e.g., simulation results, statistical measures, etc.) is displayed on the screen. The most crucial aspect of this concept is that users not familiar with the computational modeling can visually interact with these models and conduct analysis scenarios by themselves. Hence, analysts and decision makers can specify objectives and constraints via a visual interface (see Figure 3.8, on the left). They can 'experiment' with different settings, and generate alternative model outputs (see Figure 3.8, on the right). Furthermore, due to a uniform visual representation of the model data, stakeholders with different backgrounds can validate the models' utility and usability. For example, the decision maker can detect aspects not covered yet by the model that the modeling expert might include in an improved model. The communication of results is facilitated, since all stakeholders work with the same visual representation.

3.3.2. Bridging Knowledge Gaps in Policy Making

The model presented in the previous section can be easily transferred to political decision making (or policy making). Nevertheless, we introduce a second model that illustrates the idea of our approach. Therefore, we expand the five-stage policy cycle by Anderson [And75] (see Figure 3.9 a) at the policy formulation and the policy adoption stages. These stages imply the definition of policy options, their analysis, and finally the decision which one to choose (see Figure 3.9 b). Several stakeholders with different expertise are involved in these stages which result in knowledge gaps to overcome.

At the 'agenda setting and problem definition' stage public problems that shape the political agenda are identified by political decision makers (policy makers). At the 'policy formulation' stage, the problem is analyzed (policy analysis) and potential solutions to the problems are defined. The main stakeholders involved at this stage are the (policy) analyst and the modeling expert. The analyst conceptualizes the problem identified in the previous stage and consults modeling experts as external adviser. Based on the requirements specified by the analyst, the modeling experts design computational models that support the creation and analysis of policy options. The model results are communicated to the analyst who uses the extracted knowledge to define policy options. In the 'policy adoption' stage the defined policy options are communicated in a condensed way to the decision maker who decides which option to choose. The model includes two feedback loops. First, the analyst might define additional requirements based on the modeling results communicated by the modeling expert. Second, the decision makers might request additional policy options from the policy analyst by refining the problem description. The policy cycle is completed with the policy implementation and the policy evaluation stages that are not considered in this model.

Similar to the previous model, this process can be augmented with visual analytics concepts. By connecting visualization techniques to the computational model, users can visual-interactively access the model parameters and the model output (see Figure 3.9 c). In the following, we discuss how different visualization disciplines are applied to support different analysis tasks.

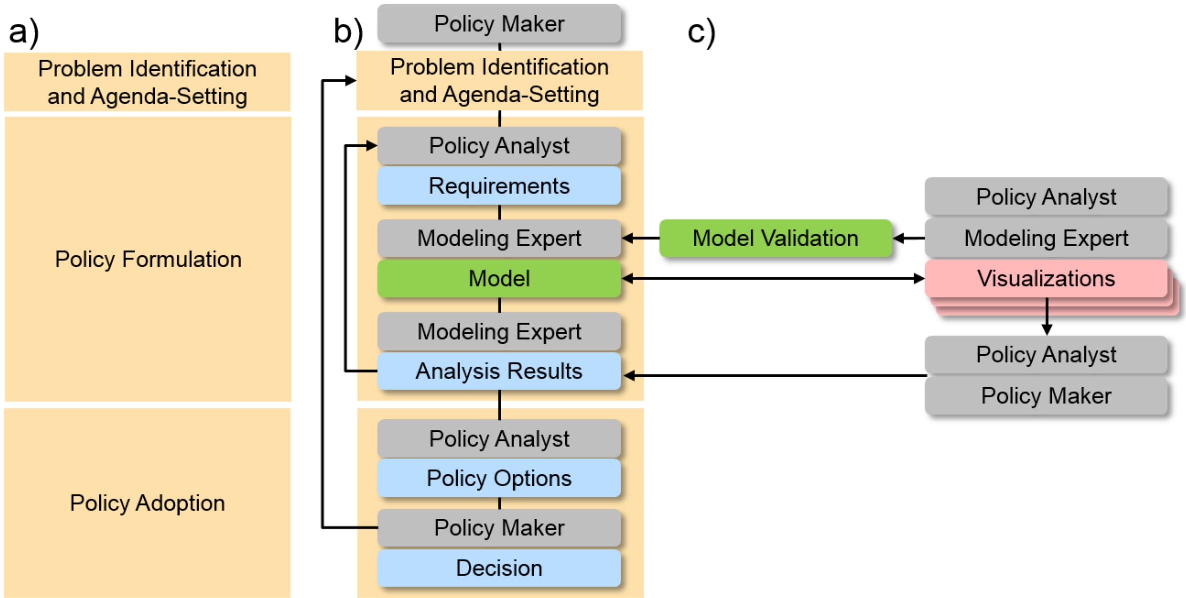

Figure 3.9.: Adaptation of policy cycle. a) Standard policy cycle [And75]. b) Adaptation of policy cycle, providing more detail in the policy formulation and policy adoption stages, and two feedback loops. c) Linking visualization to the model to bridge knowledge gaps.

3.3.3. Complexity Reduction via Appropriate Visualization Disciplines

In the previous sections, we introduced two concepts on how to include visual analytics in business and political decision making processes to bridge knowledge gaps between stakeholders. The complexity of the analysis process can be further reduced by applying appropriate visualization disciplines in the decision making process. Depending on the user and task at hand, visualization disciplines of varying complexity can be applied.

We differentiate between two main usage scenarios of visualization techniques: (1) the visual-interactive exploration and sense-making as analysis tasks and (2) the visual communication as knowledge transfer task (cf. Figure 2.4 by Few [Few09]). In the first case, the users require visual-interactive control of the computation model. Users are enabled to define input parameters and run the model to create alternative output data. The interactive control of the model is realized by using concepts from the field of visual analytics, which are coined by Bertini and Lalanne as "white-box-integration" [BL10]. The visual analytics concept connecting visualization with computational models is already applied to different application domains. We propose to introduce this concept to the decision making domain. In the second case, users require visual access to the analysis results generated with the model. The output data is presented with information visualization techniques. Users can visually explore the results of the model by search and filter operations. The information may be visualized with different visualization techniques depending on the users' expertise and knowledge. The main focus of these techniques lies on the usability of the system. The intuitive usage of the visualization has to be ensured.

In the following, we describe how the appropriate selection of visualization techniques can further support the bridging of knowledge gaps between stakeholders.

Visualization for Decision Makers (and Public stakeholders)

The visualization design for the decision maker consists of easy-to-understand interfaces only depicting the information relevant for the decision process. It enables the decision maker to get quick access to analysis results. This interface bridges knowledge gap between decision maker and the modeling expert (see competence gap). Moreover, the decision maker can give high-level feedback to the modeling expert if some information is missing, or if the model needs to be refined (see iteration gap). As another 'gap bridger', the translation of analysis results to decision options can be derived by the visualization. This bridges the knowledge gap between analyst and decision maker (see analysis and competence gap). The considered visualization techniques are mainly static visualization (or infographics), and easy-to-use information visualization techniques. As an example, the ManyEyes system enables a user friendly access to visualization techniques with the option for the users to upload and visualize their own datasets [VWvH*07].

Visualization for Analysts

The visualization design for the analyst consists of the basis functionality provided to the decision maker, and advanced interaction techniques, that offer a closer connection to the model. This interface bridges the knowledge gap between the analyst and the modeling expert (competence gap). The analyst is enabled to validate the model from the domain perspective, and refine it, e.g., by changing input parameters (iteration gap and analysis gap). The communication between the analyst and the modeling expert is supported, since both can work with the same information representation. Again, the access to the complex model is facilitated. This enables the analyst to interact with the model, gain an understanding of the model, and finally produce analysis results without the help of the modeling expert (see analysis gap). The considered visualization techniques come from both fields, information visualization and visual analytics. Examples for visual analytics systems designed to support analysts can be found in the fishery policy domain [BMPM12], and the energy domain [Hea12].

Visualization for Modeling Experts

The visualization design for the modeling expert comprises the highest functionality. Depending on the requirements of the modeling expert, a visual-interactive editing of the model may be realized. Visualizing the model input and output supports the modeling expert in refining the model and validating its functionality (see iteration gap). That way, new analysis results are produced and communicated to the analyst and the decision maker via their respective visualization design (see analysis gap, and competence gap). This bridges the knowledge gap between modeling experts, analysts, and decision makers, who collaboratively refine the functionality of the model, and validate the correctness of the model (see iteration gap). The considered visualization techniques for the modeling expert are mainly from the field of visual analytics. In [MK08], a visual system for the data-driven verification of hypothesis is provided. In [IMI*10], an interactive data analysis process is supported with visualization techniques.

3.3.4. Synergy Effects of Applying Visual Analytics to Decision Making

In order to address the challenges imposed on decision making and policy making, we proposed the incorporation of visual analytics into decision making (see Figure 3.8) and policy making (see Figure 3.9) processes. Hereby, visual analytics serves as an important component of decision support systems itself. In the following, we summarize the benefits resulting from this integration:

Communication. The communication between relevant fields, e.g., science and policy making, is facilitated. Visualization may serve as a mediator of information between two distinct environments. Through the unified visual presentation of information, different stakeholders are enabled to discuss at the same knowledge level. Thereby, the communication between scientists and decision makers in the decision making process is supported.

Complexity. Through the abstraction of user tasks and the design of visual analytics systems adapted to the expertise of the targeted users the complexity of the underlying models can be hidden. With visual analytics complex data operations can be executed on the machine, while the parameters to control their execution can be intuitively displayed on the screen. Visual interfaces provide the information on the level of detail needed by the respective user role.

Subjectivity. The aspect of subjectivity can be reduced since different stakeholders get access to the same information provided in an ‘objective’ way via visual analytics techniques. The provided information can be discussed among the stakeholders to balance subjective interpretations of the findings.

Validation. The outcomes of the decision making process can be transparently presented to all involved stakeholders including public stakeholders. That way decisions can be justified since they have been made based on an objective analysis. This can improve the trust in scientific results, and political decision making.

Transparency and reproducibility of results. If open access to the visual analytics system is provided, public stakeholders (e.g., journalists, interest groups, etc.) are enabled to generate analysis results on their own, and therefore, better understand the rational background of strategical decisions.

3.4. Outlook on Technical Contributions of this Thesis

In the remainder of this chapter, we provide an outlook on the technical contributions presented in the following chapters. These chapters address the six technical challenges identified in Section 3.1. They serve as proof of concepts for the applicability of our design methodology presented in Section 3.2. We categorize the underlying decision scenarios according to our domain characterization. Three of them target textual data, one addresses empirical data, and the last two tackle model-driven data. We also discuss which general decision making tasks are supported by the respective visual analytics system and which stakeholder have been involved in the design process.

Visual-Interactive Access to Decision Making Processes

In Chapter 4, we address the process overview challenge $\mathbf{C}_{\text{Proc}}$, presenting a visual analytics system for the creation, exploration, and analysis of decision-related content. The visual analytics system organizes relevant text documents in a timeline visualization that we named 'PolicyLine'. Documents can be distinguished based on their publication date, title, author category, and relevance to the decision process. Users are enabled to create new processes including the definition of process step, or augment existing processes with additional documents, document ratings, and comments on documents. Automatic analysis methods calculate the likelihood of alternative solutions (represented by specific documents) to succeed. In the design process, decision makers, analysts, and domain experts were involved. The contributions of this chapter are: (1) a problem characterization about providing decision making process overviews, (2) a visual analytics system for the text document-driven overview of the decision process, and (3) a design study including three cycles of design, implementation, and evaluation conducted with eighteen real-world users.

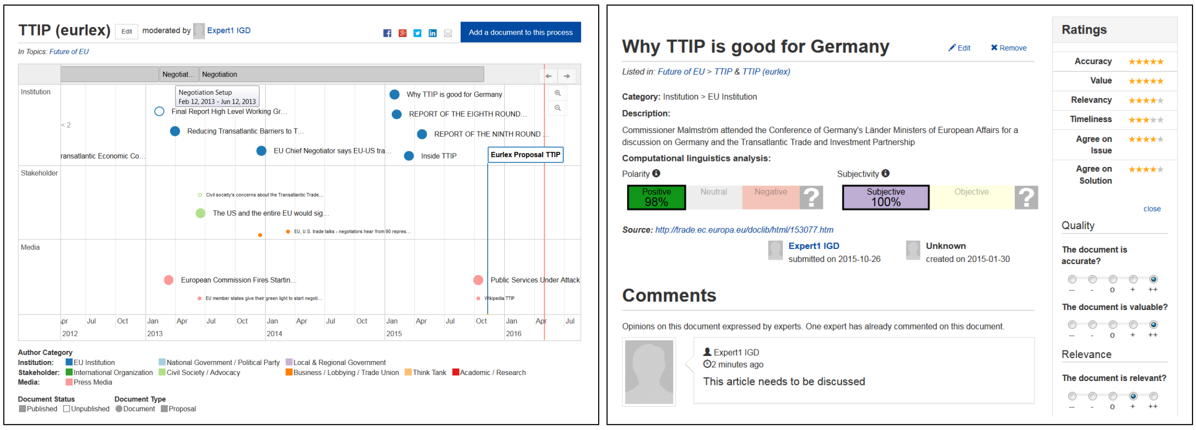

Figure 3.10.: PolicyLine: visualizing the decision process. Presented in Chapter 4.

Visual-Interactive Access to Text Document Collections

In Chapter 5, we demonstrate a visual analytics system for providing content-based overviews of text document collections and solve challenge $\mathbf{C}_{\text{Doc}}$. Text clustering as a method of choice is applied to condense large numbers of documents to a small set of representatives. The visual analytics system provides visual and interactive access to the entire text clustering process from the pre-processing of unstructured text documents, the selection of features for representing documents, and the parameterization and computation of text clusterings to the comparison of different clustering results and their validation. Due to the complexity of text clustering in general, and the provided degrees of freedom in the interaction design of our visual analytics approach, the system is targeting analyst as primary users. Nevertheless, the results generated with the system can be used to provide content-based overviews to decision makers and other stakeholders in the process. The contributions of this chapter are: (1) a visual and interactive interface for the selection of a feature vector representation required for the clustering text documents, (2) a visual and interactive interface for the creation and analysis of text clusterings adapted to the specific data, users, and task at hand, and (3) a visual and interactive interface for the comparison of different text clusterings that allows users to choose the most appropriate clustering result for the specific data, users, and task at hand.

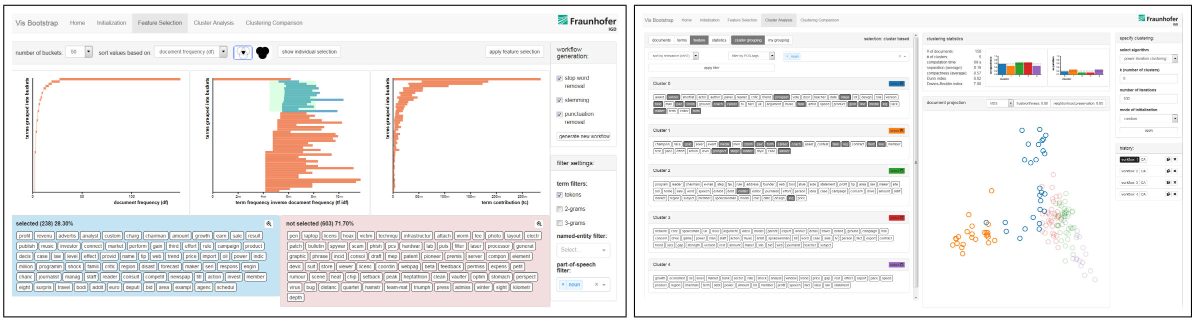

Figure 3.11.: Content-based exploration and analysis of text document collections. Presented in Chapter 5.

Visual-Interactive Access to Online Debates

In Chapter 6, we present a visual analytics system for aggregating textual content extracted from public debates in social media channels, which addresses challenge $\mathbf{C}_{\text{Deb}}$. The system provides visual and interactive access to text segments extracted from the web and mapped to pre-defined policy models and arguments. In addition, the sentiment of text segments is extracted and shown to the user. The system allows users to monitor the public debate on specific topics and identify new arguments that may inform the decision process. In addition, the user is enabled to refine the underlying model by providing feedback on the accuracy of the automatically generated text analysis results. During the design process decision makers, domain experts, analysts, and modeling experts from the text analysis domain were involved. The contributions of this chapter are: (1) the definition of a text analysis workflow adapted to the specific needs for monitoring the public debate, (2) a visual analytics system that provides visual and interactive access to the results of the text analysis workflow, and (3) a user feedback concept that allows users to improve the accuracy of the text analysis models via user interactions.

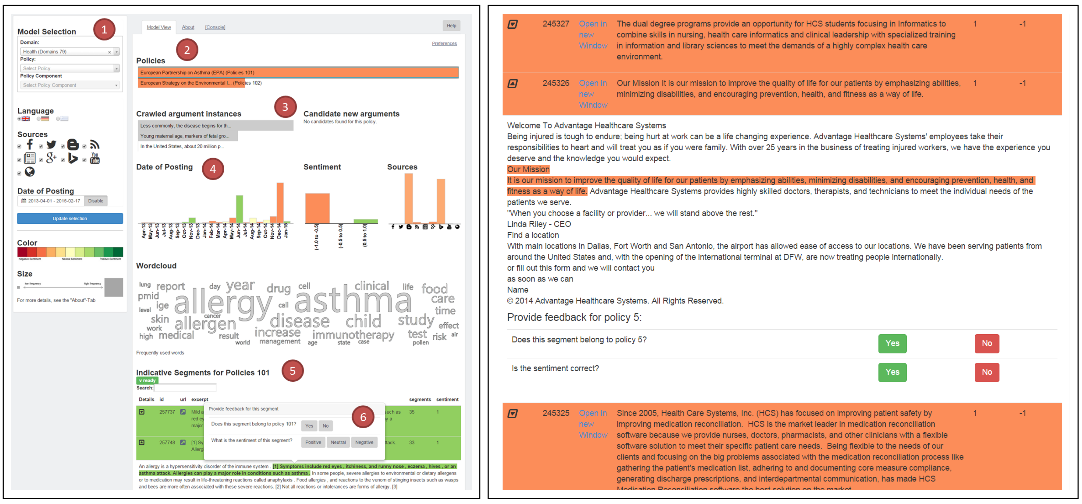

Figure 3.12.: Analysis of Arguments and Opinions in Text Documents. Presented in Chapter 6.

Visual-Interactive Access to Empirical Datasets

In Chapter 7, we address the challenge about the exploration, analysis, and comparison of empirical data $\mathbf{C}_{\text{Dat}}$. We introduce a visual and interactive system for accessing country-specific performance indicators in the mining sector. The system is designed for governmental and business (investors) decision makers as well as for public stakeholders. Domain experts provided knowledge on the governmental and legal aspects related to mining in resource-rich countries. This knowledge was transformed to qualitative performance indicators. Our visual analytics system allows stakeholders to get an overview of individual country performances, identify weaknesses and strengths of individual countries, compare the performances of multiple countries, and search for similar countries. As overarching goals, investors are supported in their investment decisions, governments are supported in their policy decisions to improve the current status and attract investments, and public stakeholders are supported in their understanding of political decisions. The contributions of this chapter are (1) a domain characterization of the mining sector as a specific decision making domain, (2) a visual analytics system for providing visual and interactive access to performance indicators, and (3) evaluation results gathered in a user workshop with domain experts from the mining sector.

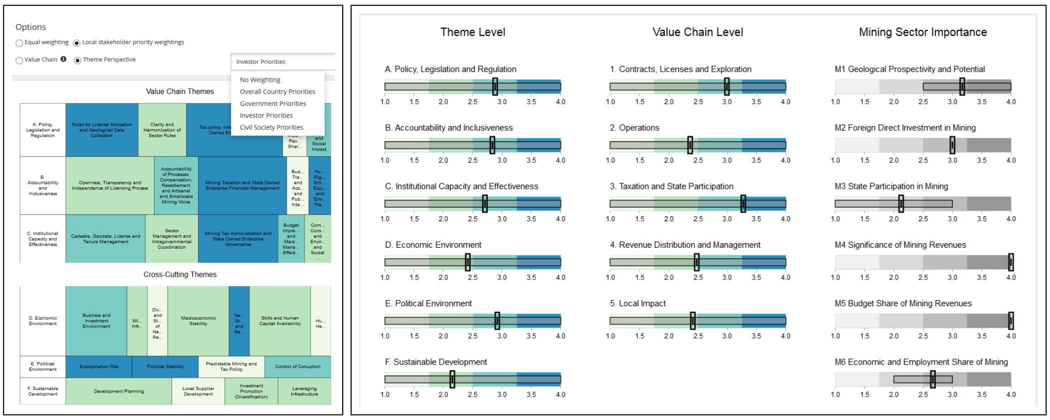

Figure 3.13.: Exploration and Comparison of Mining-Related Empirical Data. Presented in Chapter 7.

Visual-Interactive Access to Simulation Models

In Chapter 8, we combine simulation techniques with visualization methods to create a visual analytics system supporting the analysis of impacts derived from alternative solutions. This chapter addresses Challenge $\mathbf{C}_{\text{Imp}}$ on the impact assessment of decision options. The visual analytics system provides visual-interactive access to an agent-based simulation model to support political decision makers in the exploration, creation, analysis, and comparison of simulation data. With the agent-based simulation model the impact of alternative governmental subsidy strategies on the public adoption of photovoltaic plants is assessed. Our visual analytics system was designed in collaboration with decision makers, analysts, domain experts in the energy sector, and modeling experts in agent-based simulation. The contributions of the chapter are (1) a problem characterization of simulation applied to policy making in the energy sector, (2) a visual-interactive interface for the impact assessment of alternative solutions via an agent-based simulation model, (3) a novel visualization technique specialized on the exploration of dependencies between input and output variables.

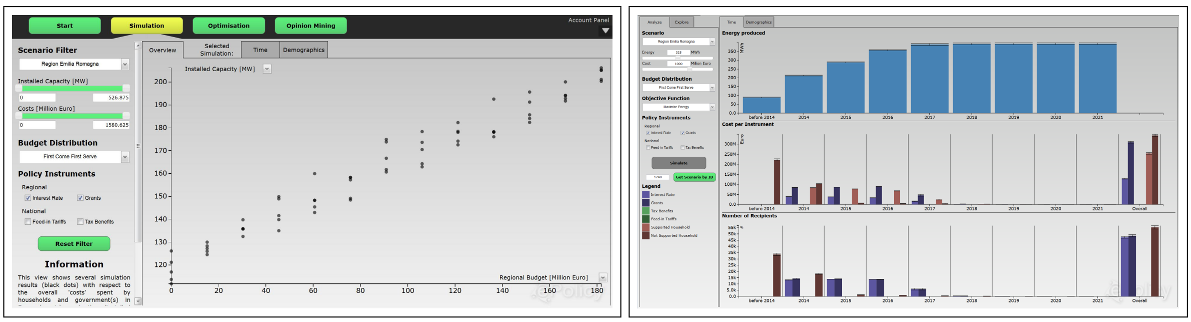

Figure 3.14.: Simulating the Impact of Decision Options. Presented in Chapter 8.

Visual-Interactive Access to Optimization Models

Finally, in Chapter 9, we combine visualization methods with optimization techniques to create a visual analytics system that supports the balancing of trade-offs, and in consequence, the generation of optimal solutions for a given problem. This chapter addresses Challenge $\mathbf{C}_{\mathrm{Opt}}$ on the creation, analysis, and comparison of optimal solutions. The optimization model enables analysts and decision makers to define an optimal energy plan based on a pre-defined target function and constraints considering environmental and economic factors. The visual analytics system was designed and evaluated in collaboration with political decision makers, analysts, domain experts, and optimization modeling experts. The contributions of the chapter are (1) a problem characterization on the strategic environmental assessment (SEA) domain targeting the development of sustainable policy decisions, (2) a visual analytics system combining visualization and optimization techniques to support the generation of optimal regional energy plan, and (3) the results of two user evaluation rounds conducted with two intermediate versions of the system.

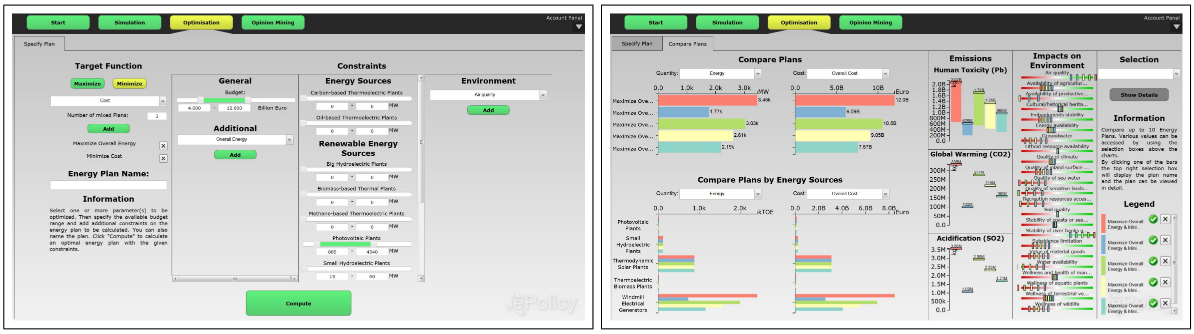

Figure 3.15.: Balancing Trade-Offs with Optimization Models. Presented in Chapter 9.

4. Visual-Interactive Access to the Decision Making Process

The visual analytics approach presented in this chapter provides visual access to the entire decision making process. It supports the exploration of the decision process, the creation of opinions and knowledge to be included in the process, the analysis of text documents relevant for the decision, and the comparison of the stakeholders' support of a textual proposal. Regarding the overall approach of this thesis, we prove the applicability of our concept to text-driven data (Challenge $\mathbf{C}_{\mathrm{Proc}}$). Figure 4.1 shows how the design methodology presented in Chapter 3 is applied in this specific scenario with text-driven data (Challenge $\mathbf{C}_{\mathrm{VDSS}}$). First, the presented approach allows users to explore the entire

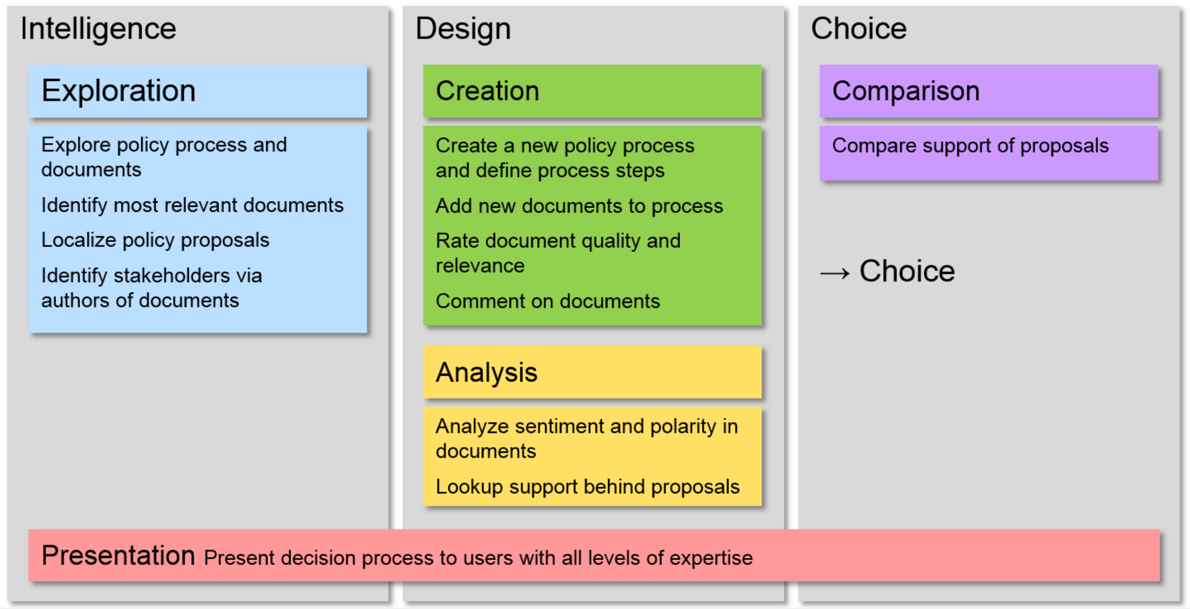

Figure 4.1.: Proof of concept for the applicability of the design methodology on **textual data**. Relevant text documents in the decision process can be explored, analyzed, and compared. Alternative solutions can be created and added to the process as proposal documents. Additional analysis and comparison tasks are supported. The intuitiveness of the visual analytics approach allows stakeholders with all levels of expertise monitor the status quo and participate in the process. Figure is adapted from Figure 3.6.

decision process, which is represented by text documents that inform or record intermediate results of the decision process. The documents are structured along a timeline, which enables uses to monitor the progress and the status quo of a decision. The documents are distinguishable via their authorship and their relevance. Moreover, specific policy proposals that serve as an alternative option to the given problem are highlighted. Second, users of the system are enabled to augment the decision process with additional information. They can create new processes including the definition of process steps. Moreover, they can augment existing processes with additional documents that bear information to be considered in the decision making. Furthermore, they can rate the document quality and relevance, and add comments to the documents to discuss their value. Third, the system automatically extracts the documents' polarity and objectiveness, which supports users in the analysis of individual documents. Fourth, based on the document ratings, users can compare the support behind the different proposals and identify the proposal most likely to succeed. Finally, the visual analytics system is designed and evaluated with a focus on its usability. This makes it usable for users with all levels of expertise. In general, the system is targeting decision makers, analysts, domain experts, modeling experts, and general stakeholders. Any of these roles can participate in augmenting the process with knowledge. Our approach bridges knowledge gaps between these stakeholders by providing a transparent view on the current status of the process and by realizing participation (Challenge $\mathbf{C}_{\mathrm{BKG}}$). This chapter is partially based on our previous work published in [RBB*16].

Contents

4.1. Introduction

The public policy process describes the cycle from the identification of a societal problem to the implementation and evaluation of a public policy. Different factors contribute to the complexity of a policy process [Mac04b] [RBK13]. First, a policy needs to pass different steps within the cycle. Second, the processes vary in their duration, some processes may last for years. Third, a policy proposal within the process may run through several iterations. Finally, a large number of stakeholders with differing expertise are involved in the process. Policy makers have to decide which policy will be finally implemented. Policy analysts and policy advisers support policy makers in making decisions. Scientific experts provide knowledge about the field of discourse. Journalists report political discussions to the public and make the policy process transparent. Lobbyists, such as environmental or social activists, try to influence the policy process. The policy process can be structured along text documents that record intermediate results and discussions within the process. In addition to official documents provided by public institutions, text documents like scientific reports, newspaper articles, and economic white papers need to be considered. Furthermore, social media sources (like Twitter, LinkedIn, etc.) are playing an increasing role in public policy making. All stakeholders involved in the policy process want to get a comprehensive overview of these documents in an efficient way due to temporal pressure.

In this design study, we worked together with stakeholders from the political EU environment. They reported on problems that impede efficient and effective policy making. Since no standardized process steps exist, the identification of the current status is complicated. Due to long process durations stakeholders need to constantly update their knowledge about the policy topic. The sheer amount and variety of relevant documents hampers stakeholders to condense the information and get the full picture. This also has an impact on the required amount of iterations until the implementation of a proposal. A further challenge remains in achieving a critical mass of expertise to be included in decision-making. Relevant experts in the field need to be identified and integrated into the process. Finally, some steps of the process are not transparent for many participants (such as the public) which decreases the trust in policy decisions. Information visualization can help solving these problems [KNRB12]. However, to the best of our knowledge the public policy process has hardly been subject to information visualization research. Based on the problems of our collaborators, we propose three core visions that can improve political decision making.

First, providing *visual access* to policy processes and associated documents can help stakeholders to deal with the complexity of policy making. As a benefit, stakeholders can intuitively access all relevant documents at once, easily identify the newest or most relevant ones, and monitor the current state of the process. Moreover, visual access to the policy process can help external stakeholders to comprehend political decisions and increase transparency.

Second, a policy process is highly collaborative. Thus, providing meaningful *interactive functionality* supports stakeholders to actively participate in the process. Users are enabled to manually contribute content to the system by creating new policy processes, adding documents, rating documents, and commenting them. Following a collaborative approach can further enrich the political debate and result in well-informed decision making.

Finally, a collaborative visualization system can benefit from *automatic algorithms* like crawlers, text analysis methods, etc. that augment the process with additional content. Related articles, blogs, and social media streams can be associated to policy processes and enrich the information density. Automatic algorithms can also help to handle the quantity of the acquired data via aggregation, content summaries, etc. A combined human-machine approach would further increase the quality of political decision making.

To support the public policy domain in addressing this vision, we conducted a design study together with experts from the public policy making field. Our contributions are:

1. We provide a *characterization of the public policy domain* and describe problems of involved stakeholders. Our domain characterization includes a comprehensive analysis of users, data, and tasks. As a result, we define requirements on a visualization system supporting the public policy process starting at the public debate.
2. We present our visualization system *PolicyLine* that provides visual access to policy processes. Users are enabled to gain an overview of existing policy processes and associated documents. Moreover, it enables users to generate content by creating new policy processes, attaching additional documents, rating documents, and providing comments. Automatic algorithms support the enrichment of policy processes via content crawling, and sentiment analysis.
3. We report on our *design study* comprising three iterations of design, implementation, and evaluation with eighteen real-world users from the public policy domain. We prove the applicability of our approach for public policy making, and reflect on lessons learned.

4.2. Related Work on Time-Oriented Text Document Overviews

Information visualization approaches in the policy or law-making domain mainly focus on the textual analysis of individual documents [CM12, ARDM11]. The main goal of our approach is to provide an overview of several documents associated to a policy process. We review related work with an emphasis on how to allocate the display axes.

Document Overviews. In information visualization a series of approaches exists reflecting structures between documents in the 2D space. A review of overview techniques for text documents is provided in [HK12]. Some approaches present both document overviews and relations between documents, e.g., in a reference map visualization [NB12]. We share the idea of presenting structures and relations between documents. However, these approaches do not reflect temporal relations, which is of utmost importance for our domain experts. Therefore, one of the most essential design decision was to assign temporal information to the horizontal axis. In the following, we survey related techniques explicitly incorporating sequential or temporal relations in combination with document overviews.

Temporal, Event, and Process Visualizations. The book of Aigner et al. provides a profound survey of visualization techniques for time-oriented data in general [AMST11]. More specifically, in the context of textual documents a prominent approach is to visualize topic evolutions over time, example techniques are ThemeRiver [HHWN02] or Tiara [WLS*10]. Our approach differs by using single documents as preferred level of granularity. Single documents in the process are arranged in their order of occurrence, similar to the PlanningLines approach [AMTB05]. However, the time primitives visualized in PlanningLines are intervals, while our approach represents documents as a series of instants. In this regard, our approach borrows features from process and event-based visualizations, such as LifeLines2 [WPS*09] or LifeFlow [WGGP*11]. Rind et al. [RWA*13] compare fourteen visualization techniques with an emphasis on discrete event data, many of them applied for the visualization of electronic health records. Furthermore, information visualization research contributes to various other domains using event-based and process visualization, such as energy distribution networks [RHF05] or web log analysis [LRTM07]. In many cases node-link diagrams or techniques using metaphors similar to Sankey diagrams are used to visualize processes. We share the idea to encode temporal information along the x-axis. In the following, we complement the approach with a review on strategies using the y-axis.

Facets and Categories. The vertical axis of our document overview is structured by different types of authors contributing documents. The idea of encoding a categorical attribute vertically is inspired by techniques using facets and categories for document layouts. An overview of facets and faceted search interfaces is provided in [Hea09]. Facets consist of different concepts of a particular category which are typically represented in list-based or hierarchical structures. Since the authors are the predominant attribute for structuring the document overview, we adopt the idea for the vertical axis. In addition, we use color to indicate sub-categories. This is inspired by the FacetMap tool where hierarchical structures are represented with nested facet structures [SCM*06]. With the horizontal time axis and a y-axis representing categories of authors, we implement a best-practice concept from a pioneer approach in the digital library domain [FHN*93].

4.3. Visual Analytics Design – The PolicyLine Approach

4.3.1. Policy Process Background

Our approach attempts to provide visual-interactive access to policy processes that are starting from public debates, continued in legislative procedures, and concluded as implementations and evaluations of public policies. A reference model for the policy cycle divides the process into five core steps: problem identification & agenda setting, policy formulation, policy adoption, policy implementation, and policy evaluation [Mac04a]. First, information about the given problem is collected and analyzed. Second, alternative solutions (or: policy proposals) are designed. These are compared or refined until they are adopted in the third step. After the official implementation of a policy in the fourth step, the policy impacts are monitored and evaluated in the fifth step. Without calling the policy cycle into question, in practice many processes vary with respect to detail and order. Therefore, each process needs to be handled differently requiring expert knowledge and careful documentation.

To record the progress of individual steps, *documents* are created and maintained within the entire policy process. For example, official documents accompanying legislative policy processes in the EU environment are stored at the EUR-Lex repository [Eur16]. Stakeholders involved in the process need to access these documents, interact with them, and possibly extend them on demand. However, most repositories lack non-legislative documents that shape the discussions in earlier steps of the process. As an example, scientific reports providing evidence on policy decisions are rarely considered. Moreover, our domain experts report that documents from new media types would supplement the political decision making process, if the content could be exploited in an effective way. To the best of our knowledge, no solution exists that unifies the broad range of textual information which finally influences policy process outcomes.

To address the described problems, we designed a visual-interactive interface that takes the goals of different *user* groups into account. In our approach we categorize the main stakeholders involved in policy processes into three groups The first group includes *institutional* stakeholders directly involved in decision making, e.g., government members. The second group combines *media* stakeholders, e.g., journalists from different online and offline media. Finally, *general* stakeholders define the third group. Examples are international organizations, business or lobbying organizations, political think tanks, or civil society.

4.3.2. Data and Task Abstraction

In our design process we built on well-received approaches describing the main ingredients for designing and evaluating information visualization systems, i.e., *users*, *data*, and *tasks* [vW13,Mun09]. After describing the user domain and characterizing their problems in the previous section, we now focus on the data and task abstraction. We compiled these abstractions in an iterative process involving workshops conducted together with real-world experts from the political EU environment. The workshops were organized as informal discussions including interviews and questionnaires in which the domain experts were invited to state their problems and challenges with respect to knowledge acquisition during ongoing policy processes. From these discussions we derived the characterization of data and tasks.

Data. The data relevant for our approach can be structured into a hierarchy with *policy topics* on the top level. These topics reflect thematic organizations, e.g., along ministries. The policy topics group associated *policy processes*, which are the core entities of our approach. The temporal expansion of a policy process is reflected by subdividing it into process steps. These steps are essential to describe the status of an ongoing political decision making process. Furthermore, a policy process is shaped by *policy documents*. In our approach any textual content relevant for a policy process can be defined as a document. Examples are official documents from the EU, tweets about public debates, online articles, etc. Table 4.1 provides an overview on the most relevant attributes of a document. Most noticeable metadata attributes are title, author, description, etc., as well as categorizations of authors that we extracted with our domain experts. In addition, official proposal documents published by legal entities are tagged.

Tasks. At a glance, the domain characterization revealed that domain experts focus on two primary goals; the *analysis* of existing content and the *creation* of new content. In the following, we further subdivide these primary goals leading to a characterization of concrete user tasks (see Table 4.2).

First, the visualization system should support the visual access to policy processes and associated content which we refer to as *analysis tasks*. Users have to explore policy processes, identify most relevant documents and proposals, explore the authorship and the support behind a document, and lookup the classified sentiment.

doc. attribute	description
title	title of document
author	person or organization
description	textual content of document
category	institution / stakeholder / media
subcategory	**institutional**: EU Institution, National Government/Political Party, Local/Regional Government
	stakeholder: International Organization, Civil Society/Advocacy, Think Tank, Business/Lobbying/Trade Union
	media: Press Media
publishing date	time stamp
proposal	proposal classification (yes or no)

Table 4.1.: Document attributes

Second, the system has to support collaborative decision making by enabling users to contribute content about policy processes (*content creation tasks*). Users need to create new processes, define process steps, associate documents to the process, rate the quality and relevance of documents, endorse specific proposal documents, and provide comments.

Finally, we identified the need to distinguish between tasks per user group. We distinguish between expert and general users. Experts in a specific policy field should have higher access rights to curate policy processes, and therefore, ensure the quality of the provided content. Moreover, these topic-based expertises have to be considered during the calculation of the documents' relevance. Since most of the PolicyLine users do not have an IT-background and cannot spend much time on learning a new application, attention has to be laid on the usability and learnability of the system. Complex visualization techniques would distract users. Moreover, their work requires a high mobility. Therefore, the solution had to be implemented as a web application.

4.3.3. PolicyLine Interface

In this section we describe our visualization system PolicyLine. After a brief overview of the technical background, we explain the overall structure of PolicyLine, including the visual encodings, the interaction design, and relations to the user tasks summarized in Table 4.2. PolicyLine is accessible online[1].

Technical Background. Our approach is implemented as a web application to enable mobile access (**G**$_2$). Metadata about the policy process and the documents added by the users is stored in a persistence

[1] https://policyline.eu

category	task
analysis	$\mathbf{A}_1$ explore policy process and documents
	$\mathbf{A}_2$ identify most relevant documents in process
	$\mathbf{A}_3$ localize policy proposals
	$\mathbf{A}_4$ lookup author of document
	$\mathbf{A}_5$ lookup support behind a proposal
	$\mathbf{A}_6$ lookup sentiments
content creation	$\mathbf{C}_1$ create a new policy process (experts only)
	$\mathbf{C}_2$ define policy process steps (experts only)
	$\mathbf{C}_3$ add documents to the process
	$\mathbf{C}_4$ rate quality and relevance of document
	$\mathbf{C}_5$ incorporate the rater's reputation
	$\mathbf{C}_6$ comment on documents
general	$\mathbf{G}_1$ distinguish between stakeholder groups
	$\mathbf{G}_2$ assure usability and mobility

Table 4.2.: User tasks

layer. The main objective of our approach is to integrate documents from different sources and organize them into policy processes. Users can directly access the original documents via the URLs provided by PolicyLine. The raw text from the original website can be crawled and analyzed with text analysis methods. For example, the results of sentiment and objectivity classification are shown at the Document page (see Section 4.3.3).

In our task analysis we identified the need for a specific user management to distinguish expert users from general users ($\mathbf{G}_1$). Therefore, PolicyLine is connected to the online service EurActory[2]. Besides social media functionalities, the system offers users to rate their own and other users' expertise in pre-defined policy topics. With this information and further reputation management methods a topic-based expertise score is computed for each user, which is used by our approach.

The relevance of a document is calculated based on a weighted sum of the author's reputation and the users' document rating. To rate a document, users have to answer six statements on a 5-level-Likert scale. The statements are: "the document is accurate", "the document is valuable", "the document is relevant", "the document is timely", "do you agree with the document's issues", "do you agree with the proposed solutions". The weights for the author reputation and the Likert-scale-items in the weighted sum were defined together with the domain experts.

In the following, we describe the visual appearance of PolicyLine, including the individual interfaces.

[2] https://euractory.eu

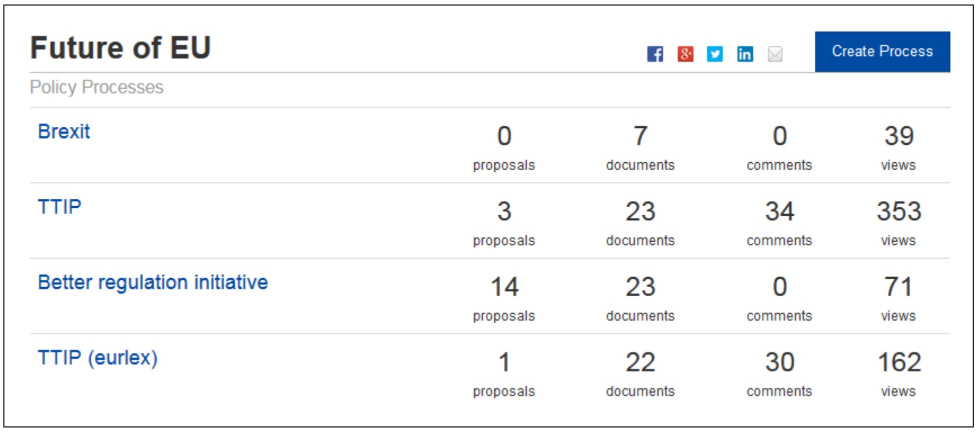

Figure 4.2.: Policy Topic View. Policy processes associated to the topic "Future of EU" are shown. The overview provides statistical information, e.g., about the number of underlying proposal documents.

Policy Topic View. The Policy Topic View (see Figure 4.2) allows the user to select a policy process from a list of processes associated to the policy topic. Besides the title, the number of associated documents classified as proposal, the total number of documents, the number of comments, and the number of views are shown. By selecting one of the policy processes, the user is guided to the respective Policy Process View. Alternatively, expert users can select "Create Process" to access the process creation form, and define a new process.

Policy Process View. The Policy Process View (see Figure 4.3) provides an overview of a policy process, including its process step, and associated documents ($\mathbf{A}_1$). The underlying timeline visualization maps the temporal dimension on the horizontal axis, and the categorical dimension on the vertical axis. The timeline is vertically separated into three rows, comprising the three main document author categories *institutions*, *stakeholders*, and *media*. On top of the timeline, the policy process steps are visualized as areas covering the respective time intervals. In the timeline, documents are represented by colored dots. The colors help users to differentiate sub-categories of authors. Textual labels show the title of the respective document. The dot sizes represent the calculated relevance of the document ($\mathbf{A}_2$). Documents that are classified as policy proposals are represented as rectangles enclosing the title of the document ($\mathbf{A}_3$). Zooming and panning interaction enables handling overplotting in the case of large document collections. On the lowest zoom level, covering the total policy process lifespan, only the most relevant documents are shown. The users can zoom into specific time intervals to display more documents. Further documents can be accessed via an additional list view that shows all documents sorted by relevance. By selecting a document, the user is forwarded to the Document View.

We briefly describe the design decisions made. We selected a scatterplot-based timeline as visualization technique providing an overview of a policy process. During the design process we had to carefully balance between the ease of use and the display of as many attributes as possible. We prioritized tasks to be supported by the visualization technique. These were in the order of relevance:

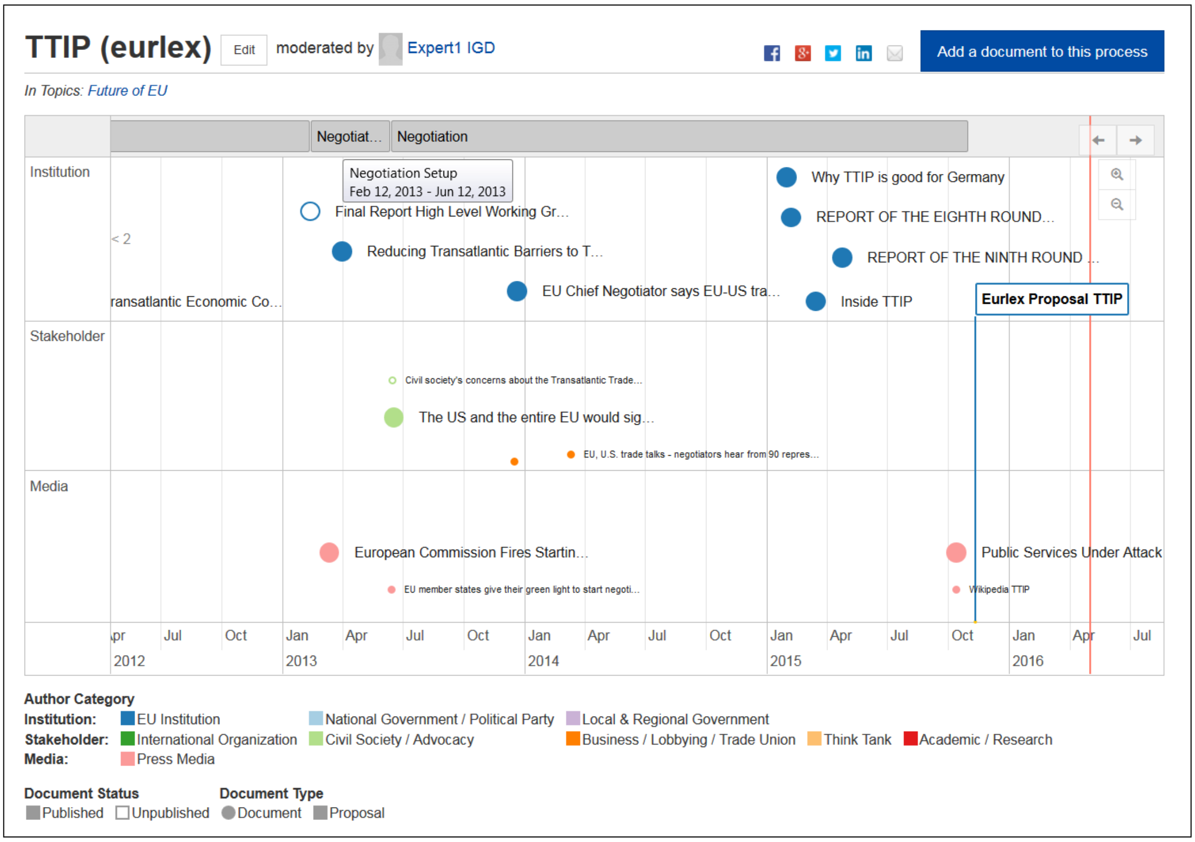

Figure 4.3.: Policy process visualization. Time is mapped on the horizontal axis, author categories on the vertical axis. Documents are represented as circles, proposal documents as rectangles. The objects' colors represent the author sub-categories. The object size reflects the relevance of the document to the process.

1. Quickly derive the timing in a series of publications
2. Get an impression which channel (institutional, media or NGOs) is most active at which time
3. Identify the authorship, or at least the type of author
4. Derive the relevance of each document
5. Access as much meta-data per document as possible

As a result, the metaphor of a timeline was advantageous. Furthermore, since many attributes need to be displayed, a scatterplot-based visualization technique was selected as it features many visual attributes. Applying the perceptual order of importance for visual attributes by Ware [War13], we derived the following visual mappings:

1. position: x-axis for time, y-axis for author categories,
2. object color: to reflect authorship of documents,

3. object size: for relevance of document,
4. object shape: filled objects for documents with a link, empty objects for upcoming documents; rectangles for proposal documents, circles for others

We also discussed the transposition of the axes, mapping time on the vertical axis. However, the given choice is superior in reducing possible overplotting of text. As an alternative to the scatterplot, a simple table could have been used for showing the documents of a policy process. However, a table only allows for sorting the documents by one attribute, e.g., time or relevance. The scatterplot enables the user to view five aspects at the same time.

Together with our domain experts, we also discussed the inclusion of more complex visualization techniques like radar charts, or projection-based techniques. However, during the first evaluation cycle we learned that users were partially overcharged with the complexity of the given visualization techniques. Additionally, the need for content creation features (create process, add document, rate documents, etc.) was identified. This fact shifted our focus from advanced visualization techniques to an intuitive user interface design.

Document View. The document view (see Figure 4.4) shows detailed information about a document ($\mathbf{A}_4$, $\mathbf{A}_5$, $\mathbf{A}_6$), and offers users to provide feedback on a document ($\mathbf{C}_4$,$\mathbf{C}_6$). Besides the title, and the author ($\mathbf{A}_4$) a short description about the content of the document is given. Additionally, meta-information, like author category and subcategory, the creation date, the PolicyLine submitter and submission date, and a link to the original source is embedded. In a box on the right-hand side, a document rating form is shown. Here, the user can rate the quality and the relevance of the document ($\mathbf{C}_4$). As described in Section 4.3.3, this is covered by six statements to be answered by the user on a 5-level-Likert scale. The user can also comment on the document ($\mathbf{C}_6$). The input provided is used for calculating the relevance of the document ($\mathbf{A}_2$) and measuring the support behind a proposal ($\mathbf{A}_5$).

The document view also categorizes the document's content by sentiment and polarity (E6). Based on the raw text crawled from the original source, a computational linguistics analysis component automatically classifies the document as positive, neutral, or negative, and as subjective or objective. In addition to the class a confidence score (in %) reflecting the accuracy of the classification is presented.

Process & Document Creation Forms. Two simple input forms are included at PolicyLine to enable users to create policy processes ($\mathbf{C}_1$) and steps ($\mathbf{C}_2$), and augment them with documents ($\mathbf{C}_3$). To create a process, users need to provide a process title, a short description of the process, and the process steps with step title, start and end date. In addition, if available, a link to a legislative procedure documented in EUR-Lex [Eur16] can be added. This link allows crawlers to add official documents stored in the EUR-Lex repository to the policy process.

After the creation of a policy process, and the respective policy process steps, the user can add relevant documents to the process. In the respective form users need to add the web link to the document, the document title, the official publishing date of the document, and a short description. Moreover, the users can manually tag the document as a policy proposal. Finally, one out of nine author categories has to be selected (cf. Table 4.1).

Figure 4.4.: Document View. Details about a document are shown. Users can read a short description or follow the link to the original source. On the right-hand side users can view the average user rating and provide their own rating.

4.4. Design Process and Evaluation

In our approach, we followed the design study methodology by Sedlmair et al. [SMM12]. Our design process was conducted in a two-years effort, comprising three iterations of design (DESI), implementation (IMPL), and evaluation (EVAL) (see Figure 4.7). The focus of each round was chosen in line with the suggestions provided by Munzner et al. [Mun09]. Based on the *domain characterization* (cf. Section 4.3.1) and the *data and task abstraction* (cf. Section 4.3.2), we implemented an early prototype that was shown to selected expert users. In interviews with the experts we validated the *problem characterization* and the *data/operation abstraction design*. In the second round, we focused on the *visual encoding* and the *interaction design*. Feedback on the usability of PolicyLine was collected. Finally, in the third round, we repeated the evaluation design of the second round to measure the progress made. In the following, we briefly describe each of the evaluation stages and present the achieved results.

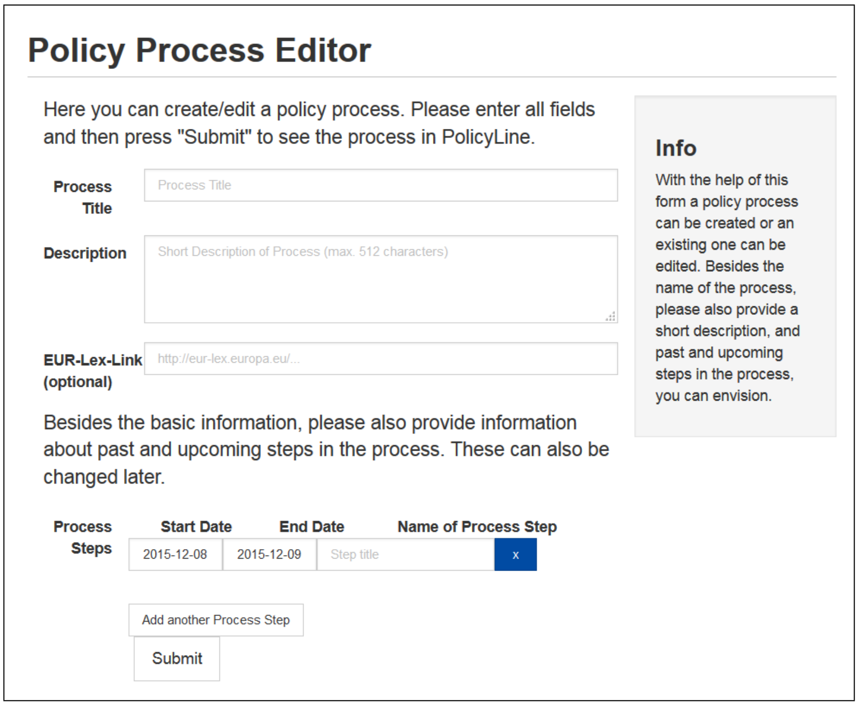

Figure 4.5.: Process Creation Form.

4.4.1. Pre-Evaluation - Expert Interviews

Initial user feedback on the first version of PolicyLine was gathered in a pre-evaluation round which has been conducted in form of informal interviews with four selected expert users highly connected to political EU institutions. The purpose was to understand whether the design of our approach de facto addresses the experts' problems. Feedback about the usefulness of PolicyLine helped us to extract requirements on a refined version. The main requirement was to improve and expand the functionalities regarding the manual creation of content. The experts preferred manually provided content instead of using automatic approaches for extracting policy processes and process steps from external sources like EUR-Lex. Also the manual classification of the documents' author categories was preferred over automatic approaches. This feedback further strengthened our vision of a collaborative system approach. Still, the experts promoted the usage of automatic approaches to augment the policy processes and analyze its textual content. Moreover, the experts shared detailed perspectives on a refined structure of the timeline visualization. A clear separation of the high level author categories on the y-axis, a decoupled display of non-overlapping process steps on top of the timeline, and an intuitive zooming functionality were requested. The detailed feedback on creating policy processes, on adding manually curated documents, and on the timeline visualization enabled us to improve and refine the design of PolicyLine significantly.

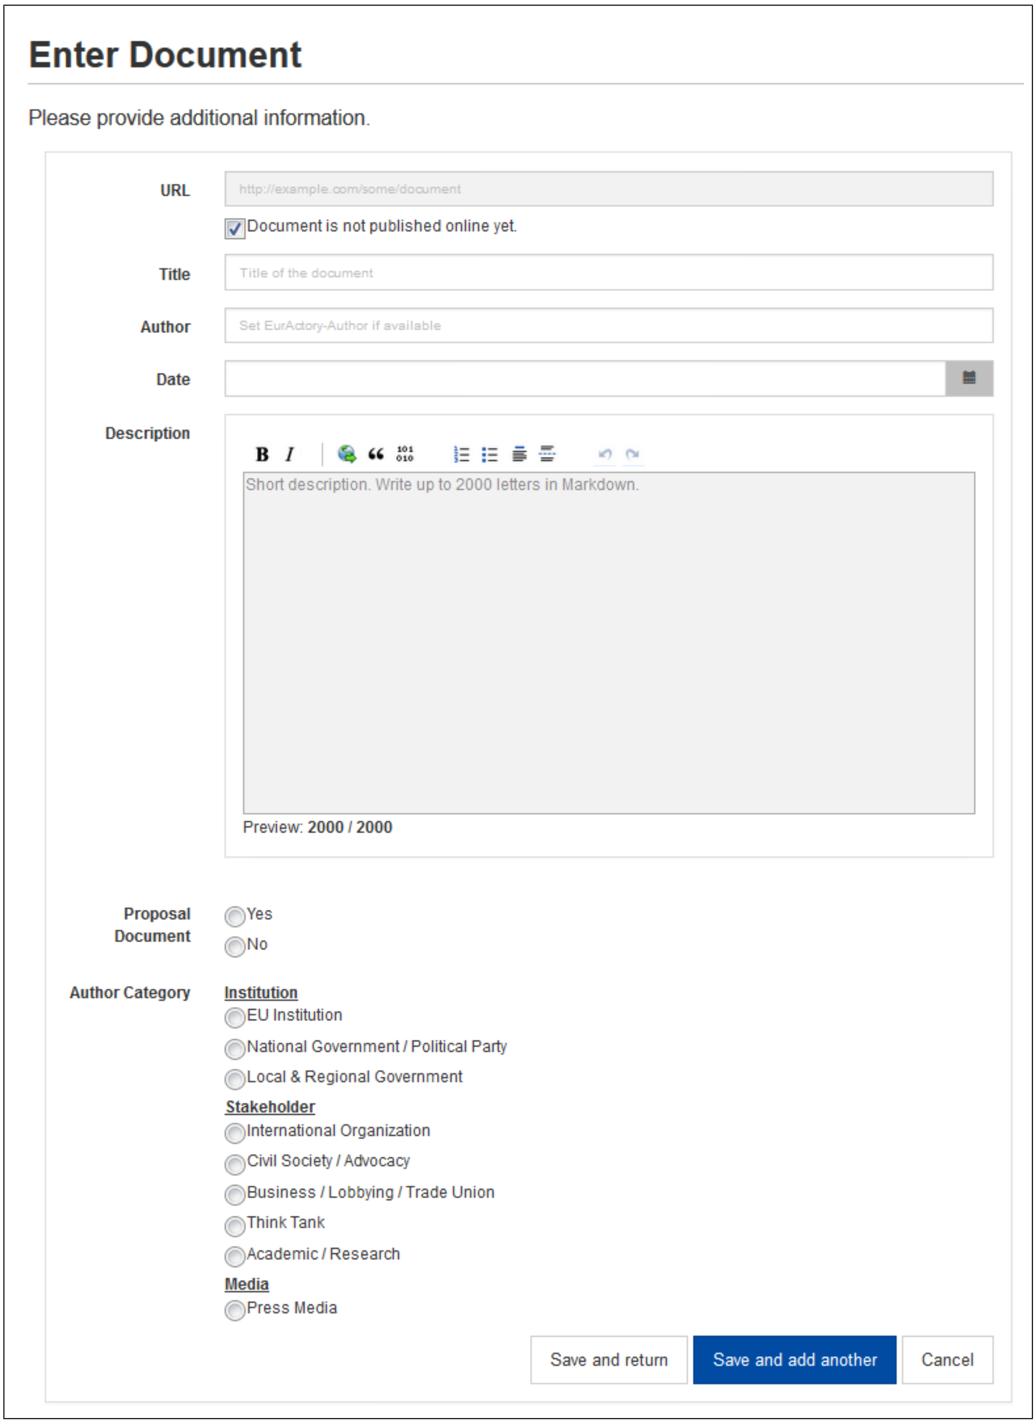

Figure 4.6.: Document Creation Form.

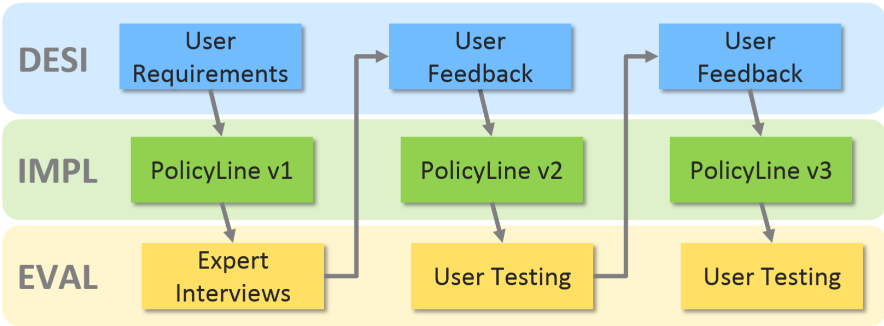

Figure 4.7.: Overall Evaluation Methodology.

4.4.2. Two Qualitative Evaluation Rounds

We tested the usefulness and the usability of PolicyLine in two consecutive evaluation rounds. We applied the same experiment design in both rounds which allowed us to (a) receive feedback and suggestions for improving the application (formative evaluation), and (b) to measure the progress between the first and the second version of PolicyLine (summative evaluation). The two rounds were conducted with fifteen and ten real-world users from the EU environment.

Evaluation Design. Both evaluation rounds were divided into four parts: (1) a personal information form, (2) a task completion test, (3) short usability questionnaires on the individual interfaces, and (4) a final questionnaire on the overall approach. Through the personal information form, we acquired information about the users' IT expertise. The task completion test was designed to cover the main functionalities of PolicyLine. The users were asked to (a) create a policy process, (b) add documents to the process, (c) explore the process with the timeline visualization, and (d) explore the information provided about a single document (only applicable for the second round). After each task, they had to respond whether they were successful in completing the task. After the task completion test, the users had to answer separate questionnaires on the respective interfaces. For each visual interface we designed four open questions, and four (closed) statements, which had to be rated on a 5-level-Likert scale. For each interface, the open questions 1-3 were covered by the control questions 2-4, while the first control question covered the task completion test. These questions/statements concerned the intuitiveness, the organization, and missing information of the interfaces. The final questionnaire contained five open questions for sharing ideas and suggestions, followed by ten closed statements about the overall impression of our approach. The evaluation was conducted with an online tool developed by Nazemi et al. [NBH*15].

Results of First Round. Given the fact that future users of PolicyLine expect the system to be self-explanatory, the evaluation was conducted without introduction. In the first round, we recorded difficulties to grasp the general concept of a policy process and the process steps. As a result, we improved the introductory aspects of the system. Therefore, we designed an infographic that explains the overall concept of PolicyLine in a compact way. Another issue concerned the EUR-Lex repository [Eur16], the main source for European legislative documents. While some experts appreciated the inclusion of a EUR-Lex link as source for automatic document crawling, others were not aware of this information source. They reported that they still use Google for searching relevant policy documents, which once more supports the idea of our approach to provide a "one stop shop decision making aid" as one participant stated. The general intuitiveness and usability of all interfaces was described very positive. Moreover, the participants welcomed our collaborative approach by requesting further manual content creation facilities, like providing the owner/initiator of a process, enabling longer textual descriptions for documents and processes, etc. Moreover, they proposed the augmentation of the policy process with events. Several participants requested refinements on the author categories, e.g., a category to identify peer-reviewed research papers was requested. As an overall feedback, most of the users stated that our approach is very innovative, and would help them in their daily work. They also provided high ratings for the learnability of PolicyLine. Surprisingly, some participants also commented on the aesthetics of the interface. While we were focusing on the functionality of the interfaces in the first implementation round, they demanded to improve the look and feel of the application in order to attract more users.

Results of Second Round. In the second evaluation round, experts identified several improvements on the process creation form. Still, most of them commented that this is a very important task that should be carefully executed by expert users only. The success of our approach would heavily depend on the quality of the curation process. Hence, once more the separation of expert users curating processes and associated documents, and general users with the intent of getting an overview of ongoing policy processes was emphasized. The refined author categories were widely accepted. The linguistics component at the document page was not well understood by the experts and required more explanations. The timeline was appreciated by most experts. Still, some participants stated that it looks "busy". This obvious trade-off between the visual information density and the subjective usability will be subject to possible future work. As a final remark, we want to emphasize the dissent between some users having problems in handling simple visualization techniques and others requesting more innovative visualization techniques. Exemplary statements on our overall approach were: the tool is "multi-stakeholder and multi-purpose", "a useful decision making tool that creates better set of options", "practically no training needed", "the timeline is useful once it is well created and all relevant documents have been uploaded", "a dynamic information channel".

4.4.3. Collaborative Usage Scenario

PolicyLine's overarching goal is to increase the transparency of the policy process among all stakeholders involved. To illustrate its power, we demonstrate a collaborative usage scenario. Policy analyst *Alice* is engaged by the European Commission (EC) to initiate and analyze the public discussion on TTIP. Using PolicyLine, Alice simply creates a new process, enters a few details and starts adding the relevant institutional documents (like white papers or policy proposals) written by the European Commission and published via the online portal EUR-Lex. Then, she shares the PolicyLine process to initiate a public discussion and becomes an observer monitoring the process. Multiple stakeholders start to contribute to the process. One of them is lobbyist *Lawrence*, who works for a large multi-national corporation. He is interested in influencing the legislative decision in the direction of his employers' favor. For this, he adds documents that argue pro TTIP. Finally, he rates existing documents in line with the vision of his employer. Alice also reached *Chloe* who works for an Anti-TTIP NGO. Chloe adds documents to the process, rates further documents against TTIP, and also adds her own proposal. As this flow is repeated by different stakeholders, the public discussion intensifies. Alice monitors the evolution closely and observes that the Anti-TTIP proposal by Chloe has a high support from many credible experts in the topic. Therefore, she puts that proposal on the agenda for one of her next meetings at the EC in order to evaluate the raised key issues. Obviously, she uses PolicyLine as a visual means to communicate the evolution of the process.

4.4.4. Lessons Learned

Providing a policy process visualization system to political stakeholders was appreciated by experts from the domain. There was an obvious lack of such a system. The content creation functionality was well accepted by most of the experts, although there was a discussion, who should be able to curate policy processes. They emphasized that the process creation cannot be automated due to the absence of standardized structures. Still, the automatic augmentation of official documents from existing repositories (e.g., EUR-Lex) was appreciated. Further automatic techniques like linguistic analysis methods were initially requested by the domain experts. However, the functioning of such techniques needs to be carefully explained to the users to increase the acceptance of text analysis results. From this, we learned that *trust* and the *awareness of uncertainty* in the data needs to be carefully considered during the design of visualization systems for policy domain users.

From the visualization perspective, we learned that a visualization system supporting the derived tasks needs to be carefully designed. We learned that the attempt to structure policy processes that are not necessarily structured and providing an intuitive access to the derived structure are difficult tasks. Meeting the expectations of different stakeholder types is cumbersome. However, the involved stakeholders are very enthusiastic about the benefits of such an approach. Seeing the whole process lifespan including the most relevant documents at a glance was a requested key feature. However, the definition of relevance is difficult due to conflicting user expectations. Therefore, drill-down (by zooming), and search functionality to explore less relevant documents is an essential functionality for knowledge acquisition.

Finally, our initial design process plan consisted of two cycles including design, implementation, and evaluation within each cycle. During the first design phase, we realized the need for an intermediate evaluation to validate whether the user expectations were met. From this, we learned that the policy domain varies from other application domains. First, computer expertise strongly varies in this field. Both highly skilled technicians and seniors with little to no computer expertise collaborate in this domain, which makes it difficult to derive clear requirements from the users. Second, due to time pressure, political stakeholders are difficult to reach. Hence, it was of key importance to collaborate with partners that had close connections to EU stakeholders.

4.5. Summary

In this chapter, we presented a visual analytics approach that serves as a proof of concept for the applicability of our design methodology illustrated in Chapter 3 on textual data. The main purpose of our visualization design was to structure political processes, shaped by underlying core documents. We presented the intermediate results of our design study that guided us in designing, implementing, and evaluating the visualization system. The design process was structured along our concept presented in Section 3.2.3. During the first intermediate design and evaluation cycle, we included a pre-evaluation activity by presenting the first version of PolicyLine to four selected policy experts. According to their feedback we refined PolicyLine focusing on the amelioration of the interfaces. On the basis of the refined system, we conducted two evaluation rounds focusing on the usability and usefulness of PolicyLine. The evaluation questionnaire included both Likert scale statements and open questions to receive qualitative feedback. The results of both rounds helped us to further improve PolicyLine towards a usable and useful system, which was illustrated in a usage scenario. Additionally, we provided ideas on how to address the shortcomings identified by the test participants.

Future work will incorporate additional technical features. Depending on the targeted use case and application domain, enhanced sentiment analysis, topic modeling, simulation, or prediction models may further add to the process. Similarly, mechanisms for the enrichment of the process with automatically crawled documents are subject to future work.

5. Visual-Interactive Access to Document Collections

In this chapter, we present a visual analytics approach that targets the provision of visual access to text document collections to be considered in the decision making process. Regarding the overall approach of this thesis, we prove the applicability of our concept to explore and analyze document collections, Challenge $\mathbf{C}_{\text{Doc}}$. The approach applies the design methodology presented in Chapter 3 in two ways (Challenge $\mathbf{C}_{\text{VDSS}}$). On an abstract level the visual analytics approach supports the exploration and analysis of text document collections that inform the decision process via text clustering. On an alternative perspective the approach supports the decision to find an adequate grouping (or clustering) of documents adapted to the respective users and tasks at hand. That way the design study targets the creation, exploration, analysis, and comparison of text document clustering results to generate content-

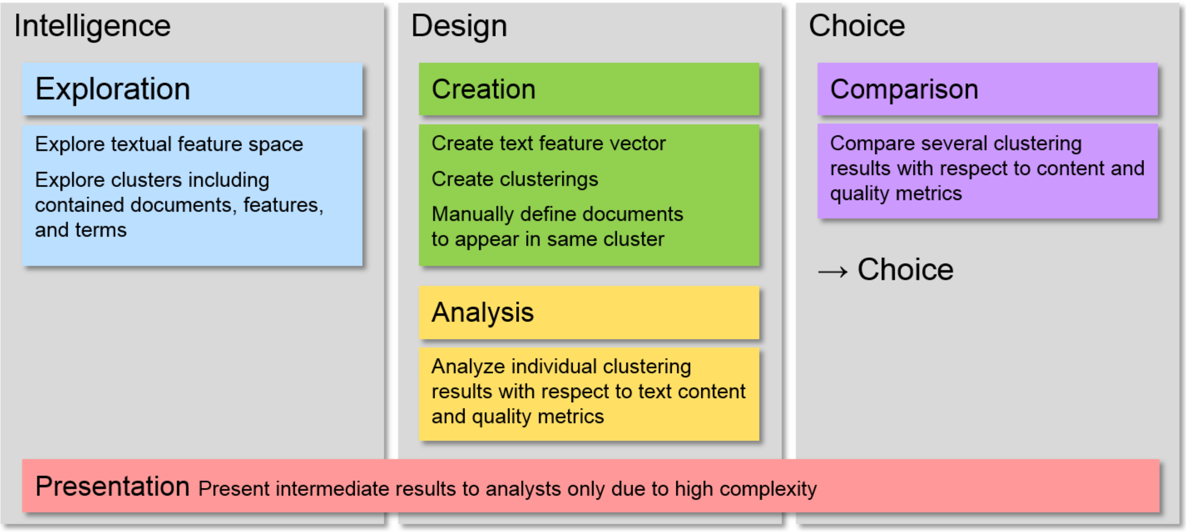

Figure 5.1.: A second proof of concept for the applicability of our design methodology on **textual data**. The visual analytics approach supports the creation and validation of content-based document clusterings. Clusters are used to create overviews on large document collections. Different clusterings can be created analyzed and compared. The complex analysis workflow is mainly usable for analysts. Figure is adapted from Figure 3.6.

based overviews of large document collections (see Figure 5.1). First, users are enabled to explore the document collections via contained textual terms. This includes the filtering based on specific parts-of-speech, named-entities, or 2- and 3-grams. Second users can explore clustering results with respect to contained documents, and frequent and discriminative cluster terms. Third, users can manually create document groupings of documents that have to appear in the same cluster. These clusters are used as a ground truth for the cluster analysis. Third, users are enabled to define a text feature vector and specify the parameters for an automatic cluster algorithm. The resulting clusters can be analyzed with respect to their content and specific cluster quality metrics. Fourth, different clustering results obtained via different feature vectors and parameterizations can be compared based on their content and quality metrics. As a final result, the user can select an optimal clustering fitting to the respective task at hand. The visual analytics system is designed to support analysts without prior knowledge in text analysis to create and analyze document clusters. These document clusters support the analyst in providing overview of the textual content of large document collections. However, the system has a high degree functionality which also increases its complexity. Therefore, the tool is not usable for decision makers or other stakeholders. Challenge $\mathbf{C}_{\text{BKG}}$, the bridging of knowledge gaps, is only addressed between the modeling expert in text analysis and the analyst as a user of the system. This chapter is partially based on our previous work published in [RSB*17].

Contents

5.1. Introduction

The volume of digitally available textual data is continuously increasing. Examples for document collections include newspaper articles, scientific papers, technical reports, patents, legislative documents or social media entries like tweets, blog posts or customer reviews. These documents are highly relevant for many types of stakeholders like journalists, researchers, political decision makers, and online-shop customers. Methods from information retrieval are the means of choice, if stakeholders can specify their information need precisely, e.g., by formulating a *search* query. However, these fact retrieval or known-item search techniques often become ineffective, if document collections are large, complex, or unknown. In such scenarios, the goal to gain an overview of the document collection can be achieved via the *exploration* of structural information within the collection.

The mechanisms needed to enable the exploration of document collections strongly differ from classical search methods [WR09]. Among others, data aggregation methods support the generation of content-based overviews, by condensing large numbers of documents into a small set of representatives. One of the most prominent classes of aggregation methods is data clustering with its plethora of techniques and its ability to solve various real-world problems. However, analysis approaches based on clustering are confronted with a variety of challenges.

First, clustering algorithms require numerical feature vectors as input, they cannot process unstructured text documents. The definition of an appropriate feature vector representing text documents is a non-trivial task. As a common practice in text analysis, a feature represents a term that occurs in a document and the feature value describes the relevance of the term to the document (cf. vector space model). Using the entire vocabulary of the document collection would result in large feature vectors that are sensitive to noise and inefficient to process. Thus, the size of the feature vector needs to be reduced by selecting a content-preserving feature subset as representatives of the documents. The effective selection of relevant features can be supported by several metrics. However, different metrics may produce different feature rankings. Moreover, the definition of appropriate thresholds for selecting or deselecting features is challenging.

A second challenge is the choice of a suitable clustering algorithm. Different clustering algorithms produce different groupings of objects owed to the fact that they are designed for different problems. Additionally, the results depend on the parameterization of the clustering algorithm. Multiple cluster quality measures exist that allow to quantify the internal quality of a clustering result (e.g., compactness and separation of clusters). However, these measures focus and different characteristics and some of them may even contradict each other. Since, there is no ground truth to measure against, a "best" clustering method does not exist [Jai10]. It is up to the user to evaluate the quality of a clustering depending on the document collection and the analytical task at hand.

This leads to the third core challenge, the comparison of different clustering results. Since no best clustering method exists, users need to be supported in the choice of the most appropriate among several clustering results. Comparing the clustering results based on internal quality metrics is reasonable. However, the comparison of document-cluster affiliations of several clustering results is a non-trivial task, since the analysts are confronted with interesting document subsets distributed over different clusters over different clusterings.

The rationale of this research approach is to formalize the design space for text document clustering processes. The resulting framework builds the basis for the design of text clustering workflows to be applied to document collections in strong accordance to the involved users, data, and tasks. In particular, we aim at enabling analysts without prior knowledge on text analysis to create analytical document clustering workflows. Visualization and interaction techniques from information visualization and visual analytics have proven to ease the access to complex data spaces and analytical models, respectively. Visual comparison and guidance concepts can be applied to make meaningful decisions in the choice of algorithms and parameters. The contributions of our approach are as follows:

1. Feature selection: we present a visual interface for the selection of textual features (terms) from a document collection with the goal to reduce the size of the feature space. Ranking based on the feature selection metrics, and filtering based on the feature type support the users during the selection process. Intermediate feedback is provided to the users by directly displaying selected and deselected features.
2. Cluster analysis: we present a visual interface for the analysis of document clustering results. The clustering can be analyzed based on the cluster contents and cluster quality measures. The content-based perspective is defined by the documents assigned to a cluster and the prevalent features and terms in a cluster. In addition, cluster quality measures support the user in evaluating the cluster's compactness and separation.
3. Cluster comparison: we present a visual interface for the comparison of clustering results. A content-based and quality-metrics-based perspective is provided. Users can identify intersecting subsets that appear throughout several clusterings and inspect the documents and terms contained in theses subsets. Additionally, the cluster quality measures of different clusterings can be compared and the F-measures of the clusterings to a manually defined reference clustering can be inspected.

5.2. Related Work on Visual Text Clustering

We review visual analytics and information visualization approaches related to document clustering. First, we discuss visual and interactive approaches that support the feature selection process, mostly executed prior to clustering. Second, we review visualization systems that apply clustering or other aggregation techniques to derive structural information from document collections to generate overviews. Third, we provide a short summary on related work about the comparison of clustering results. And finally, we review uncertainty visualization techniques that raise the users' awareness of projection errors.

Feature Selection. Several approaches exist, that address the visual and interactive selection of features. Examples are the work by Guo [Guo03], SmartStripes [MBD*11], INFUSE [KPB14], or the Rank-by-Feature framework [SS05]. We share the idea of defining ranking criteria to enable the reduction of the multi-dimensional feature space. However, these systems differ from our approach, since all of them work on numerical data, while we are focusing on textual features. Most textual features selection approaches only allow users to define thresholds for the feature selection metrics. Features with values beyond these thresholds are excluded from further analysis steps in the clustering workflow (e.g., [CLL*13]). Other text clustering approaches use the entire feature space without applying any feature selection mechanism. To the best of our knowledge, no system exists that allows user to visually select textual features.

Document Collection Overviews. Several approaches from the field of visual analytics target the exploration and/or analysis of document collections. Some approaches use meta-information like author, publication year, citations, etc. to group documents. Examples include SurVis [BKW16], CiteRivers [HHKE16], PolicyLine [RBB*16] and an approach by Oelke et al. [OSR*14]. We do not use any meta-information except the document title (if available) but focus on the content data. A class of content-based approaches use a vector space model (consisting of feature term weights) to represent documents. From these vectors, topic models can be extracted in order to structure the document collection. Among others, Latent Dirichlet Allocation (LDA) is a prominent topic modeling approach [BNJ03]. Each document is represented by a mixture of topics. The topics are represented by a weighted set of terms. Examples of interactive visualization systems that apply LDA to provide overviews of document collections are ParallelTopics [DWCR11], TIARA [WLS*10], and TextFlow [CLT*11]. All of these approaches are limited to one single clustering algorithm. The experimentation with different clustering techniques and the comparison of differing results is not provided. Moreover, none of the approaches support the visual and interactive refinement of workflow steps, which was one of the goals of our system.

The iVisClustering approach [LKC*12] and the UTOPIAN system [CLRP13] allow the refinement of topic models via user interaction and visual feedback. In addition, both approaches project documents to the display space. Cluster affiliations are represented by categorical color maps, and weighted topic keywords can be adjusted by the user. While iVisClustering incorporates an enhanced LDA model, the UTOPIAN system introduces an alternative approach for the interactive refinement of topic models, non-negative matrix factorization. Although both approaches support the interactive refinement of the underlying models, they are restricted to only one model. They do not address the comparison of results coming from different models.

We highlight two approaches that provide content-based overviews of document collections via clustering: IN-SPIRE [Wis99], and Overview [BISM14]. IN-SPIRE generates thematic document landscapes by combining document clustering, projection, and keyword extraction. The Overview system organizes document collections in a tree structure based on the results of a hierarchical clustering. While IN-SPIRE has limited interaction and refinement capabilities, the Overview system allows users to document findings by manual tagging. However, both systems rely on a single clustering method, the comparison of clustering results is not addressed.

The Jigsaw visual analytics system supports the exploration of a document collection by extracting entities in documents and analyzing their co-occurrences [GLK*13]. In addition, document clustering and document summarization techniques are incorporated. However, the approach differs from ours since it neither addresses the feature selection process nor the comparison of different clustering algorithms.

Visual Cluster Comparison. Our approach is related to techniques supporting the visual comparison of multiple clustering results. We selected the parallel set visualization to compare document affiliations to clusters from different clustering results [KBH06]. Further visualization techniques that support the comparison of sets are presented in a survey by Alsallakh et al. [AMA*16]. The clustering comparison component in our approach is also inspired by XCluSim, a visual analytics tool applied in the area of bioinformatics [LKS*15]. The tool allows the comparison of clustering results coming from different clustering algorithms. It uses an enhanced parallel sets visualization that incorporates a tree color map to allow the identification of related clusters coming from different clustering results. The clusters are colored according to their similarity. Documents are depicted via gray bands between the clusters. Since the main target in our clustering comparison is to identify stable subsets, we follow the parallel set visualization approach (see above).

The paper most related to our approach was presented by Choo et al. [CLL*13]. It introduces an interactive visual testbed system that allows the definition of dimension reduction and clustering workflows. While the paper primarily focuses on the integration of different data types, our system targets textual data. We provide content overviews, summarizing most relevant features in clusters, and cluster subsets. In addition, we also support users in the feature selection phase.

Uncertainty Visualization. Finally, we draw a connection to uncertainty visualization. In our approach documents are projected on the display space and represented as circles to analyze their similarities. In a recent publication by Sacha et al. the role of uncertainty, awareness, and trust in visual analytics is discussed [SSK*16]. A comprehensive overview about visualizing geospatial uncertainty is provided by MacEachren et al. [MRG*05]. The work summarizes sources of uncertainty and possible techniques for visualizing them. Among others, saturation can be used to depict the uncertainty of objects. This technique is also called pseudo-coloring in a survey about depicting uncertainty in scientific visualization approaches by Pang et al. [PWL97]. We adopt this concept to represent projection errors using a sequential color map.

5.3. Requirements

In our approach, we introduce a visual interface for the creation of text clustering workflows with the goal to structure and explore document collections. The targeted *user* group are data analysts. The approach aims at opening up the design space for text clustering workflows, and making them accessible for data analysts. The resulting system should also be applied by users without a specific expertise in data mining, NLP, or statistics. However, prior knowledge about the applied methods is beneficial for the selection of algorithms and the interpretation of results. In a realistic scenario a data analyst will use the system to design an optimal text clustering workflow for a stakeholder with a specific interest in a text document collection. The resulting clustering workflows should answer several questions that the stakeholder might have specified prior to the design. Examples include: What is the collection about? Which groups of documents emerge? What are the groups about? Are there alternative groupings? How do these groupings differ? Which documents are similar? Why are the documents similar? What is a document about? Keeping these questions in mind, we defined some concrete requirements that should be considered during the design of a visual text clustering system.

First, the desired clustering workflow should heavily rely on the underlying data and task. Therefore, users have to get *access* to the entire clustering workflow. That way, domain knowledge can be incorporated into the analysis process. Workflow steps include text pre-processing, feature selection, clustering specification, analysis of a single clustering result, and comparison of multiple clustering results.

Second, since there is no “best” clustering workflow, users should *analyze* the quality of a clustering result, depending on the underlying data and task at hand. The quality assessment can be supported in two ways: (i) by providing the user overviews on the clusters’ content (showing prevalent documents and features/terms), and (ii) by incorporating cluster quality measures in the overviews.

Third, users should be enabled to *compare* several clustering results. This requirement is of key importance to allow the user to decide upon the most appropriate clustering result for the task and data at hand. Similar to the analysis of a single clustering result, the comparison of several clustering results should be based on (i) the cluster content and (ii) the cluster quality measures. To simplify the comparison, the users should be supported in identifying subsets of documents that are constantly grouped together in a cluster across many clustering workflows.

Fourth, to allow the iterative refinement and comparison of clustering results, intermediate results in the workflow should be stored and made accessible for the user in a *history*. This allows to recall and/or refine previous results for comparisons and/or improvements, respectively. The benefits and purposes for data provenance have been presented recently by Ragan et al. [RESC16].

The requirements can be summarized as follows:

- $\mathbf{R}_1$ Access: each workflow step needs to be made accessible to the user. The parameterization of workflow steps should be controlled by the user. Interim results of each workflow step should be presented.
- $\mathbf{R}_2$ Analysis: quality of a clustering result should be evaluated by the user. The cluster assignment of documents, the prevalence of feature terms, and the cluster quality documented by cluster measures should help users in their judgment
- $\mathbf{R}_3$ Comparison: users should be able to compare different clustering results based on the resulting clusters, and the respective quality measures
- $\mathbf{R}_4$ History: to enable an iterative workflow, intermediate analysis results should be stored in a workflow history.

5.4. Text Analysis & Clustering Methods

For the realization of our approach, we incorporated techniques from the field of data mining, natural language processing (NLP), and statistics. An overview of the applied techniques is given in Table 5.1.

We use a vector space model, representing each document with a weight vector [Liu07]. The dimensions in the vector represent unique terms (features), the weights are calculated based on the term frequency-inverse document frequency (tf-idf) in the underlying document. The resulting vectors are used in the text clustering workflow, e.g., to calculate the similarity between documents.

Pre-processing. To generate the vector space model for a document collection the originally unstructured texts is preprocessed. Pre-processing includes (a) optional stop word removal, (b) optional

feature selection	clust. specification	cluster representation	doc. projection
feature selection metrics:	**clustering methods:**	**content-based cluster**	**projection method:**
document frequency (df),	k-means++,	**representation:**	MDS,
term frequency-inverse	hierarchical	document affiliation,	Sammon mapping
document frequency	clustering,	frequent (tf) or correlated (χ^2)	**projection error:**
(tf-idf),	power iteration	cluster terms or features,	neighborhood
term contribution (tc)	clustering,	feature type-filters:	preservation,
feature type-filters:	bisecting k-means	POS, named entities	trustworthiness
Part-Of-Speech (POS)		**cluster quality measures:**	
tagging;		compactness, separation	
named-entity recognition;		Dunn & Davies-Bouldin index	
token, 2-gram, 3-gram extraction			

Table 5.1.: Integrated methods and metrics: NLP, clustering, and projection methods are incorporated in the text clustering workflow. Additional feature selection, cluster quality, and projection error metrics support the user in the creation and validation of the workflow.

punctuation removal, (c) optional stemming, (d) the extraction of single terms, 2-grams, and 3-grams, (e) part-of-speech tagging (POS), and (f) named entity recognition.

Feature selection. Since the vector space model contains the entire vocabulary of the document collection, the resulting feature vector might be very large. Therefore, feature selection is applied to reduce the dimensionality of the model. The features are ranked based on feature selection metrics to support the user in the selection. We included three commonly applied metrics: term frequency-inverse document frequency (tf-idf), document frequency (df), and term contribution (tc) [LLCM03]. In addition to the ranking, filters can be applied to the features. In our approach, we incorporated (a) POS filtering, (b) named entity filtering, and (c) token, 2-gram, and 3-gram filtering based on respective extraction techniques. After the feature selection step, a document is represented by the reduced feature vector with weights defined by the tf-idf.

Clustering. Clustering algorithms require the definition of document similarity. We apply the cosine distance between the documents' feature vectors. We incorporated four clustering methods which are often applied to cluster documents: k-means++ [AV07], hierarchical clustering, bisecting k-means [SKK00], and power iteration clustering (PIC) [LC10]. For the evaluation of the clustering results, we apply four cluster quality measures: compactness, separation, Dunn- and Davies-Bouldin-index [IPRE08].

Feature extraction. To represent the content of a cluster we extract the most relevant and the most frequent cluster terms (or features). The cluster-wide term frequency (tf) is applied to extract the most frequent terms in a cluster. However, terms that are frequent in several clusters are not discriminative. Therefore, we include a second measure to extract terms that highly correlate with a cluster, the χ^2 statistics. To ensure that terms are extracted that occur in the cluster, we only take the positively correlated terms into account. The two measures (tf and χ^2) can be applied to both the reduced feature space (derived from the feature selection step), and the full document vocabulary.

Projection. Finally, we incorporated two projection and layout techniques to provide a visual overview of the documents space: multi-dimensional scaling (MDS) [CC00] and Sammon Mapping [Sam69]. Due to the curse of dimensionality, the projection error might induce a misinterpretation of the vector space similarities between documents. To make the users aware of these effects, we included two measures representing the projection error: trustworthiness and neighborhood preservation [KNO*03].

5.5. Visual Analytics Design

Our visualization system supports the creation and validation of text clustering workflows to explore document collections. The system was designed based on the requirements presented in the previous sections. The standard text clustering workflow comprising the stages pre-processing, feature selection, clustering method selection and parameterization, and cluster analysis was expanded by an additional stage, which allows users to compare several clustering results (see Figure 5.2). In our approach, the text clustering workflow is grouped into three stages of which each is presented in a separate view: the Feature Selection View, the Cluster Analysis View, and the Clustering Comparison View. Details about these interfaces will be provided in the following sections.

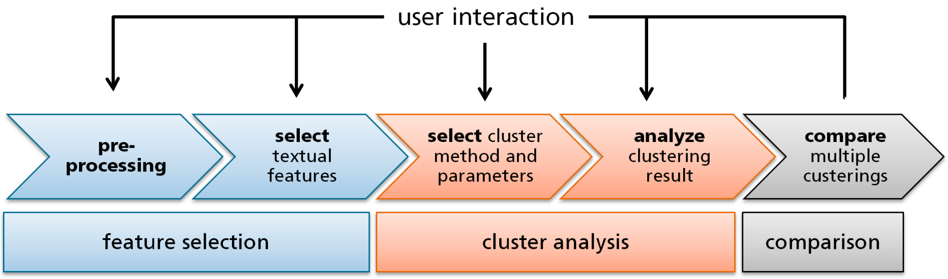

Figure 5.2.: Text clustering workflow: the standard workflow comprising pre-processing, feature selection, clustering selection, and cluster analysis is expanded by the clustering comparison step. The complete workflow is covered by three views: the Feature Selection View, the Cluster Analysis view, and the Clustering Comparison View.

Figure 5.3.: Feature Selection View: **Workflow generation panel** (top right): pre-processing is executed to generate a new workflow; optional pre-processing steps include stop word removal, punctuation removal, and stemming. **Feature selection metrics panel** (top): three metrics are shown (df, tf-idf, tc), features are grouped into buckets on the vertical axis, for each bucket the value range of the underlying features is shown on the horizontal axis; users can adjust bucket size and sorting of features (each chart sorted individually, or all charts sorted based on single metric); intersection or union of selected features can be applied. **Feature panel** (bottom): shows samples of selected and not selected features. **Filter panel** (middle right): user can filter features based on length (token, 2-gram, 3-gram), named entity types, or parts-of-speech.

5.5.1. Feature Selection View

The Feature Selection View (see Figure 5.3) supports the user in the visual selection of features ($\mathbf{R}_1$). To prepare the feature selection, pre-processing on the original documents needs to be executed. Therefore, the user has to generate a new workflow in the *workflow generation panel*. Users can choose whether the pre-processing should include stop word removal, punctuation removal, and stemming. After the pre-processing the new workflow is shown in the *history panel* (see an example in Figure 5.4 (bottom right)). Here, each workflow is represented by a quadruplet: the workflow ID, the latest workflow step that was successfully executed in this workflow, a copy button that allows users to duplicate existing processes, and the delete button to remove a workflow from the history. The *history panel* is shown in each view, and allows users to switch between different workflows. Moreover, it allows users to define several workflows with differing parameterizations, which can be compared in the Cluster Comparison View ($\mathbf{R}_4$). In the *feature selection metrics panel* range bar charts represent the feature metrics extracted in the pre-processing step. Three metrics are incorporated: document frequency (df), term frequency-inverse document frequency (tf-idf; here, the tf in the entire document collection is used), and term contribution (tc). To enable the visualization of large vocabularies, the features are grouped into buckets and mapped on the vertical axis. The number of buckets (resolution) can be selected by the user. Each bar represents a bucket by depicting the value range of the feature selection metric on the horizontal axis, from the minimal to the maximal feature value in the bucket. The features on the horizontal axis can be sorted individually, or based on one of the feature selection metrics. Users can select buckets in the chart via a rectangular rubber-band. Intermediate feedback is provided to the user by showing samples of the selected and (unselected) features in the *feature panel* below the respective chart ($\mathbf{R}_1$). Users may select features from any of the three charts, and decide whether the actual feature subset is defined as a union or intersection. The resulting feature space is shown in the *feature panel*, represented by "selected" and "not selected" features ($\mathbf{R}_1$). The *filter panel* (below the *workflow generation panel*) allows users to apply additional filters on the features. The filtering of tokens, 2-grams, 3-grams, named entities (e.g., locations, persons, etc.), and parts-of-speech (e.g., nouns, adjectives, punctuations, etc.) are supported. The selected features are used for the representation of documents, and the calculation of document similarities in the subsequent workflow steps.

5.5.2. Cluster Analysis View

In the Cluster Analysis View (see Figure 5.4), users can define and validate a clustering based on the feature representation selected in the previous view ($\mathbf{R}_2$). The view is divided into three panels. The *clustering specification panel* allows users to select one of four clustering methods and set its parameterization. After the clustering is executed, the clustering results are shown in the *cluster panel* and the *document projection panel*. In the *cluster panel* the clusters are represented by distinct colors from a categorical color map. The user has several options to explore the content of a cluster by showing (a) the documents in the cluster, (b) the most frequent terms or features in the cluster (tf) (again POS and named entity filters can be applied), (c) the terms or features most correlated to the cluster (χ^2), and (d) the cluster quality, represented by the compactness and separation of the cluster. Users can switch between these options on top of the *cluster panel*. The overall quality of the clustering can be derived from the average compactness and separation, and the overall Dunn and Davies-Bouldin indexes shown in the *clustering statistics panel*. To help the user to examine the clustering result with respect to the similarity of documents, we incorporated an additional visualization. The *document projection panel* in the middle of the Cluster Analysis View shows documents represented as circles projected onto a 2D plane. The projection is derived by executing an MDS on the document's feature vectors. The colors of the dots reflect the cluster affiliation. The projection attempts to minimize the error between the feature vector distances and the Euclidean distances in the 2D panel. Therefore, similar documents are shown close to each other in the visualization. Due to the reduction of a high-dimensional vector to a 2D vector a projection error is introduced to the visualization. Neighborhood preservation and trustworthiness measures shown at the top of the *document projection panel* make the users aware of this fact. By selecting one of the measures, the individual scores are mapped on the color of the document dots via a sequential grayscale color map. The *document projection panel* and the *cluster panel* are linked, documents highlighted or selected in one view are also highlighted in the

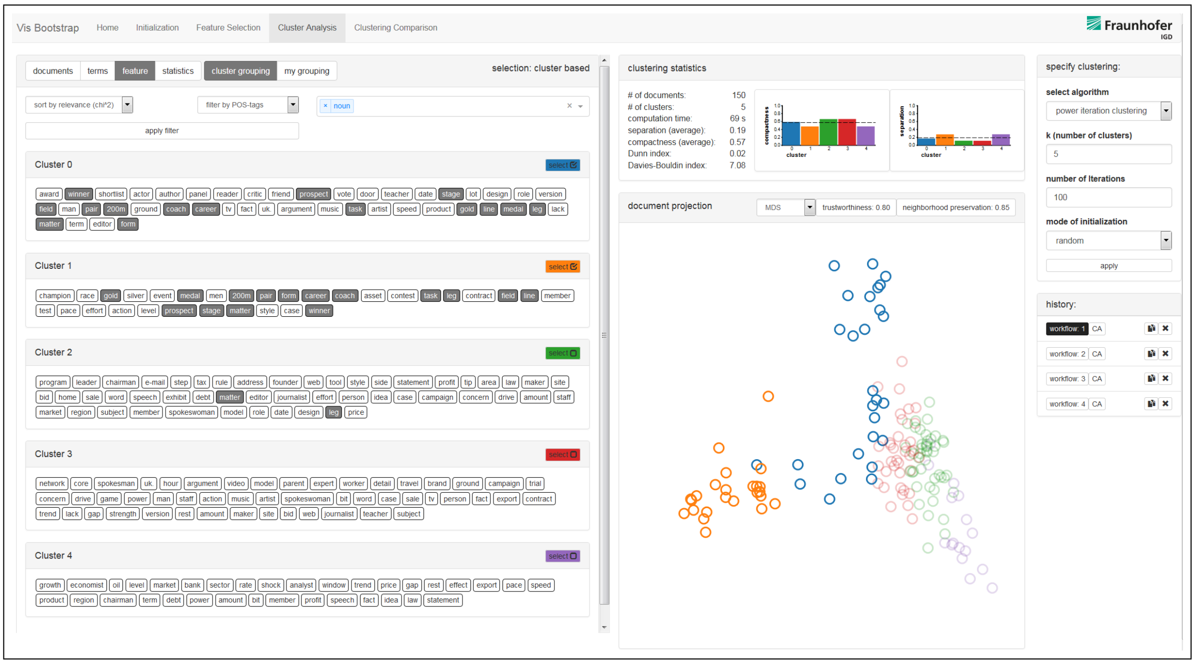

Figure 5.4.: Cluster Analysis View: **Clustering specification panel** (right): clustering algorithm and parameters can be defined. **Clustering statistics panel** (top): cluster quality measures, computation time, and number of documents and clusters are shown; two bar charts show the separation, and compactness per cluster. **Document projection panel** (bottom): documents are projected on display space based on selected projection method; colors represent cluster affiliation; trustworthiness and neighborhood preservation values are shown. **Cluster panel** (left): clusters can be represented by affiliated documents, by most frequent (or most correlated) terms, or by most frequent (or most correlated) features; POS- and named entity-filters can be applied to terms and features. Document and term highlighting supports the comparison of clusters. Here, blue and orange clusters are highlighted, most correlated features are shown.

other view. In Figure 5.4, the "blue" and "orange" clusters are selected. The respective document dots are highlighted in the *document projection panel*. Moreover, feature terms that occur in both clusters are highlighted. The Cluster Analysis View allows user to evaluate the quality of the clustering with respect to the clusters' content and based on the shown cluster quality measures ($\mathbf{R}_2$).

As an additional method to evaluate the quality of a clustering result, users can manually group documents into clusters. An example is given in on the left side of Figure 5.5. Here, all documents on "sports" are grouped into Cluster 0. The remaining documents are grouped into an "unassigned" cluster. Users can add further clusters by clicking on the plus button, and add documents to these clusters. The manual document grouping ("my grouping" tab) is used as a ground truth representing the mental model of a user. The quality of a clustering can also be evaluated by measuring how adequate the clustering represents the manual grouping of the user. Therefore, an F-measure is calculated and presented in the Clustering Comparison View, discussed in the following section.

Figure 5.5.: Clustering Comparison View: **Workflow statistics panel** (right): F-measure (if reference clustering is available), computation time, average separation, average compactness, Dunn and Davies-Bouldin indexes of different clusterings can be compared. **Cluster statistics panel** (top): clustering statistics of workflow on top of cluster comparison panel are shown. **Workflow comparison panel** (middle): clustering results (rows) are separated into clusters according to underlying document counts; order of clusterings can be adapted via drag-and-drop; color of bands represent clusters of clustering on top; band width represents number of documents in band. **Cluster intersection panel** (left): shows documents within a band selected in the cluster comparison panel; alternatively, most frequent or correlated terms can be shown. **History panel** (bottom right): created workflows are represented by ID, latest workflow step, copy and delete buttons; current workflow is highlighted, via the check boxes users can select workflows to be compared.

5.5.3. Clustering Comparison View

Finally, the Clustering Comparison View (see Figure 5.5 and Figure 5.7) allows the user to compare several clustering results generated by different clustering workflows ($\mathbf{R}_3$). In Figure 5.5 the clustering workflow results are compared to the reference clustering, in Figure 5.7 only the clustering workflow results are shown. The view is divided into three panels. In the *workflow statistics panel*, users can compare the quantitative cluster quality scores of different clustering results. Six bar charts show (i) the F-measure between the respective clustering and the user-defined “my grouping”, (ii) the computation time, (iii) the average compactness, (iv) the average separation, (v) the Dunn index, and (vi) the Davies Bouldin index of the different clustering results. The *workflow comparison panel* in the middle of the view shows a parallel sets visualization [KBH06] that allows users to analyze document-cluster affiliations over several clusterings. In the visualization, each row represents a clustering. The rows are divided into sections of differing lengths, representing the clusters and their respective sizes (number of associated documents). The order of the clusterings can be adapted via drag-and-drop. The clusters of different clusterings are connected via colored bands. The colors of the bands are defined by the clusters of the workflow shown on top of the view (in Figure 5.7, workflow 1). The colored bands between two rows represent a subset of documents that appear in both clusterings in a single cluster. This helps the user to see overlapping sub-clusters coming from different clusterings. The analysis of stable document subsets that appear in one single cluster in each clustering is supported. Details about these subsets are shown in the *cluster intersection panel*. Subsets can be selected by clicking on the respective band. The user can choose whether the documents or the most frequent (or correlated) terms in this document subset are shown.

5.5.4. Discussion on Design Decisions

In the previous section, we described our visual-interactive system including visualization techniques to support the exploration of document collections. In this section, we want to briefly discuss the design decisions and chosen visualization techniques. In the Feature Selection View, we needed visualization techniques to (a) show the quantitative feature selection metrics, and (b) the resulting features. We chose a range bar chart and a word list to address these issues, respectively.

Range bar chart (to visualize feature selection metrics): The purpose of the feature selection is to reduce the dimensionality of the vector space model while retaining the quality of the representation. Feature selection metrics can support users in removing non-informative features from the feature space. A visualization technique for the visualization of features selection metrics needs to fulfill the following requirements: (a) features should be sortable based on their feature selection metrics, (b) features should easily be selected or deselected, (c) the visualization should be scalable, since the vocabulary of the entire document collection needs to be represented, (d) the combination of different metrics should be supported. A simple sortable table that shows the features and their metrics is not scalable due to the high dimensionality of the vocabulary. Still, we want to keep the metaphor of a list. Hence, an aggregation of the features is needed. We aggregate the features into buckets of equal size. To visualize the scores, we decided to show the user the range (min-max) of the prevalent features selection metrics within a bucket. We also discussed alternative representations like box plots, or dot plots. However, for box plots and dot plots overplotting becomes an issue if the user increases the resolution to a large numbers of buckets. Dot plots with only one dot representing the average score in a bucket are a further alternative. However, if the resolution is high, it is difficult to spot them in the chart. Moreover, outliers are not covered by the average. For example, users cannot grasp the minimal score within a bucket, which could be important for defining score thresholds. We also discussed alternative aggregation methods. For example, by grouping features based on their feature selection metrics, histograms of the metrics' distributions could be shown. However, this would impede users to estimate the ratio of selected features, while our aggregation method explicitly shows the proportions on the vertical axis. Moreover, our aggregation method allows the combination of feature selection metrics, e.g., by sorting a metrics chart based on another metric. Other aggregation methods would not allow comparisons due to varying bucket sizes. The results of the feature selection in the range bar charts are presented in the word list.

Word list (to visualize features/terms): In several views, we needed a visualization technique to represent relevant features or terms: selected and not selected features in the Feature Selection View; most frequent terms and most correlated features in the Cluster Analysis View, and named entities within cluster intersections in the Cluster Comparison View. We needed a compact representation, since the available space in the view was limited. Moreover, the most relevant terms should be identified quickly. Although word clouds are a popular tool to visualize text, research has shown that simple tables are the better choice for identifying (a) the presence or absence of terms, and (b) most and least relevant terms (e.g., [OC10]). Due to the matter of space, we attempt to combine the benefits of both tables and word clouds. Therefore, we sort the terms based on their relevance, and show all terms with the same size to keep the structure of a list, but display them like a space-filling word cloud to reduce white space.

Bar chart (to visualize cluster quality metrics): In the Cluster Analysis View and the Clustering Comparison View, we needed a visualization technique to compare the cluster quality scores of different clusters and clusterings. Research has proven that bar charts are most appropriate for the comparison of quantitative data (e.g., [Few09]).

Scatterplot projection (to visualize similarity and cluster affiliation of documents): In the Cluster Analysis View, we needed a visualization that intuitively represents the similarity and the cluster affiliation of documents. For this purpose, a similarity matrix or a projection-based scatterplot visualization are appropriate choices. We decided to apply a projection-based scatterplot that shows the similarity between documents by their spatial distance, and the cluster affiliation via a color coding. While the precision of a matrix visualization is higher, a projection-based scatterplot offers a global perspective on the distances. Patterns in a matrix visualization are more difficult to interpret. In a projection view, the user can directly inspect the similarity between documents via the spatial distance. Moreover, it is easier to spot outliers in a cluster.

Parallel sets visualization (to visualize document intersections between clusters from different clusterings): Finally, for comparing different clusterings, we needed a visualization that was capable of showing stable cluster subsets that always appear in the same cluster, independently of the chosen clustering method. Related research can be found in the visualization and comparison of sets. A comprehensive overview of set visualization techniques has been recently published [AMA*16]. Out of the proposed techniques, we selected the parallel sets visualization, introduced by Kosara et al. [KBH06], since it fits best to our purpose and is still easy to comprehend. We also discussed the alternative of a heatmap matrix representing clusters of one clustering as rows, and clusters of a second clustering as columns. The cells could represent the number of documents in the intersection via a color map [AAMH13]. However, this view would only be suitable for comparing two clusterings. Finally, as discussed in the related work section Lyi et al. present an enhanced parallel sets visualization coloring the clusters across all clusterings and depicting stable subsets via edges between the clusterings [LKS*15]. However, we prefer to color the bands instead of the rows in order to better comprehend where specific cluster subsets can be found in the different clusterings.

5.6. Usage Scenario

In the following, we demonstrate the usefulness of our visualization system with a real-world dataset. The BBC dataset containing news articles in five topics from 2004-2005 serves as a test case [GC06]. We selected a random sample of 150 documents (30 per topic) from the collection. The topic labels are business, entertainment, politics, sport, and tech. Each document title consists of the underlying topic label and a document ID (e.g., tech22). Our system only uses the document content for clustering. In this usage scenario, the document titles and labels will help us to validate the clustering results. In an unlabeled dataset, the user will use his previous knowledge to assess the topics of the documents via their titles or content.

We will structure our analysis process as follows. First, we will select a small subset of the features extracted from the document collection in the pre-processing step. Second, we will run several clustering algorithms on the resulting feature space and analyze the derived clusterings in the document projection panel, observing whether similar documents are associated to the same cluster. We will choose the most promising clustering result and analyze the clusters via most frequent cluster terms and documents in the clusters. Third, we will compare the performance of our feature vector with alternative feature vectors. Therefore, the respective clustering results will be compared based on their internal cluster quality, and towards a manually defined reference clustering. The usage scenario just illustrates one possible way to use the presented system. The order of the analysis steps might be adapted and the process may have more or less iterations.

Feature Selection. As a first step in our usage scenario, we create a new workflow in the Feature Selection View as shown in Figure 5.3. We keep the pre-processing routines stemming, stop word removal, and punctuation removal activated, since they already help to decrease the feature space. To further reduce the size of the feature vector, we only use nouns as features by applying a part-of-speech filter on the features. The resulting feature quality metrics are shown in Figure 5.3. Instead of sorting all metric charts individually, we sort them based on the features' document frequencies (df). It can be seen that the document frequency and the term contribution charts show similar shapes. Still, the upper bucket in the term contribution chart also contains small values (as in the tf-idf chart). To focus on features with high scores throughout all metrics, we select the upper buckets in the tf-idf chart excluding the top bucket that also contains low scores. Here, the combination of metrics helps to remove features with low scores from the feature vector. The selection results in 238 selected features which is less than 10% of the total vocabulary with 2482 features. We will evaluate in a later analysis step whether this relatively small feature vector sufficiently represents the document collection. In the *feature panel*, we can get a first notion about the content of the dataset. Features ranked highest are "profit", "revenue", "advertisement", "analyst", "custom", which gives us the notion of a business dataset. However, it is difficult to estimate the coverage of a document collection by looking at a small

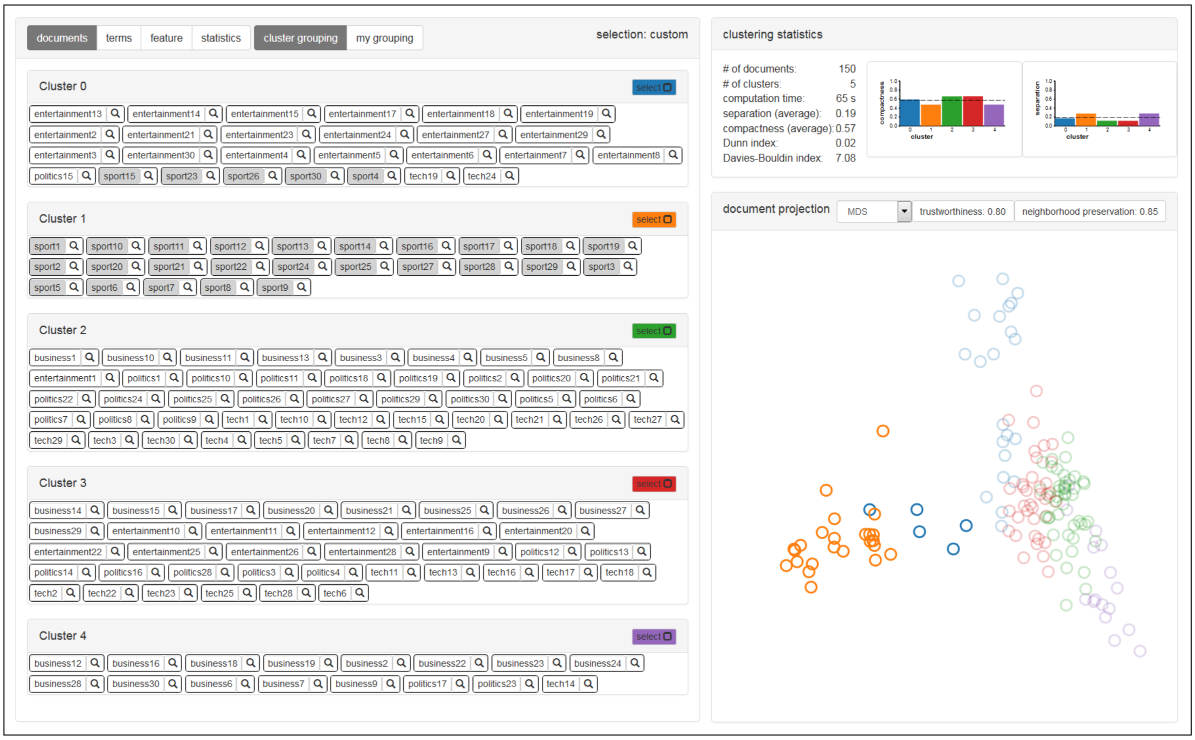

Figure 5.6.: Cluster Analysis View: in the cluster panel (left), the documents per cluster are represented by their title (in this dataset a combination of genre and an ID). In the document projection panel, documents horizontally separated from the others have been selected. These document are highlighted in the document projection panel (right) and the cluster panel (left). All selected documents are about sports. The “orange” cluster only contains documents on sport. The “blue” is a mixture of topics. The other clusters do not contain documents on sports.

excerpt of the vocabulary. To get a better overview of the dataset by grouping documents with similar content, we proceed to the cluster analysis step.

Cluster Analysis. We switch to the Cluster Analysis View (Figure 5.4), and execute several clustering algorithms, since we cannot say yet, which clustering algorithm will perform well with the selected features. In the *clustering statistics panel* and in the *document projection panel*, we can assess the separation and compactness of the resulting clusters. In this specific scenario, we prefer the clustering results achieved with the power iteration clustering, as shown in Figure 5.4. In the *document projection panel* the five clusters are depicted via five colors. The orange cluster seems to be well separated from the other clusters. Only a few “blue” documents are close to the “orange” documents. In the *cluster panel*, we inspect the most relevant features per cluster (derived from the χ^2 measure). By selecting the orange and the blue cluster, the underlying documents are highlighted in the *document projection panel*. The relevant features shared by both clusters are highlighted in the *cluster panel*. The orange and the blue cluster have several features in common. Relevant feature terms include “gold”, “medal”, "200m", and "coach". Further features occurring in the orange cluster are "champion", "race", and "contest". Without knowing anything about the dataset in advance, this gives us the notion that the orange cluster mainly contains documents about "sports". The "blue" cluster contains sportive terms, too, but also terms like from arts and entertainment like "award", "actor", "author", and "tv". From this, we learn that the blue cluster represents more than one theme. The "green" cluster contains political and business terms like "program", "leader", "chairman", "tax", "law". The "red" cluster contains a mixture of terms like "network", "model", "worker", "artist", "music", while the "purple" cluster includes several business terms like "growth", "economist", "oil", "market", "bank". We conclude that the orange and the purple clusters provide coherent topics, while the other clusters contain a mixture of themes. This can also be validated by looking at the separation scores of the respective clusters in the *clustering statistics panel*.

As a next step, to confirm our analysis results, we want to inspect the documents per cluster. Therefore, we switch from features to documents in the *cluster panel* (see Figure 5.6). All documents contained in the clusters are shown represented by their titles (in this scenario including topic labels). Our analysis results presented in the previous paragraph can be confirmed. The "orange" cluster only contains documents about sports. The purple cluster mainly contains documents about business. The blue cluster contains a mixture of sports and entertainment documents.

To further experiment with the dataset, we want to create a "sports" cluster that can be used as a ground truth for alternative clustering workflows. We use the interface depicted in Figure 5.5 to create a manual grouping with all sports documents in one group and the remaining documents in another group.

ID	description	size
1	noun features with best df scores w/o 1^{st} bucket	238
2	all nouns in vocabulary	841
3	all features in vocabulary	2482
4	intersection of best noun features per metric	243

Table 5.2.: Usage scenario: Defined workflows.

Clustering Comparison. Finally, we want to compare the performance of our selected vector space model (cf. Figure 5.3) with alternative models. Therefore, we create additional workflows by selecting different feature subsets in the Feature Selection View. Table 5.2 shows an overview of the selected features. In addition to the initial vector space model (workflow 1), we created alternative models using all nouns (2), all single token features (3), and an intersection of the noun features with the best values per feature selection metric (4).

For all of these vector space models we run a power iteration clustering with five clusters, which allows us to analyze the variation induced by different vector space models. In Figure 5.7, we compare the individual cluster groupings. The clustering results are rather stable. Hence, our initial feature selection including only 238 features (Workflow 1) produces similar results like the full vector space

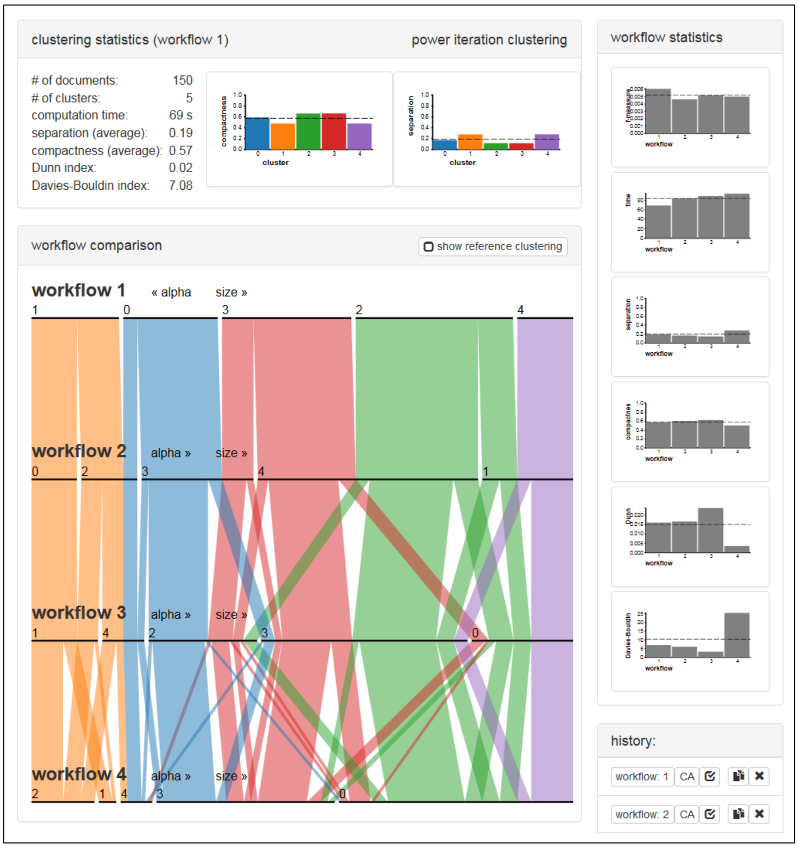

Figure 5.7.: Clustering Comparison View: The results of four different clustering workflows based on four vector space models are shown. The clustering results have several overlappings in the different clusters. We can assume that the clusterings are already rather stable. The documents contained in the first two orange and the first blue cluster bands can be inspected in Figures 5.8, 5.9, and 5.10. All of them only include documents about "sports".

model including all single term tokens (Workflow 3). Moreover, our feature vector consumes less computation time than the full feature vector (see second bar char on the right). By clicking on the first two orange and the first blue bands, we can inspect the documents contained (see Figures 5.8, 5.9, and 5.10, respectively). All of these documents are about "sports".

In Figure 5.5, we compare the clustering results to the reference clustering that we manually created in the previous analysis step. The orange band represents the documents of the reference cluster - documents on "sports". We can see that in our initial cluster workflow these documents are distributed over two clusters, while the other workflows dispense them into three clusters. By comparing the F-measures in the *workflow statistics panel* (see Figure 5.5 (top right)), we can conclude that Workflow 1 represents our reference cluster best. Therefore, we conclude that our selected feature vector represents the reference clustering best, while consuming less computation effort and memory space. To finalize our usage scenario, we report the following findings:

1. By combining POS-filtering and multiple feature selection metrics, we were able to select a relatively small feature vector that already provides satisfying clustering results.
2. Inspecting a sample of the selected features provided us a high-level view on the dataset. Deeper insights could only be gained by exploring the clusters' content.
3. The document projection panel helped us to get a first overview on the quality of the clusters. Overlapping clusters and well separated clusters were found immediately.
4. The workflow comparison panel supported the identification of stable document subsets throughout several clusterings. The clustering was effectively compared with the reference clusterings.

Figure 5.8.: First cluster intersection band (orange)

Figure 5.9.: Second cluster intersection band (orange)

Figure 5.10.: Third cluster intersection band (blue)

5.7. Summary

In this chapter, we presented a visual analytics system for the visual-interactive creation and validation of text clustering workflows with the purpose to explore document collections. The system serves as a proof of concept for the applicability of our general domain characterization (presented in Section 3.2.2) on textual data. First, we discussed current challenges in the exploration of document collections via document clustering. Second, we presented requirements on the design of a visualization system that addresses these challenges. The requirements were: (1) open up the design space for text clustering workflows by providing visual and interactive access to all workflow steps; (2) support the users' evaluation of the clustering results based on the clusters' content (relevant terms and documents) and cluster quality metrics; (3) also support the comparison of different clustering results on these levels; (4) provide a workflow history that allows to switch back and forth between different clustering workflows. Based on these requirements, we introduced a visualization system structured into three views to support each step of the text clustering workflow: (a) pre-processing and features selection, (b) cluster definition and analysis, and (c) clustering comparison. We also provided our design rationale discussing the choice of the visualization techniques integrated in the system. To underline the usefulness of the system, we showed its applicability in a usage scenario that targeted the exploration of BBC news articles. The system was mainly designed for analysts without prior knowledge in text analysis who need to create and analyze document clusters which are later used generate overviews of large document collections. Therefore, we identify three main challenges with potential for future improvements of our approach.

Usability. The presented visualization system was designed for analysts that have some experience in NLP, data mining, and statistics. The main goal was to open up the design space for text clustering workflows by providing visual access to crucial steps in the process. However, it would be desirable to reduce the system's complexity to allow a larger user group the exploration of document collections. As a possible solution to address this challenge, we could incorporate workflow patterns into the system that users can select and adapt to their specific scenario. The patterns could be extracted by replicating best practices in the configuration of workflow steps from related research. Additional meta-information should explain the specific characteristics and targets of the underlying configuration to the users.

Scalability. First, so far we tested our system mainly on document collections containing around two hundred documents. For the exploration of larger collections, the scalability of the *document projection panel* needs to be addressed. As an option, the documents could be represented by their cluster centroids. By zooming into the view, the documents contained in the cluster could be shown. Second, if the document collections become larger, the computation time will also increase. To improve the efficiency, and realize user interaction more calculations could be shifted into the pre-processing phase (if possible). This could also be achieved by calculating the workflow patterns mentioned in the previous paragraph in the pre-processing step.

Expandability. Finally, the visualization system could be expanded by further feature selection metrics, clustering and projection algorithms, cluster quality metrics, projection error measures, and visualization techniques. Moreover, an interface for integrating external algorithms into the system could be envisioned. This would allow researchers to test and evaluate new algorithms with the system.

6. Visual-Interactive Access to the Public Debate

In this chapter, we present a visual analytics approach for monitoring and analyzing online discussions on predefined topics (Challenge $\mathbf{C}_{\text{Deb}}$). The system allows the exploration, analysis, and comparison of textual statements on specific topics and underlying arguments with respect to their relevance. Regarding the overall approach of this thesis, we prove the applicability of our concept to stakeholder opinions and arguments in the form of textual data. Figure 6.1 shows how the design methodology presented in Chapter 3 is applied in this specific scenario with textual data (Challenge $\mathbf{C}_{\text{VDSS}}$). First, the approach supports the decision makers, analysts, and stakeholders in the exploration of predefined policies, arguments, and policy terms with respect to their relevance represented by the number of occurrences in online media. Users can explore total numbers and their temporal distribution. Second, the visual analytics system allows to extract new arguments that can be included into the domain model. These arguments are re-used for supporting pre-defined policies, or for creating new policies. Third, the automatic extraction of sentiments from text documents and the drill-down to the original data support the analysis of individual policies and arguments in detail. Fourth, users can compare the relevance

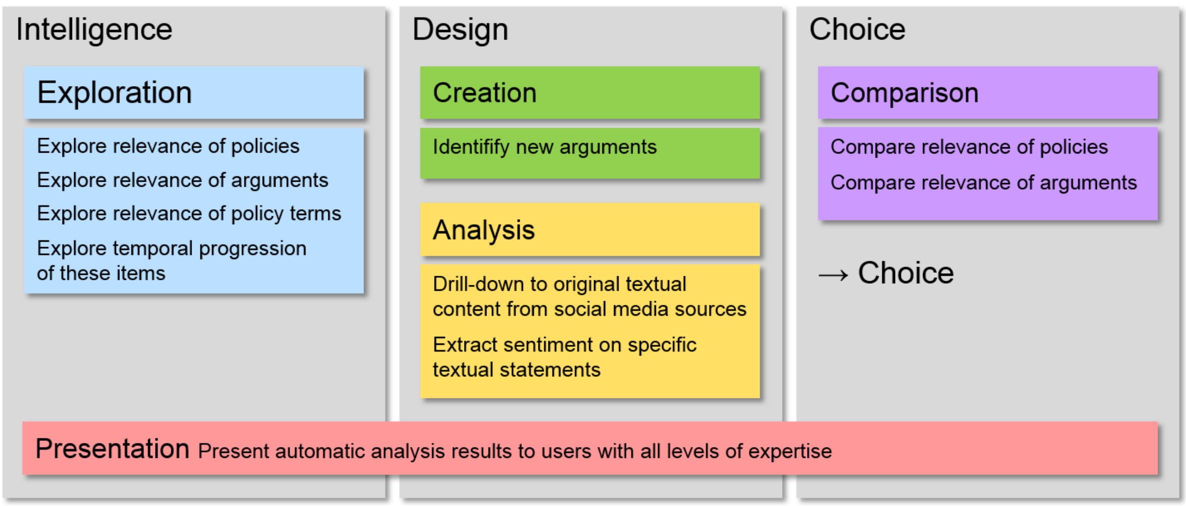

Figure 6.1.: Design methodology applied to **textual data** as third proof of concept. Visual analytics system supports the exploration, analysis, and comparison of relevance of alternative solutions, arguments, and sentiments in the public debate. New arguments can be identified and included in the decision making process. Visual access to the public debate for all expertise levels. Figure is adapted from Figure 3.6.

of several policies and arguments in order to decide which policy to choose and support with which arguments. Finally, the described datasets are presented in an intuitive way, which allows users with all levels of expertise to incorporate data in their decision process. Besides the providing visual access to the public debate extracted from online media, our system bridges knowledge gaps between different stakeholders involved in the process (Challenge $\mathbf{C}_{\text{BKG}}$). In general, the visual analytics system is targeting decision makers and analysts who can use the extracted information in the decision process. In addition, the underlying text analysis model, designed by a modeling expert is refined by direct user feedback through interaction with the system. This enables the bridging of knowledge gaps between users and modeling experts. Moreover, the pre-defined policies and arguments are modeled by a domain expert into an ontology. The visual analytics system allows these domain experts to analyze the performance of the text analysis components on their model and refine the model on demand. This bridges knowledge gaps between modeling expert and domain expert. This chapter is partially based on our previous work published in [RBLT*15].

Contents

6.1. Introduction

Today, policy makers are requested to integrate large amounts of external knowledge and public opinions in their decision making processes. Large parts of the information to be considered are available on the web. However, the manual monitoring and analysis of this data is time-consuming and therefore not applicable in real-world scenarios. Automatic methods for the mining and analysis of textual content exist. To make use of these powerful tools, some challenges need to be tackled. First, the methods need to be combined to a workflow and adapted to the addressed domain. Second, the results of the workflow need to be presented to policy makers in an intuitive way. Third, since the accuracy of text analysis methods is not necessarily satisfying the users' expectations, concepts for feedback loops need to be considered.

Our contributions to tackle these challenges are:

1. As a baseline for our approach, we describe a text analysis workflow targeted on the domain of political decision making. The workflow applies text analysis methods that extract segments

from a crawled document collection and associates them with predefined political concepts – policy models and arguments.

2. We present a visualization dashboard designed for the presentation of text analysis results. The dashboard helps policy makers to access the results in an intuitive way.
3. We extend the workflow with visual-interactive feedback concepts that enable the users to improve the accuracy of the text analysis results. As a result, we increase the credibility of the system.
4. We implemented our approach in a real-world environment during a European research project to prove its applicability for political decision making.

6.2. Related Work on Document-Level Text Analysis

In the following, we discuss related work in the fields of text analysis and visual text analysis that target the analysis of individual text documents.

Text Analysis. A general introduction of text mining (also including some visualization examples) is presented by Feldman and Sanger [FS06]. Liu provides a comprehensive work about data mining techniques for the extraction and analysis of textual data from the web. This includes topics like crawling, opinion mining, and sentiment analysis [Liu07]. A general introduction to opinion mining and sentiment analysis is provided by Pang et al. [PL08]. Argumentation mining as a research field is relatively young. Relevant works can be found in Teufel's thesis [T*00], and further approaches presented by Palau et al. [PM09], and Feng and Hirst [FH11]. The state-of-the-art report by Jones describes recent approaches about the automatic summarization of text documents [SJ07].

Visual Text Analysis. During the design of our visualization dashboard we followed Few's suggestions for the visualization of quantitative data [Few09]. Prominent examples for visual analytics approaches related to sentiment analysis are proposed by Liu et al. [LHC05] and Chen et al. [CISSW06]. Oelke et al. visualize feature-based opinion clusters for Amazon products [OHR*09]. The Document Cards approach describes a technique for the summarization of single documents [SOR*09]. A comprehensive work about practical techniques for visualizing arguments is introduced by Kirschner et al. [KBSC03]. A visual analytics approach for analyzing social media content from Twitter and Youtube is presented by Diakopoulos [DNKS10]. Although these approaches are related to ours, most of them are mainly monitoring tools, feedback concepts are not considered.

6.3. Visual Analytics Design

Our approach tackles tasks that evolved during a European research project together with political decision makers. The tasks can be summarized as follows:

R_1 Extract the relevance of policy models, policy arguments, and policy terms in online discussions.

R$_2$ Analyze their relevance over time and per source.

R$_3$ Analyze the sentiment of extracted text segments.

R$_4$ Get access to the original textual content and sources.

R$_5$ Identify new arguments.

Our approach operates on individual text segments extracted from a large document collection. A text segment comprises one to several consecutive sentences. A query for extracting the documents and the segments from the collection is created from a concept graph representing user interests. This predefined graph relates policy domains (e.g., energy, transport, etc.), policy models (e.g., renewable energy directive, etc.), and arguments. The role of the concept graph is twofold. Firstly, it implicitly defines a document query by using search keywords from the political concepts. Secondly, it is used to structure the text segments, based upon the user's understanding of the policy domain. For more details about the graph and its editing process we refer to Spiliotopoulos et al. [SDK14]. The document collection is crawled from the web and comprises textual statements from newspaper sites, social media platforms, blogs, etc.

6.3.1. Text Analysis Workflow

We present a text analysis workflow (see Figure 6.2) that was designed and implemented to tackle the tasks described in the previous section. The individual text analysis modules that constitute the workflow are explained in the following. More details about the underlying linguistic pipeline are provided by Komourtzis et al. [KGP*14]. As described above a crawled document collection and a predefined concept graph serve as a prerequisite for our approach.

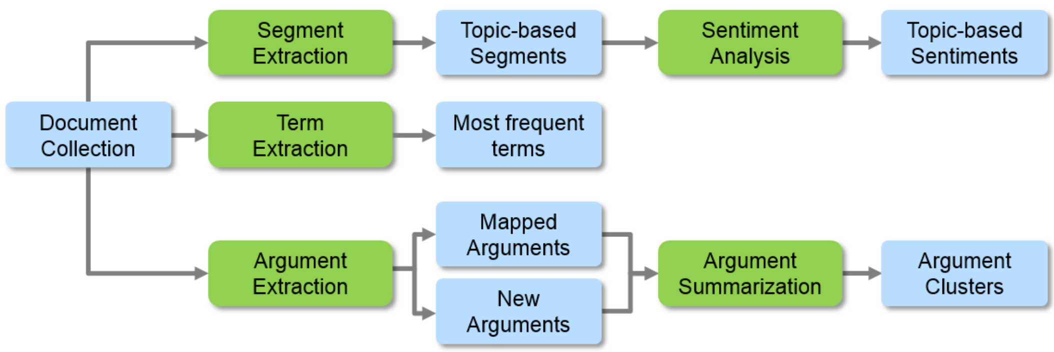

Figure 6.2.: Text analysis workflow including text analysis modules (green) and intermediate results (blue).

Segment Extraction. The segment extraction module analyses the document collection and extracts text segments related to the definitions provided by a given policy model, or argument (**R**$_1$). The module classifies text segments as associated or not associated with the given political concept.

Sentiment Analysis. The sentiment analysis module operates on two distinct modes. In the first case, the sentiment score of a text document is calculated. In the second case, the module extracts the sentiment from individual segments with respect to the associated political concept. Hence, for an entire document the general sentiment score is calculated. For the individual segments the topic-based sentiment scores are calculated ($\mathbf{R}_3$). The applied sentiment analysis methods were introduced by Petasis et al. [PSTT14].

Argument Extraction. The argument extraction module identifies argumentative sentences ($\mathbf{R}_1$,$\mathbf{R}_5$). Text segments that pertain an argument structure, e.g., containing claims and premises, are extracted from the document collection and classified as arguments. In a second step the extracted arguments are mapped on the predefined arguments in the concept graph. The argument extraction methods applied in this approach are presented by Goudas et al. [GLPK14].

Argument Summarization. The main purpose of the argument summarization is to extract arguments that are not yet defined in the concept graph ($\mathbf{R}_5$). It shall help the users to improve the existing policy models. The module forms argument clusters either based on (1) existing mappings of extracted arguments on predefined arguments (cf. argument extraction), or (2) based on textual similarities. In the latter case, a representative is chosen that summarizes the content of the new argument cluster.

Term Extraction. The term extraction module discovers the most frequent terms found in a document collection ($\mathbf{R}_1$). The module is not restricted to terms directly relevant to the category (as these are more useful for classifying the content), but rather discovers and presents terms that are frequently used within the context. (e.g., 'wind farm' is a term related to the domain 'energy', while 'noise', or 'efficiency' are terms that are common in discussions under the 'energy' category, thus they denote issues that must be taken into account when constructing a policy).

In summary, the text analysis workflow extracts text segments from a large document collection and associates these segments with the predefined policy models, and arguments. Moreover, argumentative segments are identified, mapped on predefined arguments, or clustered and marked as potentially new arguments. For all documents the general sentiment scores are calculated. For extracted text segments the sentiment score towards the associated policy model, or argument is calculated. Finally, for all subgroups of the document collection the most frequent terms are calculated.

6.3.2. Visualizing Text Analysis Results

The visualization module of our visual analytics system was designed as a dashboard with the goal to present the text analysis results to the users in an intuitive way (see Figure 6.3). Since most of the users do not have an IT background, we chose familiar visualization techniques. The dashboard is divided into three areas: a navigation panel (Figure 6.3(1)), a statistics panel (Figure 6.3(2)(3)(4)), and a text segment panel (Figure 6.3(5)). As denoted by the legend, in all views the color reflects the sentiment score (green = positive sentiment, yellow = neutral sentiment, red = negative sentiment) ($\mathbf{R}_3$). The size of the visual objects reflects the number of extracted segments, hence, the concept's relevance ($\mathbf{R}_1$). The navigation panel consists of a hierarchical topic selection menu. The menu represents the political concept graph described in Section 6.3. The user can select (a) a policy domain to get details about un-

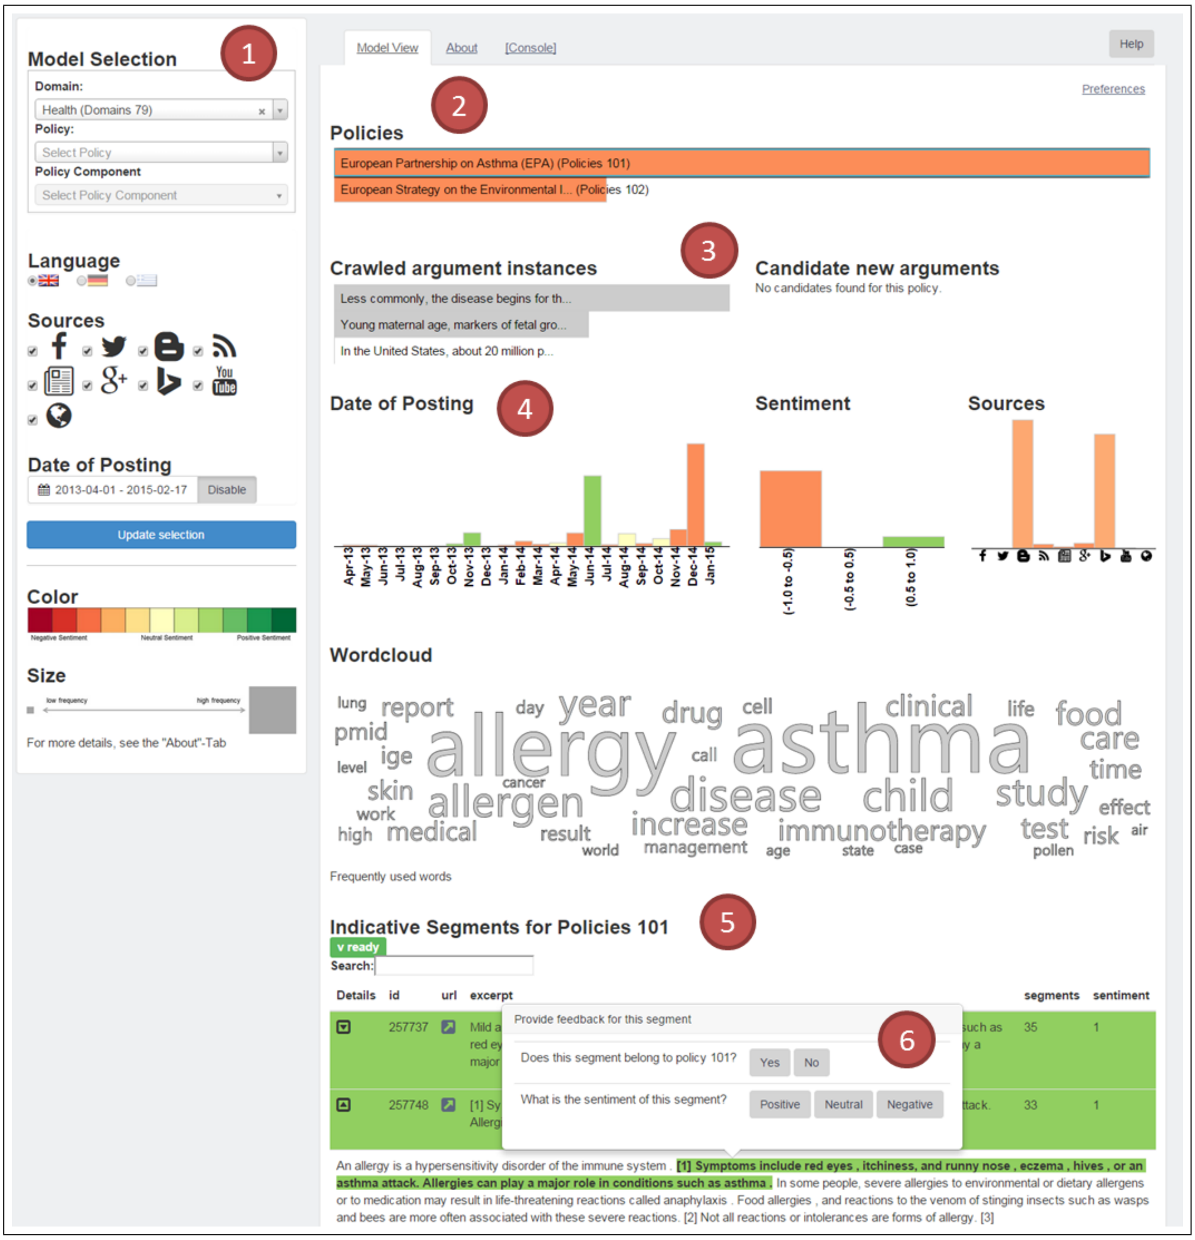

Figure 6.3.: Visualization dashboard.

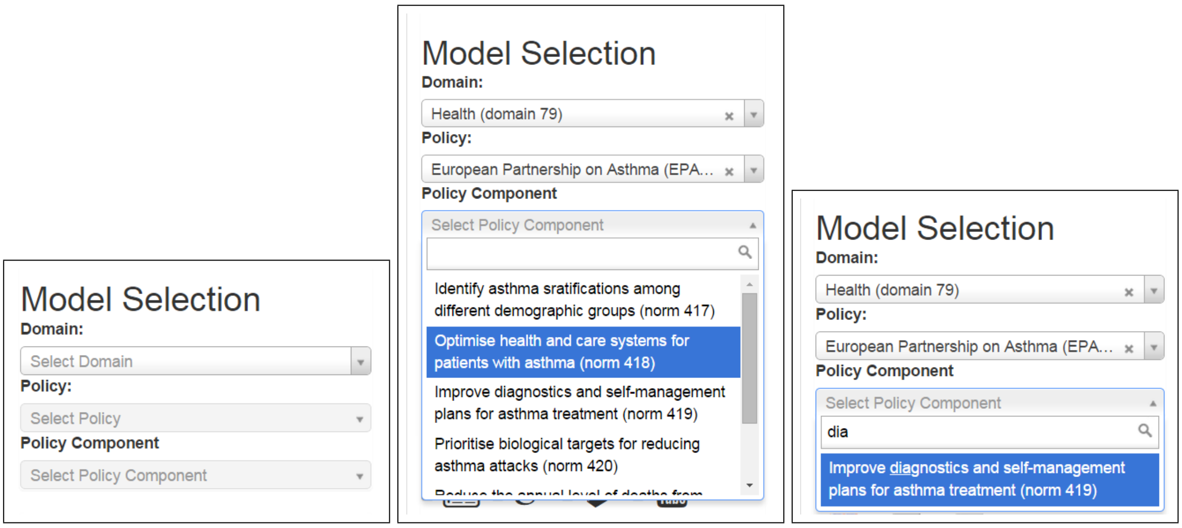

Figure 6.4.: Model Selection. Searching models can be done "top-down" by first specifying the domain, then the policy and optionally the policy component. This is the entry point for the analysis of results about the specified topic.

derlying policy models, (b) a policy model to get details about underlying policy components (parts of a policy model), and (c) a policy component to get details about underlying arguments (see also Figure 6.4). These details are shown in the statistics panel. The underlying policy models, policy components, or arguments are displayed in a sorted bar chart (see Figure 6.3(2) or Figure 6.5). This enables the user to get a quick overview about the relevance of the political concepts ($\mathbf{R}_1$) and the overall sentiment ($\mathbf{R}_3$). An additional bar chart shows the extracted argument clusters separated into predefined (left) and potentially new clusters (right) (Figure 6.3(3))($\mathbf{R}_5$). Further statistical information includes the temporal distribution of underlying text segments, the distribution per web source, and a sentiment distribution (Figure 6.3(4))($\mathbf{R}_2$,$\mathbf{R}_3$). An additional word cloud provides users an idea about the discussed textual content. A tabular view provides the original text documents including the extracted and highlighted segments to the user. (Figure 6.3(5))($\mathbf{R}_4$). Finally, the queries can be refined based on language, web source, and date of posting filters.

6.3.3. User Feedback

In general, it cannot be assumed that the results of text analysis processes perfectly match with human understanding of a domain. To mitigate this problem, the user is able to refine the results by giving incremental feedback on policies, arguments, segments, sentiments and their proposed relations. Feedback is generally triggered from a sub-menu by selecting a corresponding visual representation. For all modules feedback can be given in at least two ways: (1) through a general validation or approval of

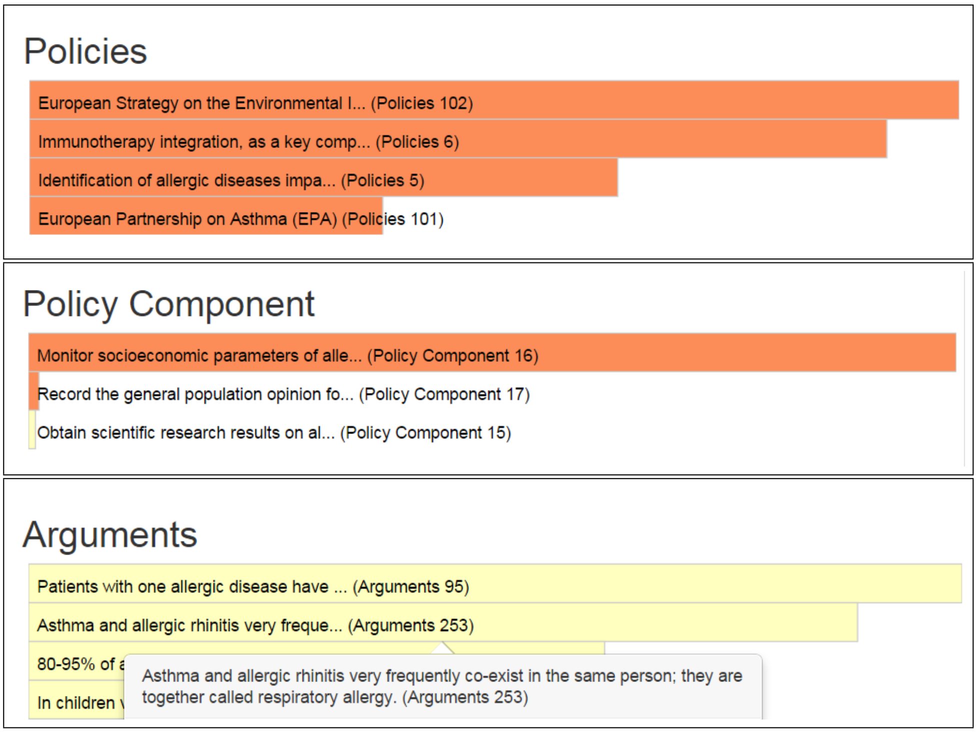

Figure 6.5.: Model View. Shows the results of the selections made in Figure 6.4.

the implied relationship. (2) through a manual correction. In any case, the user feedback is collected in a database for the refinement of the analytical models. Because the text corpus is too large for interactive adaption, the actual modification of the models is done in an offline process on a regular basis. For the ongoing session, the feedback only changes the specified modules. In the following, we will describe the type of feedback for every text analysis module presented in Section 6.3.1. An exemplary visual-interactive feedback concept is shown in Figure 6.3(6).

Segment Extraction. For the validation of the segment extraction, a user can indicate that a text segment is, in fact, relevant for a policy model (cf. Figure 6.3(6)). For a correction, a user may attach the text segment to another policy model.

Sentiment Analysis. For the validation of the sentiment analysis module, the user is able to feed back whether the sentiment scores are correct or not. If this is not the case, the user may adjust the scores (cf. Figure 6.3(6)). Corrected sentiment scores for documents or segment-topic-pairs are included into the training corpus of the sentiment analysis module.

Argument Extraction. Concerning the argument extraction module, a user has three options for feedback: Because not all segments might in fact be arguments, the feedback includes the validation whether a specific segment can be accounted for an argument at all. Arguments are identified by their similarity to predefined ‘template’ arguments. A user may specify whether this association is valid or possibly suggest another predefined argument. Finally, the user can rephrase a new argument and add it to the corpus of predefined arguments to capture a new aspect or to better distinguish between different predefined arguments.

Argument Summarization. With respect to the argument summarization module the user may approve the grouping of extracted arguments or remove outliers from their respective groups. In addition, the user may associate a similarity-based argument cluster with an existing predefined argument, or phrase a new argument, that describe the argument cluster and add it to the corpus of predefined arguments.

Term Extraction. As a possible user feedback, terms that are automatically extracted from a textual corpus can be excluded from the display. This might be feasible for terms that are obvious for a given domain and should not be highlighted anymore. As an example the term ‘energy’ does not provide any helpful insights in the energy domain, while the term ‘efficiency’ would. Therefore, the exclusion of terms from the most frequent term list would be a valuable user feedback that can improve the quality of the term extraction module. From a technical perspective, the terms to be deleted could be added to a user-defined stop word list.

6.4. Summary

In this chapter, we presented a visual text analysis system applied to the political decision making domain. The system extracts text segments from the web and associates them with predefined policy models and arguments. It combines a text analysis workflow with a visualization dashboard with the focus to facilitate the access to text analysis results. In addition, we introduced concepts that enable the user to provide direct feedback on results. These concepts help to improve the accuracy of individual text analysis modules and increase the credibility of our system. The system showed how to apply the concept of this thesis (Chapter 3) on textual data.

7. Visual-Interactive Access to Performance Indicators in the Mining Sector

In this chapter, we prove the applicability of our concept to empirical data. We present a visual analytics approach targeting the provision of visual access to empirical performance indicators in the mining sector to different stakeholder groups (Challenge $\mathbf{C}_{\text{Dat}}$). First, governmental decision makers are supported in political decisions on how to improve the governmental factors in mining. Second, investors are supported in their decision in which country's mining sector to invest. And third, the governmental deficits are made transparent to the civil society which supports public stakeholders' decisions at future elections. Figure 7.1 shows how the design methodology presented in Chapter 3 is applied in this specific scenario with empirical data (Challenge $\mathbf{C}_{\text{VDSS}}$). Our approach explicitly supports the tasks exploration, analysis, comparison, and presentation, and implicitly the creation task in the decision making workflow. Users are enabled to explore the performance indicators of several countries. Ana-

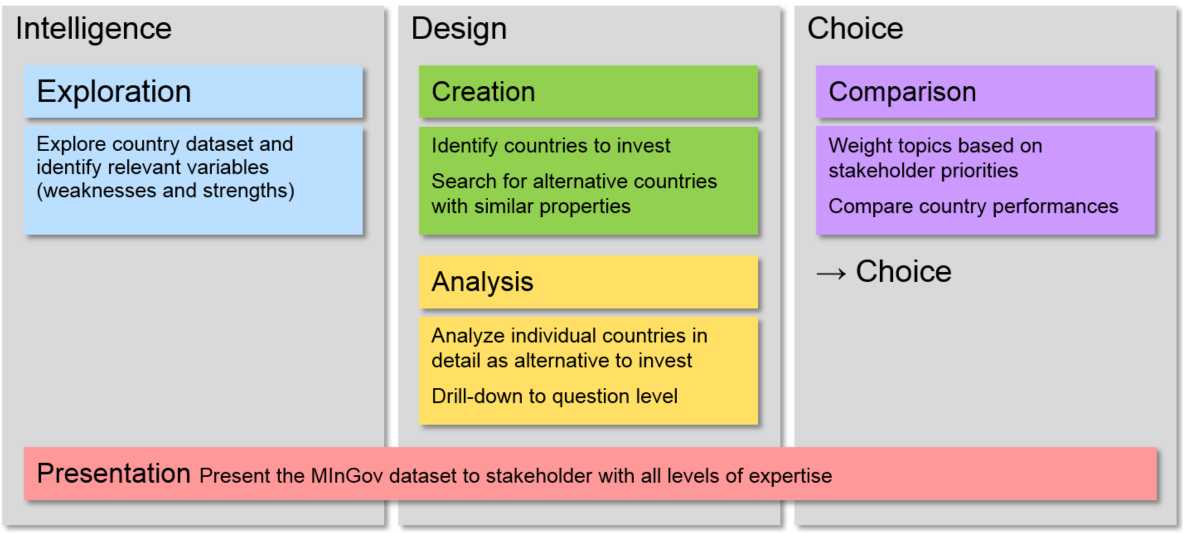

Figure 7.1.: Proof of concept for the applicability of our design methodology to **empirical data**. Exploration, creation, analysis, and comparison tasks are supported. Via the presentation task stakeholders with all levels of expertise are enabled to access the data. The system supports investment-related decisions in the mining sector. Figure is adapted from Figure 3.6.

lytical functionality like drill-downs and similarity search are supported. Moreover, users can compare several countries based on individual country scores and stakeholder based priorities. Finally, the visual design is evaluated by real-world users who attested the usability of the visualization system. This enables different stakeholders to use our approach for the presentation of the underlying data. In fact, the incorporated export functionality allows the inclusion of visualizations from the visualization system into written country reports. The visual analytics system mainly targets decision makers and analysts as users. However, the underlying dataset is collected partially via interviewing domain experts in the mining sector. Moreover, public stakeholders are enabled to access country datasets, and thereby, monitor the government performances towards the mining sector. Our approach implicitly supports the bridging of knowledge gaps between these stakeholders which addresses Challenge $\mathbf{C}_{\text{BKG}}$. This chapter is partially based on our previous work published in [RBB*17].

Contents

7.1. Introduction

The mining sector is one of the key industries ensuring economic growth in many resource-rich developing countries. As Martin Lokanc from the World Bank Group says, “Mining projects, when well-managed, offer an opportunity to transform resource wealth into sustainable development in many poor countries” [Wor17]. Several stakeholders play an important role for realizing a sustainable development. *Governments* need to provide regulatory conditions in the mining sector that attract *investors*, while in parallel serving the wealth of the country’s *civil society*. As a response to the growing demand of these stakeholders for understanding country-specific mining conditions, the World Bank Group has designed a methodology for collecting a dataset describing these conditions: The Mining Investment and Governance Review (MInGov) [Wor17]. The resulting dataset allows investors, governments and the civil society to assess the governance quality as well as the sector competitiveness and attractiveness of the respective country. In 2016, MInGov was conducted in seven countries. The target is to cover all mining-related countries and periodically collect data to monitor progress. The results are published in individual country reports and as open datasets [Wor17]. However, stakeholders requested an intuitive access to the complex datasets to make them usable. Moreover, the charts to be included in the country reports have to be generated manually in a tedious process.

The goal of this design study is to support decision making in the mining sector by providing visual-interactive access to MInGov data. The domain characterization with the involved user group revealed three core requirements: The visualization system has to allow (a) the exploration of individual country datasets, (b) the comparison of different country datasets, and (c) the export of charts for the inclusion in individual country reports. The contributions of our approach are threefold. First, we provide a domain characterization describing the underlying data, relevant user groups, and the tasks to be addressed with the data. Second, we introduce a visualization system that allows the identified users to conduct the described tasks in an effective and efficient way. Third, we present the results of a usability workshop conducted with real-world users demonstrating the efficiency and effectiveness of the system.

7.2. Background on Commercial Visualization Systems

We reviewed commercial visualization systems towards their suitability for our approach. Examples include Tableau, Tibco Spotfire, Microsoft Power BI, QlikView, and SAS Visual Analytics. An extensive review of these systems was already provided by Mittelstädt et al. [MBW*12]. However, the specific structure of the dataset and the tasks to be supported required some customized visualization techniques. This includes techniques providing (a) topic overviews in a matrix-like structure, (b) drill-down functionality to the question level, and (c) overviews of topic scores and weights via a specific treemap layout. None of the reviewed visualization systems includes these techniques.

7.3. Domain and Problem Characterization

We describe the three ingredients to be analyzed prior to the design of a visualization system: data, users, and tasks [vW13]. The results built the basis for our user-centered design approach.

Data. The foundation of the MInGov dataset is a questionnaire containing 350 questions about a country's performance in the mining sector. These questions are answered by domain experts, desktop research, or secondary sources. The question answers are aggregated and categorized into one of four ordinal classes (very low, low, high, very high) representing the country's performance on the respective question. On top of the questions a hierarchy was defined, which we exploit for the design of overview visualizations (see Figure 7.2). Questions are grouped into indicators, indicators are grouped into topics, and topics are grouped into themes. Some topics are additionally organized into groups along the mining value chain. This allows us to map them in a matrix-like structure, with themes on the vertical, and value chain levels on the horizontal axis (see Figure 7.3 upper matrix). Finally, four stakeholder-specific topic weights are provided by the respective domain experts to reflect (1) government priorities, (2) investor priorities, (3) civil society priorities, and (4) country-specific priorities. More details on the dataset are given at the MInGov website [Wor17].

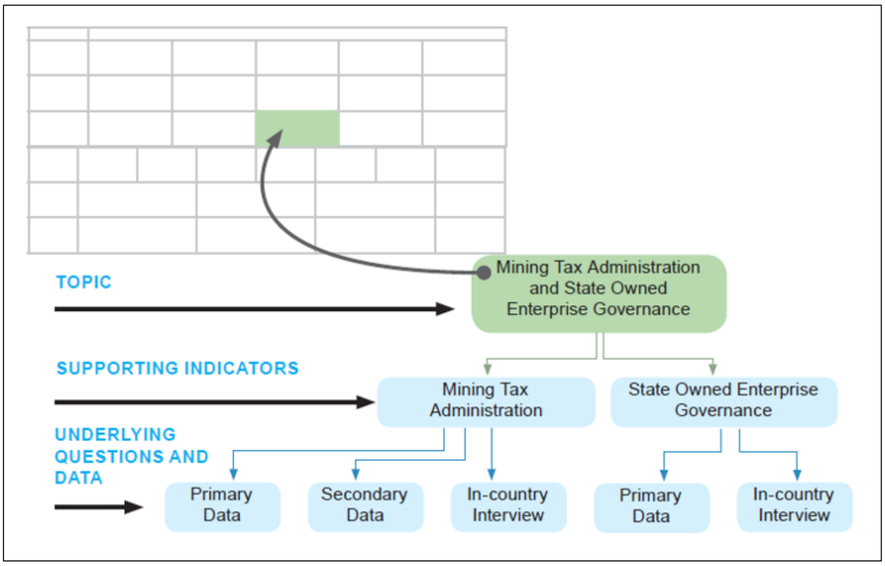

Figure 7.2.: The MInGov dataset structure extracted from the MInGov website [Wor17].

Users. We differentiate between three stakeholder groups: governments, investors, and civil society. The main goal of the government is to attract investors in order to generate economic growth. Therefore, they have to identify weak indicators that motivate the adaptation of regulations or the design of new policies. The main objective of investors is to make better investment decisions. They have to identify negative indicators to anticipate challenges prior to investments. Finally, civil society organizations aim at monitoring the mining sector performance. The provided transparency increases the understanding of critical government decisions. Moreover, civil projects can help to attract mining sector investments.

Tasks. Together with domain experts from the World Bank Group, we identified several tasks to be conducted with the MInGov dataset. The tasks and their order match the decision making process described in the concept of this thesis (Section 3).

R_1 Explore a country dataset and identify relevant variables (weaknesses and strengths) (*intelligence*)

R_2 Analyze individual countries in detail (e.g., as alternative to invest); allow drill-down to question level (*design*)

R_3 Search for alternatives with similar properties (*design*)

R_4 Weight topics based on stakeholder priorities (*choice*)

R_5 Compare countries and make a choice (*choice*)

7.4. Visual Analytics Design – Visual MInGov

Our visualization system was designed to support decision making in the mining sector in an efficient and effective way. Two views support the exploration of individual country datasets ($\mathbf{R}_1$): the Matrix View (Section 7.4.2) and the Statistics View (Section 7.4.4). Both allow a drill-down to the question level ($\mathbf{R}_2$). A Similarity Search (Section 7.4.7) enables users to search for similar countries ($\mathbf{R}_3$). The Weighted Matrix View (Section 7.4.3) incorporates topic weightings based on stakeholder priorities ($\mathbf{R}_4$). Different countries are compared with the Country Comparison View (Section 7.4.6) ($\mathbf{R}_5$).

7.4.1. Color Map

In the visualization system, we map the country performance scores on a discrete color map representing the four performance categories very low, low, high, and very high. According to the domain experts, two sequential, colorblind, and print-friendly color maps were needed - one for the Mining Sector Importance topics and one for the other topics. Since, The Mining Sector Importance topics do not provide insights in the country's performance, they need to be distinguishable from other topics. Based on these requirements, we selected two color maps with the ColorBrewer [HB03]: one multi-hue for the performance scores and one single hue for the Mining Sector Importance scores (see Figure 7.3 (legends at the bottom)).

7.4.2. Matrix View

The Matrix View (Figure 7.3) gives an overview of the topic scores of an individual country dataset ($\mathbf{R}_1$). According to the domain experts the topics have a "natural" order which we consider in the visualization design. The questions in the dataset are grouped into 36 topics (colored cells and gray-scale bars on the right). The topics are grouped into seven themes (A-F, M). Some of the topics are additionally associated to the five levels of the mining value chain (1-5). These topics are organized in a matrix-like structure with the themes (A-C) as rows and the levels in the value chain (1-5) as columns. The remaining themes (D-F) are displayed separately below the matrix. The domain experts requested a distinguishable design for the "Mining Sector Importance" theme, since the underlying data does not provide information on the country's performance. Mining Sector importance topics are visualized via range charts depicting the minimal, maximal, and average question value within a topic (Figure 7.3 right). The "Mining Sector Importance" scores are depicted via a gray-scale color map, while the other topic scores are depicted via a multi-hue color map (see Section 7.4.1). The Matrix View replicates the structure of the dataset defined by the domain experts. This structure does not allow an alternative display of the MInGov data. However, by removing one dimension in the upper matrix (value chain or themes), we obtain a tree that allows the application of alternative visualization methods. We integrated an icicle plot into our system. Users could choose between the value chain and the theme perspective, and then drill-down from the value chain (or theme) to the question level ($\mathbf{R}_2$). Individual scores were

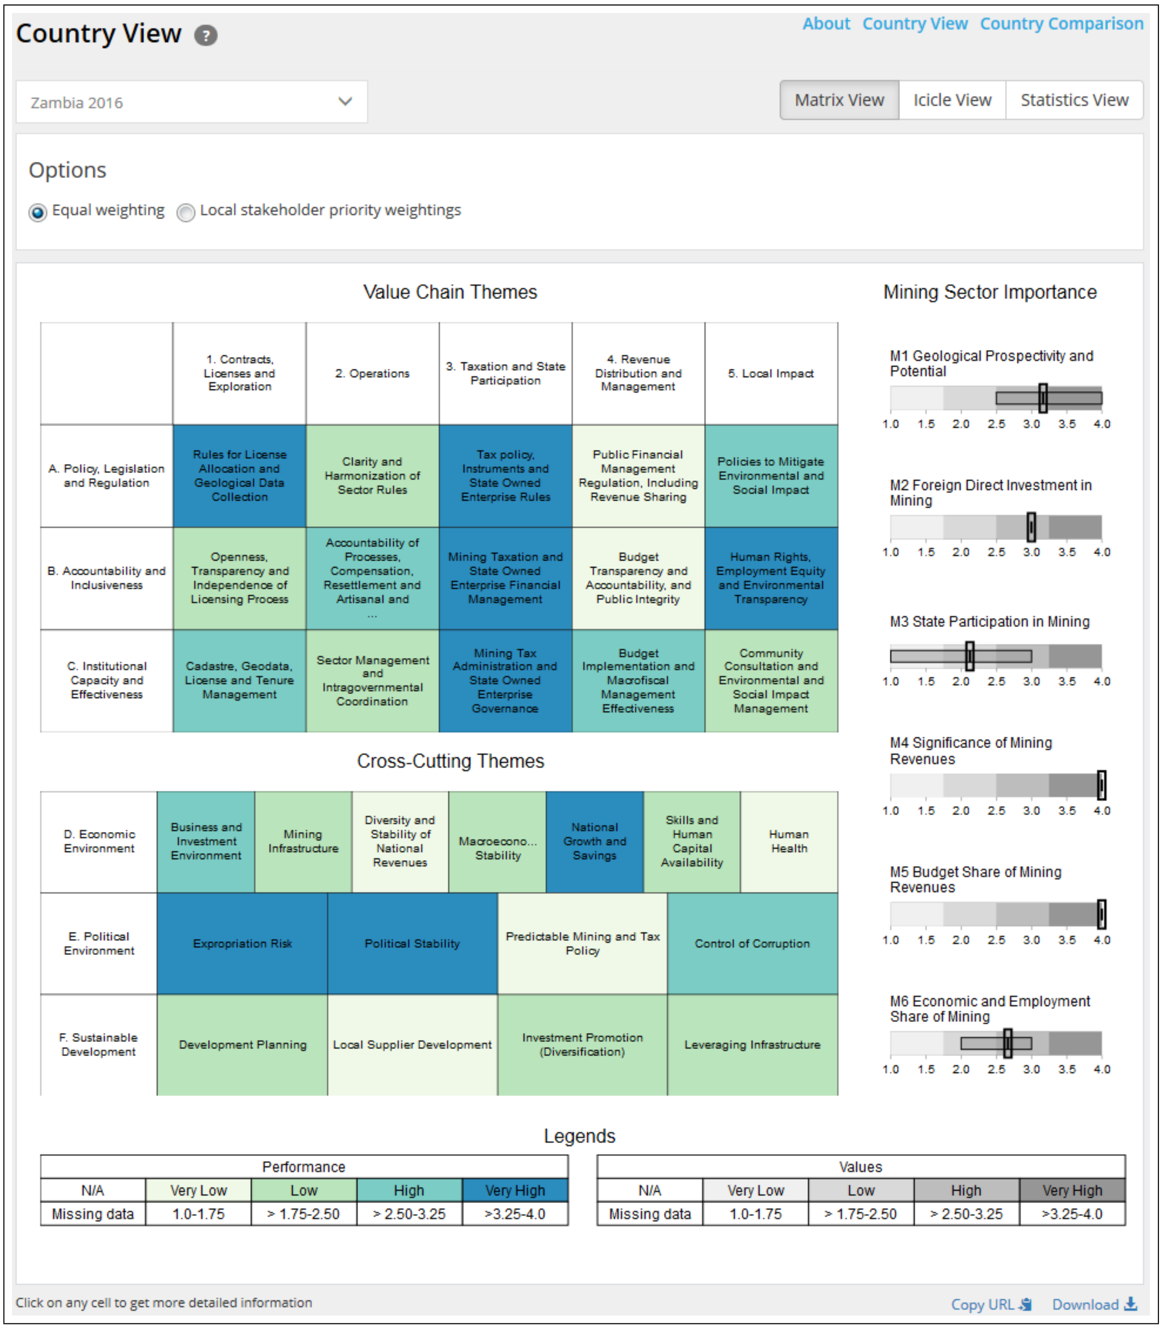

Figure 7.3.: Matrix View: topics are grouped into the theme-value chain matrix (top left), cross-cutting themes (lower left), and the Mining Sector Importance theme (right). The visualization can be exported via the 'export image' link (bottom right).

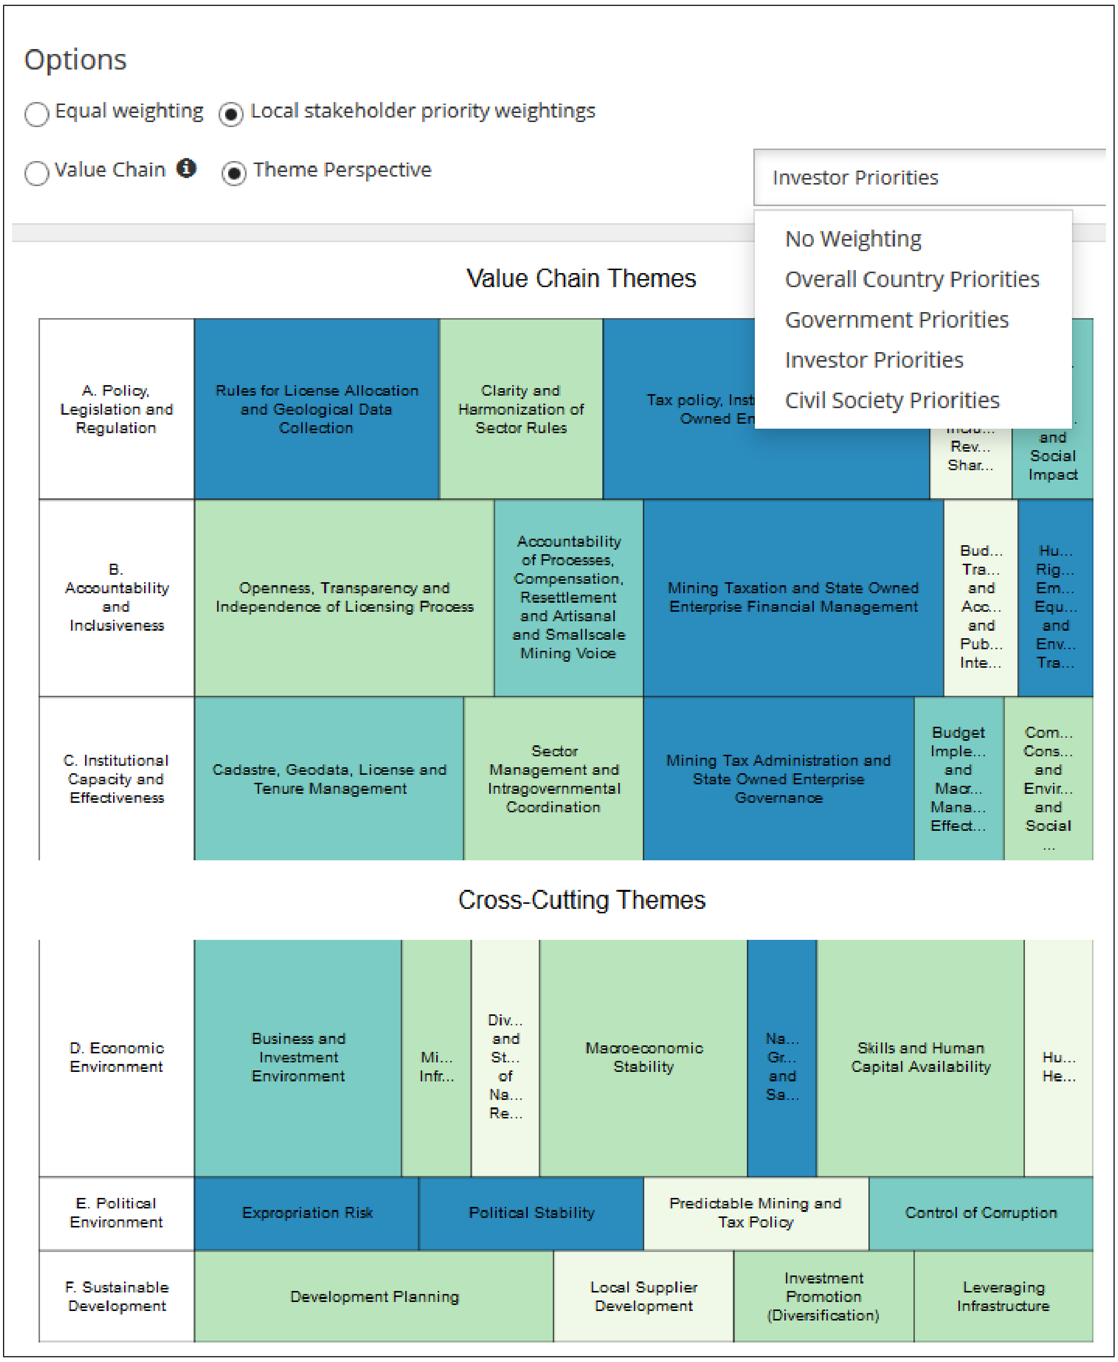

Figure 7.4.: Weighted Matrix View: Depending on the selected perspective (value chain or theme) and selected stakeholder priorities (here: country priorities) the sizes of the matrix cells are adapted.

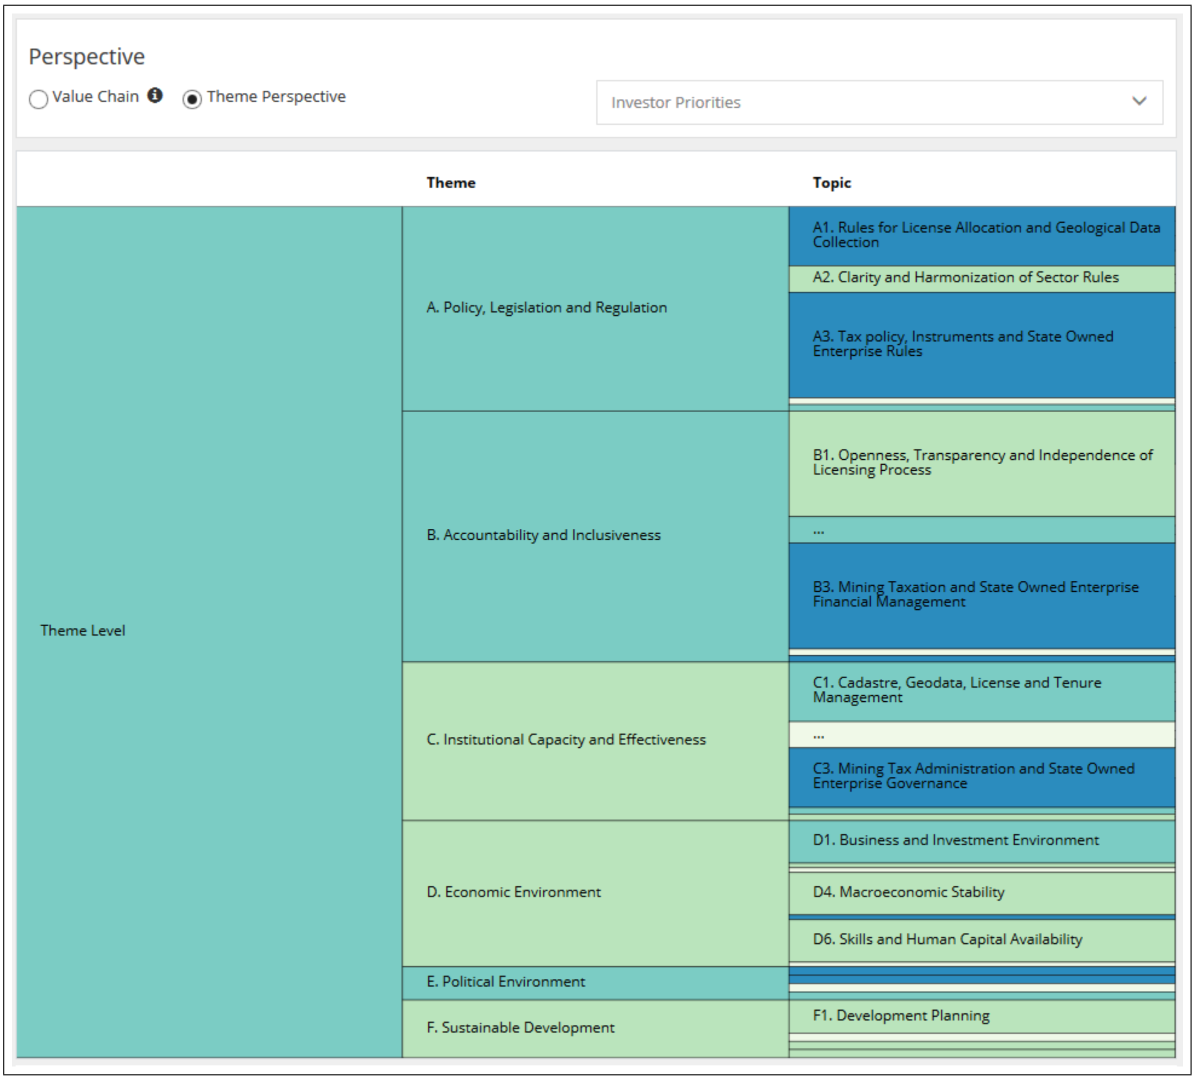

Figure 7.5.: Icicle View: Value Chain perspective. Value chain levels and underlying topics are shown. Users can select a cell to display the underlying hierarchy below this cell.

depicted via a color map. The weighting was mapped on the cell height ($\mathbf{R}_4$). We will not discuss this view in more detail, since the users preferred the Matrix View (see Section 7.5).

7.4.3. Weighted Matrix View

The domain experts also requested the visualization of topic weights ($\mathbf{R}_4$). This is realized in the Weighted Matrix View (Figure 7.4). The area size of the topic cells is used to depict the topic weights. The mapping as shown in Figure 7.4 is only applicable on two level hierarchies. Users can select whether they want to inspect the upper matrix from the value chain or the theme perspective. Based on the selection, the theme labels, or the value chain labels are removed. During the design phase, we also elaborated a visualization technique on the basis of a tree map, as an alternative to the weighted matrix (e.g., [JS91], [GACOR05]). Finally, together with the domain experts, we decided to withdraw the tree map since the reordering of topics caused by the tree map layout confused the users.

7.4.4. Statistics View

To get a high-level overview of the themes and the value chain levels, we included an additional Statistics View (Figure 7.6). It aggregates the questions based on the overarching themes (left) and value chain levels (middle). To provide the full picture the mining sector importance topics are also shown (right).

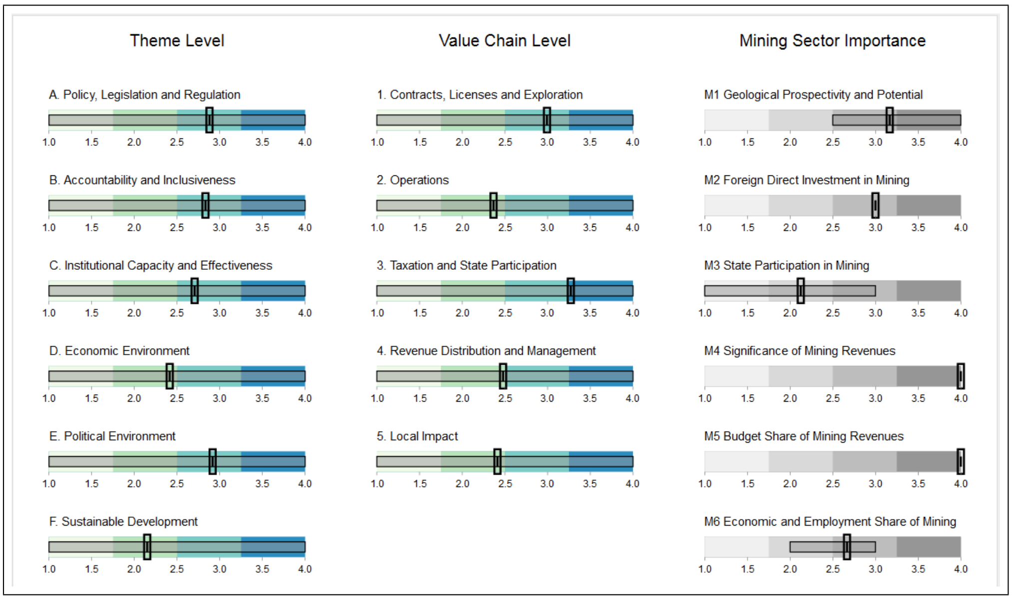

Figure 7.6.: Statistics View: questions aggregated to theme, value chain, or mining sector importance level.

7.4.5. Detail View

The Detail View (Figure 7.8 (left)) allows the drill-down from topics to indicators to questions ($\mathbf{R}_2$). The Detail View is accessed, whenever a user clicks on a cell or a bar in the Matrix or the Statistics View. A headline indicates the selected topic (e.g., “mining tax administration and state owned enterprise governance”). Below a list of indicators is presented, colored with respect to the average of the underlying question scores. Users can select one of the indicators to see the distribution of the question scores in the bar chart below. By selecting one of the bars in the bar chart, the underlying questions, their performance score, and their question type is shown.

7.4.6. Country Comparison View

The Country Comparison View (Figure 7.7) raises the granularity of the analysis from single countries to the comparison of multiple countries ($\mathbf{R}_5$). Users can compare several country datasets with respect to their value chain level, theme or topic scores. The comparison is realized by a radar chart. The domain experts preferred this compact representation over a grouped bar chart.

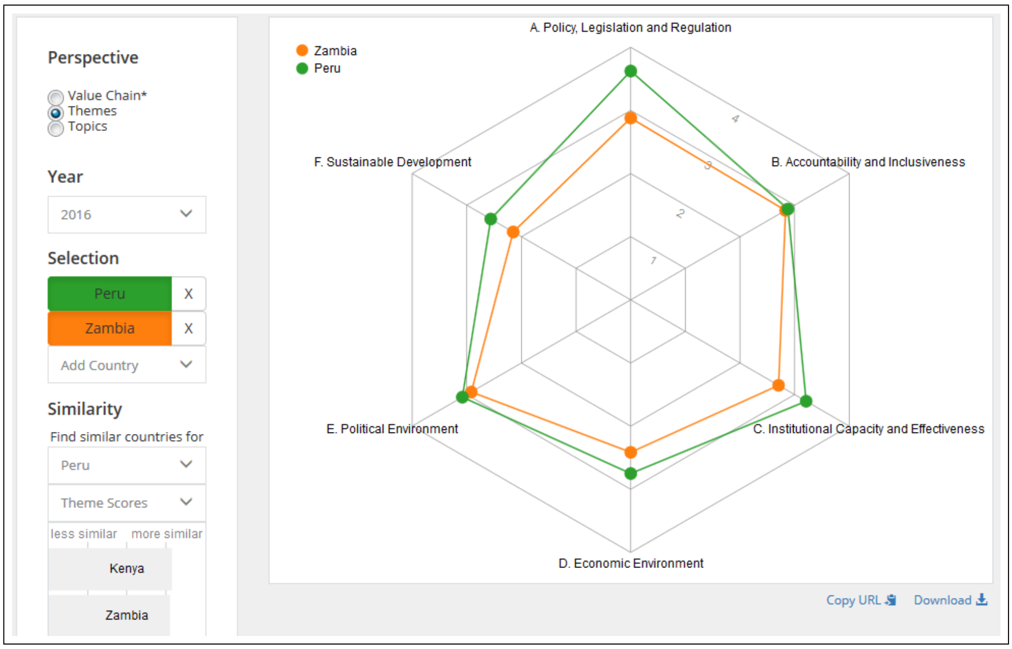

Figure 7.7.: Country Comparison. Countries are compared in radar chart. Value chain, theme, or topic level can be selected. Similar countries can be found with the Similarity Search (bottom left).

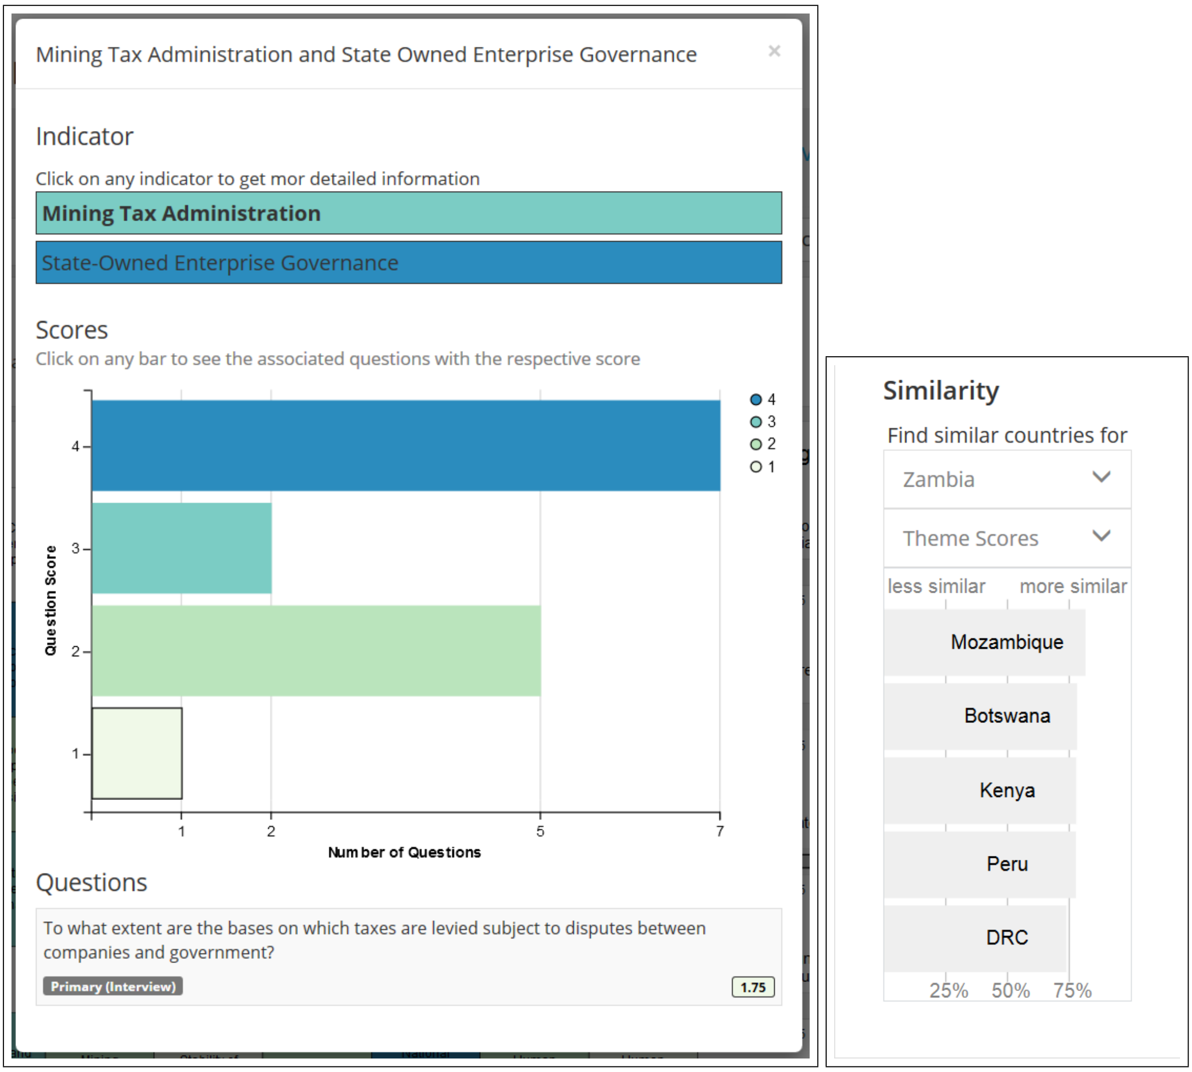

Figure 7.8.: Left: Detail View (as shown when selecting the topic "Mining Tax Administration..."): Topic label, indicators assigned to the topic (in bars), and questions assigned to the indicators (in bar chart) are shown. Questions can be accessed by selecting a bar. **Right:** Similarity Search: For a given country, similar countries can be searched by specifying similarity measure. Here, similarity is calculated based on the theme scores of the respective countries.

7.4.7. Similar Country Search

Finally, we incorporated a retrieval of similar countries in the Country Comparison View (Figure 7.7 bottom left) ($\mathbf{R}_3$). Based on the selected country and similarity measure the most similar countries are retrieved and depicted in a bar chart. The measure of similarity produces percentage scores, i.e., a perfect match has a score of 100%. Based on a discussion with domain experts, seven similarity measures were incorporated. The Euclidean distance between the following vectors is used: (1) all topic scores (excluding mining sector importance topics), (2) all value chain topic scores (excluding mining sector importance topics), (3) the mining sector topic scores, (4) all country specific weights, (5) all government weights, (6) all investor weights, (7) all civil society weights.

7.5. Evaluation - Usability Testing

We conducted a usability testing workshop with five mining experts at the Indaba conference 2017 [Ind17]. Users had to execute seven tasks with the visualization dashboard. After the task completion test an additional usability questionnaire allowed users to provide qualitative (open questions) and quantitative (Likert scale) feedback on the individual views and the overall system. In the second half of the workshop, we discussed our approach in the group.

In general, the designed visualization system was appreciated by the participants. As a user stated the “substance and data display are adequate for guiding governments in their agenda-setting”. All participants were able to efficiently complete the tasks. While the users accepted the mapping of topic scores on color, the mapping of weights on the cell size was questioned. Participants noted that they could easily distinguish large from small cells. However, the identification of cells with the same size but different aspect ratios was difficult for the users. Only the icicle plot received mainly negative ratings. One participant questioned the added value provided by this view. Another stated that is was “difficult to view”. Since the view was provided as an alternative to the Matrix View, we might remove it from the system. Although the Statistics View was rated positive, we had to explain the participants that the bars are clickable for drill-down. The participants’ missing awareness of interactivity in the views was a lesson we learned from the workshop. The Similarity Search was also rated very positive by the participants. One participant even requested more flexibility in the similarity measures, the seven similarity scores provided by the view should be augmented. Although the dashboard was designed for larger displays, one user was successfully conducting the test with a mobile phone. Simply the height of the bars in the Detail View would have to be decreased to fulfill the usability on small screens.

7.6. Summary

We presented a visual analytics approach that provides visual access to country-based performance indicators in the mining sector. First, based on our general decision making domain characterization (see Section 3.2.2) we gave a domain characterization of the mining sector, discussing the structure of the MInGov dataset, the relevant stakeholders, and the tasks to be supported along mining-related decision making processes. Second, we introduced our visualization system that supports these tasks by providing an intuitive access to the categorical questionnaire data, organized in a hierarchical structure. Finally, we discussed the results of a user evaluation conducted with five domain experts at an international mining conference. The design process was structured along the guidelines presented in the concept of this thesis (Section 3.2.3). Moreover, the system serves as a proof of the applicability of our concept (Chapter 3) on empirical data.

8. Visual-Interactive Access to Simulation Models

In this chapter, we present a visual analytics that realizes the visual access to simulation results developed for estimating the impact of decisions. The visual analytics system supports the exploration of multiple simulation scenarios with respect to dependencies between input parameters and generated output, the creation of alternative simulation scenarios by specifying alternative input parameters, the detailed analysis of simulation results originating from individual scenarios, and the comparison of several simulation scenarios based on cost and benefit (energy production). Regarding the overall approach of this thesis, we prove the applicability of our concept to model-driven data with a focus on estimating the impact of possible solutions to the given problem (Challenge $\mathbf{C}_{\text{Imp}}$). Figure 8.1 shows how the design methodology presented in Chapter 3 is applied in this specific scenario with model-driven data (Challenge $\mathbf{C}_{\text{VDSS}}$). First, the visual analytics system allows users to explore dependencies between input and output variables resulting from several simulation scenarios. That way, analysts and decision makers can explore how much the potential impact (output) depends on the chosen de-

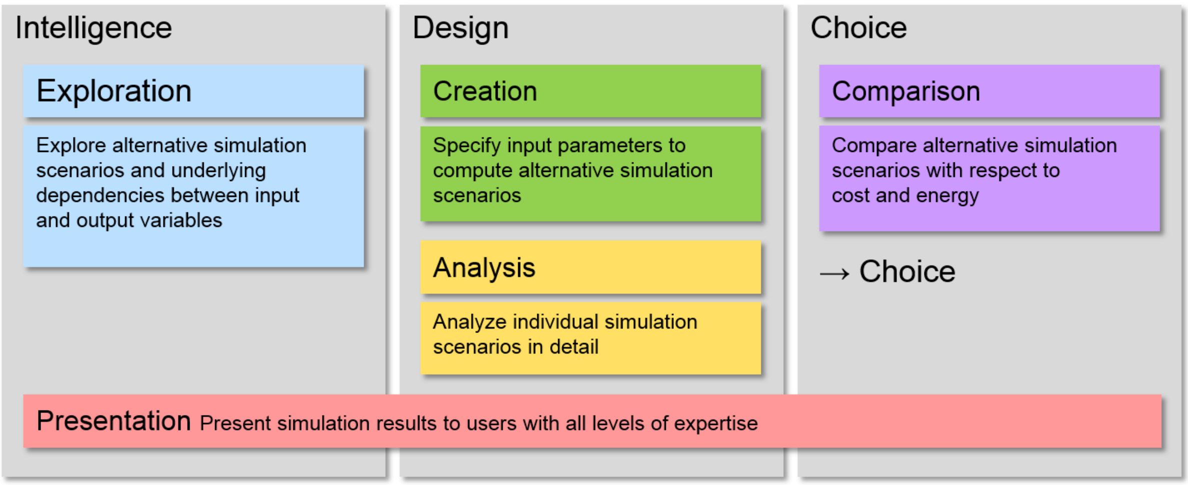

Figure 8.1.: Proof of concept for the applicability of the design methodology to visual analytics approaches supporting **model-driven data**. Visualization techniques are connected to a simulation model to analyze the impact of alternative solutions. The exploration of dependencies between scenario and impact, the creation of new scenarios, and the comparison of impacts resulting from alternatives are supported. Visual access is provided for all levels of expertise. Figure is adapted from Figure 3.6.

cisions (input). Second, users can create alternative simulation scenarios by specifying an alternative strategy (input). Third, the results of a simulation scenario can be analyzed in detail. In our approach, we apply an agent-based simulation model with the agents representing a population of a region in Italy. The detailed analysis allows the drill-down to demographic characteristics of the agents, and the time-dependent analysis of the simulation results. Finally, our visual analytics approach enables users to compare several simulation scenarios with respect to the underlying cost and benefit. This supports analysts in proposing a subset of alternative solutions and decision makers in making a final decision. Our system is useful for several user roles involved in the decision process. Modeling experts responsible for the design of the simulation model are enabled to analyze the simulation results, and thereby, validate their model. During the design process, the modeling experts were applying our visual analytics system to identify failures and fine-tune the simulation model. Moreover, domain experts used the system to judge on the validity of the simulation model. With their expertise and knowledge, they could, first, inform the modeling experts during the simulation model design, and second, provide feedback on the plausibility of the results. Analysts were enabled to use the simulation model without having to understand the internal complexity of the model. They could explore existing simulation scenarios, define new parameters for creating alternative scenarios, analyze the results in detail, and finally select a number of alternatives to be presented to the decision maker. The decision maker can finally use the system to compare alternative strategies presented by the analyst. Therefore, we emphasize that knowledge gaps between different stakeholders in the decision process were bridged by providing visual-interactive access to the simulation model. This chapter is partially based on our previous work published in [RBMK14] [RBU*14].

Contents

8.1. Introduction

Political decision making is a complex task. Policy makers are responsible for providing solutions to societal problems on the political agenda. In most cases several alternative solutions, called policy options exist. As an example, policy makers might want to shift from conventional to regenerative energy production. Therefore, they have to decide upon a subsidy strategy for renewable energy sources. Depending on economic, environmental and social determinants different policy options for reaching a renewable energy target can be defined. These policy options are provided by policy advisers or policy analysts. In order to evaluate the effectiveness of policy options it is crucial to estimate their potential impact on society *before* their implementation into practice. The social behavior of the public is one important aspect that needs to be considered when measuring the potential impact of policies. A common choice for assessing the impact of a policy on the population is agent-based simulation [GMHO12]. Agent-based simulation attempts to model the actions and interactions between autonomous agents in an environment in order to assess the impacts on the system as a whole. These agent-based models are developed by modeling experts from the social sciences and computational engineering. The policy analysts use the condensed simulation results to generate policy options and communicate them to the policy maker. Finally, the policy maker puts a single policy into practice.

In the described policy making process (cf. [HRP09]) both competences and responsibilities are distributed over different stakeholders. We identify two potentials which might improve future policy making. Firstly, we assume that an intuitive *visual access* to simulation models would enhance the creation and evaluation of new policy options by policy makers (and analysts). Secondly, the possibility to *visually explore* a multitude of policy options at a glance would enable policy makers (and analysts) to better understand the decision space and help them in selecting the most appropriate policy option. These potentials of information visualization and visual analytics for supporting the decision making process have already been identified in various domains [KKEM10]. However, to the best of our knowledge, only few visualization approaches in the policy making domain exist.

In our approach we introduce a visual-interactive interface for an agent-based simulation model to support the political decision making process in a regional energy planning scenario. The visual interface addresses two analytical tasks: a) the definition of the simulation input and analysis of its output; and b) the exploration of multiple simulation scenarios based on distinct parameter setups. With our visual interface a political decision maker is enabled to define input parameters for a specific simulation scenario and analyze the simulated impacts in an effective and efficient way. We thereby support policy makers as a new type of stakeholders to work with complex simulation models. Moreover, we introduce a new visualization technique for the exploration of different input-output dependencies and verify its functionality on a synthetic data set. Finally, we apply our visual interface to a real-world example to show its usefulness. Therefore, we connect it to an agent-based simulation model designed for simulating the adoption of photovoltaic (PV) panels on a household level in the region Emilia Romagna in Italy. Our three main contributions are:

1) We contribute to the introduction of policy making as an application field for information visualization and visual analytics. Policy makers need to include information coming from multiple sources

into their decision process. Here, visualization can serve as a mediator providing an intuitive access to this information.

2) We motivate the connection between the research fields agent-based simulation and information visualization. This combination offers valuable potential to both research fields. On the one hand, visualization support simulation experts in validating the agent-based simulation model by providing a detailed view on the simulation results at the agent-level. On the other hand information visualization and visual analytics benefit from addressing research problems from another application area.

3) We introduce a visualization technique specialized for the visual analysis of modeling approaches with the characteristic of using data variables that serve for both as input and output. The visualization enables users to get an overview of the data space, to detect deviations between input and output, and explore the distribution of input variables in the output space and vice versa.

8.2. Related Work on Simulation and Visualization

In the following we review information visualization and visual analytics approaches related to our approach. These include interdisciplinary approaches between agent-based simulation and visualization and related works that address the analysis of input-output parameter spaces. A more general introduction to the field of agent-based simulation is outlined below where we explain the simulation model serving as our use case.

The combination of visualization with simulation approaches in the policy making domain was, to the best of our knowledge, published just a few times. In particular, the 'Vismon' approach by Booshehrian et al. enables scientists from the fishery domain to apply visual data analysis on simulation models in order to perform sensitivity analysis, and global trade-off analysis. The tool was introduced to facilitate the communication between fishery scientists and policy makers in Alaska [BMPM12]. Afzal et al. introduce an interactive decision support environment applied to epidemic simulation models. The developed visual analytics tools help researchers, analysts and public health officials to evaluate epidemic scenarios and observe the impacts of their decisions over time [AME11]. Two visualization tools for analyzing agent-based simulations in the political context are introduced by Crouser et al. [CK12]. The first view applies a projection method to the high-dimensional simulation output, the second view is focused on the agent-level information. Although this work is highly related to ours regarding the simulation model and the application field, it does not allow to compare simulation scenarios with differing input parameters. A visual support system for the understanding of stochastic simulation processes is presented by Unger et al. [US09]. The authors structure a simulation process into three levels: model, experiment, and multi-run simulation data, and provide visual interfaces for all levels. The two latter levels are also considered in our approach. However, the underlying data structure, biochemical reaction networks, and a missing dedication to a specific user group differs from our approach. General design guidelines for agent-based model visualization are presented by Kornhauser et al. [KWR09]. However, the authors focus on the scientific visualization aspects. Spatial and temporal aspects are depicted via the visual variables position and animation, while glyphs are used to represent the agents' current states. In our approach, we focus on information visualization aspects like the statistical analysis of the simulation output to provide the user an overview of the system as a whole.

Information visualization and visual analytics techniques are a popular means for both: to create and to validate varieties of models. From a visual perspective the most similar interface to our approach was presented by Bremm et al. [BvLBS11]. The technique based on Self-Organizing-Maps (SOM) compares different descriptors by aligning the grid-based outputs of respective SOM models. Similar to our interface the visual variables position and color are used to facilitate the comparison of model results. The approach differs in the targeted application task which aims to select appropriate descriptors for the feature-based representation of objects. A related approach, but with a focus on spatio-temporal data is shown by Andrienko et al. [AAB*10]. Here, the authors include temporal information in a geographical view, and geographical information in a time-dependent view, by replacing the absent data type with a color map. This work is related from a methodological perspective. Further examples of parameter space projections, where position is also represented by a color map can be found in the works of Guo et al. [GGMZ05] and Bernard et al. [BRG*12].

The comparison of input and output spaces that cover the same dimensions can be (visually) represented and analyzed by vector field visualizations, where vectors connect the input and output projected into 2D. Telea et al. presented a visualization for vector fields not being limited to input-output representation [TvW99]. Similar to our approach a simplification technique is applied to reduce the visual complexity of the 'display' space. The chosen aggregation technique is based on hierarchical clustering of the flow vectors while our visual interface realizes a grid-based aggregation technique. This choice offers a higher visual robustness, since heterogeneous vector fields would generate overplotting. Another vector-based visualization was presented by Pölzlbauer et al. [PDR06]. Here, the output of SOMs is analyzed to reveal cluster structures. Similar to our visual interface the output of the applied SOM model is visually represented in a 2D grid. However, the input space of the neural network is not incorporated visually.

8.3. Background on the Agent-Based Simulation Model

Agent-based modeling is a method that is frequently used in the political context in order to simulate the impact of political instruments on the public behavior. According to Gavanelli et al. a large number of policy models rely on agent-based simulation [GMHO12]. Gilbert defines agent-based modeling as "a computational method that enables a researcher to create, analyze, and experiment with models composed of *agents* that interact within an *environment*" [Gil08]. The agent's actions and interactions with other agents are driven by autonomous decision rules. As an example these interactions can consist of the transfer of data from one agent to another agent typically located close by in the simulated environment. The environment may be constituted by a geographical region. Further examples include knowledge networks between researchers, etc. Although each individual agent acts independently within the environment, collective group behavior may emerge. That way the individual heterogeneity of real societies can be considered, while analyzing the effects on the system as a whole, which distinguishes agent-based simulation from other modeling approaches.

8.3.1. Simulation Model

In this section, we describe the agent-based simulation model that serves our concrete use case. For more information about the model from a social scientist's perspective we refer to Johnson et al. who explain it in more detail [JBK14].

The model is designed to simulate the public adoption of photovoltaic (PV) panels supported by political subsidies. Its purpose is to evaluate different policy options and to support the policy makers with the selection of the most promising one. In our model the agents represent households that interact in the geographical region Emilia Romagna in Italy. The agent's environment is modeled based on the characteristics of the region. The model is designed based on information that social scientists acquired through surveys in the respective region. The agent model reflects demographic variables of the population like age, education and income aggregated on a household level. Further encoded variables are, e.g., the household's awareness of PV subsidies, the type of household, etc. Moreover, the agent's behavior - their action and interaction rules - is modeled. Based on the agent's individual state, it decides whether to install PV panels or not. Further model variables influencing the agent behavior are policy instruments (e.g., financial subsidies) that may be applied in the simulation to support households in their investments.

In our use case, a policy option is defined as a set of input parameters for the simulation that constitutes a simulation scenario. These input parameters can be set by the user, e.g., a policy analyst. The main input variables of our considered model are the *cost* to be spent for the PV panel installations and the targeted *energy*. Further input parameters to the simulation model and their options are:

- the policy instruments to be applied: regional incentives (e.g., grants, interest rate), national incentives (e.g., feed-in tariffs, tax benefits).
- the budget distribution over years: no distribution rule (first-come-first-serve), even distribution, increasing distribution, decreasing distribution
- the objective of the scenario: maximize energy, minimize budget, maximize participation

The two main output variables of the simulation are the effectively spent cost and produced energy. Please note that these variables are used as input and output. Both output variables - the cost spent and the energy produced - can be analyzed on an agent-level. Each agent is categorized according to its demographic characteristics.

Hence, for every agent, we extract three simulation outputs: the energy produced per year, the cost spent by the household and the financial support received through the policy instruments. Moreover, the number of subsidy recipients is reported. Each simulation covers several years, in our use case from 2014 to 2021. This adds another dimension to the output data space. In addition, each simulation scenario can be run multiple times in order to measure the variance of the simulation output.

In our approach we introduce a visual-interactive interface to agent-based simulation models. These agent-based simulation models can be used to evaluate the impact that policy instruments might induce on society. In our use case the simulation model is designed to simulate the public adoption of PV panels supported by governmental subsidies. In the following, we firstly summarize the user requirements formulated as tasks that policy makers wish to address with the simulation model introduced in the previous Section 8.3. Finally, we present our visual interface providing access to the simulation model.

8.3.2. Requirements for the Visual Interface

The overall goal of our use case is to detect the optimal policy option (by means of a subsidy strategy) to reach an energy target with a given budget. Thereby the public behavior regarding the PV adoption should be observed. The requirements for our system were collected through questionnaires and interviews conducted with policy makers from the region Emilia Romagna in Italy. Based on the overall objective, the extracted qualitative feedback and informal suggestions, we derived the following concrete tasks to be addressed with the visual interface:

Req.	**Description**	**Challenge**
$\mathbf{R}_1$	Specification of input parameters to define a simulation scenario	creation
$\mathbf{R}_2$	Visual analysis of output variables. Focus on cost and effect	analysis
$\mathbf{R}_3$	Drill-down into agent-specific output information	analysis
$\mathbf{R}_4$	Comparison of multiple simulation outputs with respect to dependencies between input and output of simulation runs	comparison

Table 8.1.: Functional requirements for the visual-interactive simulation system.

8.4. Visual Analytics Designs

The visual interface is based on the two most important analytical tasks that were derived in the requirement phase: analyzing a single simulation scenario (*analysis*, $\mathbf{R}_1$, $\mathbf{R}_2$, $\mathbf{R}_3$) and explore different simulation scenarios (*exploration*, $\mathbf{R}_4$). Both tasks are addressed with distinct visual modes.

Analysis. The analysis mode enables the user to set the input parameters and run the respective simulation scenario once or several times with the same input. The user can inspect the results of the simulation runs including the mean values and standard deviation for all output variables.

Exploration. The exploration mode enables the user to explore the output of multiple simulation runs generated with different scenarios. The user can observe dependencies between input and output variables. For example, she might explore how an increase of the maximal costs (input) might affect the cost effectively spent (output).

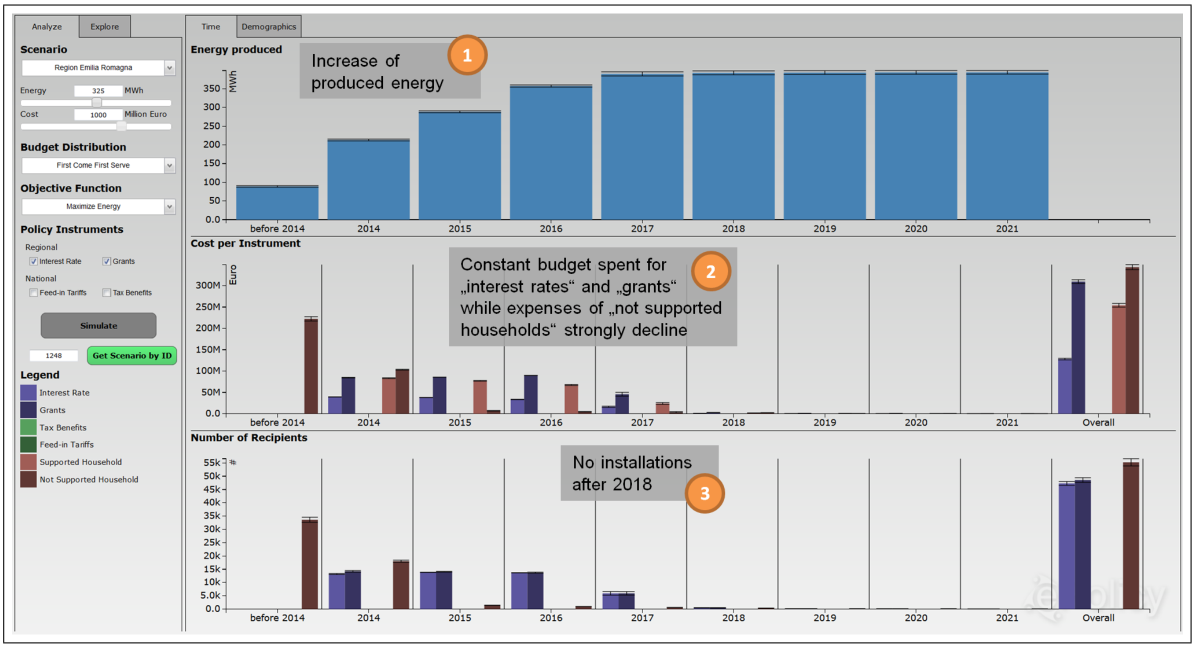

Figure 8.2.: Time View: Analysis mode with time-oriented simulation output. Left: input parameters to be defined. Right: Time-oriented visualization of simulation output. Three bar charts are illustrating temporal progression of energy produced (top), costs spent per policy instrument (middle) and number of subsidy recipients including non-supported households that installed photovoltaic panels (brown bars) (bottom).

From an analytical perspective, these two modes complement each other. Following the visual information-seeking mantra [Shn96], the user can get an *overview first* with the exploration mode, *zoom and filter* into interesting regions of the parameter space, and get *details-on-demand* in the analysis mode. Moreover, in the analysis mode, the user can enlarge the data set by generating new simulation scenario outputs. In the following, we describe these two visual modes in more detail.

8.4.1. Analysis of Single Simulation Scenarios

The analysis mode enables the user to (a) specify the input parameters of a single simulation scenario ($\mathbf{R}_1$), and (b) analyze the simulated output for this scenario ($\mathbf{R}_2$, $\mathbf{R}_3$). Figure 8.2 shows the Time View in the analysis mode view. On the left the energy target, the budget to be spent by the government, the budget distribution over time, the policy instruments to be applied and the number of simulation runs can be specified. After the specification the simulation is executed. The outputs of the simulation run(s) can be analyzed in a time-oriented view, the Time View (see Figure 8.2) and a demographics-based view, the Demographics View (see Figure 8.3).

The Time view provides details about the energy produced, the cost spent, and the number of households that installed PV panels (recipients) over the years 2014 to 2021. The visualization provides

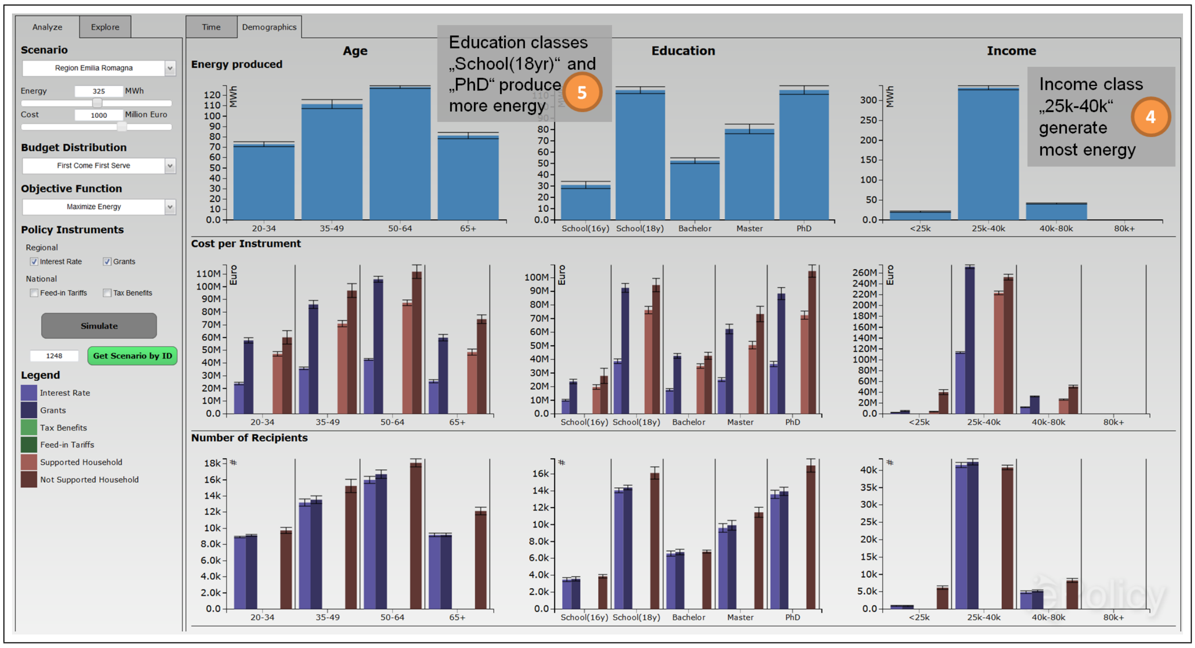

Figure 8.3.: Demographic View: Analysis mode with demographics-based simulation output. Left: input parameters. Right: Nine bar charts illustrating energy produced, costs spent and number of subsidy recipients (vertically arranged), separated into demographic categories age, education and income (horizontally arranged).

information about the costs spent per policy instrument (interest rate (bright blue), grant (dark blue), tax (bright green), feed-in tariffs (dark green)) and by the households themselves. The latter are separated into households that received support (bright brown) and those that did not (dark brown). The same holds for the number of recipients separated into households that received funding from one or more of the four policy instruments (again in blue and green), and those households that did not receive any budget (dark brown). The Demographics View presents details about energy, cost per policy instrument and number of recipients per policy instrument sorted by the affiliation to demographic groups. In our use case these groups are separated by the agents' age, education status and income. The color mapping is identical to the one described for the time-dependent view.

Following the recommendations of Few [Few09], we choose bar charts as visualization techniques in all views of the analysis mode to depict the quantitative information calculated with the simulation. To raise the awareness of uncertainty in the data, in each view the standard deviation of the simulated runs is presented with error bars.

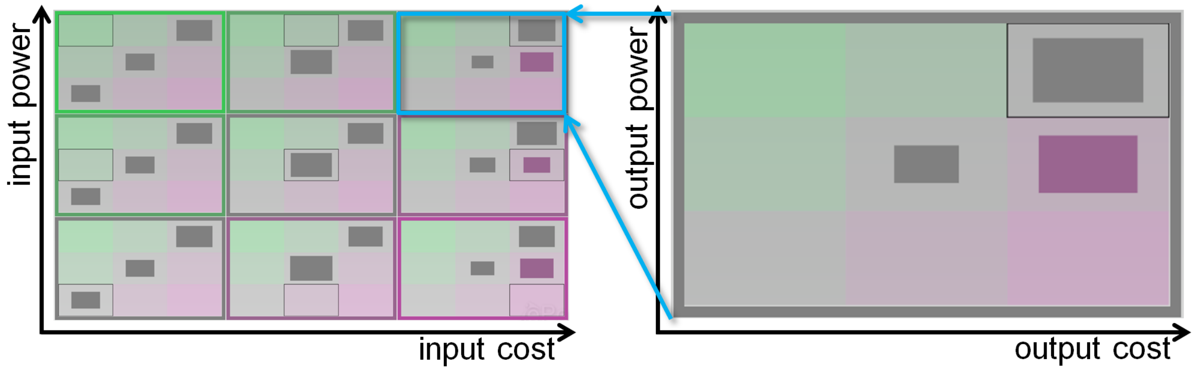

Figure 8.4.: Schematic view of the nested visualization. The global grid (left) reflects the input variables. In each cell of the global grid a local grid displays the input variables (right). The localization of the local grid within the global grid is supported by a black frame. In the local grid the number of samples per cell is depicted by the size of the filled rectangles.

8.4.2. Exploration and Comparison of Different Simulation Scenarios

The exploration mode enables the user to detect dependencies between different simulation scenarios ($\mathbf{R}_4$). One specific characteristic of this approach is the exploration of two main variables that serve as input and output. In our use case these are energy and cost. The exploration mode supports the observation of dependencies between targeted and effectively reached energy and costs.

The exploration mode is realized as a nested visualization in a 2D coordinate system. This results in one *global* grid and multiple *local* grids being displayed in each cell of the global grid (see Figure 8.4). Both grids cover the same parameter space. The value ranges of both axes are separated into areas of constant width. This results in a regular grid with fixed number of rectangles representing the respective value ranges. The input variables of the simulation are mapped on the global grid. The output variables of the simulation are mapped on the local grids. Hence, for a given input scenario depicted in a global cell the output is aligned locally. For example, in Figure 8.4 the cost is aligned at the x-axis and the power is aligned at the y-axis. The output simulated with maximum costs and maximum power as input can be seen in the upper right of the global grid.

We use the size of the shown local grid cells to visually encode the frequency for a particular output. Moreover, in each local grid the corresponding input value range is highlighted with a frame. That way the user can compare the input parameters (framed rectangle) with the output values (filled rectangles). For example, in the local grid selected in Figure 8.4, it can be seen, that no output data (green and blue rectangles) matches the defined input (framed rectangle).

In order to support the orientation in the 2D coordinate system, we apply a bipolar color map on both the global and the local grids. They reflect the position of the values in the grid. In our use case it ranges from purple (bottom right; low energy with high cost) to green (top left, high energy with low cost). The color map is applied to the local grid by filling the rectangles and to the global grid by

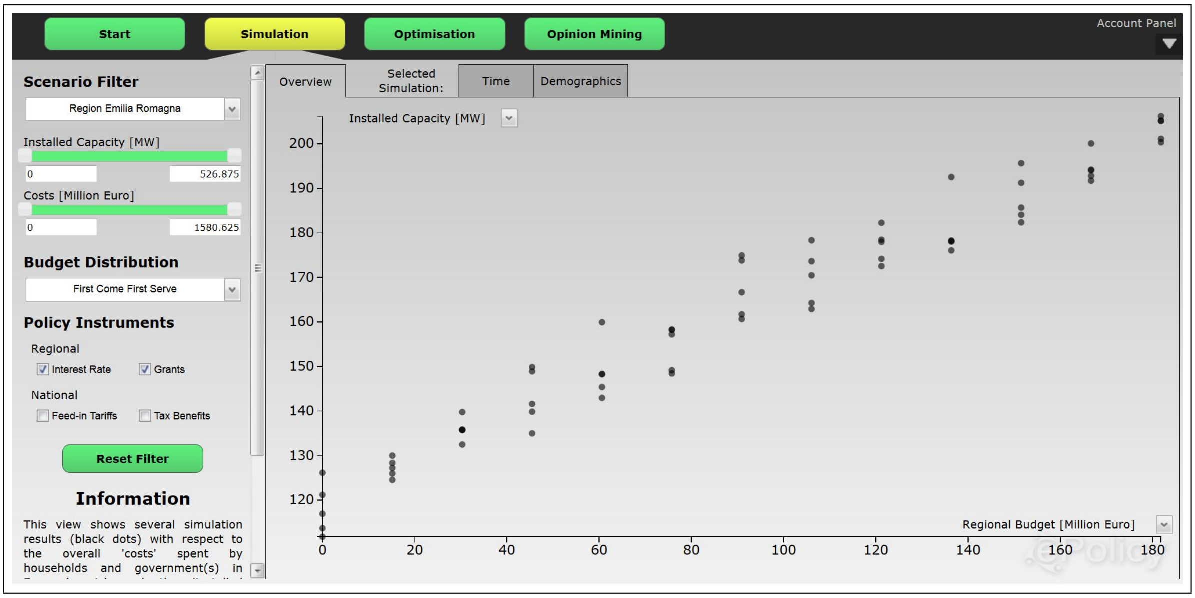

Figure 8.5.: Visual simulation interface: Overview.

coloring the frames of the global cells. Please note, that the mapping between input and output can be switched by the user.

Since this exploration view is difficult to understand, e.g., for decision makers, we included an alternative view in the visual analytics system. Figure 8.5 shows the Overview that allows the comparison of several simulation scenarios with respect to energy and cost. Users can select one of simulation scenarios represented by the dots to analyze the underlying data in the Time View and the Demographics View.

8.5. Case Study

We evaluate our visual interface in two case studies. First, we apply a synthetic model to verify our visual interface. Second, we validate the interface in a use case based on the simulation of photovoltaic (PV) plant adoption by public households.

8.5.1. Synthetic Model

To verify the visual interface, we firstly apply it to a synthetic model. The model is based on a deterministic mathematical function. By eliminating any non-deterministic behavior caused by the model we establish ‘laboratory conditions’ and thus are able to focus on the verification of the visual interface. By that means, we aim at a) verifying the correct input-output parameter representation of the interface

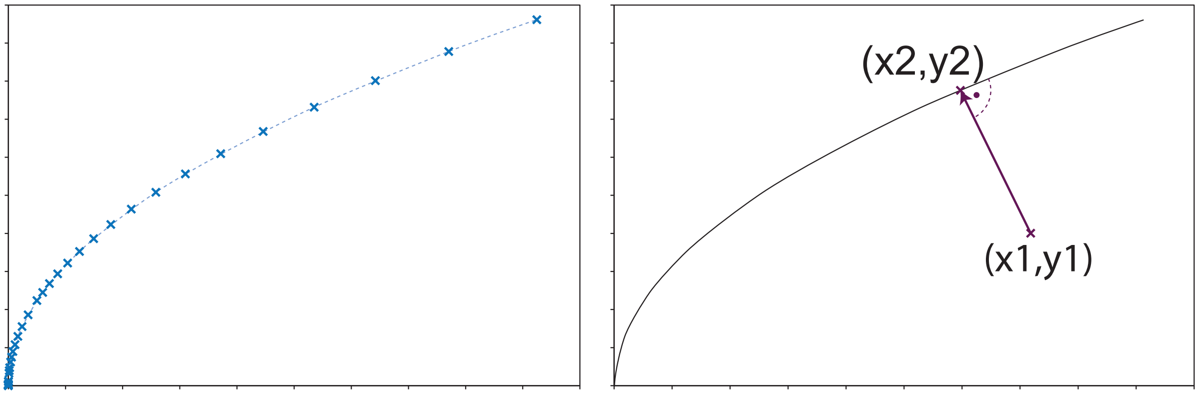

Figure 8.6.: Synthetic model that projects a given parameter pair on the function $\sqrt{x}$). Left: square root function. Right: input-output mapping.

and b) validating the usefulness of the visual interface for analytical reasoning. Our selected model has two input parameters (x_1, y_1) and calculates as output the coordinates $(x_2, y_2 = \sqrt{x_2})$ with the minimal distance to the input parameters (cf. Figure 8.6). We define two hypotheses for the visual interface:

(H_a): We expect the output parameters shown in the visual interface to match a square root distribution, similar to the function graph in Figure 8.6. This verifies the correct implementation of the visual interface.

(H_b): To prove the usefulness of the visual interface for an input-output analysis, we expect that all possible input parameters are mapped to the square root function graph. We assume that, if the visual interface is useful it must be easy to identify that all 25 input grid elements are mapped to the discretized output grid elements according to the square root function graph.

The model is executed 500 times with sampled input parameters, to guarantee generalizability. Moreover, we make sure that each of the 25 discretized input areas of the parameter space are represented at least 10 times. The result can be seen in Figure 8.7. We identify that all outputs of the model in the 25 grid elements are aligned according to the targeted square root function, which verifies H_a. Moreover, it can be seen that all input parameters that already lie on the model function are mapped to the identical local grid element which proves that the input equals the output of the model. Finally, it can be seen that parameterizations in the upper left of the grid and at the lower right of the grid are mapped onto the square root output positions. Based on the human ability to recognize patterns visually, analysts are able to identify input-output parameter shifts made by the model easily (H_b). More generally speaking the visual interface enables the analysis of the model behavior for a large variety of input parameters at a glance.

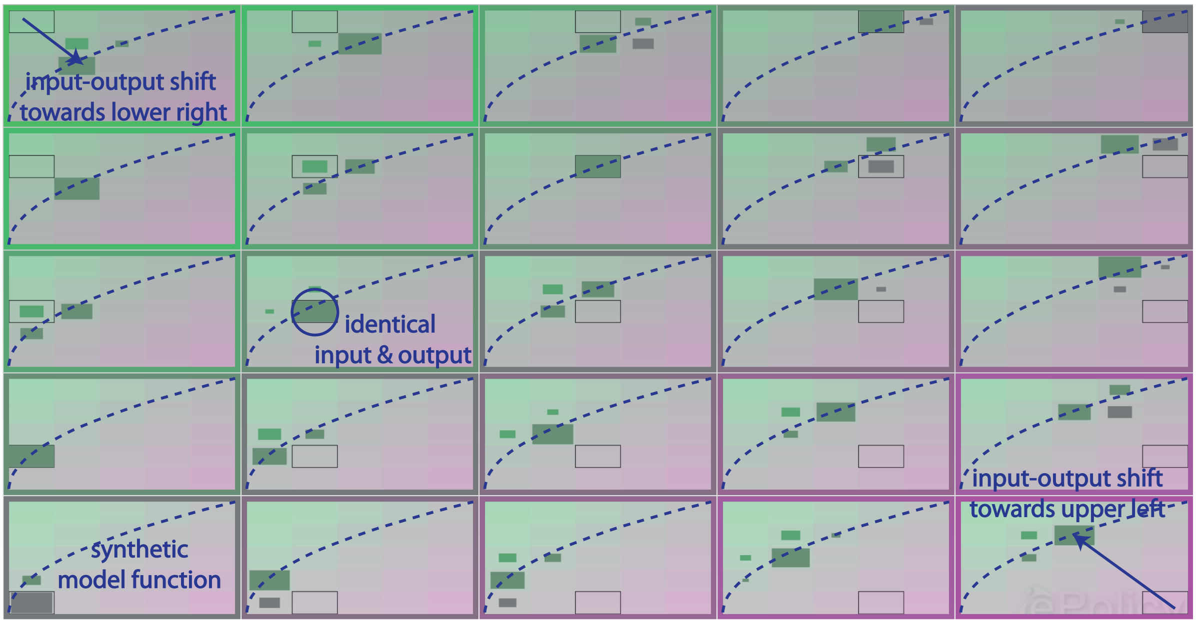

Figure 8.7.: Input-output behavior of the synthetically generated 'square root' model. All parameters are mapped towards the target function (dotted lines).

8.5.2. Real-World Example

In our real-world example we analyze and explore the simulated impact of policy options supporting the PV adoption in the region Emilia Romagna in Italy on a household level. The simulation model is discussed in detail in Section 8.3. In the following, we present findings discovered with our visual interface. In Figure 8.2 the input (left) and the output (right) of a single input parameter set calculated with the simulation model can be seen. The following input parameters have been set: region: Emilia Romagna, energy target: 325MWh, budget: 1000 million Euro, budget distribution: first-come-first-serve, policy instruments: only regional policy instruments (grants and interest rates) activated.

In the output visualization (right) the energy generated per year by the specified simulation scenario is shown. First of all, it can be seen that the annual energy target of 325MWh is reached already in 2016 with the given budget constraints (1). The highest increase of energy can be observed in the years 2014 to 2017, after this period the energy production remains stable. We next focus on the budget used per policy instrument and the costs spent by households (see middle row). It can be seen that before 2014 households spent roughly 220 million Euro on PV (left). Beginning in 2014 the region starts to subsidize PV installation by *interest rates* and *grants*. Between 2014 and 2016 the budget spent on these policy instruments remains constant. However, from 2015 the budget spent by *not supported households* declined drastically (2). One interpretation of this finding may be that the policy instruments are accepted by the region. The statistics about the *number of recipients* (lower row) also supports this interpretation. Another finding in this view is that no installations are made after 2018, although the maximum budget to be spent is not reached (3). Further conclusions can be drawn when switching to the demographics-based view (Figure 8.3). At the upper right where the produced energy is divided by the different incomes, the demographic group with a *household income between 25K and 40K* produces the most energy in the simulation (4). Moreover, in the education column, we identify that the education group of *school(18y)* produces the most energy, together with the *PhD‘s* (5). It may be important for policy makers to reach these two demographic groups since they invest a majority of the money. In contrary, the question arises why other demographic groups seem to be underrepresented. With these two views the output of a single simulation scenario can be analyzed in detail.

We now switch to the exploration view presented in Figure 8.8. The visualization shows 490 simulation runs with multiple input parameter sets at a glance. As described in Section 8.4.2, the x-axis shows the budget spent while the y-axis shows the energy produced. A first finding is that the output (local cell distributions) is very similar in vertically aligned cells of the global grid, while it changes with respect to horizontally aligned cells. Hence, for nearly all runs the simulation output (local grids) only marginally depends on the input energy while it strongly varies with the input costs. Moreover, many output values are aligned along the diagonal of both axes. According to the chosen color map those values are colored in gray. Based on that finding we can assume that the model shows linear behavior for most parameter sets of the 2D input space. As another finding in the first column of the global grid (lowest budget) the highest spread within the local grids can be spotted. This means that for low budget the uncertainty of the generated output is highest. In general, a very heterogeneous spread of output values along the diagonal for nearly all input parameter sets exists. This can be interpreted as a high variability of the simulation outcome. However, most stable solutions seem to be generated when high input values (for costs and energy) are selected (upper right of the global grid).

From these findings, we conclude that the visual interface adds benefits to different stakeholder groups of the system. Firstly, the policy maker can access the complex simulation model in an intuitive way. She can specify a simulation scenario ($\mathbf{R}_1$), analyze the output variables with a focus on energy, costs, and number of recipients ($\mathbf{R}_2$), get more detail about the agents’ output in a demographics-based view ($\mathbf{R}_3$), and compare different simulation scenarios as policy options with a focus on energy and cost of the respective scenarios ($\mathbf{R}_4$). Secondly, simulation experts can validate the consistency of their models by conducting a detailed visual-interactive analysis of the agents’ output. This has been verified through several discussions with the developers of the simulation model.

8.5.3. User Evaluation

In addition to the qualitative evaluation of the visual analytics system presented in the two previous sections, we evaluated our visual analytics approach together with the visual analytics system introduced in the following chapter (Chapter 9). This evaluation was conducted with users realized through task completion tests and usability questionnaires. The evaluation design and the results of the user testing are presented in the following chapter in Section 9.6.

Figure 8.8.: Visual exploration of 490 simulation scenarios and their output with respect to energy (y-axis) and costs (x-axis). Global grid defines input (maximum budget and targeted energy). Local grids reflect output (effectively reached energy and costs).

8.6. Discussion

We presented a visual interface for an agent-based simulation model that enables decision makers to evaluate the impacts of alternative policy options for a regional energy planning scenario. While we showed that the exploration mode of the interface is capable for the class of 2D to 2D parameter spaces, multivariate policy options exist. We showed one way to address these by selecting the two most relevant parameters, cost and power in our case, from the multivariate input and output. Another way is to incorporate more complex models, which also brings new requirements to the visual interface. We suggest an alternative use of color to extend the model to 3D. For high-dimensional parameter sets data projection techniques can be applied to map the parameter space into the 2D screen space (e.g., by mapping high-dimensional data to a regular grid with a SOM approach). However, data projection is always accompanied with a loss of information which should be indicated visually; one of the current information visualization research problems. Another point of discussion relates to scalability issues. While our visual interface is independent of the particular implementation of the applied model, it depends on the respective model scalability, though. Since we enable an online access to an agent-based model we suggest the choice of models that are based in a parallel architecture in order to simulate various scenarios simultaneously. Regarding the visual scalability, we had hardly any problems to encode the displayed information in real time. However, the web-based data communication may become a bottleneck, depending on the net speed and the amount of data to be transferred.

8.7. Summary

In this chapter, we combined an agent-based simulation model with visualization techniques to support evidence-based decision making in the policy making domain. The resulting visual analytics system proves the applicability of our concept (see Section 3) on model-driven data. Based on our visual interface, policy makers were enabled to analyze specific policy options in a more effective and efficient way. This supports the bridging of knowledge gaps between policy makers (decision makers), analysts, modeling experts (who designed the simulation model), and domain experts (who provided data to fill the model). Moreover, in our application we connected an agent-based simulation model and visualization to allow policy makers to a) analyze underlying policy options in more depth and b) explore dependencies between multiple policy options. Based on the results of a requirement analysis we implemented a system with two visual-interactive modes to provide both of the latter described analysis tasks. Finally, to achieve our objectives (extracted with the help of our abstract task characterization; Section 3.2.2), we contributed a new visualization technique for the visual input output analysis of simulation models. The visual interface provides an overview of large varieties of 2D input-output parameter sets and thus, enables analysts to efficiently capture the global optimum of the simulation model output. We verified the visual interface by means of a synthetic 2D input-output model and validated it with a real-world case study conducted at the region Emilia Romagna, Italy. Here, an agent-based simulation model of photovoltaic plant adoption by public households was applied to support the optimization of financial expense and clean energy revenue.

9. Visual-Interactive Access to Optimization Models

In this chapter, we present a visual analytics approach that combines visualization techniques with an optimization model to support decision makers in mitigating trade-offs during the decision process. The visual analytics system supports the creation of optimal solutions based on a user-defined target(s) and constraints, the detailed analysis of a created optimal solution considering multiple output parameters, and the comparison of several optimal solutions. Regarding the overall approach of this thesis, we prove the applicability of our concept to model-driven data with a focus on creating optimal solution to a given problem mitigating trade-offs between relevant parameters (Challenge $\mathbf{C}_{\mathrm{Opt}}$). Figure 9.1 shows how the design methodology presented in Chapter 3 is applied in this specific scenario with model-driven data (Challenge $\mathbf{C}_{\mathrm{VDSS}}$). First, the user is enabled to create an optimal solution by specifying one to many target functions and constraints on the input variables. In the case of multiple target functions, several

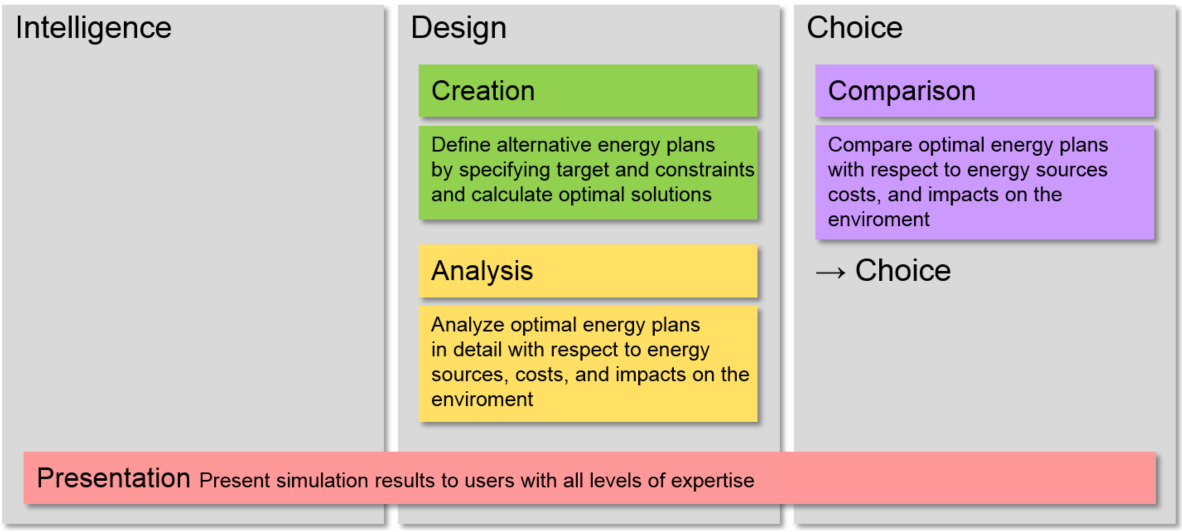

Figure 9.1.: Proof of concept for the applicability of our design methodology to **model-driven data**. Visualization techniques are combine with an optimization model to support mitigating trade-offs and creating optimal solutions. Solutions can be compared and analyzed in detail. Visual-interactive access is provided to all expertise levels. Figure is adapted from Figure 3.6.

Pareto optimal solutions are calculated and presented to the user. Second, the visual analytics approach allows the detailed analysis of the created solutions. And finally, the user is enabled to compare several optimal energy plans with respect to the underlying output variables. The system was designed with a focus on usability, which makes is usable for different stakeholders. Modeling experts designed the optimization model in cooperation with domain experts, who provide details on the dependencies of the considered variables. Both user types could use the system to validate and fine-tune the underlying model. Decision makers and analysts used the system to create, analyze, and compare the calculated optimal solutions. In conclusion, our visual analytics approach supported the bridging of knowledge gaps between the involved stakeholders in a sense that all users could access the relevant information in an intuitive way with the identical perspective (Challenge $\mathbf{C}_{\text{BKG}}$) This chapter is partially based on our previous work published in [RBU*13] [RBMK14] [RDK*15].

Contents

9.1. Introduction

Visualizing information helps humans to understand problems. Perception science has shown that the human visual system can grasp information faster if relevant aspects are visually highlighted [War13]. This is especially important when difficult decisions have to be made and a good knowledge foundation is crucial. Besides supporting informed decisions, visualization can help in communicating important information during and after the deciding making process. This makes decisions more transparent for a broad group of stakeholders and enables constructive feedback and creative discussions.

Policy making is a domain where the scope of decisions is broad and the impact can affect large parts of the society. At the same time, influencing factors like environmental, economic, and social aspects have to be considered. In recent years environmental and social impacts gained more importance so that strategic environmental assessment (SEA) was enforced by law in many countries. The goal of SEA is to analyze impacts on the environment caused by political decisions. This results in a complex set of relationships that demands a scientific analysis. Multiple authors from the field of SEA claim that the methods to analyze environmental impacts are not well integrated into the policy making cycle [Fis07] [The12] [DCS05]. One reason for this is the complexity of SEA systems that usually depend on many influencing 'dimensions'. The authors suggest that SEA concepts have to be presented more clearly and robust to the decision makers.

One possible way to address multidimensional decision problems is mathematical optimization. Established models and algorithms that support multidimensional problems exist, but they have to be transferred to the policy making process. Recent attempts in the optimization domain have created models which support multidimensional decision making with integrated environmental assessment, but most of them are lacking a visual interface to make the results accessible to non-experts [GMHO12] [YTGS12] [LSKP10]. SEA could provide even more value to the decision making process if decision makers could directly work with the models.

In our approach, we introduce a visual interface to multidimensional optimization models. Goal of this interface is to reduce the complexity of the underlying optimization models in order to provide non-IT-experts an intuitive access to advanced analysis functionality. The user can interactively adjust input parameters, and analyze the resulting alternative solutions. As a use case our visual interface is coupled with a SEA model designed for creating regional energy plans. The underlying optimization model concerns environmental, social and economic impacts of these energy plans on a regional level. The results of this optimization are interpreted by decision makers as well as domain experts, and may conclude in policy options to be considered in the policy making process. The visual interface is also designed to overcome knowledge gaps between different stakeholders. With our approach collaboration between decision makers and domain experts is facilitated due to a common information base provided by the visualization. The visual-interactive design makes use of an optimization model to compute solutions and provides clearly structured information about environmental impacts as requested by [The12].

9.2. Related Work on Strategic Environmental Assessment and Optimization

We review related work on strategic environmental assessment (SEA) at the policy level, and optimization in general as the application domains of our approach.

Strategic environmental assessment (SEA) is a proactive approach for integrating environmental concerns into early phases of decision making. The target is to anticipate and prevent environmental, economic and social damage by predicting the impacts. Especially in the policy making process where contrary options have to be evaluated SEA concepts can be applied profitably. The first SEA system was already established in 1969 by The US National Environmental Policy Act (NEPA). With the first SEA system applied by the World Bank in 1989 the acknowledgment rose, and more countries started to make use of SEA [Sad05] [DCS05]. Nowadays, SEA is an approved methodology and is used in many countries [The12]. In [Fis07], Fischer postulates three meanings of SEA: a) a systematic decision support process, b) an evidence-based instrument for scientific assistance in policy making, and c) a framework for the better consideration of alternative policy options for sustainable development. Therivel additionally recommends an increasing participation and collaboration of multiple domains in the policy making process. The target of SEA should be to deliver robust data and clearly presented information [The12]. The LEAP system implements an SEA approach in the field of energy planning [Hea12]. It provides functions to analyze energy consumption, production and resource extraction while monitoring the resulting greenhouse gas emissions. Further approaches that are concerned with energy efficiency can be found in the survey of Markovic et al. [MCM11]. A further work to be mentioned is the visual-interactive tool ComVis created by Matkovic et al. to assist the engine design process by optimizing diesel injection [MGJ*10]. The optimal parameter set up for an engine construction is analyzed via visual-interactive simulation and optimization systems.

Optimization models can describe complex decision making problems. To make the results understandable for non-domain-experts visualization can help as suggested in Jones' work [Jon96]. Multiple approaches that consider environmental, economic and social impacts were submitted by the domain of optimization, since multi-objective optimization models are able to support these problems [GGG10] [YTGS12] [Aga13]. Further authors have combined the policy making process with optimization and environmental assessment like You et al. for optimizing bio-fuel supply chains [YTGS12], or Lim et al. for water infrastructure optimization [LSKP10]. They have committed models that are able to solve multidimensional decision making problems. Yet, they lack visual interfaces that could enable the involvement of decision makers into the process.

9.3. Domain and Problem Characterization

In our approach we introduce a visual interface to provide access to multidimensional optimization models for SEA. In our use case the optimization model tackles the problem of defining an optimal energy plan on a regional level. This complex mathematical model is difficult to understand by non-modeling-experts like policy makers. To address this challenge, we connect visualizations to the model which facilitates the access for policy makers. In the following, we first briefly describe the model with its variables, dependencies, possible target functions, and constraints. Then, we summarize the user requirements coming from policy makers and domain experts. And finally, we present the visual designs to enable the visual access to the multidimensional optimization problem.

9.3.1. Modeling a Regional Energy Plan

In this section, we describe our collaborative approach with the goal to find an optimal energy plan on a regional level. The resulting energy plan consists of a set of energy sources (primary activities), that each produces a specific amount of energy. The plan also includes secondary activities needed for the installation of the respective energy sources (e.g., aerial power lines, dams, etc.). Multiple aspects have to be considered in this scenario. The government has only a limited budget to incentivize the construction of new plants. Still, a defined minimum of energy has to be produced. There are multiple types of energy sources that can be installed. Some are more efficient, others are more sustainable. Moreover, each region has geographical characteristics that restrict some types of energy sources. In addition, governmental laws have to be observed. Often, they aim at the protection of nature and prohibit the extensive use of polluting energy plants. This is also a demand of the society, which is directly affected by the impact of new policies. In summary, considering all dependencies and finding a solution that satisfies all constraints results in a multidimensional decision problem.

In our approach, we made use of an optimization model to address this multidimensional decision problem. More specifically, a linear optimization model is used, designed by modeling experts, which can be reviewed in [GMHO12]. Please note, that our visual designs only consider input parameters and output data of the optimization model. Hence, it can be easily adapted to other linear, and even non-linear optimization problems. A linear optimization problem can be mathematically described as $max(\vec{c}^T\vec{x}\,|\,A\vec{x} \leq \vec{b}, \vec{x} \geq 0)$. In our case the vector $\vec{x}$ to be optimized consists of the amount of energy to be produced by each of the energy sources included in the model. $\vec{c}^T\vec{x}$ defines the target function to be maximized. Thereby, $\vec{c}$ encodes the target of the optimization problem, e.g., overall energy produced, overall cost, impact on an environmental receptor, etc. $A\vec{x} \leq b$ encodes the constraints on the problem. Thereby, similar to $\vec{c}$, each row of the matrix A together with the boundary value comprised by $\vec{b}$ describes a constraint. After the definition of the optimization model, an optimal solution vector $\vec{x}$ can be calculated, if a solution exists. This vector comprises the optimal amount of energy to be produced by each energy source.

9.3.2. User Requirements

At the beginning of our approach, requirements were identified with the user groups of policy makers, domain experts and modeling experts. Moreover, a questionnaire was sent to the potential user groups to confirm the identified requirements, and determine further refinements. As a result of this requirements analysis, the final requirements for this approach are shown in Table 9.1.

9.4. Visual Analytics Design

Based on the results of the requirements analysis phase, we present a web-based system for the visual access to multidimensional optimization models in the application field of strategic environmental assessment. The various input parameters of the model can be defined in a visual-interactive manner.

Req.	Description	Challenge
$\mathbf{R}_1$	Visual definition of target function and constraints	creation
$\mathbf{R}_2$	Analysis of individual solution	analysis
$\mathbf{R}_3$	Comparison of multiple (possibly pre-calculated) solutions	comparison
$\mathbf{R}_4$	Consideration of environmental, economic and social impacts	analysis

Table 9.1.: Functional requirements for the visual-interactive DSS.

This encapsulation of the optimization model itself via visual interfaces helps to reduce the complexity of the input space. The output space of the model is represented in a two-stage design. Firstly, a result visualization gains insight into the output data of the model. Secondly, an interface to visually compare results of different parameterizations helps to determine optimal input parameter setups. In the following, the visual designs of the web-system are described.

9.4.1. The Input Interface

makes available all possible degrees of freedom of the optimization model (see Figure 9.2). The user is enabled to define target function and constraints to specify the optimization problem (see $\mathbf{R}_1$). The visualization provides three sections to address these tasks. In the upper part the target variable to be maximized or minimized can be chosen. Below constraints on the energy sources can be specified. Please note, that in our use case as a maximum value for each energy source the available regional capacity of each respective source is set. Hence, the user can refine these constraints within the range of 0 and the maximal regional capacity. Moreover, additional constraints on the environmental, social, and economic impacts can be set. Finally, the specified parameters can be labeled as a plan, and the optimal solution can be calculated.

9.4.2. The Optimized Plan View

visualizes the output data of the calculated energy plan based on the inputs defined in the Input Interface (see Figure 9.3). This addresses requirement $\mathbf{R}_2$. It gives an overview of the amount of energy (in megawatt or kilotons of oil equivalent) and the costs to be produced by the plan (top left). Additionally, environmental, social and economic impacts are displayed (top right), which addresses requirement $\mathbf{R}_4$. Moreover, the secondary activities needed for the installation of the energy sources are depicted (top middle and bottom).

The information in this view consists of nominal and quantitative data. Hence, we chose bar charts as visualization technique, as proposed in the literature [Few09]. Moreover, this technique is easy to understand for non-IT-experts. A normalized stacked bar chart depicts the percentages of secondary activities needed for the installation of primary activities (energy sources). The impacts are depicted via a heatmap to save display space. The values of the different impact types cannot be compared

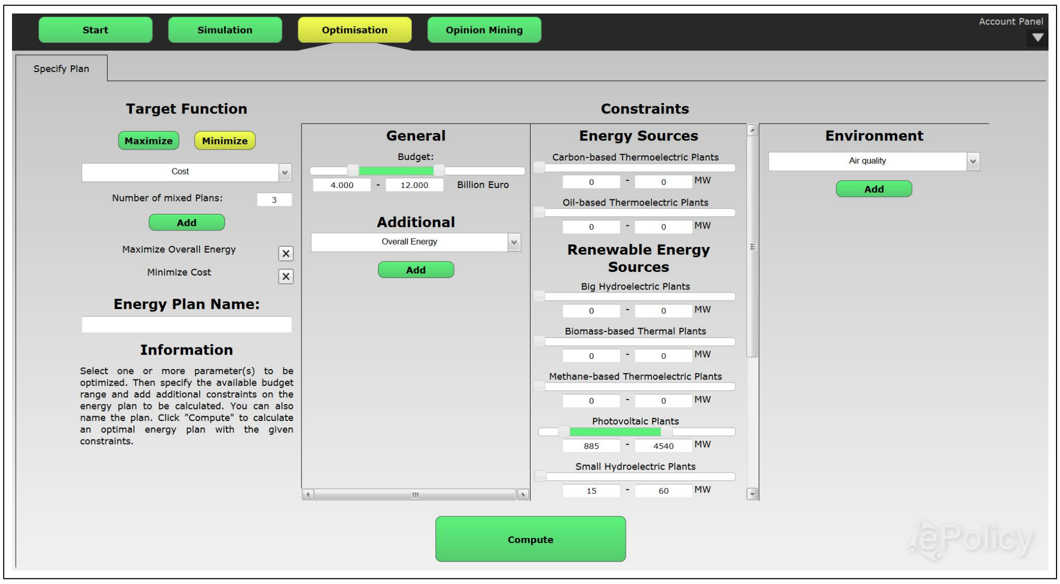

Figure 9.2.: The input interface for specifying target function (e.g., maximize energy), constraints on budget and energy source capacities (e.g., biomass plants capacities), and constraints on environmental impacts caused by the energy plan (e.g., air quality).

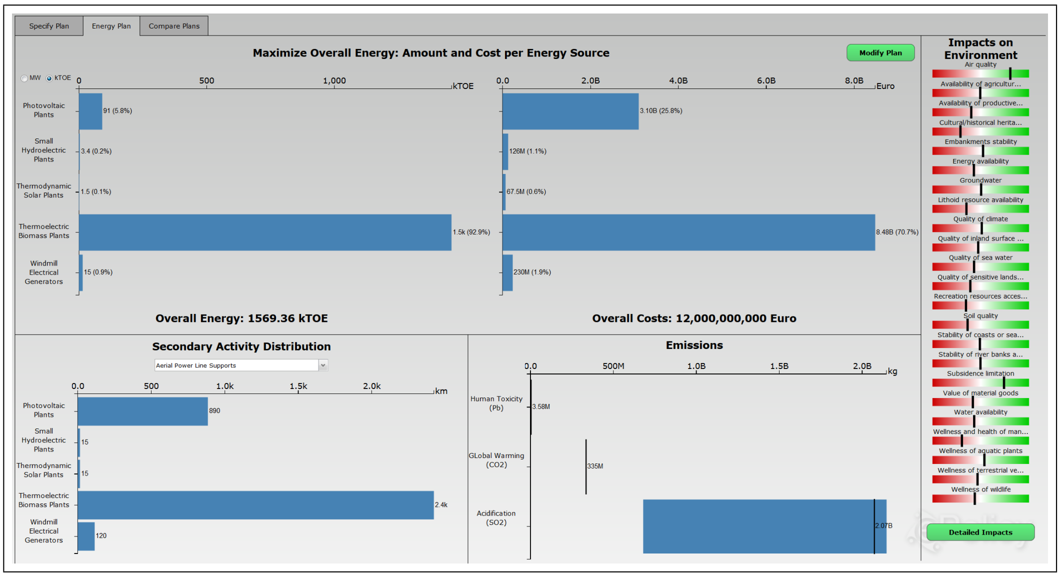

Figure 9.3.: The optimized plan view visualizes the output data of the optimization. It depicts the quantity and costs of energy to be produced per source, impacts on the environment, and additional secondary activities needed for the realization of the plan.

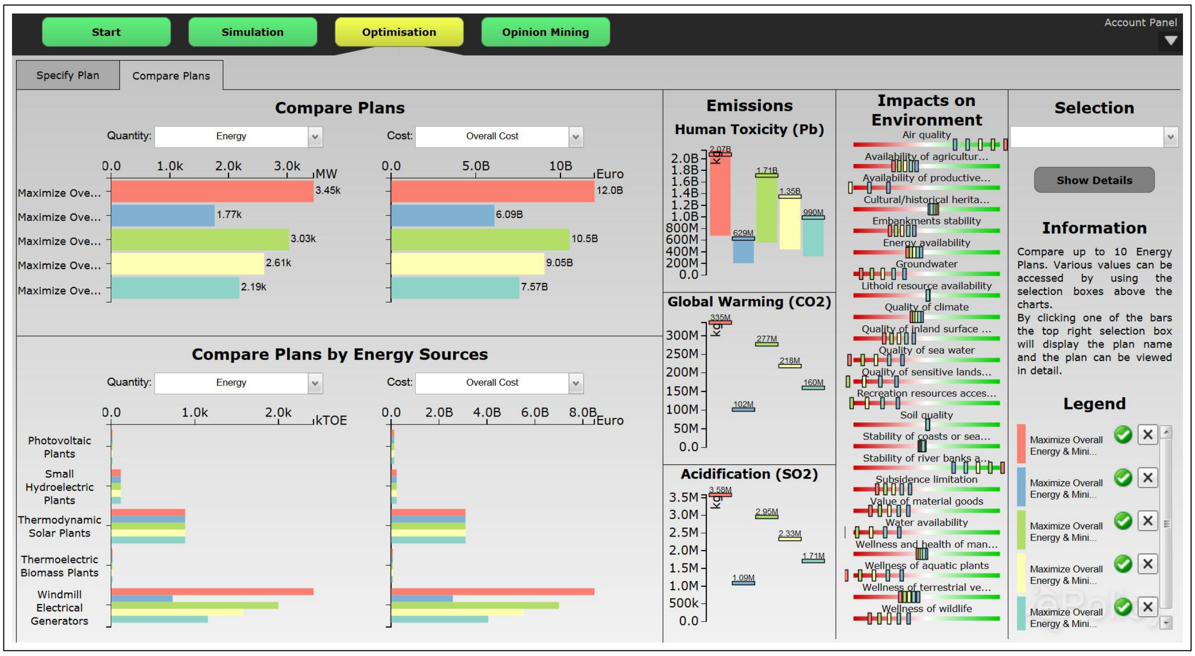

Figure 9.4.: The comparison of plans view enables the user to compare plans calculated with different input parameters. Here, four different plans are depicted. The user can compare the overall energy and costs to be produced by each plan, detailed information about the energy source mix, and impacts on the environment.

because they are measured in different units. If requested, more detail on the impacts can be viewed in the Impacts View, an additional matrix heatmap visualization mapping secondary activities (rows) on impacts (columns). For the visualizations we used an evaluated categorical color map [HB03] to depict the nominal data labeling the distinct energy sources. A diverging color map is used to depict the quantitative impact values in the heatmap; negative (red), neutral (white), or positive (green).

9.4.3. The Compare View

visualizes a set of energy plans the user wants to compare (see Figure 9.4). This view covers functional requirement $\mathbf{R}_3$. To compare the different facets of the plans each variable is visualized separately. The top layer allows the user to get a fast overview of the compared plans by presenting the overall energy, and the overall costs produced by each plan via bar charts. The middle layer of the visualization splits the energy produced and the costs into the energy sources and displays them as grouped bar charts. The heat map, also presented in the Optimized Plan View, shows the different impacts on environment, society and economy, and therefore addresses requirement $\mathbf{R}_4$.

9.5. First Evaluation Round

We conducted a case study to evaluate the outcome of our approach. In the following, we describe the methodology and the results of our evaluation.

9.5.1. Experimental Design

The system evaluation was hampered by the fact that, to the best of our knowledge, no comparable visual interfaces for SEA could serve as a baseline to evaluate against. Neither the targeted policy analysis domain, nor the applied model to solve multidimensional optimization problems have been presented in a comparable manner. For this reason, the experiment was designed towards an evaluation of the targeted purpose of this approach. We considered a) the validation of the functional user requirements and b) the verification of the visual encodings of the system as the two important factors to prove the success of the system.

We conducted an informal summative case study based on task-completion tests and a user questionnaire. The field-based case study was conducted with our web-based application to enable a real-world setup. The unsupervised environment enabled the participants to perform the tests without being influenced, in order to give credible feedback. Model parameterizations for the respective tasks were carefully selected and tested in a previously applied test-run to control the resulting data values and their visual representations. The dependent variable of the case study was the task-completion. Additionally, qualitative feedback based on user preference was gathered within the user questionnaire.

Task-Completion Test. We designed the tasks to validate the functional user requirements (cf. Section 9.3.2). The tasks were chosen by means of covering each functional component. By defining the task setup, attention was paid to ensure that each task covers an 'atomic' functional unit of the analysis workflow. This enabled the identification of bottlenecks in the analytical workflow and a target-oriented enhancement thereof. We were interested in what way non-expert users would be able to comprehend the analytical context of the domain-specific model. In this way, statements regarding the 'generalizability' of the case study became possible. First of all, the user had to calculate an energy plan with the default input parameters. Then, she had to solve task 1 and task 2 with the Optimized Plan View ($\mathbf{R}_2$). The first task should test, whether she is able to find the relevant information displayed. The second task should test whether the user discovered the additional information in the heatmap by using the mouse-over tool-tip. Next, the user had to apply the detailed impact view to solve task 3 ($\mathbf{R}_4$). Task 4 could be solved by understanding the remaining visualization of the Optimized Plan View ($\mathbf{R}_2$). Afterwards, the user had to specify an alternative plan in the Input Interface by changing some of the default constraints ($\mathbf{R}_1$). The two calculated plans were compared by the user in the Compare View to solve tasks 5 and 6 ($\mathbf{R}_3$). The six tasks to cover the functional requirements were:

1. Which energy source in the plan produces the lowest amount of energy?
2. On which receptor the most negative impact is induced? What is its value?
3. What has the most negative impact on 'quality of sensitive landscapes'?

4. Which source needs the highest amount of 'aerial power line supports'?
5. Which energy plan has a more positive impact on the air quality?
6. Which of the two plans produces more energy?

Questionnaire. In addition to the task-completion tests, we provided a questionnaire that enabled the participants to give qualitative feedback. It targeted the usability of the system. Questions of concern were the verification of the visual encodings and interactive capabilities. First, questions about the participants' profession, their domain of expertise and their common analysis tasks, were asked. Additionally, the questionnaire incorporated eleven closed questions about the visual encodings and the usefulness of each view. Finally, we provided open questions to gather informal feedback of individual user preferences.

Participants. Ten non-domain experts agreed to participate voluntarily in the task-completion test. All of them had a profound background in information visualization. None of them had expertise in the energy domain. Thus, the results of the task-completion test were not influenced by domain knowledge which allows the generalization of the results. The questionnaire was answered by all non-experts, and two additionally recruited experts from the energy domain with no expertise in visualization. Most of the participants were male, none of them reported color blindness. The age of the participants reached from 22 years to 38 years with a median of 28 years. All of the participants had an academic degree.

9.5.2. Evaluation Results

Results of the Task-Completion Tests. Figure 9.5 shows the results of the task-completion tests. Tasks one, three and four were completed correctly by most of the participants while tasks two, five and six had higher error rates. The tasks with the low error rates were completed by using the Input Interface, the Optimized Plan View and the Impacts View. For the tasks with the higher error rates the users had to make use of an additional Overview and the Compare View. The Overview consisted of a scatterplot visualizing all computed plans with respect to their overall energy produced and overall costs. In the Overview the user had to select a set of plans for comparison. The usability test shows that on average 90% of the first four tasks were completed correctly but only 65% of the last two tasks. This concludes that the users were able to use the Input Interface and calculate optimized energy plans. Thus requirement $\mathbf{R}_1$ is fulfilled. The Optimized Plan View supports the analysis of energy plans and the Impacts View displays the environmental impacts. The tasks 1 – 4 related to these views were completed correctly by most of the users. Thus, requirements $\mathbf{R}_2$ and $\mathbf{R}_4$ are met. The results of the last two tasks indicated that the exploration of multiple energy plans and the following comparison is not as easy as the previous tasks. To solve the task, the Overview and the Compare View had to be used. The error rates of the last two tasks showed that the visual-interactive design of these views did not support the user as well as the other views. Thus the requirements $\mathbf{R}_3$ and $\mathbf{R}_4$ are met but the solutions have to be improved.

Results of the Questionnaire. The questionnaire was designed to test the usability of the system with respect to the visual encodings and interactive capabilities. For each view the users were asked

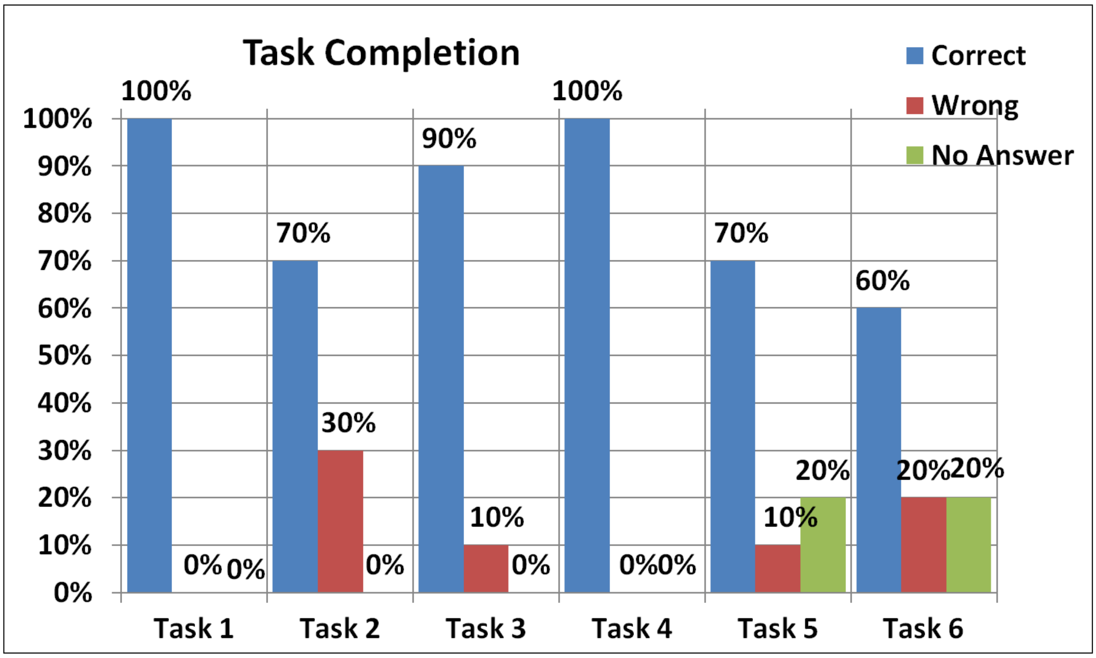

Figure 9.5.: Results of the task-completion tests.

regarding its intuitiveness and usefulness. For all views, except for the Overview, above 80% of the users stated that they are easy to understand. Above 70% of the users argued that the views are useful, the Compare View even gained 100% of the users' approval. The Overview was only intuitively understood by 50% of the users. Its usefulness was only affirmed by half of the users. Hence, we assume that the lower task-completion rates of tasks 5 and 6 are caused by the inappropriate design of the Overview, not the Compare View. As a result, the Overview was excluded from the application after the test. Calculated energy plans are now directly added to the Compare View.

9.6. Second Evaluation Round

Goal and Task. Our visual analytics approach was tested in a second evaluation round together with the visual analytics approach presented in the previous chapter (Chapter 8), and an additional component addressing the visual-interactive access to opinion mining data (see [RBMK14]). In the second evaluation, we considered a) the validation of the functional user requirements (usefulness) and b) the verification of the visual encodings and the interaction design of the system (usability). Therefore, we conducted an informal summative case study based on task completion tests and a usability questionnaire. The field-based case study was conducted with our web-based application to enable a real-world setup. The unsupervised environment enabled the participants to perform the tests without being influenced, in order to give credible feedback. We carefully selected and tested the model parameterizations for the respective tasks in a previously applied test-run to control the resulting data values and their visual representations. The dependent variable of the case study was the task-completion. Additionally, we gathered qualitative feedback based on user preferences within the usability questionnaire.

Task-Completion Test and Questionnaire. To validate the functional user requirements presented in Table 8.1 (on the simulation component) and Table 9.1 (on the optimization component), we designed six tasks (cf. Table 9.2). The test covers all technical components of the system with an emphasis on the 'atomic' functionalities of the decision making workflow described in the use case. Although, we provided domain-specific information about the real-world problem, we did not introduce the participants to the system before the test. Our goal was to observe whether non-expert users would be able to visual-interactively access the domain-specific models, which would support the generalizability of the evaluation results. In addition to the task-completion tests, we gathered qualitative feedback via a usability questionnaire. It included closed questions about the visual encodings and the interactive features of each view. Finally, we provided open questions to gather informal feedback about individual user preferences.

Task	View	Description
1	Opinion Mining (see [RBU*14])	The user has to drill-down to the most positive comment within a specific time period and extract the name of the author and exact date of post in the table view.
2	Overview (Simulation)	The user has to apply filters on the input parameters. From the remaining simulation outputs he has to state the maximal capacity and the exact costs not exceeding 700M Euros.
3	Time View (Simulation)	The overall cost on a specific subsidy instrument has to be named.
4	Demographics View (Simulation)	The income group receiving the largest amount of grant support hast to be named
5	Input View (Optimization)	The user has to specify input parameters with defined target functions and constraints. From the calculated plans the one that installs the largest capacity has to be found.
6	Output View (Optimization)	The most positively affected energy plan has to be named.

Table 9.2.: Task completion test.

Participants. The evaluation was conducted with 23 participants (7 female and 16 male) between the age of 19 and 67 years (median: 30 years). All participants had an academic degree. 11 participants had a background in energy, 12 had a background in politics.

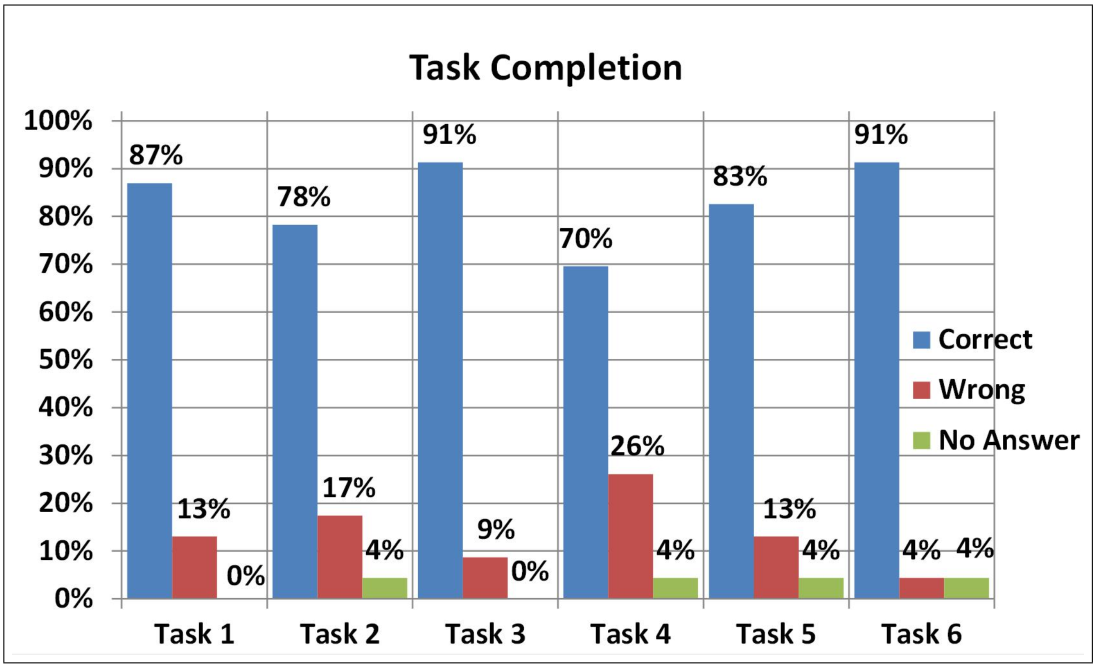

Figure 9.6.: Task completion test: results may not sum to 100% due to rounding

Results of the Evaluation. For all tasks the percentage of correct answers was above or equal to 70% (cf. Figure 9.6). Hence, we observe that the majority of users is able to use the system with respect to the requirements tested through the tasks. This underlines the usefulness of the system. However, even during the test the participants suggested possible improvements on the usability of the system. In the usability questionnaire, more than 70% of the test participants approved the usability of all visual interfaces except the visual opinion mining interface. The reason for this, as stated by the participants, was an interaction design for the selection of multiple dots in the scatterplot which was not easy to understand. This interaction scheme has to be improved or textual explanations for its usage have to be included. In general, more textual descriptions could be added to the pages to help and guide the users through their workflow. Another usability issue was to improve the adaption of the visual interface to the screen size. For example, during the test some labels overlapped, or visual features were not visible. Since decision makers are often on the move, an adaptation of the system to smaller screens, e.g., those of tablets or laptops was desired. Solving these usability issues would also result in better task completion test results, and thus, in an increased usefulness. However, we want to recall that no introduction to the system has been provided to the participants before the test. Participants stated that an introductory tour or video could also help in mitigating the remaining usability issues. In summary, most of the test participants were able to solve the tasks of the task completion test, although, they had no prior knowledge and expertise in the respective domain and the underlying models. Some minor usability issues have been detected and will be addressed in the future. Finally, from the open questions we summarize that most of the users have a positive impression of the overall system. Nearly all of them stated that this tool could be very useful for decision making.

9.7. Lessons Learned

In this section, we summarize lessons learned during the case study and the collaboration with policy makers and domain experts. We share Few's view with regard to his choice of techniques [Few09]. All views designed with bar charts were easy understood by all user groups. That way, our system is applicable by all relevant stakeholders. Still, we learned that different stakeholders need different visual designs with regard to the amount of information presented. While domain experts were satisfied with several visualizations in one view, the policy makers only want selected output data to be presented. Further variables should only be shown, if they exceed critical values (e.g., a scatterplot matrix depicting relations between variables was considered too complex). This is caused by the fact, that policy makers do not have time to learn complex applications, while domain experts need higher functionalities and would adopt to more informative visualizations.

As future work, a user-adaptive visual interface for policy makers with a focused functionality could be designed. The user will be supported by customized views, depicting only selected information. Still, the policy maker view and the domain expert view have to build a coherent basis for communication, which has to be considered during the design process. Moreover, the generalizability of the system could be proved by applying it to different use cases and non-linear optimization problems. As another issue, sensitivity analysis may be included in the interface, since politicians are often interested in adjusting optimal solutions. The complexity introduced to the system may be further reduced by user-guidance techniques.

9.8. Summary

In this approach, we combined the capabilities of information visualization and optimization in order to support the policy making process with methods from SEA. Firstly, we addressed the demand to ensure that SEA is more effective in policy making, by presenting the information clearly to the decision makers, as stated in [The12]. Secondly, we addressed the demand to provide access to an optimization model for non-IT-experts. Therefore, we developed a visualization-tool that connects to an optimization module [GMHO12], using approved concepts of information visualization [Few09] to reduce the complexity of the data created by the optimization. As a result, modeling experts could spot errors in the model and energy experts were able to validate the model. Especially the policy makers, who had to rely on the knowledge of experts, were able to understand the process behind creating an energy plan, and thus could make better founded decisions. The visual analytics system was designed applying re-using the domain characterization and the design process as presented in the concept of this thesis (Section 3.2.2 and Section 3.2.3, respectively). Moreover, we showed how to bridge knowledge gaps between stakeholders in the decision making process (as conceptually presented in Section 3.3).

10. Conclusions and Future Work

To conclude this thesis, we summarize the research challenges and scientific contributions presented in this thesis. Finally, we discuss future research directions that build upon our work.

10.1. Conclusion

In this thesis, we have investigated how visual analytics technology can support evidence-based decision making. In the introduction of our approach, we described the problems and challenges that decision makers at the business and the political level face in modern decision making (Chapter 1). First, the complexity of decision making increases, since besides economic factors, social and environmental aspects are taken into consideration. At the political level this triple is augmented by the public opinion. Second, despite the digital availability of data, information hidden in these datasets needs to be extracted before it becomes useful for the decision process. Third, additional complexity is added to modern decision making through the participation of several stakeholders that differ with respect to their expertise. Fourth, although visual analytics technology bears great potential to support the extraction of knowledge from data informing decision making, the research field lacks theoretical foundations on how to incorporate visual analytics in the decision process. The identified challenges motivated our research and this thesis as a summary of our research results.

As a baseline for our thesis, we reviewed scientific approaches on decision support systems and computational policy making support (Chapter 2). We re-used Simon's decision making model containing the three stages intelligence, design, and choice for the task taxonomy of our concept [Sim60]. Power's decision support categories motivated the definition of data categories relevant for decision making [Pow02]. From policy making research, we adapted the incorporation of eParticipation concepts and stakeholder opinions into the decision process [DMRI16]. Furthermore, we identified specific challenges evolving through the participation of multiple stakeholders in the decision making process, especially between scientists and policy makers [Shu99]. After having reviewed decision support system and policy making support theory, we shed light on foundations in visual analytics design and technologies. The main insight from this review was the lack of a design methodology targeting visual analytics support for decision making. However, we re-used and adapted the visualization design and validation concepts from Munzner [Mun09] and Andrews [And08] as inspiration for our approach.

As a result of the review of scientific foundations related to decision making, policy making, and visual analytics, we defined research challenges to be addressed in order to incorporate visual analytics in decision making processes (Chapter 3). The challenges are summarized in Figure 10.1 (on the left).

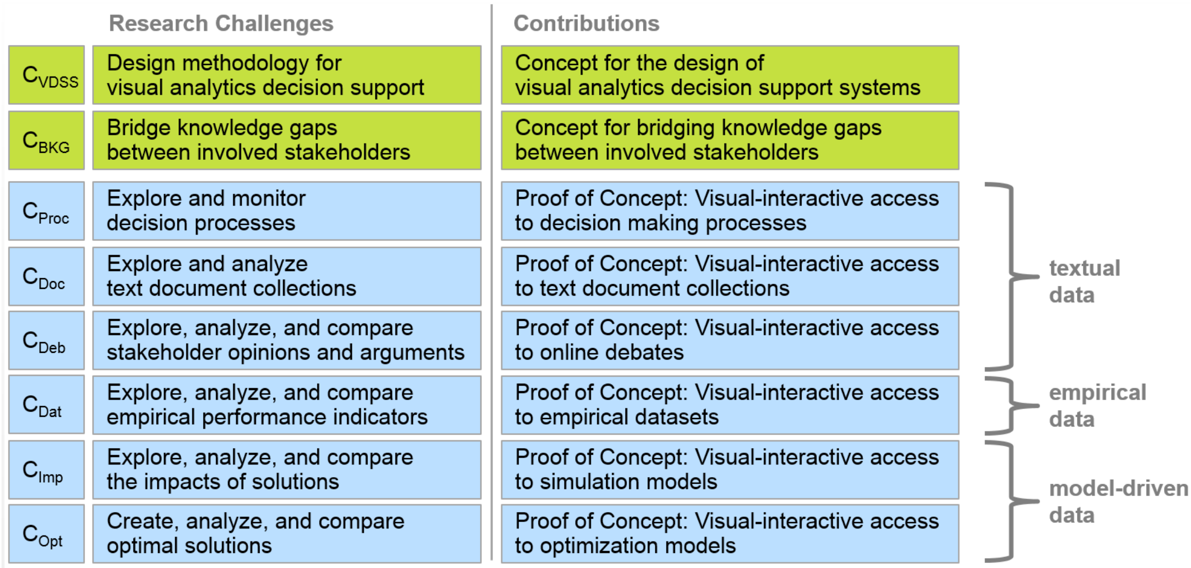

Figure 10.1.: Research challenges and contributions: Two conceptual (in green) and six technical (in blue) research challenges are depicted on the left. Each challenge is addressed by one of our scientific contributions (depicted on the right). We introduced a concept for the design of visual analytics decision support systems and a concept for bridging knowledge gaps between stakeholders involved in the decision making process. In addition, three proof of concepts related to textual data, one proof of concept related to empirical data, and two proofs of concept related to model-driven data were contributed.

We formulated two conceptual challenges (in green): first, the need for a methodology for the design of visual analytics solutions that support decision making, and second, the need for a methodology on how to simplify communication between stakeholders involved in decision making by bridging knowledge gaps between them. In addition to the conceptual challenges, we formulated six technical challenges (in blue) related to specific tasks and datasets relevant in the decision making process. These relate to the extraction of knowledge from (1) documents framing the decision process, (2) text document collections in general, (3) the public debate in form of opinions' and arguments' relevancies, (4) empirical datasets, (5) simulation model data for decision impact assessment, and (6) optimization model data for balancing trade-offs.

Conceptual Contributions

Throughout this thesis, we addressed all challenges (see Figure 10.1 (on the right)). In Chapter 3.2, we presented a concept for the design of visual analytics decision support systems. The concept comprises a general decision making domain characterization and a description of a process for the design and evaluation of visual analytics decision support systems. First, after having described the decision making process from the perspective of a visual analytics expert, we introduced a domain characterization for decision making. Therefore, we analyzed data, user, and task categories relevant for decision making in general. From the data perspective, we distinguished between unstructured textual data, structured empirical data, and structured model-driven data. From the user perspective, we promoted the differentiation between decision makers, analysts, domain experts, modeling experts, and stakeholders. From the task perspective, we organized exploration, creation, analysis, comparison, and presentation tasks along the decision making process. The domain characterization is beneficial in two ways: (a) it provides guidelines for the design of visual analytics solutions supporting the decision making process and (b) it supports the categorization of existing visual analytics solutions and allows the identification of potential extensions. With the domain characterization additional stakeholders to be involved into the decision making process can be identified, alternative data categories can be taken into consideration, and missing tasks to be supported by the visual analytics solution are brought to mind. As a second part of our design methodology, we introduced a process for the design and evaluation of visual analytics decision support systems. Adapting Andrews' concept [And08], our design process contains four consecutive stages: phase before design, phase before implementation, phase during implementation, and phase after implementation. For every stage, we propose goals to be achieved and potential methods for the validation of their achievement. Our design process follows the principles of user-centered design promoting a close collaboration with the potential users of the targeted visual analytics system. As a main benefit, our process definition introduces guidelines for the user-centered design of visual analytics solutions aiming at decision making support.

As a second conceptual contribution of this thesis, we defined a concept for the bridging of knowledge gaps between stakeholders involved in the decision making process (Chapter 3.3). Our concept is based on the incorporation of visual analytics technology in the decision making process. We proposed two models, one for organizational decision making and one for political decision making. The main point of our concept is that visualization can be used for the presentation and communication of information extracted from data. Depending on the expertise of the respective user, different visualization disciplines can be applied to reduce the complexity of data. We identified several benefits of these concepts. First, the communication between stakeholders involved in the decision making process is facilitated by visually sharing intermediate analysis results. Second, depending on the applied visualization discipline the complexity of the presented information can be controlled and adapted to the respective user. Third, by showing the same information to all stakeholders the subjectivity introduced into decision making is reduced. Fourth, visualizing data analysis results also facilitates the validation of the analysis process. For example, modeling experts can test and calibrate their models based on the visual comparison of results. And finally, through the visual presentation of information transparency is introduced into the process.

Technical Contributions

The conceptual contributions presented in the previous section built the foundation for six additional technical contributions (Chapters 4 – 9). The concepts presented in Chapter 3 are applied in each of the technical contributions. Hence, besides their contributions to visual analytics research, they served as a proof of concept for the applicability of our two conceptual contributions. Each contribution contained a visual analytics decision support system targeting a specific aspect in the decision making process. The visual analytics design was conducted based on our design methodology and we explained how the systems supported the bridging of knowledge gaps between stakeholders involved in the decision making process.

In our first technical contribution (Chapter 4) we presented a visual analytics system providing visual-interactive access to the decision making process. The system provides a temporal overview of text documents framing the decision making process. Among other functionalities, users are enabled to create a decision process, add text documents to the process, rate the quality and relevance of documents, comment on them, and compare the support behind specific proposal documents. Furthermore, the expertise of the logged-in users is considered in the computation of relevance. We could prove the system's usability and usefulness within a real-world scenario with domain experts. The system serves as a proof of concept for the applicability of our design methodology on textual data with a focus on exploration, analysis, creation, comparison and presentation tasks. Furthermore, we could bridge knowledge gaps by providing to all stakeholders involved in the process an overview on relevant documents and letting them participate.

To create a temporal overview of text documents published during the decision process is essential to improve evidence-based decision making. However, in most cases the sheer amount of documents is difficult to assess. Getting an overview on the underlying content is cumbersome and time-consuming. Our second technical contribution (Chapter 5) addressed this challenge. We presented a visual analytics system for the creation of document collection overviews via content-based text clustering. The system allows analysts the visual-interactive creation of text clustering workflows from feature selection over algorithm parameterization and cluster analysis to the comparison of multiple clustering results. As a result, analysts can create document clusters optimized for the underlying data and tasks at hand. The complexity of the workflows makes the system usable mainly for analysts. The reduction of its functionality resulting in a lower complexity remains future work. However, the system allows the bridging of knowledge gaps between modeling experts (in text analysis) and analysts. Moreover, it serves as a proof of concept for applying our design concept on textual data and the decision making tasks exploration, creation, analysis, and comparison.

The public debate is another important source of knowledge to be considered in modern decision making. Assessing the content of numerous online discussions is challenging, reading the entire digitally available content is not possible. Our third technical contribution aimed at measuring the relevance of decision alternatives, arguments, and opinions via social media sources (Chapter 6). In collaboration with modeling experts from text mining research, we designed a visual analytics system that allows the exploration, analysis, and comparison in online discussions. The underlying text analysis workflow analyzes the frequency of online statements related to user-defined policies, arguments, and policy terms and uses these scores for estimating their relevance. The resulting information is presented to the users who can support the improvement of the algorithms' accuracy via interactive feedback. The presented visual analytics system is the third proof of concept for applying our design methodology on textual data. The system bridges knowledge gaps between unexperienced users like decision makers and stakeholders and modeling experts.

Besides textual data, empirical data was identified as an important data category to be incorporated into the decision making process. As a proof of concept for applying our design methodology on empirical data, we presented an approach that provides visual-interactive access to empirical performance indicators in the mining sector (Chapter 7). The underlying dataset was collected via desktop research and domain expert interviews. The visual analytics system allows the exploration and analysis of country-specific data, the creation of alternatives via similarity-based search (alternative countries to invest), and the comparison of country datasets. Decision makers from governments, investment companies, and the civil society are supported in investment-related decisions in the mining sector. The system was tested with domain experts at an African mining conference. It bridges knowledge gaps between domain experts and decision makers involved in mining-related decisions.

Another important task during the decision making process is the estimation of impacts of alternative solutions to a given problem. We addressed this task by combining visualization and simulation techniques in a visual analytics decision support system (Chapter 8). The resulting system allows users to simulate the impact of alternative subsidy strategies on the adoption of photovoltaic plants in private households. Decision makers and analysts can define alternative subsidy strategies and simulate the energy capacity being installed with an agent-based simulation model. The approach serves as a proof of concept for the application of our design methodology on model-driven data. The system supports exploration, creation, analysis, and comparison tasks. Moreover, knowledge gaps between domain experts on photovoltaic, simulation modeling experts, analysts, and decision makers are bridged. We proved the usability and usefulness of the system through a user evaluation.

Finally, we presented a second proof of concept for the applicability of our design methodology on model-driven data. The numerous factors to be considered in today's decision making processes often provoke conflicting characteristics. Hence, the balancing of trade-offs remains a challenge to be addressed by the decision maker. In Chapter 9, we contributed a visual analytics system that provides visual-interactive access to an optimization model. Users can specify targets and constraints and calculate optimal solutions, in this specific case, a regional energy plan that defines an optimal mixture of renewable energy sources to be installed by the local government. The visual analytics system allows the creation of optimal solutions, the analysis of these solutions including the consideration of environmental impacts, the comparison of optimal energy plans and the presentation of the results. That way, the system bridges knowledge gaps between domain experts providing knowledge on environmental impacts, modeling experts designing the optimization model, and analysts creating alternative solutions with the tool and presenting the results to decision makers. The visual analytics system was tested with users towards its usability and usefulness.

Summary of Contributions

In summary, we were able to address all research challenges formulated in this thesis (see Figure 10.1). The presented design process and the concept for bridging knowledge gaps was successfully applied to six scenarios, in which we proved its applicability with respect to the identified data, user, and task categories. We validated the efficiency and effectiveness of our visual analytics systems through (heuristic) justification and immense user testing. In total, three of our systems were tested with 57 users including 24 domain experts in 5 user testing rounds. In addition, the presented visual analytics approaches also address the general challenges specified in the introduction. Our solutions provide visual-interactive access to various datasets relevant for the decision making process. They reduce the complexity of the data through the application of basic visualization techniques and user-centered design. The quantity of the information is handled by providing overviews of the data, supporting filtering and zooming, and allowing users to drill-down to basic data entities. If applicable, our approaches also address the transparent assessment of data quality. The trust in the data analysis process is increased by making the analysis process transparent through the visualization of intermediate results. And finally, the usability of data and our visual analytics systems has been proven via design justifications or immense user testing. In the remainder of this chapter, we define future research directions that build upon this thesis.

10.2. Future Work

The presented contributions of this thesis serve as a baseline for several extensions to be addressed as future work. Future research directions on the individual visual analytics systems presented as technical contributions of this thesis are discussed in the respective chapters (Chapter 4 – 9)

Applying the design methodology to additional scenarios: In this thesis, we already successfully applied our design methodology on six decision making-related scenarios. These served as a proof of concept. In the future, we plan to apply the methodology to additional scenarios. In addition, we promote other researchers to apply our methodology. This has two positive effects. First, the methodology can be further tested towards its suitability. Second, if shortcomings are identified it can be improved. In our approach, we presented three visual analytics systems targeting textual data. Text as unstructured data category is difficult to handle. Text analysis methods that support humans in extracting knowledge from text are difficult to comprehend. Therefore, we see great potential for combining text analysis with visualization techniques to support decision making. Methods like text summarization, topic modeling, text classification, etc. provide great potential for informing the decision making process. The same holds for empirical datasets. Most of the related work on visual analytics addresses this data category, but only few explicitly target decision support. Furthermore, the combination of modeling approaches with visualization bears great potential for future research. For example, recently, tremendous research on deep neural networks has been conducted. Resulting approaches are also applied to the business context. However, the underlying methods are difficult to comprehend. We promote further research in combining visualization and neural networks to reduce the complexity of the algorithms and make them applicable in decision making scenarios.

Extending the design methodology to other types of unstructured data: The focus of this thesis has been laid on the analysis of textual, empirical, and model-driven data. However, additional data categories exist. For example, unstructured audio, video, and image data play an increasingly important role in decision making. Therefore, in the future it would be interesting to apply our methodology to these data categories. As an example existing algorithms allow the extraction of text from speech in video and audio data. This text can be further processed as described in this thesis.

Extending the concept for bridging gaps to other domains: In this thesis, we introduced a concept for bridging knowledge gaps between stakeholders involved in decision making. Collaborative visual analytics has already been targeted by related research approaches, e.g. by Isenberg et al. [IES*11]. However, besides time (asynchronous vs. synchronous) and space (co-located vs. distributed) the level of expertise (expert vs. non-expert) should be added to their model in order to support the communication, e.g., between scientists and decision makers. In this thesis, we presented a concept specifically targeting decision making. However, collaboration is critical for several domains. Therefore, we promote to further generalize the presented concept, or to define additional knowledge bridging concepts adapted to other domains.

Defining analysis workflow patterns for complexity reduction and result presentation: In this thesis, we realized the bridging of knowledge gaps via interactive visualization. We focused on basic visualization techniques in order to allow stakeholders with different expertise levels to communicate on the same information basis. In most cases, this implies a trade-off between analysis functionality and usability of the visualization. For example, in Chapter 5, we presented a visual-interactive text analysis workflow system that integrates multiple analysis features but can mainly be used by analysts with a background in data analysis. To mitigate this trade-off, we propose the design of workflow patterns that can be used to encode best practice workflows including parameterization, filter settings, etc. (e.g., as presented by the DimStiller approach [IMI*10]). These best practices can be recommended to novice users to simplify their access to the visual analytics systems. In addition, analytics process capturing methods can support analysts in the presentation of results to decision makers. Research in visual analytics story-telling can contribute to the dissemination of information and improve the transparency of the analysis process [KM13].

Integrating individual visual analytics components to single framework: In our approach, we presented different visual analytics systems that target different datasets and tasks. As a future task, these components could be integrated into a single system. This would allow using analysis results from one component as input for another. We have already published an integrated version that combines the simulation, the optimization, and an additional sentiment analysis component to a visual-interactive system [RBMK14]. However, we promote further research in integrating different datasets and tasks into a single system.

Realizing device-adaptive interfaces. Mobility plays in important role for decision makers. Access to information and analysis functionalities needs to be granted from different locations and devices. This poses challenges to visual analytics, since solutions need to be provided as web applications and on mobiles with small screens. From the technical perspective, the easiest solution is the realization of responsive design as commonly realized for web applications. However, since visualization design targets the exploitation of the available screen space, this is not always a trivial task. An alternative solution would be the reduction of the presented information. The analysis results can be aggregated to the most relevant information and provide simple infographics as output of the visual analytics process.

Acceleration of evaluation processes. Finally, as already discussed, decision making processes are time-critical. At the same time, the user-centered design process that ensures the usability and usefulness of visual analytics solutions is time-consuming. In order to provide useful and usable visual analytics solutions for decision making scenarios, evaluation processes need to be accelerated. As an option, current usability heuristics may be adapted to the decision making domain. Moreover, user interaction logging may support evaluation. Additionally, we recommend the evaluation of generic visualization techniques with different user groups, including non-expert users, in order to provide guidelines for the design of easy-to-use visual analytics systems.

A. Publications and Talks

The thesis is partially based on the following publications and talks:

A.1. Journal Publications and Book Chapters

1. Ruppert T., Dambruch J., Krämer M., Balke T., Gavanelli M., Bragaglia S., Chesani F., Milano M., Kohlhammer J.: Visual decision support for policy making: Advancing policy analysis with visualization. In: Policy Practice and Digital Science. Springer, 2015.
2. Bernard J., Daberkow D., Fellner D., Fischer K., Koepler O., Kohlhammer J., Runnwerth M., Ruppert T., Schreck T., Sens I.: Visinfo: a digital library system for time series research data based on exploratory search - a user-centered design approach. International Journal on Digital Libraries (IJDL) 16, 1 (2015), Springer.
3. Kohlhammer J., Nazemi K., Ruppert T., Burkhardt D.: Toward visualization in policy modeling. IEEE Computer Graphics and Applications 32, 5 (2012), IEEE Computer Society.
4. Bernard J., Brase J., Fellner D., Koepler O., Kohlhammer J., Ruppert T., Schreck T., Sens I.: A visual digital library approach for time-oriented scientific primary data. International Journal on Digital Libraries (IJDL) 11, 2 (2011), Springer.
5. Ruppert T., May T., Kohlhammer J., Schreck T.: Visuelle Analysen des Datensatzes: Wie versteckte Zusammenhänge sichtbar werden. In: Allgemeinbildung in Deutschland: Erkenntnisse aus dem SPIEGEL-Studentenpisa-Test. Springer, 2010.

A.2. Conference Papers

1. Ruppert T., Bannach A., Bernard J., Lokanc, M., Kohlhammer J.: Visual Access to Performance Indicators in the Mining Sector. In: Conference on Visualization (EuroVis) (2017), The Eurographics Association.
2. Ruppert T., Staab M., Bannach A., Lücke-Tieke H., Bernard J., Kuijper A., Kohlhammer J.: Visual interactive creation and validation of text clustering workflows to explore document collections. In: Visualization and Data Analysis (VDA), Electronic Imaging (2017), IS&T.
3. Ruppert T., Bannach A., Bernard J., Lücke-Tieke H., Ulmer A., Kohlhammer J.: Supporting collaborative political decision making - an interactive policy process visualization system. In: In-

ternational Symposium on Visual Information Communication and Interaction (VINCI) (2016), ACM. [Honorable Mention - “Selected Readings in Computer Graphics” at Fraunhofer IGD]

4. Ruppert T., Bernard J., Lücke-Tieke H., May T., Kohlhammer J.: Visual-interactive text analysis to support political decision making - from sentiments to arguments to policies. In: International Workshop on Visual Analytics (EuroVA) (2015), The Eurographics Association.
5. Ramfos A., Kiousi A., Kokkonidis M., Leclercq C., Mekkaoui D., Sattonnay M., Maragoudakis M., Androutsopoulou A., Charalabidis Y., Kohlhammer J., Ruppert T., Lücke-Tieke H., Dimakopoulos N., Kallipolitis L., Nikodem P., Madlenak T., Mureddu F., Pyrenis D., Protonotarios M., Ipektsidis C.: The “EU Community” Project - coupling the power of data with community expertise. In: Workshop on Enabling Effective Policy Making 2015 - Coupling the Power of Data with the Wisdom of the Crowd (EEPM) co-located with IFIP Electronic Government and Electronic Participation Conference (eGOV) (2015), CEUR-WS.
6. Koldijk S., Bernard J., Ruppert T., Kohlhammer J., Neerincx M., Kraaij W.: Visual analytics of work behavior data - insights on individual differences. In: Conference on Visualization (EuroVis) (2015), The Eurographics Association.
7. Ruppert T., Bernard J., May T., Kohlhammer J.: Combining computational models and interactive visualization to support rational decision making. In: International Symposium on Visual Computing (ISVC) (2014), Springer.
8. Ruppert T., Bernard J., Ulmer A., Lücke-Tieke H., Kohlhammer J.: Visual access to an agent-based simulation model to support political decision making. In: International Conference on Knowledge Management and Knowledge Technologies (i-KNOW) (2014), ACM. [BEST PAPER AWARD]
9. Bernard J., Sessler D., Ruppert T., Davey J., Kuijper A., Kohlhammer J.: User-based visual interactive similarity definition for mixed data objects - concept and first implementation. In: Journal of Winter School of Computer Graphics (WSCG) (2014)
10. Ruppert T., Bernard J., Ulmer A., Kuijper A., Kohlhammer J.: Visual access to optimization problems in strategic environmental assessment. In: International Symposium on Visual Computing (ISVC) (2013), Springer.
11. Ruppert T., Bernard J., Kohlhammer J.: Bridging knowledge gaps in policy analysis with information visualization. In: Electronic Government and Electronic Participation: Joint Proceedings of Ongoing Research of IFIP EGOV and IFIP ePart (2013), GI-Edition - Lecture Notes in Informatics (LNI).
12. Steiger M., Krämer M., Ruppert T., Kohlhammer J.: Visualizing uncertain underground information for urban management. In: International Conference Computer Graphics, Visualization, Computer Vision and Image Processing (CGVCVIP) (2013), IADIS Press.
13. Bernard J., Ruppert T., Goroll O., May T., Kohlhammer J.: Visual-interactive preprocessing of time series data. In: SIGRAD, Interactive Visual Analysis of Data (2012), Linköping University Electronic Press.

14. Bernard J., Ruppert T., Scherer M., Kohlhammer J., Schreck T.: Content-based layouts for exploratory metadata search in scientific research data. In: Joint Conference on Digital Libraries (JCDL) (2012), ACM. [BEST STUDENT PAPER AWARD]
15. Bernard J., Ruppert T., Scherer M., Schreck T., Kohlhammer J.: Guided discovery of interesting relationships between time series clusters and metadata properties. In: International Conference on Knowledge Management and Knowledge Technologies (i-KNOW) (2012), ACM.
16. Burkhardt D., Ruppert T., Nazemi K.: Towards process-oriented information visualization for supporting users. In: International Conference on Interactive Collaborative Learning (ICL) (2012), IEEE Computer Society.
17. May T., Bannach A., Davey J., Ruppert T., Kohlhammer J.: Guiding feature subset selection with an interactive visualization. In: IEEE Conference on Visual Analytics Science and Technology (VAST) (2011), IEEE Computer Society.
18. May T., Davey J., Ruppert T.: Smartstripes - looking under the hood of feature subset selection methods. In: International Workshop on Visual Analytics (EuroVA) (2011), The Eurographics Association.
19. Bernard J., Brase J., Fellner D., Koepler O., Kohlhammer J., Ruppert T., Schreck T., Sens I.: A visual digital library approach for time-oriented scientific primary data. In: Research and Advanced Technology for Digital Libraries: European Conference on Digital Libraries (ECDL) (2010), Springer.
20. Kohlhammer J., Ruppert T., Davey J., Mansmann F., Keim D.: Information visualisation and visual analytics for governance and policy modelling. In: CROSSROAD Call for Contributions on Future Internet on ICT for Governance and Policy Modelling (2010)
21. Krämer M., Ruppert T., Klien E., Kohlhammer J.: DeepCity3D: Integration von 3D-Stadtmodellen und Untergrundinformationen. In: Geoinformatik 2010. Die Welt im Netz. Konferenzband (2010), Akademische Verlagsgesellschaft AKA.
22. Heinze T., Ruppert T., Korn P., Bonaventura L.: Velocity reconstruction by radial basis functions in a triangular staggered c grid. In: Working Group on Numerical Experimentation (WGNE) Blue Book: Research Activities in Atmospheric and Ocean Modelling (2007), World Climate Research Programme (WCRP)

A.3. Posters

1. Ruppert T., Bernard J., Lücke-Tieke H., Kohlhammer J.: Towards a tighter coupling of visualization and public policy making. In: IEEE Conference on Visual Analytics Science and Technology (VAST) (2014), IEEE Computer Society.
2. Schreck T., Sharalieva L., Wanner F., Bernard J., Ruppert T., von Landesberger T., Bustos B.: Visual exploration of local interest points in sets of time series. In: IEEE Conference on Visual Analytics Science and Technology (VAST) (2012), IEEE Computer Society.

3. Dummer M., Krämer M., Ruppert T., Kohlhammer J.: Visualizing uncertain underground information for urban management. In: Working with Uncertainty Workshop (IEEE VisWeek 2011) (2011), IEEE Computer Society.
4. Krämer M., Dummer M., Ruppert T., Kohlhammer J.: Tackling Uncertainty in Combined Visualizations of Underground Information and 3D City Models. In: GeoViz Hamburg 2011 Workshop: Linking Geovisualization with Spatial Analysis and Modeling (2011)
5. Ruppert T., Kohlhammer J.: A radial visualization tool for depicting hierarchically structured video content. In: IEEE Symposium on Visual Analytics Science and Technology (VAST) (2010), IEEE Computer Society. [BEST POSTER AWARD]
6. Ruppert T., Bremm S., Schreck T., Kohlhammer J.: VisServer - a visual analytics framework. In Eurographics/IEEE Symposium on Visualization (EuroVis) (2009), The Eurographics Association.

A.4. Talks

1. Tobias Ruppert, "Information Visualisation & Visual Analytics for Policy Making", at 1st international Summer School on Open & Collaborative Governance, https://egov2013.pns.aegean.gr, 1st – 6th July, Samos, Greece, 2013.

B. Supervising Activities

The following list summarizes the student bachelor, diploma and master thesis supervised by the author. The results of these works were partially used as an input into the thesis.

B.1. Master Theses

1. Vishal Lakhani, Tobias Ruppert (Supervisor): Visual-Interactive Document Classification, TU Darmstadt, 2016.
2. Michael Staab, Tobias Ruppert (Supervisor): Visuell-Interaktive Exploration von Text Clustering Ergebnissen, TU Darmstadt, 2016.
3. Maximilian Pohst, Tobias Ruppert (Supervisor), James Davey (Supervisor): Definition eines Evaluierungs- prozesses für Visual-Analytics-Expertenlösungen, TU Darmstadt, 2013.

B.2. Bachelor Theses

1. Michael Maus, Tobias Ruppert (Supervisor), Hendrik Lücke-Tieke (Supervisor): Definition und Visualisierung von zoombaren 2D-Projektionen im Web, TU Darmstadt, 2016.
2. Alex Ulmer, Tobias Ruppert (Supervisor): Visual Analysis of Multidimensional Optimization Problems, TU Darmstadt, 2013.

B.3. Internships

1. Alex Ulmer, Tobias Ruppert (Supervisor): Visueller Zugang zu Agenten-basierten Simulations-modellen, TU Darmstadt, 2014.
2. Alex Ulmer, Dennis Basgier, Marcel Weiler, Tobias Ruppert (Supervisor): Webapplikationen zur visuell-interaktiven Zeitserienanalyse, TU Darmstadt, 2012.

C. Curriculum Vitae

Personal Data

Name	Tobias Ruppert
Birth date	30th October, 1981
Birth place	Flörsheim am Main, Germany
Family status	Unmarried
Nationality	German

Education

2012 – 2017	PhD Student at Interactive Systems Group Faculty of Computer Science, Technische Universität Darmstadt, Germany Focus: Visual Analytics for Decision Support Publications: Google Scholar Profile
2008	Graduation (Diploma) in Mathematics Faculty of Mathematics, Technische Universität Darmstadt, Germany
2007	Diploma Thesis "Vector Field Reconstruction by Radial Basis Functions" Deutscher Wetterdienst (DWD), Offenbach, Germany
2001 – 2008	Study of Mathematics Faculty of Mathematics, Technische Universität Darmstadt, Germany

Work Experience

2008 – 2017	Researcher, Department Information Visualization and Visual Analytics Fraunhofer Institute for Computer Graphics Research IGD, Germany Focus: Information Visualization and Visual Analytics

Darmstadt, 14^{th} of September, 2017

Bibliography

[AA03] ANDRIENKO N., ANDRIENKO G.: Informed spatial decisions through coordinated views. *Information Visualization 2*, 4 (2003). doi:10.1057/palgrave.ivs.9500058. 41

[AAB*10] ANDRIENKO G., ANDRIENKO N., BREMM S., SCHRECK T., VON LANDESBERGER T., BAK P., KEIM D.: Space-in-time and time-in-space self-organizing maps for exploring spatiotemporal patterns. In *Conference on Visualization (EuroVis)* (2010), The Eurographics Association. doi:10.1111/j.1467-8659.2009.01664.x. 143

[AAJ*07] ANDRIENKO G., ANDRIENKO N., JANKOWSKI P., KEIM D., KRAAK M. J., MACEACHREN A., WROBEL S.: Geovisual analytics for spatial decision support: Setting the research agenda. *International Journal of Geographical Information Science 21*, 8 (2007). doi:10.1080/13658810701349011. 17, 38, 41

[AAMH13] ALSALLAKH B., AIGNER W., MIKSCH S., HAUSER H.: Radial sets: Interactive visual analysis of large overlapping sets. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 19*, 12 (2013). doi:10.1109/TVCG.2013.184. 107

[ABD08] ALTHAUS C., BRIDGMAN P., DAVIS G.: *The Australian Policy Handbook*. Allen and Unwin, 2008. 20

[AES05] AMAR R., EAGAN J., STASKO J.: Low-level components of analytic activity in information visualization. In *IEEE Symposium on Information Visualization (InfoVis)* (2005), IEEE Computer Society. doi:10.1109/infvis.2005.1532136. 39

[Aga13] AGARWAL R. K.: Assessment and optimization of an airplane's environmental impact. *Aircraft Engineering and Aerospace 85* (2013). doi:10.1108/00022661311313632. 158

[Alt80] ALTER S.: *Decision support systems: current practice and continuing challenges*. Addison-Wesley Publ., 1980. 12, 13, 14

[AMA*16] ALSALLAKH B., MICALLEF L., AIGNER W., HAUSER H., MIKSCH S., RODGERS P.: The state-of-the-art of set visualization. *Computer Graphics Forum 35*, 1 (2016). doi:10.1111/cgf.12722. 97, 107

[AME11] AFZAL S., MACIEJEWSKI R., EBERT D. S.: Visual analytics decision support environment for epidemic modeling and response evaluation. In *IEEE Conference on Visual Analytics Science and Technology (VAST)* (2011), IEEE Computer Society. doi:10.1109/VAST.2011.6102457. 18, 142

[AMST11] Aigner W., Miksch S., Schumann H., Tominski C.: *Visualization of Time-Oriented Data*. Human-Computer Interaction. Springer, 2011. doi:10.1007/978-0-85729-079-3. 78

[AMTB05] Aigner W., Miksch S., Thurnher B., Biffl S.: Planninglines: Novel glyphs for representing temporal uncertainties and their evaluation. In *IEEE Conference on Information Visualization (IV)* (2005), IEEE Computer Society. doi:10.1109/IV.2005.97. 78

[And75] Anderson J. E.: *Public Policymaking*, 7 (2010) ed. Wadsworth Publishing, 1975. 19, 20, 66, 67

[And08] Andrews K.: Evaluation comes in many guises. In *Workshop on BEyond time and errors: novel evaLuation methods for Information Visualization (BELIV)* (2008), ACM. doi:10.1145/1168149.1168151. 32, 61, 62, 64, 169, 171

[Ant65] Anthony R. N.: *Planning and Control Systems: A Framework for Analysis*. Harvard University Press, 1965. 13

[ARDM11] Assogba Y., Ros I., DiMicco J., McKeon M.: Many bills: engaging citizens through visualizations of congressional legislation. In *SIGCHI Conference on Human Factors in Computing Systems (CHI)* (2011), ACM, ACM. doi:10.1145/1978942.1979004. 78

[AS04] Amar R., Stasko J.: A knowledge task-based framework for design and evaluation of information visualizations. In *IEEE Symposium on Information Visualization (InfoVis)* (2004), IEEE Computer Society. doi:10.1109/infvis.2004.10. 39

[AV07] Arthur D., Vassilvitskii S.: k-means++: The advantages of careful seeding. In *ACM-SIAM Symposium on Discrete Algorithms (SODA)* (2007), Society for Industrial and Applied Mathematics, ACM. 100

[Bel09] Bell G.: Foreword. In *The fourth paradigm: data-intensive scientific discovery*, Hey T., Tansley S., Tolle K., (Eds.). Microsoft Corporation, 2009. 2

[Ber83] Bertin J.: *Semiology of Graphics: Diagrams, Networks, Maps [orig. Semiologie Graphique]*. University of Wisconsin Press, 1983. 27, 35

[Ber15] Bernard J.: *Exploratory Search in Time-Oriented Primary Data*. PhD thesis, Technische Universität Darmstadt, 2015. URL: http://tuprints.ulb.tu-darmstadt.de/5173/. 38

[BISM14] Brehmer M., Ingram S., Stray J., Munzner T.: Overview: The design, adoption, and analysis of a visual document mining tool for investigative journalists. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 20*, 12 (2014). doi:10.1109/TVCG.2014.234631. 97

[BKW16] Beck F., Koch S., Weiskopf D.: Visual analysis and dissemination of scientific literature collections with survis. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 22*, 1 (2016). doi:10.1109/TVCG.2015.2467757. 97

[BL10] Bertini E., Lalanne D.: Investigating and reflecting on the integration of automatic data analysis and visualization in knowledge discovery. *SIGKDD Explorations Newsletter 11* (2010). doi:10.1145/1809400.1809404. 67

[BM13] Brehmer M., Munzner T.: A multi-level typology of abstract visualization tasks. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 19*, 12 (2013). doi:10.1109/TVCG.2013.124. 4, 39

[BMPM12] Booshehrian M., Möller T., Peterman R. M., Munzner T.: Vismon: Facilitating analysis of trade-offs, uncertainty, and sensitivity in fisheries management decision making. *Computer Graphics Forum 31* (2012). doi:10.1111/j.1467-8659.2012.03116.x. 41, 68, 142

[BNJ03] Blei D. M., Ng A. Y., Jordan M. I.: Latent dirichlet allocation. *Journal of Machine Learning Research 3* (2003). 97

[BRBF14] Boy J., Rensink R. A., Bertini E., Fekete J. D.: A principled way of assessing visualization literacy. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 20*, 12 (2014). doi:10.1109/TVCG.2014.2346984. 4

[Bre74] Brewer G. D.: The policy sciences emerge: To nurture and structure a discipline. *Policy Sciences 5*, 3 (1974). doi:10.1007/BF00144283. 20

[BRG*12] Bernard J., Ruppert T., Goroll O., May T., Kohlhammer J.: Visual-interactive preprocessing of time series data. In *SIGRAD, Interactive Visual Analysis of Data* (2012), Linköping University Electronic Press. URL: http://www.ep.liu.se/ecp/article.asp?issue=081&article=006. 143

[BvLBS11] Bremm S., von Landesberger T., Bernard J., Schreck T.: Assisted descriptor selection based on visual comparative data analysis. In *Conference on Visualization (EuroVis)* (2011), The Eurographics Association. doi:10.1111/j.1467-8659.2011.01938.x. 143

[CC00] Cox T., Cox M.: *Multidimensional scaling*. CRC Press, 2000. doi:10.1007/978-3-540-33037-0_14. 101

[CCS12] Chen H., Chiang R., Storey V.: Business intelligence and analytics: From big data to big impact. *MIS Quarterly: Management Information Systems 36*, 4 (2012). 2, 16, 24, 47

[CISSW06] Chen C., Ibekwe-SanJuan F., SanJuan E., Weaver C.: Visual analysis of conflicting opinions. In *IEEE Symposium on Visual Analytics Science and Technology (VAST)* (2006), IEEE Computer Society. doi:10.1109/VAST.2006.261431. 117

[CK12] Crouser R. J., Kee D.: Two visualization tools for analyzing agent-based simulations in political science. *IEEE Computer Graphics and Applications* (2012). doi:10.1109/MCG.2011.90. 142

[CL04] Carenini G., Lloyd J.: ValueCharts: Analyzing Linear Models Expressing Preferences and Evaluations. *Working Conference on Advanced Visual Interfaces (AVI)* (2004).

doi:10.1145/989863.989885. 41

[CLC14] CHARLES L. COCHRAN E. F. M.: *Public policy: perspectives and choices*, 5 ed. Lynne Rienner Publishers, 2014. 19

[CLL*13] CHOO J., LEE H., LIU Z., STASKO J., PARK H.: An interactive visual testbed system for dimension reduction and clustering of large-scale high-dimensional data. In *Proceedings of SPIE, Visualization and Data Analysis (VDA)* (2013), vol. 8654, SPIE. doi:10.1117/12.2007316. 96, 98

[CLRP13] CHOO J., LEE C., REDDY C. K., PARK H.: UTOPIAN: User-driven topic modeling based on interactive nonnegative matrix factorization. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 19*, 12 (2013). doi:10.1109/TVCG.2013.212. 97

[CLT*11] CUI W., LIU S., TAN L., SHI C., SONG Y., GAO Z., QU H., TONG X.: Textflow: Towards better understanding of evolving topics in text. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 17*, 12 (2011). doi:10.1109/TVCG.2011.239. 97

[CM12] CURTOTTI M., MCCREATH E.: Enhancing the visualization of law. In *Law via the Internet Twentieth Anniversary Conference, Cornell University* (2012). doi:10.2139/ssrn.2160614. 78

[CMS99] CARD S., MACKINLAY J., SHNEIDERMAN B. (Eds.): *Readings in information visualization: using vision to think*. Morgan Kaufmann Publishers, San Franc. CA, USA, 1999. 27, 28, 29, 35, 38, 39, 40, 44

[Cou01] COURTNEY J. F.: Decision making and knowledge management in inquiring organizations: toward a new decision-making paradigm for DSS. *Decision Support Systems 31*, 1 (2001). doi:10.1016/S0167-9236(00)00117-2. 2, 17

[DCS05] DALAL-CLAYTON D., SADLER B.: *Strategic Environmental Assessment: A Sourcebook and Reference Guide to International Experience*. Earthscan, 2005. 157, 158

[DLT16] DE MARCHI G., LUCERTINI G., TSOUKIAS A.: From evidence-based policy making to policy analytics. *Annals of Operations Research 236*, 15 (2016). doi:10.1007/s10479-014-1578-6. 24, 37

[DMRI16] DANIELL K. A., MORTON A., RÍOS INSUA D.: Policy analysis and policy analytics. *Annals of Operations Research 236*, 1 (2016). doi:10.1007/s10479-015-1902-9. 1, 24, 169

[DNKS10] DIAKOPOULOS N., NAAMAN M., KIVRAN-SWAINE F.: Diamonds in the rough: Social media visual analytics for journalistic inquiry. In *IEEE Symposium on Visual Analytics Science and Technology (VAST)* (2010), IEEE Computer Society. doi:10.1109/VAST.2010.5652922. 117

[DWCR11] DOU W., WANG X., CHANG R., RIBARSKY W.: Paralleltopics: A probabilistic approach to exploring document collections. In *IEEE Conference on Visual Analytics Sci-*

ence and Technology (VAST) (2011), IEEE Computer Society. doi:10.1109/VAST.2011.6102461. 97

[Elk97] ELKINGTON J.: *Cannibals with Forks: The Triple Bottom Line of 21st Century Business (The Conscientious Commerce Series)*. Capstone Publishing, 1997. 1, 53

[Eng05] ENGELS A.: The science-policy interface. *The Integrated Assessment Journal 5*, 1 (2005). URL: http://journals.sfu.ca/int_assess/index.php/iaj/article/viewArticle/165. 23

[Eur10] EUROPEAN COMMISSION: 7th framework programme for research and technological development. Online, 2010. URL: http://ec.europa.eu/research/fp7/index_en.cfm. 30

[Eur16] EUROPEAN UNION: EUR-Lex - Access to European Union Law. Online, accessed in 2016. URL: http://eur-lex.europa.eu. 79, 85, 90

[Few09] FEW S.: *Now You See it: Simple Visualization Techniques for Quantitative Analysis*. Analytics Press, 2009. 26, 27, 67, 107, 117, 147, 160, 168

[FH11] FENG V. W., HIRST G.: Classifying arguments by scheme. In *Annual Meeting of the Association for Computational Linguistics: Human Language Technologies (HLT)* (2011), vol. 1, Association for Computational Linguistics. 117

[FHN*93] FOX E. A., HIX D., NOWELL L. T., BRUENI D. J., RAO D., WAKE W. C., HEATH L. S.: Users, user interfaces, and objects: Envision, a digital library. *Journal of the American Society for Information Science 44*, 8 (1993). doi:10.1002/(SICI)1097-4571(199309)44:8<480::AID-ASI7>3.0.CO;2-B. 79

[Fis07] FISCHER T. B.: *The theory and practice of strategic environmental assessment: towards a more systematic approach*. Earthscan, 2007. 22, 157, 158

[FPSS96] FAYYAD U., PIATETSKY-SHAPIRO G., SMYTH P.: From data mining to knowledge discovery in databases. *AI Magazine 17* (1996). 29, 31

[FS06] FELDMAN R., SANGER J.: *Text Mining Handbook: Advanced Approaches in Analyzing Unstructured Data*. Cambridge University Press, New York, NY, USA, 2006. 117

[GACOR05] GAMON M., AUE A., CORSTON-OLIVER S., RINGGER E.: Pulse: Mining customer opinions from free text. In *International Conference on Advances in Intelligent Data Analysis (IDA)* (2005), Springer. doi:10.1007/11552253_12. 133

[GC06] GREENE D., CUNNINGHAM P.: Practical solutions to the problem of diagonal dominance in kernel document clustering. In *International Conference on Machine Learning (ICML)* (2006). doi:10.1145/1143844.1143892. 108

[GGG10] GUILLEN-GOSALBEZ G., GROSSMANN I.: A global optimization strategy for the environmentally conscious design of chemical supply chains under uncertainty in the damage assessment model. *Computers & Chemical Engineering 34*, 1 (2010). doi:10.1016/j.compchemeng.2009.09.003. 158

[GGMZ05] GUO D., GAHEGAN M., MACEACHREN A. M., ZHOU B.: Multivariate analysis and geovisualization with an integrated geographic knowledge discovery approach. *Cartography and Geographic Information Science 32*, 2 (2005). arXiv:http://www.tandfonline.com/doi/pdf/10.1559/1523040053722150, doi:10.1559/1523040053722150. 143

[Gil08] GILBERT G. N.: *Agent-Based Models*. Quantitative applications in the social sciences. SAGE Publications, 2008. 143

[GLG*13] GRATZL S., LEX A., GEHLENBORG N., PFISTER H.-P., STREIT M.: LineUp: Visual analysis of multi-attribute rankings. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 19*, 12 (2013). doi:10.1109/TVCG.2013.173. 41

[GLK*13] GÖRG C., LIU Z., KIHM J., CHOO J., PARK H., STASKO J.: Combining computational analyses and interactive visualization for document exploration and sensemaking in jigsaw. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 19*, 10 (2013). doi:10.1109/TVCG.2012.324. 97

[GLPK14] GOUDAS T., LOUIZOS C., PETASIS G., KARKALETSIS V.: Argument extraction from news, blogs, and social media. In *Artificial Intelligence: Methods and Applications - 8th Hellenic Conference on AI, SETN* (2014), Likas A., Blekas K., Kalles D., (Eds.), Springer. doi:10.1007/978-3-319-07064-3_23. 119

[GMHO12] GAVANELLI M., MILANO M., HOLLAND A., O'SULLIVAN B.: What-if analysis through simulation - optimization hybrids. In *European Conference on Modelling and Simulation (ECMS)* (2012), Troitzsch K. G., Möhring M., Lotzmann U., (Eds.). 141, 143, 157, 159, 168

[GSM71] GORRY G. A., SCOTT MORTON M. S.: A Framework for Management Information Systems. *MIT Sloan Management Review 13*, 1 (1971). 12, 13

[Guo03] GUO D.: Coordinating computational and visual approaches for interactive feature selection and multivariate clustering. *Information Visualization 2*, 4 (2003). doi:10.1057/palgrave.ivs.9500053. 96

[HB03] HARROWER M., BREWER C. A.: Colorbrewer.org: An online tool for selecting colour schemes for maps. *The Cartographic Journal 40*, 1 (2003). doi:10.1179/000870403235002. 129, 162

[HCSG01] HILL L. L., CROSIER S. J., SMITH T. R., GOODCHILD M.: A content standard for computational models. *D-Lib Magazine 7*, 6 (2001). URL: http://www.dlib.org/dlib/june01/hill/06hill.html. 15, 55

[Hea09] HEARST M. A.: *Search User Interfaces*. Cambridge University Press, 2009. 79

[Hea12] HEAPS C. G.: Long-range Energy Alternatives Planning (LEAP) system. Stockholm Environment Institute. Somerville, MA, USA., accessed in 2012. URL: www.energycommunity.org. 68, 158

[HHKE16] HEIMERL F., HAN Q., KOCH S., ERTL T.: CiteRivers: visual analytics of citation patterns. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 22*, 1 (2016). doi:10.1109/TVCG.2015.2467621. 97

[HHWN02] HAVRE S., HETZLER E., WHITNEY P., NOWELL L.: Themeriver: Visualizing thematic changes in large document collections. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 8*, 1 (2002). doi:10.1109/2945.981848. 78

[HK12] HERRMANNOVA D., KNOTH P.: Visual search for supporting content exploration in large document collections. *D-Lib Magazine 18*, 7/8 (2012). doi:10.1045/july2012-herrmannova. 78

[Hop99] HOPPE R.: Policy analysis, science and politics: from 'speaking truth to power' to 'making sense together'. *Science and Public Policy 26*, 3 (1999). doi:10.3152/147154399781782482. 21

[Hov07] HOVE, SYBILLE VAN DEN: A rationale for science-policy interfaces. *Futures - The journal of policy, planning, and futures studies 39* (2007). doi:10.1016/j.futures.2006.12.004. 2, 23

[How09] HOWLETT M.: Policy analytical capacity and evidence-based policy-making: Lessons from canada. *Canadian Public Administration 52*, 2 (2009). doi:10.1111/j.1754-7121.2009.00070_1.x. 21, 22

[HPS15] HÖCHTL J., PARYCEK P., SCHÖLLHAMMER R.: Big data in the policy cycle: Policy decision making in the digital era. *Journal of Organizational Computing and Electronic Commerce 26*, 1-2 (2015). doi:10.1080/10919392.2015.1125187. 21, 24

[HRP09] HOWLETT M., RAMESH M., PERL A.: *Studying Public Policy: Policy Cycles and Policy Subsystems*, 3 ed. Oxford University Press, 2009. 20, 21, 37, 141

[HW09] HOWLETT M., WELLSTEAD A.: Re-Visiting Meltsner: Policy Advice Systems and the Multi-Dimensional Nature of Professional Policy Analysis. *Available at SSRN 1546251* (2009). doi:10.2139/ssrn.1546251. 23, 37

[IES*11] ISENBERG P., ELMQVIST N., SCHOLTZ J., CERNEA D., MA K.-L., HAGEN H.: Collaborative visualization: Definition, challenges, and research agenda. *Information Visualization 10*, 4 (2011). doi:10.1177/1473871611412817. 4, 64, 175

[IH01] IVORY M. Y., HEARST M. A.: The state of the art in automating usability evaluation of user interfaces. *ACM Computing Surveys 33*, 4 (2001). doi:10.1145/503112.503114. 63

[IMI*10] INGRAM S., MUNZNER T., IRVINE V., TORY M., BERGNER S., MÖLLER T.: Dimstiller: Workflows for dimensional analysis and reduction. In *IEEE Symposium on Visual Analytics Science and Technology (VAST)* (2010), IEEE Computer Society. doi:10.1109/VAST.2010.5652392. 69, 175

[Ind17] INDABA M.: Investing in African Mining Indaba. Online, 2017. URL: (https://www.miningindaba.com/. 136

[IPRE08] INGARAMO D., PINTO D., ROSSO P., ERRECALDE M.: Evaluation of internal validity measures in short-text corpora. In *International Conference on Computational Linguistics and Intelligent Text Processing (CICLing)* (2008), Springer. doi:10.1007/978-3-540-78135-6_48. 101

[IZCC08] ISENBERG P., ZUK T., COLLINS C., CARPENDALE S.: Grounded evaluation of information visualizations. In *Workshop on BEyond time and errors: novel evaLuation methods for Information Visualization (BELIV)* (2008), ACM. doi:10.1145/1377966.1377974. 32, 62

[Jai10] JAIN A. K.: Data clustering: 50 years beyond k-means. *Pattern recognition letters 31*, 8 (2010). doi:10.1007/978-3-540-87479-9_3. 95

[JBK14] JOHNSON P. G., BALKE T., KOTTHOF L.: Integrating optimisation and agent-based modelling. In *European Conference on Modelling and Simulation (ECMS)* (2014). doi:10.1016/j.procs.2014.03.097. 144

[Jon96] JONES C.: Visualization and optimization. *Interactive Transactions on ORMS (INFORMS) 2*, 1 (1996). URL: http://www.informs.org/Pubs/ITORMS/Archive/Volume-2/No.-1-Jones. 158

[JS90] JENKINS-SMITH H. C.: *Democratic politics and policy analysis*. Brooks/Cole Pacific Grove, 1990. 21

[JS91] JOHNSON B., SHNEIDERMAN B.: Tree-maps: A space-filling approach to the visualization of hierarchical information structures. In *IEEE Conference on Visualization (VIS)* (1991), IEEE Computer Society. doi:10.1109/VISUAL.1991.175815. 133

[JW07] JANN W., WEGRICH K.: Chapter 4: Theories of the policy cycle. In *Handbook of Public Policy Analysis: Theory, Politics, and Methods*, Fischer F., Miller G., Sidney M., (Eds.). CRC Press, 2007. doi:10.1201/9781420017007.pt2. 20

[KAF*08] KEIM D., ANDRIENKO G., FEKETE J.-D., GÖRG C., KOHLHAMMER J., MELANÇON G.: Visual analytics: Definition, process, and challenges. In *Information Visualization*, Lecture Notes in Computer Science (LNCS). Springer, 2008. doi:10.1007/978-3-540-70956-5_7. 17, 29, 30, 38, 44

[KBH06] KOSARA R., BENDIX F., HAUSER H.: Parallel sets: Interactive exploration and visual analysis of categorical data. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 12*, 4 (2006). doi:10.1109/TVCG.2006.76. 97, 105, 107

[KBSC03] KIRSCHNER P., BUCKINGHAM-SHUM S., CARR C.: *Visualizing Argumentation: Software Tools for Collaborative and Educational Sense-Making*. Computer Supported Cooperative Work. Springer, 2003. 117

[Kei02] KEIM D. A.: Information visualization and visual data mining. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 7*, 1 (2002). doi:10.1109/2945.981847. 35, 36

[KGP*14] KIOMOURTZIS G., GIANNAKOPOULOS G., PETASIS G., KARAMPIPERIS P., KARKALETSIS V.: NOMAD: Linguistic Resources and Tools Aimed at Policy Formulation and Validation. In *International Conference on Language Resources and Evaluation* (2014), European Language Resources Association (ELRA). 118

[KKEM10] KEIM D., KOHLHAMMER J., ELLIS G., MANSMANN F.: *Mastering the Information Age - Solving Problems with Visual Analytics*. The Eurographics Association, 2010. 17, 29, 40, 44, 141

[KM13] KOSARA R., MACKINLAY J.: Storytelling: The next step for visualization. *IEEE Computer 46*, 5 (2013). doi:10.1109/MC.2013.36. 4, 30, 175

[KMH09] KOHLHAMMER J., MAY T., HOFFMANN M.: Visual analytics for the strategic decision making process. In *GeoSpatial Visual Analytics: Geographical Information Processing and Visual Analytics for Environmental Security*, Amicis R. D., Stojanovic R., Conti G., (Eds.). Springer, 2009. doi:10.1007/978-90-481-2899-0_23. 18

[KNO*03] KASKI S., NIKKILÄ J., OJA M., VENNA J., TÖRÖNEN P., CASTRÉN E.: Trustworthiness and metrics in visualizing similarity of gene expression. *BMC Bioinformatics 4*, 1 (2003). doi:10.1186/1471-2105-4-48. 101

[KNRB12] KOHLHAMMER J., NAZEMI K., RUPPERT T., BURKHARDT D.: Toward visualization in policy modeling. *IEEE Computer Graphics and Applications 32*, 5 (2012). doi:10.1109/MCG.2012.107. 21, 30, 41, 43, 77

[KOCZ93] KLEIN G., ORASANU J., CALDERWOOD R., ZSAMBOK C.: *Decision Making in Action: Models and Methods*. Ablex Publishing Corporation, 1993. 17

[KPB14] KRAUSE J., PERER A., BERTINI E.: Infuse: Interactive feature selection for predictive modeling of high dimensional data. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 20*, 12 (2014). doi:10.1109/TVCG.2014.2346482. 96

[KPHH12] KANDEL S., PAEPCKE A., HELLERSTEIN J. M., HEER J.: Enterprise data analysis and visualization: An interview study. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 18* (2012). doi:10.1109/TVCG.2012.219. 37, 40

[KPW13] KOHLHAMMER J., PROFF D. U., WIENER A.: *Visual Business Analytics: Effektiver Zugang zu Daten und Informationen*. dpunkt.verlag, 2013. 17, 54

[KSM78] KEEN G. W., SCOTT MORTON M. S.: *Decision support systems: an organizational perspective*. Addison-Wesley Publ., 1978. 15

[KWR09] KORNHAUSER D., WILENSKY U., RAND W.: Design guidelines for agent based model visualization. *Journal of Artificial Societies and Social Simulation 12*, 2 (2009). URL: http://jasss.soc.surrey.ac.uk/12/2/1.html. 142

[KWS*14] KONEV A., WASER J., SADRANSKY B., CORNEL D., PERDIGAO R. A. P., HORVATH Z., GROELLER M. E.: Run Watchers: Automatic simulation-based decision support in flood management. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 20*, 12 (2014). doi:10.1109/TVCG.2014.2346930. 41

[Las56] LASSWELL H. D.: *The Decision Process: Seven Categories of Functional Analysis*. University of Maryland, College Park, 1956. 19

[LBI*11] LAM H., BERTINI E., ISENBERG P., PLAISANT C., CARPENDALE S.: Empirical Studies in Information Visualization: Seven Scenarios. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 18*, 9 (2011). doi:10.1109/TVCG.2011.279. 32, 33, 44

[LC10] LIN F., COHEN W. W.: Power iteration clustering. In *International Conference on Machine Learning (ICML)* (2010). 100

[LDWB10] LIU S., DUFFY A. H. B., WHITFIELD R. I., BOYLE I. M.: Integration of decision support systems to improve decision support performance. *Knowledge and Information Systems (KAIS) 22*, 3 (2010). doi:10.1007/s10115-009-0192-4. 17

[LHC05] LIU B., HU M., CHENG J.: Opinion observer: analyzing and comparing opinions on the web. In *International Conference on World Wide Web* (2005), ACM, ACM. doi:10.1145/1060745.1060797. 117

[Lin11a] LINDQUIST E.: Grappling with Complex Policy Challenges: exploring the potential of visualization for analysis, advising and engagement. *HC Coombs Policy Forum, The Australian National University* (2011). URL: https://coombs-forum.crawford.anu.edu.au/publication/hc-coombs-policy-forum/2497/grappling-complex-policy-challenges-exploring-potential. 25

[Lin11b] LINDQUIST E.: Surveying the World of Visualization. *HC Coombs Policy Forum, The Australian National University* (2011). URL: https://coombs-forum.crawford.anu.edu.au/publication/hc-coombs-policy-forum/2498/surveying-world-visualization. 25

[Liu07] LIU B.: *Web data mining: exploring hyperlinks, contents, and usage data*. Springer, 2007. doi:10.1007/978-3-540-37882-2. 99, 117

[LKC*12] LEE H., KIHM J., CHOO J., STASKO J., PARK H.: iVisClustering : An interactive visual document clustering via topic modeling. *Computer Graphics Forum 31*, 3 (2012). doi:10.1111/j.1467-8659.2012.03108.x. 97

[LKS*15] L'YI S., KO B., SHIN D., CHO Y.-J., LEE J., KIM B., SEO J.: XCluSim: a visual analytics tool for interactively comparing multiple clustering results of bioinformatics data. *BMC Bioinformatics 16*, 11 (2015). doi:10.1186/1471-2105-16-S11-S5. 98, 108

[LL51] LASSWELL H. D., LERNER D.: *The Policy Sciences*. Stanford University Press, 1951. 21

[LL10] LAUDON K. C., LAUDON J. P.: *Management information systems: Managing the digital firm*, 11 ed. Pearson Education, 2010. 12

[LLCM03] Liu T., Liu S., Chen Z., Ma W.-Y.: An evaluation on feature selection for text clustering. In *International Conference on Machine Learning (ICML)* (2003). 100

[LRTM07] Lam H., Russell D., Tang D., Munzner T.: Session viewer: Visual exploratory analysis of web session logs. In *IEEE Symposium on Visual Analytics Science and Technology (VAST)* (2007), IEEE Computer Society. doi:10.1109/VAST.2007.4389008. 79

[LSKP10] Lim S.-R., Suh S., Kim J.-H., Park H. S.: Urban water infrastructure optimization to reduce environmental impacts and costs. *Journal of Environmental Management*, 3 (2010). doi:10.1016/j.jenvman.2009.09.026. 157, 158

[MA14] Miksch S., Aigner W.: A matter of time: Applying a data-users-tasks design triangle to visual analytics of time-oriented data. *Computers & Graphics 38* (2014). doi:10.1016/j.cag.2013.11.002. 33, 48

[Mac04a] Macintosh A.: Characterizing e-participation in policy-making. In *Annual Hawaii International Conference on System Sciences* (2004), IEEE Computer Society. doi:10.1109/HICSS.2004.1265300. 21, 79

[Mac04b] Macintosh A.: Using information and communication technologies to enhance citizen engagement in the policy process. In *Promise and Problems of E-democracy*. OECD, 2004. doi:10.1787/9789264019492-en. 76

[MBD*11] May T., Bannach A., Davey J., Ruppert T., Kohlhammer J.: Guiding feature subset selection with an interactive visualization. In *IEEE Conference on Visual Analytics Science and Technology (VAST)* (2011), IEEE Computer Society. doi:10.1109/VAST.2011.6102448. 96

[MBW*12] Mittelstädt S., Behrisch M., Weber S., Schreck T., Stoffel A., Pompl R., Keim D., Last H., Zhang L.: Visual analytics for the big data era - a comparative review of state-of-the-art commercial systems. In *IEEE Conference on Visual Analytics Science and Technology (VAST)* (2012), IEEE Computer Society. doi:10.1109/VAST.2012.6400554. 127

[MCF*14] McInerny G. J., Chen M., Freeman R., Gavaghan D., Meyer M., Rowland F., Spiegelhalter D. J., Stefaner M., Tessarolo G., Hortal J.: Information visualisation for science and policy: Engaging users and avoiding bias. *Trends in Ecology and Evolution 29*, 3 (2014). doi:10.1016/j.tree.2014.01.003. 25

[MCM11] Markovic D., Cvetkovic D., Masic B.: Survey of software tools for energy efficiency in a community. *Renewable and Sustainable Energy Reviews 15*, 9 (2011). doi:10.1016/j.rser.2011.06.014. 158

[MGJ*10] Matkovic K., Gracanin D., Jelovic M., Ammer A., Lez A., Hauser H.: Interactive visual analysis of multiple simulation runs using the simulation model view: Understanding and tuning of an electronic unit injector. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 16*, 6 (2010). doi:10.1109/TVCG.2010.171.

158

[MJR*11] MacEachren A. M., Jaiswal A., Robinson A. C., Pezanowski S., Savelyev A., Mitra P., Zhang X., Blanford J.: Senseplace2: Geotwitter analytics support for situational awareness. In *IEEE Conference on Visual Analytics Science and Technology (VAST)* (2011), IEEE Computer Society. doi:10.1109/VAST.2011.6102456. 18

[MK08] May T., Kohlhammer J.: Visual verification of hypotheses. In *International Symposium on Visual Computing (ISVC)* (2008), Springer. doi:10.1007/978-3-540-89646-3_4. 69

[MMT*14] Malik A., Maciejewski R., Towers S., McCullough S., Ebert D. S.: Proactive spatiotemporal resource allocation and predictive visual analytics for community policing and law enforcement. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 20*, 12 (2014). doi:10.1109/TVCG.2014.2346926. 41

[MRG*05] MacEachren A. M., Robinson A., Gardner S., Murray R., Gahegan M., Hetzler E.: Visualizing geospatial information uncertainty: What we know and what we need to know. cartography and geographic. *Cartography and Geographic Information Science 32*, 3 (2005). doi:10.1559/1523040054738936. 98

[MSQM15] Meyer M., Sedlmair M., Quinan P. S., Munzner T.: The nested blocks and guidelines model. *Information Visualization 14*, 3 (2015). doi:10.1177/1473871613510429. 35, 44

[Mun08] Munzner T.: Process and pitfalls in writing information visualization research papers. In *Information Visualization: Human-Centered Issues and Perspectives*. Springer, 2008. doi:10.1007/978-3-540-70956-5_6. 34

[Mun09] Munzner T.: A nested model for visualization design and validation. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 15*, 6 (2009). doi:10.1109/TVCG.2009.111. 32, 33, 44, 48, 61, 62, 63, 64, 80, 86, 169

[Mun14] Munzner T.: *Visualization Analysis and Design*. CRC Press, 2014. 34, 35, 50

[MvDB13] Mayer I. S., van Daalen C. E., Bots P. W. G.: Perspectives on policy analysis: A framework for understanding and design. In *Public Policy Analysis: New Developments*, Thissen W. A. H., Walker W. E., (Eds.). Springer, 2013. doi:10.1007/978-1-4614-4602-6_3. 23, 24

[NB12] Nocaj A., Brandes U.: Organizing search results with a reference map. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 18*, 12 (2012). doi:10.1109/TVCG.2012.250. 78

[NBH*15] Nazemi K., Burkhardt D., Hoppe D., Nazemi M., Kohlhammer J.: Web-based evaluation of information visualization. *Procedia Manufacturing 3* (2015). International Conference on Applied Human Factors and Ergonomics (AHFE). doi:10.1016/j.promfg.2015.07.718. 89

[ND86] NORMAN D. A., DRAPER S. W.: *User-Centered System Design: New Perspectives on Human-Computer Interaction*. Lawrence Erlbaum Associates, 1986. 31

[OC93] ORASANU J., CONNOLLY T.: The reinvention of decision making. In *Decision Making in Action: Models and Methods*, Klein G., Orasanu J., Calderwood R., Zsambok C., (Eds.). Ablex Publishing Corporation, 1993. 17

[OC10] OOSTERMAN J., COCKBURN A.: An empirical comparison of tag clouds and tables. In *Conference of the Computer-Human Interaction Special Interest Group of Australia on Computer-Human Interaction (OZCHI)* (2010), ACM. doi:10.1145/1952222.1952284. 107

[OHR*09] OELKE D., HAO M., ROHRDANTZ C., KEIM D. A., DAYAL U., HAUG L., JANETZKO H.: Visual opinion analysis of customer feedback data. In *IEEE Conference on Visual Analytics Science and Technology (VAST)* (2009), IEEE Computer Society. doi:10.1109/VAST.2009.5333919. 117

[Org04] ORGANIZATION N. I. S.: *Understanding Metadata*. Tech. rep., National Information Standards Organization (NISO), 2004. URL: http://www.niso.org/publications/press/. 51

[OSR*14] OELKE D., STROBELT H., ROHRDANTZ C., GUREVYCH I., DEUSSEN O.: Comparative exploration of document collections: A visual analytics approach. *Computer Graphics Forum 33*, 3 (2014). doi:10.1111/cgf.12376. 97

[OSS*17] ORTNER T., SORGER J., STEINLECHNER H., HESINA G., PIRINGER H., GROLLER E.: Vis-A-Ware: Integrating Spatial and Non-Spatial Visualization for Visibility-Aware Urban Planning. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 23* (2017). doi:10.1109/TVCG.2016.2520920. 41

[PA04] POMEROL J.-C., ADAM F.: Practical decision making: From the legacy of Herbert Simon to decision support systems. In *IFIP TC8/WG8.3 International Conference* (2004). 13

[PC05] PIROLLI P., CARD S.: The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *International Conference on Intelligence Analysis* (2005). 44

[PDR06] PÖLZLBAUER G., DITTENBACH M., RAUBER A.: Advanced visualization of self-organizing maps with vector fields. *Neural Networks 19*, 6 (2006). doi:10.1016/j.neunet.2006.05.013. 143

[PL08] PANG B., LEE L.: Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval 2*, 1-2 (2008). doi:10.1561/1500000011. 117

[PM09] PALAU R. M., MOENS M.-F.: Argumentation mining: the detection, classification and structure of arguments in text. In *International Conference on Artificial Intelligence and Law* (2009), ACM, ACM. doi:10.1145/1568234.1568246. 117

[Pow02] POWER D. J.: *Decision Support Systems: Concepts and Resources for Managers*. Quorum Books, 2002. 4, 14, 36, 50, 169

[Pow03] POWER D. J.: A brief history of decision support systems. Online, 2003. URL: http://dssresources.com/history/dsshistoryv28.html. 15, 16

[Pow13] POWER D. J.: *Decision support, analytics, and business intelligence*, 2 ed. Business Expert Press, 2013. 15, 16

[PS83] PATTON C. V., SAWICKI D. S.: *Basic Methods of Policy Analysis and Planning*. Pearson Education, 1983. 20

[PS07] POWER D. J., SHARDA R.: Model-driven decision support systems: Concepts and research directions. *Decision Support Systems 43*, 3 (2007). doi:10.1016/j.dss.2005.05.030. 2, 15, 17

[PST*15] POEL M., SCHROEDER R., TREPERMAN J., RUBINSTEIN M., MEYER E., MAHIEU B., SCHOLTEN C., SVETACHOVA M.: Data for Policy: A study of big data and other innovative data-driven approaches for evidence-informed policymaking. Online, 2015. URL: https://ofti.org/wp-content/uploads/2015/05/dataforpolicy.pdf. 25

[PSTT14] PETASIS G., SPILIOTOPOULOS D., TSIRAKIS N., TSANTILAS P.: Sentiment analysis for reputation management: Mining the Greek web. In *Artificial Intelligence: Methods and Applications - 8th Hellenic Conference on AI, SETN* (2014), Likas A., Blekas K., Kalles D., (Eds.), vol. 8445 of *LNCS*, Springer. doi:10.1007/978-3-319-07064-3_26. 119

[PWL97] PANG A. T., WITTENBRINK C. M., LODHA S. K.: Approaches to uncertainty visualization. *The Visual Computer 13*, 8 (1997). doi:10.1007/s003710050111. 98

[Qua15] QUACK K.: Vorausschauendes Analytics-System bei der DB: Predictive Maintenance spart Geld. *Computerwoche* (2015). URL: https://www.computerwoche.de/a/predictive-maintenance-spart-geld,3210427,2. 1

[RBB*16] RUPPERT T., BANNACH A., BERNARD J., LÜCKE-TIEKE H., ULMER A., KOHLHAMMER J.: Supporting collaborative political decision making - an interactive policy process visualization system. In *International Symposium on Visual Information Communication and Interaction (VINCI)* (2016), ACM. doi:10.1145/2968220.2968223. 76, 97

[RBB*17] RUPPERT T., BANNACH A., BERNARD J., LOKANC M., KOHLHAMMER J.: Visual access to performance indicators in the mining sector. In *Conference on Visualization (EuroVis)* (2017), The Eurographics Association. doi:10.2312/eurovisshort.20171150. 126

[RBK13] RUPPERT T., BERNARD J., KOHLHAMMER J.: Bridging knowledge gaps in policy analysis with information visualization. In *Electronic Government and Electronic Participation: Joint Proceedings of Ongoing Research of IFIP EGOV and IFIP ePart* (2013), GI-Edition - Lecture Notes in Informatics (LNI). URL: http://subs.emis.de/

LNI/Proceedings/Proceedings221/article13.html. 3, 43, 76

[RBLT*15] RUPPERT T., BERNARD J., LÜCKE-TIEKE H., MAY T., KOHLHAMMER J.: Visual-interactive text analysis to support political decision making - from sentiments to arguments to policies. In *International Workshop on Visual Analytics (EuroVA)* (2015), The Eurographics Association. doi:10.2312/eurova.20151101. 116

[RBLTK14] RUPPERT T., BERNARD J., LÜCKE-TIEKE H., KOHLHAMMER J.: Towards a tighter coupling of visualization and public policy making. In *IEEE Conference on Visual Analytics Science and Technology (VAST)* (2014), IEEE Computer Society. doi:10.1109/VAST.2014.7042525. 43

[RBMK14] RUPPERT T., BERNARD J., MAY T., KOHLHAMMER J.: Combining computational models and interactive visualization to support rational decision making. In *International Symposium on Visual Computing (ISVC)* (2014), Springer. doi:10.1007/978-3-319-14249-4_33. 140, 156, 165, 175

[RBU*13] RUPPERT T., BERNARD J., ULMER A., KUIJPER A., KOHLHAMMER J.: Visual access to optimization problems in strategic environmental assessment. In *International Symposium on Visual Computing (ISVC)* (2013), Springer. doi:10.1007/978-3-642-41939-3_35. 156

[RBU*14] RUPPERT T., BERNARD J., ULMER A., LÜCKE-TIEKE H., KOHLHAMMER J.: Visual access to an agent-based simulation model to support political decision making. In *International Conference on Knowledge Management and Knowledge Technologies (i-KNOW)* (2014), ACM. doi:10.1145/2637748.2638410. 140, 166

[RDK*15] RUPPERT T., DAMBRUCH J., KRÄMER M., BALKE T., GAVANELLI M., BRAGAGLIA S., CHESANI F., MILANO M., KOHLHAMMER J.: Visual decision support for policy making: Advancing policy analysis with visualization. In *Policy Practice and Digital Science*. Springer, 2015. doi:10.1007/978-3-319-12784-2_15. 43, 156

[RESC16] RAGAN E. D., ENDERT A., SANYAL J., CHEN J.: Characterizing provenance in visualization and data analysis: an organizational framework of provenance types and purposes. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 22*, 1 (2016). doi:10.1109/TVCG.2015.2467551. 99

[RHF05] RIEHMANN P., HANFLER M., FROEHLICH B.: Interactive sankey diagrams. In *IEEE Symposium on Information Visualization (INFOVIS)* (2005), IEEE Computer Society. doi:10.1109/INFVIS.2005.1532152. 79

[RSB*17] RUPPERT T., STAAB M., BANNACH A., LÜCKE-TIEKE H., BERNARD J., KUIJPER A., KOHLHAMMER J.: Visual interactive creation and validation of text clustering workflows to explore document collections. In *Visualization and Data Analysis (VDA)* (2017), Electronic Imaging, IS&T. doi:10.2352/ISSN.2470-1173.2017.1.VDA-388. 94

[RWA*13] RIND A., WANG T. D., AIGNER W., MIKSCH S., WONGSUPHASAWAT K., PLAISANT C., SHNEIDERMAN B.: Interactive information visualization to explore and query electronic health records. *Foundations and Trends in Human-Computer Interaction 5*, 3 (2013). doi:10.1561/1100000039. 79

[Sad05] Sadler B.: *Strategic Environmental Assessment at the Policy Level: Recent Progress, Current Status and Future Prospects*. Ministry of the Environment, Czech Republic, 2005. 158

[Sam69] Sammon J. W.: A nonlinear mapping for data structure analysis. *IEEE Transactions on Computers 18*, 5 (1969). doi:10.1109/T-C.1969.222678. 101

[SB03] Shneiderman B., Bederson B.: *The Craft of Information Visualization: Readings and Reflections*. Morgan Kaufmann Publishers, San Francisco, CA, USA, 2003. 28

[SC82] Sprague R. H., Carlson E. D.: *Building effective decision support systems*. Prentice Hall, 1982. 17

[SCM*06] Smith G., Czerwinski M., Meyers B., Robbins D., Robertson G., Tan D.: Facetmap: A scalable search and browse visualization. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 12*, 5 (2006). doi:10.1109/TVCG.2006.142. 79

[SDK14] Spiliotopoulos D., Dalianis A., Kouroupetroglou G.: Accessibility driven design for policy argumentation modelling. In *Universal Access in Human-Computer Interaction. Design for All and Accessibility Practice*, Stephanidis C., Antona M., (Eds.), vol. 8516 of *Lecture Notes in Computer Science*. Springer, 2014. doi:10.1007/978-3-319-07509-9_10. 118

[Shn96] Shneiderman B.: The eyes have it: a task by data type taxonomy for information visualizations. In *IEEE Symposium on Visual Languages* (1996), IEEE Computer Society. doi:10.1109/VL.1996.545307. 35, 36, 38, 146

[Shu99] Shulock N.: The paradox of policy analysis: If it is not used, why do we produce so much of it? *Journal of Policy Analysis and Management 18*, 2 (1999). doi:10.1002/(SICI)1520-6688(199921)18:2<226::AID-PAM2>3.0.CO;2-J. 22, 23, 169

[SIBB11] Sedlmair M., Isenberg P., Baur D., Butz A.: Information visualization evaluation in large companies: Challenges, experiences and recommendations. *Information Visualization 10*, 3 (2011). doi:10.1177/1473871611413099. 32

[Sid07] Sidney M. S.: Chapter 6 - policy formulation: Design and tools. In *Handbook of Public Policy Analysis: Theory, Politics, and Methods*, Fischer F., Miller G., Sidney M., (Eds.). CRC Press, 2007. doi:10.1201/9781420017007.pt2. 21

[Sim60] Simon H. A.: *The New Science of Management Decision*. Harper & Brothers, 1960. 7, 12, 13, 48, 169

[SJ07] Spärck Jones K.: Automatic summarising: The state of the art. *Information Processing and Management 43*, 6 (Nov. 2007). doi:10.1016/j.ipm.2007.03.009. 117

[SKK00] STEINBACH M., KARYPIS G., KUMAR V.: A comparison of document clustering techniques. In *KDD Workshop on Text Mining* (2000), vol. 400. 100

[SME08] SAVIKHIN A., MACIEJEWSKI R., EBERT D. S.: Applied visual analytics for economic decision-making. In *IEEE Symposium on Visual Analytics Science and Technology (VAST)* (2008), IEEE Computer Society. doi:10.1109/VAST.2008.4677363. 41

[SMM12] SEDLMAIR M., MEYER M., MUNZNER T.: Design study methodology: Reflections from the trenches and the stacks. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 18*, 12 (2012). doi:10.1109/TVCG.2012.213. 32, 34, 37, 44, 62, 63, 64, 86

[SNHS13] SCHULZ H. J., NOCKE T., HEITZLER M., SCHUMANN H.: A design space of visualization tasks. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 19*, 12 (2013). doi:10.1109/TVCG.2013.120. 39, 40

[SOR*09] STROBELT H., OELKE D., ROHRDANTZ C., STOFFEL A., KEIM D., DEUSSEN O.: Document cards: A top trumps visualization for documents. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 15*, 6 (2009). doi:10.1109/TVCG.2009.139. 117

[Spr80] SPRAGUE R. H.: A framework for the development of decision support systems. *MIS Quarterly: Management information systems 4*, 4 (1980). doi:10.2307/248957. 37

[SS05] SEO J., SHNEIDERMAN B.: A rank-by-feature framework for interactive exploration of multidimensional data. *Information Visualization 4*, 2 (2005). doi:10.1057/palgrave.ivs.9500091. 96

[SSK*16] SACHA D., SENARATNE H., KWON B. C., ELLIS G., KEIM D. A.: The role of uncertainty, awareness, and trust in visual analytics. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 22*, 1 (2016). doi:10.1109/TVCG.2015.2467591. 2, 98

[SSS*14] SACHA D., STOFFEL A., STOFFEL F., KWON B. C., ELLIS G., KEIM D. A.: Knowledge generation model for visual analytics. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 20*, 12 (2014). doi:10.1109/tvcg.2014.2346481. 44

[Ste46] STEVENS S. S.: On the theory of scales of measurement. *Science 103*, 2684 (1946). URL: http://www.jstor.org/stable/1671815. 34

[SWC*02] SHIM J., WARKENTIN M., COURTNEY J. F., POWER D. J., SHARDA R., CARLSSON C.: Past, present, and future of decision support technology. *Decision Support Systems 33*, 2 (2002). Decision Support System: Directions for the Next Decade. doi:10.1016/S0167-9236(01)00139-7. 14

[T*00] TEUFEL S., ET AL.: *Argumentative zoning: Information extraction from scientific text.* PhD thesis, University of Edinburgh, 2000. 117

[TC05] THOMAS J. J., COOK K. A.: *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. IEEE Computer Society, 2005. 17, 28, 29, 35, 36

[The12] THERIVEL R.: *Strategic Environmental Assessment in Action*. Taylor & Francis, 2012. 157, 158, 168

[TMLB13] TSOUKIAS A., MONTIBELLER G., LUCERTINI G., BELTON V.: Policy analytics: an agenda for research and practice. *EURO Journal on Decision Processes 1*, 1-2 (2013). doi:10.1007/s40070-013-0008-3. 25

[Tom06] TOMINSKI C.: *Event-Based Visualization for User-Centered Visual Analysis*. PhD thesis, University of Rostock, 2006. 59

[TSD14] TURBAN E., SHARDA R., DURSUN D.: *Decision support and business intelligence systems*, 9 ed. Pearson Education, 2014. 13, 15

[TvW99] TELEA A., VAN WIJK J. J.: Simplified representation of vector fields. In *IEEE Conference on Visualization (VIS)* (1999), IEEE Computer Society. doi:10.1109/VISUAL.1999.809865. 143

[US09] UNGER A., SCHUMANN H.: Visual support for the understanding of simulation processes. In *IEEE Pacific Visualization Symposium (PacificVis)* (2009), IEEE Computer Society. doi:10.1109/PACIFICVIS.2009.4906838. 142

[VW93] VELLEMAN P. F., WILKINSON L.: Nominal, ordinal, interval, and ratio typologies are misleading. *The American Statistician 47*, 1 (1993). 35

[vW05] VAN WIJK J. J.: The value of visualization. In *IEEE Visualization* (2005), IEEE Computer Society, IEEE Computer Society. doi:10.1109/VISUAL.2005.1532781. 44

[vW13] VAN WIJK J. J.: Evaluation: A challenge for visual analytics. *IEEE Computer 46*, 7 (2013). doi:10.1109/MC.2013.151. 31, 33, 48, 80, 127

[VWvH*07] VIEGAS F. B., WATTENBERG M., VAN HAM F., KRISS J., MCKEON M.: ManyEyes: a site for visualization at internet scale. *IEEE Transactions on Visualization and Computer Graphics (TVCG)* (2007). doi:10.1109/TVCG.2007.70577. 68

[War13] WARE C.: *Information visualization: perception for design*, 3 ed. Morgan Kaufmann series in interactive technologies. Elsevier, 2013. 35, 36, 84, 156

[WGGP*11] WONGSUPHASAWAT K., GUERRA GÓMEZ J. A., PLAISANT C., WANG T. D., TAIEB-MAIMON M., SHNEIDERMAN B.: Lifeflow: Visualizing an overview of event sequences. In *SIGCHI Conference on Human Factors in Computing Systems (CHI)* (2011), ACM. doi:10.1145/1978942.1979196. 79

[Wij06] WIJK J. J. V.: Bridging the gaps. *IEEE Computer Graphics and Applications 26*, 6 (2006). doi:10.1109/MCG.2006.120. 37

[Wis99] WISE J. A.: The ecological approach to text visualization. *Journal of the Association for Information Science and Technology 50*, 13 (1999). doi:10.1002/(SICI)1097-4571(1999)50:13<1224::AID-ASI8>3.0.CO;2-4. 97

[WL90] WEHREND S., LEWIS C.: A problem-oriented classification of visualization techniques. In *IEEE Conference on Visualization* (1990), IEEE Computer Society. doi:10.1109/VISUAL.1990.146375. 39

[WLS*10] WEI F., LIU S., SONG Y., PAN S., ZHOU M. X., QIAN W., SHI L., TAN L., ZHANG Q.: Tiara: A visual exploratory text analytic system. In *SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)* (2010), ACM. doi:10.1145/1835804.1835827. 78, 97

[Wor17] WORLD BANK: The Mining Investment and Governance Review (MInGov). Online, accessed in 2017. URL: http://www.worldbank.org/mingov. 126, 127, 128

[WPS*09] WANG T., PLAISANT C., SHNEIDERMAN B., SPRING N., ROSEMAN D., MARCHAND G., MUKHERJEE V., SMITH M.: Temporal summaries: Supporting temporal categorical searching, aggregation and comparison. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 15*, 6 (2009). doi:10.1109/TVCG.2009.187. 79

[WR09] WHITE R. W., ROTH R. A.: Exploratory search: Beyond the query-response paradigm. *Synthesis Lectures on Information Concepts, Retrieval, and Services 1*, 1 (2009). doi:10.2200/S00174ED1V01Y200901ICR003. 2, 46, 95

[WV05] WEIMER D. L., VINING A. R.: *Policy analysis: Concepts and practice*. Prentice Hall, 2005. 22

[YKSJ07] YI S., KANG Y. A., STASKO J. T., JACKO J. A.: Toward a Deeper Understanding of the Role of Interaction in Information Visualization. *IEEE Transactions on Visualization and Computer Graphics (TVCG) 13*, 6 (2007). doi:10.1109/TVCG.2007.70515. 39

[YTGS12] YOU F., TAO L., GRAZIANO D. J., SNYDER S. W.: Optimal design of sustainable cellulosic biofuel supply chains: Multiobjective optimization coupled with life cycle assessment and input - output analysis. *AIChE Journal 58*, 4 (2012). doi:10.1002/aic.12637. 157, 158