CYBERNETIC AUTONOMY:
AN ANALYSIS AND CRITIQUE OF ADAPTIVE LEARNING SYSTEMS

A Thesis
Presented To
The Academic Faculty
by

Michael A. Madaio

In Partial Fulfillment
of the Requirements for the Degree of
Master of Science in the
School of Literature, Media, and Communication.

Georgia Institute of Technology
May, 2015

CYBERNETIC AUTONOMY:
AN ANALYSIS AND CRITIQUE OF ADAPTIVE LEARNING SYSTEMS

Dr. Ian Bogost
School of Literature, Media, and Communication
Georgia Institute of Technology

Dr. Brian Magerko
School of Literature, Media, and Communication
Georgia Institute of Technology

Dr. Nassim Jafarinaimi
School of Literature, Media, and Communication
Georgia Institute of Technology

Dr. Wendy Newstetter
School of Interactive Computing
Georgia Institute of Technology

Date Approved: April 22nd, 2015

TABLE OF CONTENTS

LIST OF FIGURES

Chapter 1 - Introduction

1.1 - Overview

Though personalized learning has been a goal of educators since the days of Aristotle and private tutors, it is only relatively recently that technological and socio-cultural drivers have made personalized learning at scale possible. The development of digital technologies that allow for the analysis of large volumes of student data, combined with greater accessibility to large processing power on cloud-based servers, has led to the increasing feasibility of adaptive learning systems. Furthermore, socio-cultural drivers such as an ever-increasing rate of skill obsolescence and a greater demand for job retraining have led to commercial and corporate interest in adaptive learning. More individuals are also increasingly using digital technology to track or quantify many aspects of their lives, as seen in the cultural phenomenon of the so-called “quantified self” movement.

As adaptive learning systems become more accessible to educational institutions, corporations, and individuals, there yet remain serious questions about the conceptual model that informs their design, and the implications that model has for the users of the system. Though the rhetoric of many adaptive learning companies attempts to situate their design philosophy in a tradition of liberal humanism, their conceptual models are in fact better understood as exemplars of a cybernetic design framework. The values of liberal humanism, including autonomy, agency, and freedom of choice are given only lip service in the design of many adaptive learning systems, and the principles of authority, command and control, and automation are, in reality, a much greater influence on their design.

It is clear that, with the speculation about the possibilities for adaptive learning systems, coupled with the challenges and risks presented by their use, more research is needed on how they work, and on how they shape the students and teachers that use them. Many companies claim the mantle of “personalized” or adaptive, but few do it in quite the same way, and with quite the same effectiveness. In the rest of this thesis, I will explicate the various components and functions of adaptive learning systems and develop a taxonomy for understanding the various methods used by such systems, in order to better understand how those methods affect the kinds of teaching and learning they make possible.

First, in Chapter 1, I will provide a background by which to understand the movement towards personalized learning that has led to the recent interest in and adoption of adaptive learning technologies. I will go on to discuss the motivating factors in culture, economics, and education that has led to this moment. Then, I will discuss the educational implications of the possibilities for these systems. Next, in Chapter 2, I will analyze the rhetoric of adaptive learning system providers and proponents, in order to understand how their desire to position themselves in the tradition of liberal humanist education differs from the actual use of those same ideas by education philosophers. Then, in Chapter 3, I will explain the specific components and functions of adaptive systems, in light of the concepts discussed in Chapter 2. Following this, I will discuss in Chapter 4 how the design of adaptive systems often enacts the values of the cybernetic tradition, and how this has manifested itself in education more broadly. Finally, in Chapter 5, I provide a taxonomy of a representative sample of adaptive learning systems, classifying them according to the models explained in Chapter 3, and providing a set of guidelines for designers of adaptive systems and criteria for the selection of systems that are situated in a more humanist tradition. Chapter 6 will be the conclusion and suggestions for future directions of research.

1.2 - Background

Due to ever increasing class sizes in K-12 public schools since the early 20th century, and a commensurate increase in the diversity of students' background knowledge, there is a need for better tools and methods to provide differentiation of instruction at scale. Adaptive learning systems are one possible technological solution, which could provide learning materials and assessments adapted to the particular abilities, goals, and learning styles of individual learners.

Personalized learning is not a new idea. Anecdotally, one can point to Socratic and Aristotelian models for learning with small groups and private tutors. Examples from the early 20th century include the Dalton and Winnetka plans in the 1920’s, which both allowed students to progress through content only after demonstrating mastery of previous material (Mödritscher et al., 2004). However, with large numbers of students at different positions in the course content, these plans became logistically unmanageable with the available educational management systems, with more than a small number of students. In the mid-1960’s, the Keller plan was proposed for the University of Brasilia, which employed proctors to certify students’ mastery of content, allowing students to progress through a course at their own pace and through a sequence that was personalized for them (Mödritscher et al., 2004).

However, the expense of the proctors, combined with the same logistical issues faced by the Dalton and Winnetka plans, led to the discontinuation of this plan as well. Then, in the 1980’s, as more computers were introduced into schools, computer-mediated, or computer-assisted instruction (CAI or CMI) systems were developed to leverage computational power in service of personalized learning goals. However, most of those CMI systems, such as the Plato Learning Management system, were developed at a “macro-adaptation” level, where a student would receive a recommendation to retake a unit or be given access to the subsequent unit depending on performance on prior summative assessments (Mödritscher et al., 2004). It was not until later that such technologies as cognitive tutors, intelligent tutoring systems and “adaptive educational hypermedia” began to be developed and implemented to approach education from a micro-adaptation perspective (Koedinger et al., 1997; Brusilovsky, 2001). Intelligent tutoring systems are typically developed for a micro-adaptation level, offering support and feedback on individual problems rather than across a whole course of study (Mödritscher et al., 2004).

1.3 - Motivating Factors

After years of educational research on computer-supported personalized learning, recent developments in data analytics, proliferation of mobile technology, and economic drivers have begun to increase market adoption and interest in adaptive learning (Newman, 2013). As Deloitte's Center for the Edge has pointed out, the rate of skill obsolescence and job retraining is accelerating, driving the search for new modes and methods of education (Hagel, et al., 2014). As they argue, conventional educational institutions will soon be faced with unprecedented demand for access to educational resources and support, along with an increasing demand for such resources at a pace and place of the students' choosing (Hagel, et al., 2014). Indeed, if, as they estimate, the usefulness of work-related knowledge that a college student in 2020 acquires will only remain viable for five years after graduation, then individuals and corporations will need to seek new methods of developing employees' skills.

One such possibility is enabled by advances in digital technology, as argued in the mission statement for ARPA-ED, the newly proposed educational research arm of the Advanced Research Projects Agency (U.S. Department of Education, 2011). Initiated in 2012 in order to pursue long term, high-risk, high-yield technology solutions for educational "grand challenges," one of their primary mission components is to address the problem of personalized learning effectively at scale (U.S. Department of Education, 2011). Additionally, as advised by the Department of Education's 2010 National Education Technology Plan (U.S. Department of Education, 2010), the education sector should look to recent developments in other sectors, such as business and entertainment, to inform the design and adoption of new digital technologies for learning (U.S. Department of Education, 2010). As other sectors invest larger proportions of their budgets into R&D than the average in the education sector (e.g. 10% for business as opposed to .2% for education), it would be wise for educational leaders to look to the ways that advances in digital technologies have shaped the business sector, and begin to adopt those practices for education.

Specifically, developments in the accessibility of digital technology allow for increasing use of platforms and content at scales that were not possible during previous waves of enthusiasm for educational technology. In the 1980's, for example, during the movement to include personal computers in the classroom, there were significant costs for adoption, such as the purchase of computers and associated equipment and infrastructure, not to mention the costly professional development to train teachers on the use of the technology (Collins & Halverson, 2009).

Since then, larger portions of the population has access to some type of computing device, reducing the need to train teachers to use the computer, and allowing them to focus their efforts on using the software at hand. The proliferation of cheap notebook devices among school districts, in addition to the Bring Your Own Device policies adopted by some districts, represents a significant change from the implementation costs of expensive desktop computers (Department of Education, 2011). The ubiquity of such devices, along with the widespread use of learning management systems and educational software, has thus allowed for the collection of volumes of data on student learning at an unprecedented scale and granularity (National Academy of Education, 2013). The education community, however, unlike the business and entertainment sector, has struggled with how to effectively deal with the quantity of data generated, and has only recently begun to adopt technologies and practices from other such enterprises (National Academy of Education, 2013).

Due to the novelty of so-called "big" data mining, learning analytics, and data visualization in education, best practices and policies must still be developed or adopted from other domains to determine the most ethical methods for collection and handling of this data to respect student privacy and anonymity (Aspen Institute, 2014; Executive Office of the President, 2014, U.S. Department of Education, 2012). Moreover, as more educational institutions adopt technologies to collect and analyze student data, more consideration needs to be given to help teachers understand how to use that data to improve their teaching (Hill & Barber, 2014; U.S. Department of Education, 2013, U.S. Department of Education, 2012). Adaptive learning systems *seem* poised to take advantage of both the greater facility with digital technologies from students and teachers, and the greater volumes of data collected on student learning.

In addition to the technological drivers motivating schools to adopt digital learning technologies and adaptive learning systems specifically, there are also cultural drivers for the adoption of such systems. The “quantified self” movement, for instance, is representative of a larger trend in which average consumers have access to data about their personal behaviors and lifestyle characteristics at an unprecedented level, and are beginning to act on that data to track and bring about improvements in their lives. The popularity of fitness wristbands, such as Fitbit, Jawbone, and Nike’s Fuelband, demonstrate the desire of many people to quantitatively understand and track their exercise, calories, or length and quality of sleep per night (Newman, 2013). The assumption embedded in the adoption of such platforms is that, given sufficient knowledge of personal behaviors, an individual could modify and improve their behavior over time.

With adaptive learning systems, the possibilities exist for increased student awareness of the progress of their own learning, at both a small and large level of granularity (Newman 20). Though, ostensibly, a report card could do this on a broad level, for whole semesters or courses, and a learning management system (or, LMS) could do this on a slightly smaller level, neither offers the flexibility and granularity that an adaptive system could. Similarly, there has been an increased desire on the part of consumers for the media they consume to be adapted and personalized to fit their personality profile and tastes. For companies like Netflix, Amazon, Google, and others, significant resources have been expended to develop software that tracks users’ preferences and predicts which products would appeal to which users (U.S. Department of Education, 2011). As people come to expect such personalization from the corporations they interact with on a daily basis, educational technology companies, schools, and teachers should also take advantage of the large quantities of data they collect to better personalize their educational experience for students (Newman, 2013; Hagel, et al., 2014).

1.4 - Educational Implications

With such cultural and technological factors driving change in educational technology, the time is ripe for increased market adoption and interest in adaptive learning. As more and more students enter the education system at different points in their lives, lifelong learning becomes ever more common and becomes ever more necessary for individuals to respond to the demands of a rapidly changing economy (Hagel, et al., 2014; U.S. Department of Education, 2010). In many sectors, the focus has changed from a fixed set of “knowledge stocks” that incoming employees must learn, to a continuous “knowledge flow” necessary for employees to stay current in their domains (Hagel, et al., 2014).

As the needs of the “traditional learner” have diversified in tandem with the expansion of what it means to be a traditional student, educational models must now take into account an ever wider range of background knowledge, skill, and ability levels across an ever wider range of student ages and learning styles. Mastery-based learning is one such model, first proposed by Benjamin Bloom in 1968, where students progress through course content only once they have mastered previous content. Though this may seem obvious, the interventions showed the largest positive effect when used with individual tutors for students, a method that is not practical at scale (Murphy et al., 2013). Indeed, in experiments with individual tutors working with students on a mastery-based learning progression, Bloom saw improvements on the order of two sigmas difference from traditional cohort paced, whole group instruction, a huge difference (Murphy et al., 2013). In subsequent years, Bloom and others investigated the possibility for computers to assist teachers in providing individualized instruction to large numbers of students at the same time. However, at the time of their research, the software was not yet capable of monitoring student progress and adapting instructional content to students’ current levels of mastery (Murphy et al., 2013).

As the software has developed, more educational researchers have attempted to address the problem of how technology can support personalized learning for individual students, at scale. Personalized learning was always most effective when using the 1:1 tutor system, as discussed earlier. In the early 20th century, pedagogical solutions were sought that could scale to address the increasing numbers of students entering the public school system (Cuban, 1986). More recently however, there has been a tension between the altruistic desire to scale educational technologies to more students to solve more problems, and the market-driven desire to scale upwards to reach an ever-larger market.

A central tenet of mastery-based learning is that developing standardized curricula to teach the “average” student is no longer sufficient, as differences in the range of student knowledge and abilities make the variance around the mean student performance ever larger (American Institute for Research, 2013). Therefore, granting educational credit based on a fixed curriculum and fixed “seat time” requirements can be replaced in a mastery learning model by granting credit when students have demonstrated competence or mastery over a given course element (American Institute for Research, 2013). With adaptive learning systems, the possibility now exists, some claim, to begin to implement mastery-based or competency-based learning at scale, without the need for individual human tutors for every student. Recently, digital tutors have begun to take on that role with increasing effectiveness, as seen in the recent DARPA-developed digital tutor that has trained Navy IT specialists more effectively (ie: faster and less costly, with the same performance) than human tutors (U.S. Department of Education, 2013).

Though technology to adapt instruction for mastery-based learning is beginning to become publicly available to schools and researchers, new pedagogical models need to be in place to take advantage of such technology. The implementation of mastery-based adaptive learning typically occurs either entirely online, through learning management systems in online courses, or through instructional platforms such as Coursera or Khan Academy, or, as is increasingly the case, through a “blended learning” model, also referred to as hybrid learning, wherein students interact with digital content on a computer or mobile device in the classroom, dividing their time between learning online and interacting with the teacher and other students face to face. In such blended learning models, an individual teacher teaching a class of roughly thirty students is much more capable of addressing individual needs and issues, with access to more information on student competencies (American Institute for Research, 2013). In addition to the support offered to the teacher, blended learning models can provide a wealth of adaptable digital content to the students, providing them content or assessments tailored to their current knowledge, ability level, or interests (Hill & Barber, 2014).

As the demands on teachers to tailor instruction and assessment to a greater range of student knowledge and ability levels increases, teachers and schools will need to take advantage of an increasing ecosystem of educational support technologies, from pre-authored lesson modules, to apps for teachers to quickly gauge student understanding, many of which are becoming increasingly personalized and adaptive (Hagel, et al., 2014). The teacher's role will necessarily shift, as a result, moving from one of direct instruction, lecturing at the front of the class, to more of a facilitator of student learning, guiding students to knowledge discovery, as students seek out information and become more autonomous in their learning, supported by technology and their teachers (American Institute for Research, 2013, National Academy of Education, 2013;, Hill & Barber, 2014; U.S. Department of Education, 2012). However, such a transition is neither inevitable nor inherently desirable, and must be actively supported by administrators and researchers to help transition from a traditional, seat-time, face-to-face classroom with whole group, direct instruction, to a more personalized, competency-based, blended classroom, with the teacher as facilitator for the students' autonomous learning.

As mentioned, a transition to an adaptive, competency-based classroom is not inevitable, despite the cultural, economic, and technological driving factors. There are, firstly, infrastructural needs to consider, such as appropriate bandwidth requirements for students to access cloud-based software, and device requirements, typically addressed either by a Bring Your Own Device (BYOD) policy or 1:1 devices supplied by the school, and finally, a secure data infrastructure that allows teachers to collect, analyze, and take action on student data (U.S. Department of Education, 2010;, Hagel, et al., 2014). Secondly, even if there is a core group of interested teachers or administrators, the transition to adaptive, competency-based learning may face resistance, as it differs so greatly from the established, traditional model of teacher-directed education that many teachers and students are accustomed to (Newman, 2013, Murphy et al., 2013). Students, in particular, may have different motivations for learning, levels of maturity, and personal needs that may place them at different points in their ability to self-regulate and thrive in a more student-directed learning environment (Murphy et al., 2013;, American Institute for Research, 2013).

As is often widely touted, adaptive learning systems may provide the opportunity to technologically support the larger movement towards personalized, competency-based learning at scale. Specifically, they present the possibility to change the traditional structure of public schooling, such as the method of grouping students in age-based cohorts, in fixed pace courses (U.S. Department of Education, 2010). When changes are made to the traditional structures, research has shown that student attrition decreases, specifically when schools change to competency-based learning models (U.S. Department of Education, 2010). This has led to much speculation about the possibility for adaptive learning technologies to alleviate the so-called “Iron Triangle” of cost, access, and quality of education, wherein gains in any one area lead to losses in the others (Newman, 2013, Hagel, et al., 2014).

Chapter 2 - Liberal Humanist Education

2.1 - Overview

Much of the promotional rhetoric written by adaptive learning companies purports to value the individual student as the primary focus of the educational process. Some even go so far as to say that their technology “can empower young children, increasing their independence and giving them control” (Helix Education, 2013). However, upon closer examination, this becomes muddled. The values these companies ostensibly endorse, such as student autonomy, learner control, and individual choice, are also values endorsed by the tradition of liberal humanism, and which have been written about by educational philosophers, often in very different ways. In reality, adaptive learning systems often have very little to do with providing true autonomy for students.

In this chapter, I will analyze the rhetoric of adaptive learning companies through themes from liberal humanist education philosophy, to better understand the differences between the values the companies claim to be influential in the design of their systems, and the way those values have been written about in educational philosophy. I will discuss the concepts of autonomy, control, and agency in order to understand students and teachers as rational actors engaged in meaningful relationships with others. For each of the concepts above, I will first analyze the rhetoric from adaptive learning companies, followed by an explanation and interpretation of that concept in educational philosophy. Finally, I will discuss the impact that these concepts have had on educational policy, reform, and the design of curricula and assessments.

2.2 - Autonomy

2.2.1 - Introduction

Autonomy is the primary value that is claimed by both the liberal humanist education philosophers as well as by adaptive learning companies. The dominant themes of liberal humanist education are the individual freedom of all people to choose what, when, and how they want to learn. According to this tradition, in order for people to develop truly autonomously, learning must be an act willfully entered into, with the freedom for students to decide what is important for them to learn. Even before digital learning technologies entered into schools, educational philosophers such as John Dewey, Paolo Freire, Ivan Illich, and others in their philosophical tradition decried what they saw as a restrictive influence on schools from capitalistic influences and oppressive government regimes.

Beginning with John Dewey in the early 20th century, educational philosophers have argued for the importance of the individual's ability to understand and be in control of the ends and goals of their education. Though this has become the thin autonomy of students choosing course electives, Dewey and others intended it to be a more far-reaching autonomy that could penetrate into the level of the course, unit, or lesson. Without that, students "will operate much as an automaton would unless [they] realize the meaning of what [they] do" (Dewey, 1916). In order for education to be considered humanist, and for students to become fully realized people, they must understand the meaning of what they are learning, and know how it is relevant for them. Similarly, Ivan Illich, an Austrian educational philosopher who wrote critically about the need to "deschool" society, argued that "compulsory learning cannot be a liberal enterprise" (Illich, 1971). Where for Dewey, the impact of compulsory learning was more private and individual, with students operating like automata, for Illich, the consequences were more social, with deep implications for the entire project of compulsory education.

However, educational philosophers were not alone in discussing the importance of student autonomy, as, more recently, educational psychologists have also joined the discussion. In the most recent publication from the conference on "Cognition and Learning in the Digital

Age," several articles were written which discuss the importance of the psychological theory of self-determination, in which autonomy is explained as "the ability to strive towards one's own goals, interests, and aptitudes free from outside influences." (Hense & Mandl, 2012) According to this explanation, much like for Dewey and Illich, students must be able to engage in learning that is personally meaningful for them, without relying on "outside influences." In this case, learning technologies that would support students in their own autonomous learning without undue influence, must "provide freedom of choice, provide freedom of action" and provide opportunities for meaningful autonomy in learning experiences. (Hense & Mandl, 2012)

2.2.2 - Rhetoric

If you believe the promotional material for adaptive learning systems, autonomy and control are phenomena that exist in the classroom solely as a result of using their learning systems. One such company, LoudCloud, goes so far as to say that their "task-centric design keeps your learners focused on what's most important for their success" (LoudCloud Systems, 2013). The assumption here is that their system, or, more accurately, algorithms embedded in the software design, can decide what is important for students to learn better than the students themselves can. More broadly, the Fred Rogers Center for Early Learning and Children's Media report on personalized learning says that "Effective uses of technology are active, hands-on, engaging, and empowering, [and] give the child control... children need tools that help them explore, problem-solve, think, make decisions, and learn with and from one another" (Radich, 2013). In this framing, although thinking, problem-solving, and decision making seem to be located within the student, it is the technology and tools that first enable those actions, "empowering" students to be independent actors only after having been enabled by the technology.

These sentiments are echoed by several other adaptive learning companies, wherein they advocate for increased student engagement in the learning process, but only when that engagement is achieved with their software. One company claims that their "technology encouraged students to be more active participants", while another argues that their system's feedback about learning "empowers students to make more informed decisions" about their learning (Helix Education, 2013). Here, students are "active participants" and "informed" learners only after using a software that encourages and empowers that participation. Though "encourage" is a soft term, it disguises the insidious idea that students can only achieve authentic participation in their learning through using a learning software. "Empowered", too, is a loaded term that typically indicates autonomy and agency, although here it is used to indicate a result that, they seem to be claiming, will only occur through the use of their technology. Similarly, as argued by Sasha Barab, an advocate for digital games in education, "digital multimedia provide a resource for children to develop a sense of autonomy… [as] children have fewer means for expressing agency than they did in the past" (Barab, et al., 2005). In the following section, I will examine to what extent autonomy is exhibited by students without technological support or resources, or whether and in what ways the introduction of technological "encouragement" reduces the authenticity of students' autonomy.

2.2.3 - Elements

What are the essential elements of autonomy, as it relates to education? First, as previously mentioned, freedom to act is a necessary precondition for autonomous learning. Illich argues that students will never become autonomous learners until they have the freedom to exercise their own competence to learn without explicit regulation by teachers (Illich, 1971). Paolo Freire, a Brazilian educational philosopher and vociferous advocate for liberal education as a democratizing force, argues that students need the freedom to seek knowledge through a process of continuous, self-driven inquiry (Freire, 2000). For him, if students' learning is not driven by their own questions and their own inquiry, then it is not authentic, and thus, is not supportive of the liberal humanist endeavor.

Though freedom of inquiry is essential, students also need the confidence to use that freedom to learn in personally meaningful ways. Confidence is not a phenomena that can be mandated by educational policy, or forced into students through technology, but one that must be fostered and cultivated through meaningful, humane interactions between students and teachers, and through providing opportunities for students to develop self-efficacy through meaningful work that interests them (Hense & Mandl, 2012). Students are not the only ones who need freedom and confidence to exercise their autonomy, however. Teachers too, in order to effectively respond to the individual learning needs and interests of their students, need to be able to autonomously decide what and how to teach, free from overbearing influence from educational policies or curricula (Hense & Mandl, 2012).

Finally, in order for students to be able to learn autonomously, as a precursor to the aforementioned confidence and freedom, they also need to have developed their cognitive abilities sufficiently to be able to provide their own structure to their learning. Lev Vygotsky, Jean Piaget, and other adolescent psychologists have each argued for the presence of cognitive stages, through which children develop into more fully capable, independent thinkers and learners. For them, such development is not strictly biological, but is instead a mutually dependent neurological and social process wherein students internalize larger portions of cognitive processing as they develop (Vygotsky, 1978). When students are less cognitively developed, they require a more active teacher or peer support to successfully perform the same tasks that, when they have developed more, they are able to perform independently. For example, providing directed writing prompts and half-completed sentences for students to complete is one method of scaffolding such student performance (Ifenthaler). However, even for younger students, generic, open-ended prompts have been found to be more efficient and effective in promoting student writing, because they provide the opportunity for greater "autonomy for self-regulative acting" (Hense & Mandl, 2012).

This area could be an opportunity for adaptive learning systems to provide support for student autonomy, rather than restricting it. The standardized model of learning derived from mass-produced textbooks and curricula is one influence on adaptive learning systems (Collins & Halverson, 2009). However, with their ability to personalize content recommendations to students based on students' current abilities, adaptive systems could provide a means to scaffold and support authentic student autonomy, by letting them choose what they learn next, from a limited range of recommended options. In Chapter 4, I will discuss in more detail the ways in which adaptive systems function, and which types of systems offer more open-ended recommendations for student learning, rather than scripted, mandatory learning paths.

2.2.4 - Who has it?

In this discussion, I have been focusing on the importance of autonomy of the students, but, equally important to consider are other levels of autonomy in the educational process. As mentioned previously, if teachers are not able to be autonomous in their curricular and pedagogical decisions, then they will be limited in their effectiveness in providing freedom of learning choices for their students (Hense & Mandl, 2012). This is just as true for a teacher given a mandatory, pre-authored curricula in a textbook as it is for a teacher provided an adaptive learning system with pre-authored content unable to be modified by the teacher.

With administrators making decisions about the use of standardized, mandatory curricula, concepts of freedom and autonomy must enter into the discussion at all levels of the educational decision making process. However, giving teachers autonomy in designing their own curricula presupposes a high level of teacher knowledge of students' cognitive development, and confidence in their own abilities to design and implement an effective curriculum (Collins & Halverson, 2009). This is an area where adaptive learning systems can support teachers' autonomy, without becoming an undue influence, if designed effectively. A learning system that allows teachers to create their own curricular material and construct the curriculum in a manner of their choosing, while supporting them in their assessment and awareness of student learning, would authentically support teachers' autonomy.

2.3 - Control

2.3.1 - Rhetoric

The concept of control is one of the particular ways in which learner autonomy is enacted, but it plays out differently for different participants in the educational process. As the promotional material from LearnSmart claims, their adaptive system “guides students - at their own pace and on their own time - through the basic knowledge and skills covered in a course.” (LearnSmart, 2011) In this case, the students would have control over the pace of learning the material, but not necessarily over which particular knowledge or skills they would learn. Such control over the pace of learning is a dramatic departure from the traditional cohort-based model for learning, in which all students in a classroom learn the same material at the same time (Collins & Halverson, 2009). Additionally, the Helix adaptive learning system claims to allow students to “choose a narrative that fits their interests so they are in control of the learning experience.” (Helix Education, 2013) They seem particularly interested in allowing students some degree of control over their learning, but even the purported free choice of student narratives is still constrained by a limited set of curricular options generated by the publishing company and presented by the software. Finally, the LoudCloud marketing department claims that their behavioral analytics “lets you personalize the authoring and experience of course content for each learner…. and personalized recommendations take charge of your schedule and the pace at which you choose to learn.” (LoudCloud Systems, 2013) However, it is not clear here who the “you” is that is being addressed, and who is able to personalize the experience of course content. Is this the individual teacher? The student? The administrator? An instructional specialist hired by the school or district? This lack of clarity obscures the locus of control over the platform, and obscures the nature of student and teacher control over creating and delivering a course on their system.

2.3.2 - Elements

Control is a controversial topic in education philosophy, with many philosophers agreeing to its necessity, but disagreeing about the role it should play in the educational process. According to Dewey, control is at best a "guiding of activity to its own ends" (Dewey, 1916). For him, control over learning should not be dictatorial or authoritarian, but a cooperative endeavor, in which teachers engage in meaningful dialogues with students to understand their educational goals and help them work towards those goals in pedagogically productive ways. Paolo Freire, on the other hand, is more dire about the consequences, arguing that "attempts to control thinking and action.. are based on a mechanistic, static, naturalistic, spatialized view of consciousness" (Freire, 2000). Freire is warning here of the danger that attempting to "control thinking and action" could lead to the dehumanization of students, if they are treated as nothing more than controllable, programmable, mechanically responsive objects. As discussed previously, it is essential that students develop a sense of confidence in their own abilities in order to become autonomous learners. However, that confidence is fostered through opportunities to develop self-efficacy, "when an individual is in a position to be in control and master a situation" (CIADIC). Without the opportunity to demonstrate mastery over a concept under their own control, students will be hindered in the development of their self-efficacy, and thus, in developing confidence in their own ability to learn.

Alan Collins, in his historical account of the emergence of various learning technologies, argues that "industrial era learning technologies are characterized by… uniformity, didacticism, teacher control. Knowledge era [learning technologies, by]… customization, interactivity, and user control" (Collins & Halverson, 2009). In his view, the learning technologies of the industrial era, such as the blackboard, the standardized textbook, and individual desks bolted to the floor, were designed to enable teacher control over students' actions.

With a rapid increase in student enrollment in public schools from the late 19th century to the early 20th century, schools and school districts responded with cohort-based classroom models and standardization curricula to enable the rapid delivery of the same content to all students at the same pace, regardless of individual students' desires or needs (Collins & Halverson, 2009). He argues that more recent, so-called "knowledge era" learning technologies like the computer, and, perhaps, the mobile phone and tablet, are characterized by their ability to provide customized, interactive, and user controlled experiences. However, this user control is not an inherent component of digital learning experiences, and simply because it may be technologically possible on those platforms does not mean that it is inevitable or likely to happen without intentional design choices, as we will see in our analysis of adaptive systems and our discussion of ideal design principles for a more humanist approach.

2.3.3 - Who has it?

Similar to the more general concept of autonomy, it is essential to understand who has control in a learning environment, and how that control is mediated through learning technologies. Though, ideally, teachers and students would both have control over the teaching and learning process, in varying degrees, it is also important to discuss the other participants in the learning process that have control, such as the administrators of the school, the educational policies that dictate what can or cannot be taught, as well as the designers of learning technologies.

First, at the large scale, there is a tension between top-down decisions over curricula and standards, and bottom-up, autonomous teacher decisions over what gets taught and how, as seen in the recent debates about the adoption of the Common Core State Standards in the public school systems of 45 states. Writing much earlier than this debate, Illich highlights this tension between, as he puts it, "social control on one hand, and free cooperation on the other." (Illich, 1971) For him, cooperation between teachers and students is essential for a humanist education that respects the goals, desires, and needs of individual students, rather than an education through standardized curricula controlled by a politician who has no knowledge of who they are. However, when faced with the desire of some humanist educators and philosophers to teach the curricula they felt appropriate, Illich noticed what he called a “resistance to separating learning from social control.” (Illich, 1971) In his writing, he points out the intertwined relationship between state-sponsored public education, and decried the system of social control that he saw as attempting to mandate what constituted appropriate knowledge and skill for each student. Freire, in his work with indigenous farmers in Brazil, even went so far as to refuse “to conduct lessons around themes selected by teachers and policy,” arguing that mandated curricula were further instances of the social oppression that he saw enacted in other aspects of their society (Illich, 1971).

Next, even if the policies of a given school district allow for a sufficient level of teacher control over curricula at a particular school, the learning technologies adopted by that school should allow for individual control and choice over the ways in which those tools are used to teach and learn. Often, according to Illich, educational research “reflects the cultural bias of a society in which technological growth has been confused with technocratic control” (Illich, 1971). For Illich, unfettered adoption of technologies that mediate social interactions should not be undertaken without consideration for how those technologies will bring with them their own forms of social control. Instead, social interactions, in this case, educational interactions, should be a result of “self-chosen personal encounters, rather than engineered values” that an overly controlling learning technologies might enforce (Illich, 1971). This will be discussed at length in the next section (2.4), on constraints to individual agency, both technological, social, and the ways in which they intertwine.

Finally, at the small scale, control over learning should reside with teachers and students, to varying degrees, and in different contexts. One model is an autocratic, didactic style of instruction, alluded to by Collins earlier, in which students are seen as “passive receptacles of teacher instruction” (Freire, 2000). If this level of teacher control is desired, then there are current examples of instructional technology which support this, such as pre-recorded video lectures or online courses which provide content to be delivered to the students. For Collins, in the “conventional school, a teacher controls the official information flow of the classroom” (Collins & Halverson, 2009). In such a teacher-directed classroom, the optimal learning technology would be one that enables complete teacher control over information flows, allowing the teacher to be the sole source of information.

In fact, many teachers believe effective control over the process of learning to be an indicator of their authority over students and a measure of their effectiveness as a teacher. Indeed, in some states and districts, teachers are judged by how well they “keep control” over the classroom, which is construed as both student behavior and student learning (Collins & Halverson, 2009). However, in line with larger social trends explored in the introduction, modern technologies are “moving control away from centralized sources” all over the economy and culture, and education is no different (Deloitte). As personal learning technologies such as computers, mobile phones, tablets spread, and the learning software that is on those devices becomes more customizable, “enhanced learner control” should be the desired outcome of that customization (Collins & Halverson, 2009).

Though control exists in every layer of the educational process in varying degrees, and is manifested in various ways, there exists an understanding of control that is faithful to the tradition of liberal humanism, and that allows for both students and teachers to learn and teach in authentic, meaningful ways. As Dewey has said, though the “giving and taking of orders modifies action and results, [it] does not effect a sharing of purpose or communication of interests” (Dewey, 1916). Such sharing and communication is, for him, an essential element of a humanist education, and will be explored further in the following section on distributed cognition and learner agency within social relationships. Control manifested as the giving of orders for thought or action is, for Dewey, as well as for Freire, a mechanism of dehumanization, that treats students as objects to be controlled, rather than as cognizing, thinking subjects in their own right.

When individual autonomy and control over learning is violated, the liberal humanist education movement will have failed, and the will and desires of larger, depersonalized powers such as social policy or technology design will control the desires and goals of individual people.

2.4 - Constraints on Individual Agency

2.4.1 - Introduction

Though autonomy and student control over their own learning seems like an ideal goal for a humanist education, when one looks closer at the cognitive processes involved in learning, it is not clear that complete student autonomy is always attainable or is always even a worthwhile goal. A useful framework here will be for us to consider the students as "agents," defined in different ways across a variety of disciplines, such as computer science, cognitive science, and philosophy. Janet Murray, for example, describes an agent as an element of code that can "has goals, preferences… and can make decisions and initiate behaviors autonomously, rather than a centrally controlled subroutine" (Murray, 1998). Similarly, Franklin and Graesser describe a computational agent as "a system situated within and a part of an environment that senses that environment and acts on it, over time, in pursuit of its own agenda and so as to effect what it senses in the future" (Franklin & Graesser, 1996).

Students too, have goals, preferences, and an agenda, with decisions and behaviors that are often shaped by and constrained by their environment. The seemingly autonomous behaviors of a computational agent, are, of course, not fully autonomous, as their underlying code has already been written, and, moreover, just like the seemingly autonomous behavior of students, their actions are also shaped by the environment in which they are operating. For students in a classroom, their behaviors will always be constrained by the norms and practices permissible in that community. Moreover, for students using educational technology, particularly a system that purports to allow for autonomy and agency, their decisions and behaviors will be shaped by the design of the technology and the choices that it makes possible.

As a counterpoint to a humanist view of learners as entirely individual, autonomous actors, at times, some humanist educational philosophers made a point to reinforce the idea that learning emerges as a social process among teachers and students. Freire and Dewey agree that teachers must be "partners with the students in a shared activity" in which each become "jointly responsible" for the growth of all (Dewey, 1916, Freire, 2000). Such social learning has been researched extensively by cognitive psychologists and anthropologists such as Lev Vygotsky and Jean Lave. They, and other cognitive scientists, have contributed to a reformulation of the concept of agency which incorporates a more holistic understanding of individual cognition and action as comprising multiple people and the tools through which they think. The social cognition of Vygotsky, and the distributed cognition of Edwin Hutchins, are both different attempts to come to terms with the fact that individual people are not always independently functioning cognitive systems in and of themselves, but have their agency mediated through their interactions with other people and technology.

Specifically, Hutchins has argued that, "the notion of the autonomous individual is profoundly misleading, and overlooks the ways in which human agency is built through and with other people and artefacts" (Hutchins, 1995). Though the concept of autonomy may have been a useful construct for humanist philosophers in the 20th century, recent work in cognitive science has argued that cognition functions across a larger cognitive system than merely a single individual processing information in isolation in their own minds. For these socio-technically oriented cognitive scientists, cognition is always "distributed" across the other people and representational tools involved in the larger system of cognition (Hutchins, 1995). In a later section, I will examine the ways in which this distributed cognitive approach challenges and mediates the interpretations of autonomy and control as seen from the perspective of liberal humanist education philosophy, and how it leads to an understanding of students as agents, to be sure, but agents constrained by the environment and tools which shape their cognition.

2.4.2 - Cognitive Constraints

One major challenge to the idea of the fully autonomous individual comes from Terry Winograd and Fernando Flores, in their critical work, *Understanding Computers and Cognition*. They argue that individuals are never wholly rational actors, which they argue is a precondition for autonomy, and go even further to argue that the very concept of objective rationality is an idealization. In order for people to be considered fully rational actors, according to Flores, they must be able to “select the course of action that leads to a desired goal” (Winograd & Flores, 1986). If students do not choose their goals for learning, and if they have limited agency in the selection of the “course of action,” then, by definition, they are exhibiting only a limited, constrained form of rationality. In Flores and Winograd’s understanding, rationality “requires a choice among all possible alternative behaviors” (Winograd & Flores, 1986).

According to them, in order for an agent to engage in rational decision-making, it must have an objective consideration of the benefits and consequences of every possible decision, and an evaluation of the optimal decision for each situation. However, as they point out, for “isolated individuals,” such objective consideration of alternatives is an idealization, because humans can only ever consider a limited set of alternative decisions, and, without effective tools for evaluating optimal consequences, can only ever imagine the potential outcomes for each scenario (Winograd & Flores, 1986). They then go on to consider the possibilities for using machine intelligence to better augment and inform human decision-making, or in our case, to support students in deciding on optimal learning paths or content elements to learn.

The implications for students using adaptive learning systems are clear. If, as Winograd and Flores say, “in actual behavior, only a very few of all these possible alternatives ever come to mind,” then students cannot be trusted to knowledgeably make independent decisions over what is best for them to learn. (Winograd & Flores, 1986). This is a space where an adaptive learning system could support students’ overall learning autonomy, though constraining their agency in the short term, by providing recommendations for appropriate content elements for students to choose from. The system could (as many do) also indicate the potential benefits or consequences for students’ learning from the various learning decisions. This will be elucidated more in Chapter 3, when I discuss the variety of models that adaptive systems use to either recommend learning paths for students, or decide a mandatory set of content elements for students to learn.

2.4.3 - Technological Constraints

One constraint to students' agency is, as was just described, the limits of unaided humans' rational decision-making abilities, though this can be augmented with a computational system to support educational decision-making (though still constraining the students' agency). This leads us to the next constraint on individual agency - the limits of cognitive support tools, and the ways that a distributed system of cognition constrains the agency of individual actors in that system.

Many critics of technology in education argue that when a student uses technology for learning, they are robbing themselves of the chance to learn the material on their own (Cuban, 1986). In fact, tool use and cognition have been intertwined since the beginnings of human language and thought (Tomasello, 1999). In Michael Tomasello's book, *The Cultural Origins of Human Cognition*, he argues that human cognition developed in tandem with the technological tools used to augment that cognition. Rather than being a late artificial intrusion into cognition, tools and technology have, since the development of language, been used to further allow humans to coordinate their actions and cognition in the service of shared goals. Such tool use has acted, he argues, as a sort of cognitive bootstrapping, in which each successive generation is able to build on and refine the tools used by their forebears in service of more beneficial coordinated action (Tomasello, 1999).

Similarly, the technology and tools the students use to think with are also considered a part and parcel of their cognitive system, supporting their individual cognition in the same way that language and verbal symbols serve to augment and extend cognition. Adaptive learning systems could be one such supplement, serving to support students' movement through a domain space at a difficulty and a pace that are appropriate to their current knowledge level and learning style. However, simply because they are supplements, and perhaps useful supplements, does not change the fact that they offer constraints to students' autonomy. Students using such systems are still agents, but semi-autonomous agents, constrained by the possible choices and actions dictated by the technical systems.

Though the education philosophers described earlier valued the integrity of the individual’s power to make decisions and choose their goals for themselves, they were not blind to the interdependence of the individual with other people and tools. Even Paolo Freire, ever the individualist, agrees that individuals need to be understood as being firmly enmeshed in the context of the world and society in which they operate. For Freire, the ideal educational approach “denies that man is abstract, isolated, independent, and unattached to the world; it also denies that the world exists as a reality apart from people” (Freire, 2000). Such an understanding of the relationship between the individual and the world would fit well with the theory of distributed cognition posited by Edwin Hutchins.

Hutchins, in his seminal study of cognition in the process of shipboard navigation, found that, in the cooperative effort of four sailors navigating a ship into harbor using their bearings, maps, and coordinates, their cognition could not be said to take place solely in the head of a single individual. As each individual was involved in one component of the navigation process, from taking the bearings, to recording them, to analyzing the logs, to comparing the current bearing to previous bearings, he argued that the cognition that was occurring was distributed across the system as a whole, across both the people and the representational tools they were using to solve the problem (Hutchins, 1995).

For adaptive learning systems, therefore, each student can be seen as a cognitive agent engaged in cognition with the other agents in the larger cognitive system, and across the representational tools they use in the classroom (Hutchins, 1995). However, if students are denied the opportunity to interact with other students, as may be the case if the students are individually using a computer-delivered learning system to the exclusion of face to face interaction, the potential borders of their cognitive system are reduced.

2.4.4 - Social Constraints

As Tomasello points out, although our tools may augment our cognition, our cognition has been, from the start, always already socially constructed. Not only is individual cognition distributed across the cognitive systems of people and the tools they use, but that cognition is always instantiated in a socially constructed setting. Cognitive psychologist Lev Vygotsky goes further to say that all individual cognition arises as a result of social processes of cognition occurring among individuals through shared discussion and activity (Vygotsky, 1978). Vygotsky may have agreed with Tomasello, had they ever met, though Vygotsky limited his analysis to the ontogenetic origins of cognition in the individual, and did not attempt to make anthropological claims about the larger cultural origins of cognition as Tomasello did.

For both Vygotsky and Tomasello, however, the central event in the beginnings of cognition is the moment of joint attention shared between a child and their parent, around a third object (Vygotsky, 1978; Tomasello, 1999). This is, in essence, the nature of the learning process, and is at the core of all interactions in schools (replacing the parent with the teacher in a school context). However, if students instead spend the majority of their time in school working individually on a computer, they may miss many of the opportunities for shared cognition in that "joint attentional frame" that is so fundamental to cognitive development.

The importance of this shared attention to individual development is also argued by social anthropologist Jean Lave, who has studied changing dynamics in communities of practice in learning ecologies. Her work has centered around the ways in which students, or apprentices, move from peripheral participation in the practices of the community, to more central roles, over time. For Lave and her colleague, Etienne Wenger, that participation 'refers not just to local events of engagement in certain activities with certain people, but to a more encompassing process of being active participants in the practices of social communities and constructing identities in relation to these communities' (Wenger, 1998; Lave & Wenger, 1999). Learners inevitably participate in communities of practice, be it that of the school, or that of the intended domain of study. Teachers and designers of learning technology should be aware of that community, and take into account what it takes to form and cultivate a community of practice within their classroom.

Therefore, in addition to the technology and nature of human cognition providing constraints to the complete agency of supposedly autonomous learners, the social nature of cognition provides another constraint to that complete agency as well. Because students think and act as members of a community, their relationship to other members of that community inevitably shapes and constrains the available choices they have over their learning, much like the environment of the computer agents described at the beginning of this section shaped the actions and choices made by those agents.

At the very least, educators should reconsider their approach to teaching and assessment to take into account the always interconnected relations of individuals with each other and with the tools with which they think. Moreover, designers of educational technologies should design tools for learning in ways that take advantage of this understanding of intelligence as technologically augmented, cognitively distributed, and socially constructed. Though such technologies will still inevitably constrain the agency of students, they should seek to shape that agency in productive ways that still respect the autonomy and control of the students. Adaptive learning systems in particular, which could become a technology that supports individual autonomy by augmenting the process of choosing what and how to learn, should be designed to authentically allow for individual student choice, coupled with an understanding of the social nature of cognition, with the ultimate goal to “help children find a principled interdependence of people and technology” (Facer, 2011).

2.5 - Implications for Curricula

Up to this point, I have been discussing the use of adaptive learning systems in the abstract, examining the values professed by the designers of such systems in light of the interpretation of those values by education philosophers and cognitive scientists. In this section, I will look at the impact those values have on our understanding of curricula, textbooks, and tracking of students in the classroom. According to Freire, many “educational plans have failed because their authors designed them according to their own personal views of reality, never once taking into account the [people] to whom their program was ostensibly directed" (Freire, 2000). As seen in the discussion of autonomy and control, students need to be able to direct their learning paths themselves, supported by the learning technology to improve the accuracy and quality of their decisions.

A curriculum, at its heart, is the division of a domain as a whole into an ordered sequence of subject matter elements. This is a necessary support for a novice in any domain, but when that curricula becomes packaged and resold to different educational contexts despite the differences in students' interests, knowledge, and culture from one school to the next, it becomes antithetical to the liberal humanist idea of education (Illich, 1971). This is one of the central challenges presented by adaptive learning software. If the content is not created by individual teachers, and is not able to be chosen with freedom individually by the students, then despite the apparent personalization of the adaptive systems, the curriculum is no more humanist than a textbook.

Textbooks were considered an educational technology when they were first developed in the late 19th century, and were originally designed to solve the problem of standardizing what large numbers of students should learn (Collins & Halverson, 2009). Despite their efficiency, textbooks made what could have been a liberating experience of learning about "oneself, about others, and about nature depend on a prepackaged process" (Illich, 1971). In their modern incarnation, the curricula developed by schools and school districts is treated as a commodity, either purchased in bulk from a curriculum provider, or developed and produced in-house by "allegedly scientific research" (Collins & Halverson, 2009). The "distributor teacher delivers this finished product to the consumer-pupil" (Illich, 1971), a process which mirrors the production and distribution of other mass-produced consumer products, and one in which students have equally little contribution or influence over the design of the curriculum.

As textbooks were developed to standardize the content students learned, so too was a method developed to place each student into an appropriate level of the course based on their knowledge and ability. Known as tracking, such a method seems like it would be in a similar way as the personalization done by adaptive learning systems. However, in its first incarnation in 1848, when students were divided into separate classes based on a test, and the whole class learned the same material at the same time, the methods for tracking and dividing students were not granular enough to be effective (Collins & Halverson, 2009). Though students were placed into classes at a rough approximation to their current ability levels, such tests were only taken once a year, and determined the placement and education of students for the next year.

At the time, this was an efficient innovation, and allowed for teachers to better "address the needs of their students" when they were grouped by age and "experience," teaching the same lesson to all students at the same time (Collins & Halverson, 2009). However, despite ostensibly being grouped by experience, the students often did not have comparable levels of knowledge or ability, due to the lack of granularity of the tests used for such tracking. The benefits of the original uses of tracking and standardized curricula were that it reduced the amount of curricular knowledge needed by teachers, homogenizing the students in a given class and making it easier to control the classes (Collins & Halverson, 2009). Such homogeneity is always an idealization, though, since students, despite their apparent similarity, will always differ in significant ways, impacting the efficacy of the whole-class instructional method (Collins & Halverson, 2009). With the use of adaptive learning systems, the representation of the learner's knowledge can be updated on a nearly real-time basis, responding to the differences in student knowledge and ability at greater levels of granularity than that provided by the tracking used in the past.

The problem with pre-authored, whole-class curricula, as discussed previously, was that it treats a domain space as a static, fixed set of elements to be consumed by the students, and not viewing knowledge as something to be co-constructed between and among the students and teachers, as seen in Lave, Vygotsky, and Hutchins. In an ideal scenario, the students would be constructing their knowledge themselves, based on their interests. However, because they do not always know the appropriate material to learn, being novices in the field, they would benefit from a progressive, hierarchical ordering of material, from which they could select the material that best suits them (Dewey, 1916, Kirschner & van Merriënboer, 2013). In such a model, the students could authentically “seek out the ties that link one problem to another” without being overwhelmed by a complex, undifferentiated knowledge space. (Freire, 2000) It is true that all domains, especially to a novice perspective, are too complex “to be assimilated as a whole,” as Dewey describes it, but the pieces they are divided into should be accessible and meaningful to the students. Such a model is one that could be supported by adaptive learning systems, where the learning process could respond to individual changes in students’ knowledge and ability, while guiding their exploration of a complex domain space.

In the next chapter, I will explore the specific ways in which adaptive systems function, and how the design of their specific components either allows for, or inhibits, student and teacher autonomy, agency, and control over the learning process.

Chapter 3 - Analysis of Adaptive Learning Systems

3.1 - Overview

In this chapter, I will analyze the technical components and system designs of adaptive learning systems, in order to understand which elements might allow for greater or lesser opportunities for liberal humanist education. This will be a general framework for the analysis of these systems, explicating three main components of their design. Though I will be providing examples from specific adaptive learning systems for clarity, the majority of the explanation will be at a conceptual or system generic level. Using this analytical framework, I have then developed a taxonomy to classify each of the 16 adaptive learning systems according to their major system components: the domain model, the learner model, and the adaptation model, to better understand how they each instantiate those models in their design. In Chapter 5, I will use that taxonomy to offer case studies of several individual adaptive systems, in order to understand how well their design instantiates the values they profess.

First, however, I will explicate here the various components of these systems in general, from the construction of the domain model, the representation of the learner model, and presentation of the adaptation model. I will then describe several considerations for their use in a school environment, and conclude with implications for research. For each of the three major components, I will discuss how their various instantiations in specific adaptive systems may allow for more opportunities for student and teacher autonomy, situating those particular systems in the liberal humanist tradition discussed in Chapter 2. Conversely, I will also be indicating ways in which the components may be designed in ways that limit student and teacher autonomy, intentionally or not, positioning the systems within the tradition of cybernetics, which I will explain in more depth in Chapter 4, following this explication of the adaptive system components.

In most adaptive learning systems, or platforms that provide adaptation across a range of content elements for an entire course, the course or **domain** must first be modeled as a knowledge map, with links between content elements with identified prerequisite relationships. Once the domain is modeled, students may progress through the content in a sequence and pace individually tailored to their current knowledge state and learner profile.

Many adaptive systems have methods for continuous granular assessment of students, which are then compiled into a **learner model**. This model may comprise such elements as learner goals for the course and current domain knowledge, as well as other elements such as cognitive and meta-cognitive ability, and non-cognitive factors such as motivation level, learning style, or preference for medium of learning.

Then, in the **adaptation model** the system offers recommendations or adaptations of the learning process based on the student's performance on assessments, as well as data on the learner model gathered through clickstream data while using the system. These recommendations are selected from among a range of possible options, such as difficulty of content, sequence of content, medium of the content, and others. In addition, these adaptations can be presented more or less visibly, and with more or less choice on the part of the student.

Finally, I will discuss the considerations for implementation of such systems in an educational environment, whether online, face to face, or blended. There are technological and infrastructural considerations, as mentioned in Chapter 1, such as adequate bandwidth, access to technology, and data analysis and privacy. In addition, there are pedagogical considerations, such as ensuring teacher buy-in, transitioning teachers and students into using the adaptive system, and ensuring that it does not interfere with current classroom pedagogical practices.

3.2 - Domain Model

3.2.1 - Introduction

The first component of an adaptive learning system I will discuss is the domain model. Before students can progress through the domain content, the domain itself must be mapped out in such a way that the relationships between the elements of the domain are clearly established. The nature of those elements, their arrangement, and the role that the teacher and course provider play in the creation of the domain model vary from adaptive system to system. Some adaptive learning providers, like SmartSparrow, allow teachers to upload their own instructional content, and arrange it into a hierarchy that makes sense for them. The course elements typically include such categories as instructional content, assessment items, links to externally produced content, as well as the metadata associating each content item to each other, and to local and national education standards.

These course elements might be created by the individual teacher, created by content experts hired by the adaptive learning provider or publisher, or discovered by the adaptive learning system through a semantic search of publicly available open education resources, which is done by LoudCloud, among other providers. Once created, they are arranged into a hierarchy, with prerequisite relationships established between linear content elements, which is either done by the individual teacher for the course or is already established by the adaptive learning provider, such as is the case with Knewton's adaptive system. Finally, the visual grammar of the domain model may vary from system to system, as some have adopted conventions from digital mapping tools, such as the ability to pan, zoom, or filter the map, or the ability for each learner to see their "position" in the domain in relation to other students' levels of mastery.

Each of those aspects of the course creation process has tradeoffs between the autonomy and control of the teachers involved, and the efficiency and efficacy of the system's recommendations. For each of the following sections, I will discuss some of the advantages and disadvantages to having the course be created by a teacher, as opposed to the adaptive provider.

3.2.2 - Knowledge Maps

Knowledge maps have a long history in education, though they have been called by a variety of names, such as graphic organizers, node-link diagrams, concept maps, and knowledge maps. Graphic organizers refer to any two-dimensional knowledge representation, such as flowcharts, timelines, and tables, which allow students to "subsume new concepts in superordinate cognitive structures" (Ausubel, 1968; Hawk, 1986; Nesbit, 2006). Node-link diagrams, of which concept maps are a subset, show a potentially more complex relationship of concepts, where the concepts are represented by nodes, connected by edges, or links, that represent a proposition about their relationship. Concept maps have been shown to be beneficial to students with low verbal ability and second language students, due to the reduced textual complexity and relative standardization of the node-edge relationship (Holliday, Brunner, & Donais, 1977; Moyer, Sowder, Threadgill-Sowder, & Moyer, 1984; Stensvold & Wilson, 1990; Nesbit, 2006).

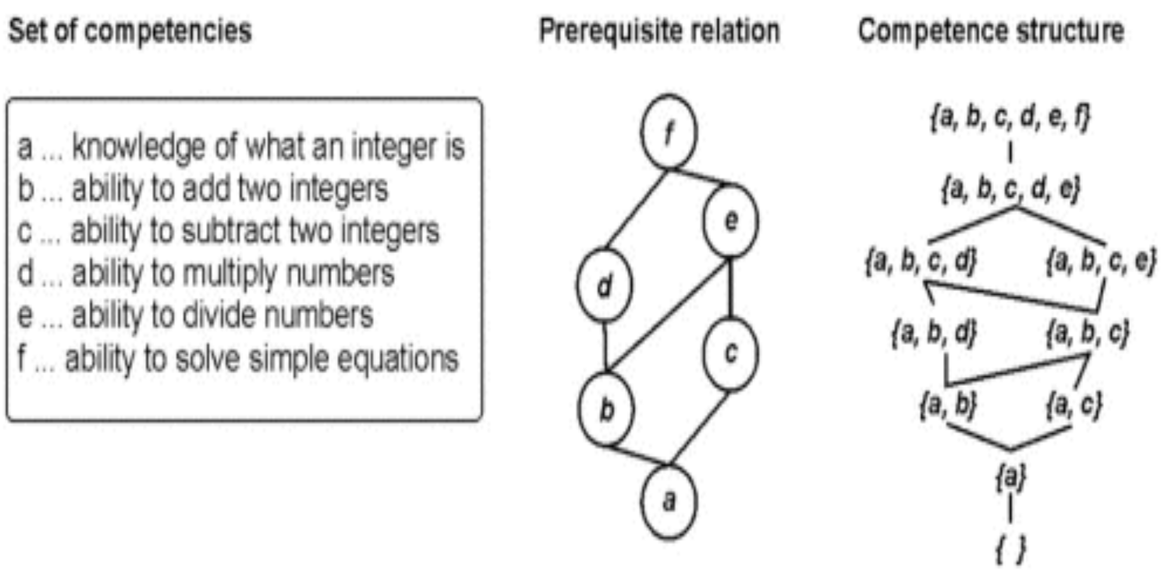

Figure 3.1 - A Simple Domain Model (Reiman, 2013)

3.2.3 - Course Elements

Although concept maps and node-link diagrams can be used to represent knowledge in a large variety of contexts, in adaptive learning systems, each node represents a course element and the edges represent their hierarchical relationship. For instance, the nodes might be instructional content, such as pre-recorded video lectures, slideshows, text articles about the content, or other types of pre-authored instructional content (Chen, 2008; Graf & Ives, 2010). They could also be assessment content, such as quizzes, problems, short response questions, or programming assignments, which would provide the system with more information about the learner, informing the learner model and the adaptations, which I will discuss in later sections of this chapter (Burgos et al., 2006; Huang et al., 2006). Finally, these course elements might be embedded external content, such as videos from Youtube, Khan Academy, or some other third party provider, or any other type of open educational resource, discovered through semantic search algorithms (Brusilovsky, 2004; Magnisalis et al., 2011). Typically, all of the above course elements are tagged with associated metadata, one element of which is the prerequisite data that informs the system which elements can be accessed by a particular student, based on how many have already been completed and mastered.

One adaptive provider, LoudCloud uses a semantic search to discover instructional video content from YouTube and Khan Academy related to the topic being learned. The student has the choice, after reading a chapter from the text, to click on the links to the related videos that are displayed in the browser. However, as with many of the related video algorithms used by YouTube and other providers, there are risks that the videos might in fact be entirely unrelated to the content being learned.

3.2.4 - Creating a Domain Model

Adaptive learning systems vary on the level of control they offer to the instructor, with some systems providing pre-authored courses, either on their own platform, or through partnership with a publishing company, while other systems offer teachers the ability to author their own course themselves. For teachers who create their own course, they must clearly establish what the overall learning goals are, and work backwards to either create or embed instructional and assessment content to lead students toward those goals (Magnisalis et al., 2011; Chung & Kim, 2012). Depending on the nature of the system, teachers may also have to manually index the content with associated metadata tags to identify which instructional standards are addressed by each content element, so the system can track which students have mastered which specific standards at any given point in accessing the system (Brusilovsky, 2004). However, though this seems like it would grant teachers more agency and control over their own teaching, in reality, the authoring programs are often unwieldy and difficult for teachers to use, if they are provided at all (Chung & Kim, 2012).

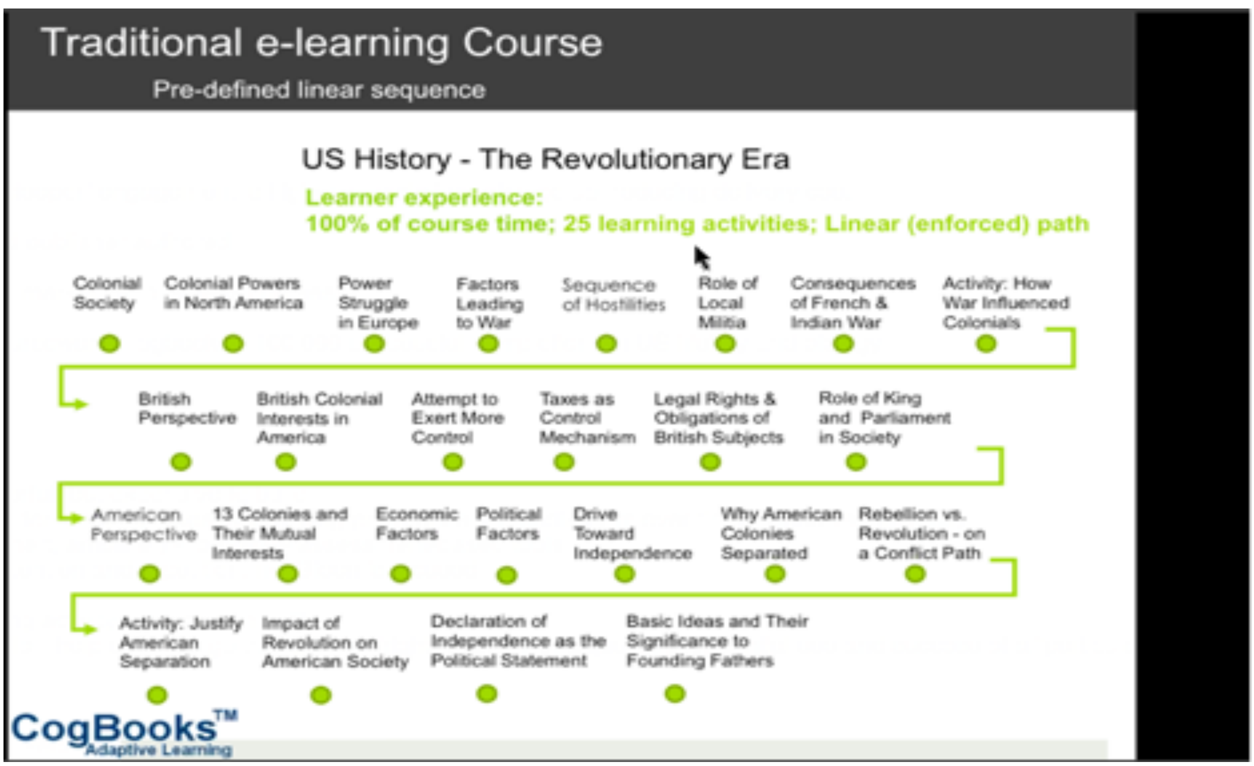

Figure 3.2 - A Traditional E-Learning Course Domain Model (CogBooks 2015)

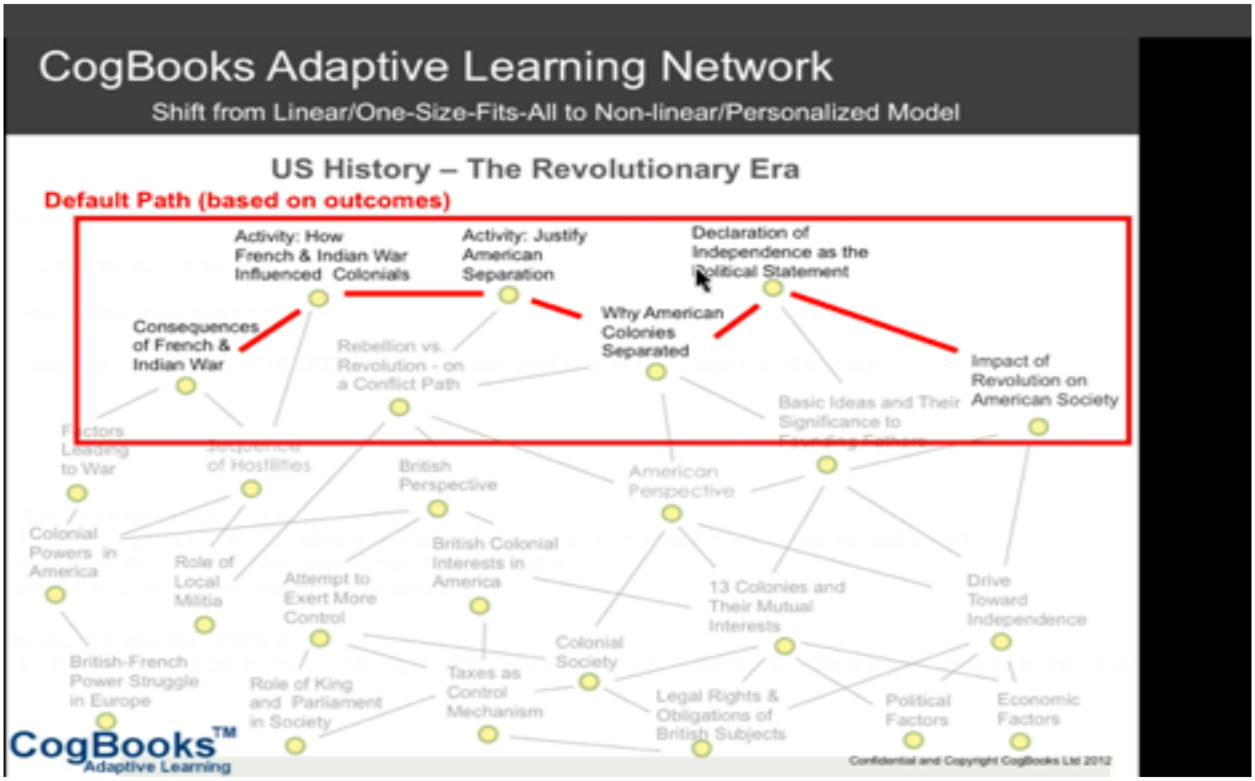

Figure 3.3 - An Adaptive Learning Domain Model (CogBooks 2015)

For systems where experts create the course, the teacher is effectively removed from having any part in the instructional design process, merely facilitating the students' progress through course content. Moreover, the adaptive system designers do not know the particular students, school, or socio-cultural context of the school where their system is being implemented, and as such, run the risk of alienating the students with content that is not relevant to their interests or life experiences, as discussed in the previous chapter. However, despite these risks, Knewton, one of the largest adaptive learning providers, uses content that is discovered, tagged, and arranged by its own employees, rather than by the teachers who will be using it. The semantic-search discovery process used by LoudCloud and others runs its own risks, particularly that of mis-identification of relevant content or of identification of relevant content that does not meet the difficulty level of the student or course (Magnisalis et al., 2011).

Though there are risks to teachers' autonomy with platform-provided content, there are equal risks associated with allowing individual teachers to create and arrange course elements, due to the complexity of the large numbers of course elements in a given domain, with interrelated prerequisites and dependencies. First, teachers may inconsistently tag the instructional content with the same kind of metadata, causing student recommendations to be inaccurate (Karampiperis & Sampson, 2005). They may also attach insufficient numbers or quality of metadata to the content , causing gaps or holes in the recommendations (Karampiperis & Sampson, 2005). If the adaptive learning system is to guide students appropriately through a learning path that both reflects the students' current knowledge state and leads to an acceptable learning outcome (ie: mastery of the appropriate competencies), then the content must be categorized consistently and sufficiently for the system to work (Brusilovsky, 2004; Karampiperis & Sampson, 2005). One possible solution presented by Doignon and Falmagne in their work on Knowledge Spaces, is to have the content experts create the initial domain model for the adaptive learning system, and then refine the available options using continuously updated user data to inform future students' options for appropriate content elements and learning sequence (Falmagne, 2011).

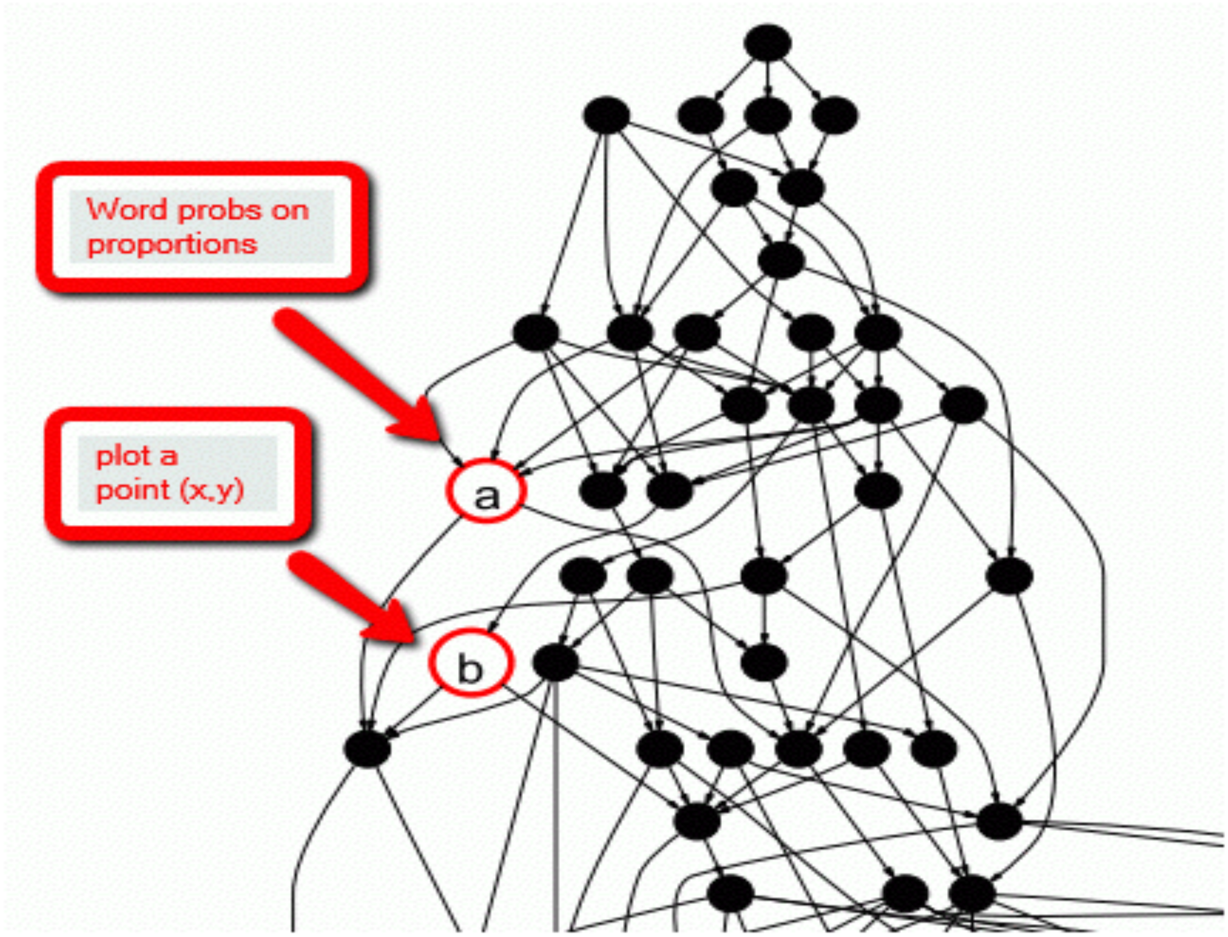

Figure 3.4 - A More Complex Domain Model (Falmagne, 2011)

3.2.5 - Representational Conventions

With all of these possibilities for construction of the domain model, there are several different possible knowledge representations at work. One, the "domain model" or "knowledge structure" is the more generalized version of the media space, which represents the concepts or standards the students should be mastering at each point, connected by the edges, which are the dependency and prerequisite information for moving from one concept to the next (Karampiperis & Sampson, 2005; Falmagne, 2011). This would be visible only to the instructional designer of the course, and perhaps to the teacher as well. In addition, there is the "media space," which contains all of the individual educational resources or content elements that the students would access, along with their associated metadata (Karampiperis & Sampson, 2005; Falmagne, 2011; Chung & Kim, 2012). Finally, there is the “knowledge state,” or, “learner model,” which is the representation of what each student “knows” at any given point in time, a model that is created through the students’ performance on assessments, among other factors (Falmagne, 2011). This will be described in more detail in the following section on the Learner Model.

Once the course elements have been created and arranged, there are different possibilities for their visualization and for interacting with those visualizations in order to orient students in a large, complex, and potentially overwhelming knowledge space. Much of the work done in this area has drawn from interaction conventions in information visualization and GIS. For instance, Bargel et al examined the use of digital map conventions such as the ability of students to zoom in and out or pan across the knowledge map, to either see the entire domain, or focus on one section that they are working on at the moment (Bargel et al, 2012). This might allow them to plan ahead for future coursework, or focus on the current content they are learning. However, this brings with it such risks as the cognitive overload that might occur when students see the entire scope of the domain, potentially leading to feelings of intimidation or paralysis (Bargel et al, 2012).

Additionally, students can typically view their individual position in the knowledge space, and, in some systems, can view their upcoming or previous “learning path” through the material, including both the course elements that they have successfully completed, and those that lay ahead, which the system has recommended that they complete (Di Bitonto et al., 2013). Some systems allow students to see their own progress through the course in relation to other students in the class, which, some argue, is beneficial for student motivation, but may lead to unnecessary or unwarranted competition (Bargel et al, 2012). One possible method for dealing with this issue of the visibility of other students, and of the potentially overwhelming path ahead of students, is by using a “fog of war” metaphor adopted from video games, to show students only the elements that are close to their position in the knowledge space (Bargel et al, 2012).

3.3 - Learner Model

The learner model is the second main component of adaptive learning systems. Once the content domain has been mapped, the system needs to be able to represent what the students know and can do, in relation to that domain. This is known as the Learner Model, also called the user model, and it is typically divided into several different components, using data collected in a variety of ways. In broad terms, these learner models can either be an overlay model or a stereotype model (Nitchot et al., 2010; Knauf et al., 2010). In an overlay model, each student's current knowledge state is represented as a subset of the domain model and overlaid onto the larger model of the domain or knowledge structure (Knauf et al., 2010). The assumption behind this model is that the domain model represents the ideal set of knowledge for all students in a given domain. In the stereotype model, student knowledge is represented in relation to clusters of similar learners, depending on the nature of the data in their learner model (Klašnja-Milićević et al., 2011). The assumption embedded in this model is that students with similar learner profiles would necessarily respond in similar ways to choices of instructional and assessment content.

3.3.1 - Method of Data Collection - Static or Dynamic

For both of those two types of learner models, the student data is collected in a variety of ways and with a variety of metrics. The data used to populate the learner model can be collected either statically or dynamically, or a combination of both, depending on the type of data. Static data is usually collected for those factors of the learner model that are assumed to remain constant over time, such as cognitive characteristics that tend to remain stable, noncognitive factors such as goals, preferences, or background knowledge students possessed before the beginning of the course (Karampiperis & Sampson, 2005). These static characteristics are usually captured through some form of survey or pre-test that the student takes before beginning their progression through the learning content (Karampiperis &

Sampson, 2005; Chen, 2008; Klašnja-Milićević et al., 2011). However, this explicit means of capturing static data runs the risk of the students choosing to not complete it, or not completing it accurately or authentically.

As opposed to static data, the data collected to inform the learner model might instead be dynamic data collected and updated on a continuous basis by the adaptive learning system. Dynamic data might comprise categories such as the learner's knowledge state, which should (ideally) be updated as the learners become more proficient and master various content elements, or the learning style of the student, which could be inferred implicitly through interactions with the adaptive learning system (Karampiperis & Sampson, 2005). A self-identified learning style survey may be filled out by a student that thinks they learn best in a certain way or in a certain medium, but, if the students' learning style is inferred through their interactions with the adaptive system, a more accurate picture might emerge (Moreno-Ger et al., 2007; Klašnja-Milićević et al., 2011). Static and dynamic data collection means are both useful, for different kinds of data, depending on what the designers of the adaptive system decide are useful kinds of data to use.

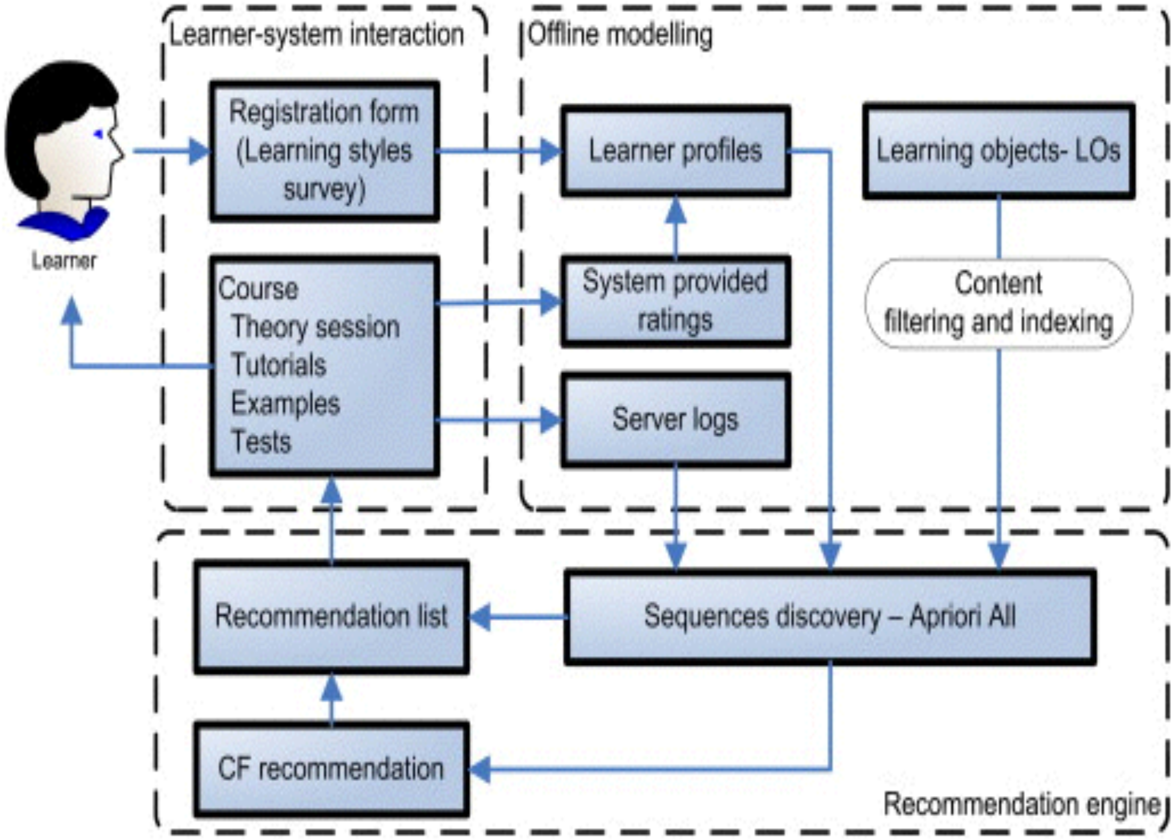

Figure 3.5 - A Learner-System Interaction Model (Klašnja-Milićević et al., 2011)

3.3.2 - Method of Learner Data Collection - Implicit or Explicit

Another way of categorizing the method of data collection is whether it is implicitly or explicitly collected. As mentioned previously, a survey or pre-test to gather user data on the learner is an explicit method of data collection, as opposed to the data gathered implicitly through clickstream data of user interactions with the system (Chen et al., 2005; Lo et al., 2012). Often, users are not motivated to answer personal questions about their demographics, and may not be as complete or accurate with self-assessment data (Joerding, 1999). Additionally, the explicit framing of a survey or a student feedback prompt after a problem interrupts the learning process, leading to fewer responses, or more difficulty in resuming the progression through course content (Koychev and Schwab, 2000; Magnisalis et al., 2011; Lo et al., 2012).

An implicit method of data collection, on the other hand, uses the interactions the user already has with the system, such as the pages the learner visits, the length of time they spend on a given page, and the links they follow, in order to infer things about the user from those actions (Klašnja-Milićević et al., 2011). Though some say that implicit methods of data collection respect the users' privacy more than explicit methods which ask for personal information, the user does not always know what data is collected from their interactions with the platform, and in what ways they will be used (Hanani, 2001; Lo et al., 2012). This seems like more of an intrusion on the privacy of the student, as they do not even have the ability to know the granular elements of their learner profile that are created, like they might in a more explicit method of data collection. The vast majority of the systems that I have analyzed for this thesis use dynamic data collection, and 15 out of 16 used implicit data collection, though some used a combination of implicit and explicit data collection.

Finally, another consideration when categorizing the method of data collection is the scope of the data being collected, whether it is on the scale of individual course elements such as problems or quiz questions, or whether the data is collected for multiple content elements across an entire course. At the micro-interaction level, or, what is known as model tracing, the system collects data on a small scale about the choices students make when solving a single problem (Koedinger et al., 1997). A subset of adaptive learning systems known as intelligent tutoring systems typically deal with model tracing and small scale user modeling to provide fine-grained support and feedback during the process of individual problem solving, much like a human tutor might do (Magnisalis et al., 2011). These systems monitor a student's progress through solving a single problem, offering feedback after the solution is presented. However, these tend to be used for well-defined problem spaces, typically in algebra courses. The other, larger scale of data collection, commonly known as knowledge tracing, monitors students' performance on multiple content elements in the domain, and can identify individual areas of difficulty or competency and recommend future learning objects or pathways through the learning space based on performance across multiple content elements (Koedinger et al., 1997; Magnisalis et al., 2011).

3.3.3 - Types of Learner Data - Performance and Interaction

While the data used to inform the learner model can be collected either explicitly or implicitly, statically or dynamically, and at either a large or a small scale, it is worth discussing here specifically which types of data are typically used to create the learner model. That data typically falls into categories such as performance data from assessments, usage and interaction data from the adaptive system, and a range of other cognitive and noncognitive factors used to better understand the learner.

First, student performance on course elements is the primary source of data used to represent the knowledge state of the learner. This may be created initially using a pre-test to determine the level of background knowledge of the student (static and explicit), though this must be updated dynamically for it to remain an accurate reflection of the students' knowledge state (Bargel et al, 2012). Some systems collect performance data at a more complex level than a merely binary determination of mastery. These data may also track the degree of correctness, understood as the speed at which the student completed the problem or quiz, as well as the relative certainty, understood as the number of attempts the student took before their performance was deemed successful, in order to construct a more robust model of the learner's knowledge (Szilagyi & Roxin, 2012; Bargel et al, 2012). In addition to student performance, another form of data collected is interaction data from student usage of the adaptive learning system. This might be aggregate time spent using the system, or granular data about time spent on individual content elements, as well as browsing behavior, such as particular videos watched before taking assessments, measured by browser clicks and other recordable actions (Karampiperis & Sampson, 2005; Szilagyi & Roxin, 2012).

3.3.4 - Types of Learner Data - Cognitive and Non-cognitive

The data collected might also be used to gather information on the students' general cognitive abilities, as well as a host of other non-cognitive data such as learning preference, learning goals, and learning style. The cognitive data is typically assumed to be relatively stable over time, and is usually collected through explicit psychometric tests, or inferred through interactions with the system collected through implicit means such as browser behavior (Brusilovsky, 2001; Lo et al., 2008; Lo et al., 2012; Lo et al., 2012). These cognitive factors may include such aspects as the students' working memory capacity, inductive reasoning ability, motivation, information processing speed, and associative learning skills, as well as the level of Bloom's taxonomy on which the assessment is operating (Karampiperis & Sampson, 2005; Essalmi et al., 2010; Akbulut & Cardak, 2012; Newman, 2013). Interestingly, in the literature on adaptive learning systems, these are typically assumed to be generalizable characteristics that can be assessed and transferred, regardless of the specific learning context or content, and assumed to be stable over "long periods of time," a fact which has yet to be definitively determined by psychological research (Brusilovsky, 2001; Karampiperis & Sampson, 2005; Poelhuber et al., 2008; Graf & Ives, 2010; Klašnja-Milićević et al., 2011; Lo et al., 2012).

The non-cognitive aspects of the learner model, on the other hand, are collected dynamically, as they are not assumed to remain stable or constant over time. These include such aspects as learner preference, either for the topic or the media form of the content, as well as learner goals and learning style. The option for media preference allows the learner to select the types of media they prefer to learn from, either textual, graphic, audio, or visual media. This might be done explicitly, where the learner manually filters content of the type they prefer to access, or implicitly, where the adaptive system forms a model of the type of content media the user performs better with, and only displays those types of media (Brusilovsky, 2001; Essalmi et al., 2010; Klašnja-Milićević et al., 2011). With the explicit filtering, the user has more choice over what and how they are learning, though, as a domain novice, they may not always be equipped to make those choices in an informed way.

Just as all learners enter a course with different background knowledge states, they also enter with different goals for their future learning path. Though those goals may be incommensurate with the required learning objectives of the course, it is possible to tailor content to appeal to the goals the learners bring with them, thus increasing their motivation for learning (Essalmi et al., 2010; Szilagyi & Roxin, 2012). These goals may be content-specific goals or goals related to the timing of completion of course content, with the system providing reminders or prompts reminding students to stay on schedule if they fall behind (Essalmi et al., 2010; Klašnja-Milićević et al., 2011). Some systems have a goal model, which comprises not just the goals, but the tasks, requirements, and workflows necessary to allow students to complete their learning goals (Natriello, 2013). More work needs to be done to mediate between students' desire for achieving their own learning goals and the required goals and outcomes mandated by the curriculum designer or teacher.

Another category of learner data collected is on the learning style of the students. An oft-disputed category, sometimes listed under cognitive characteristics, and sometimes under non-cognitive, learning styles have a long and complicated history in cognitive psychology (Gardner, 1989; Klašnja-Milićević et al., 2011; Kirschner & van Merriënboer, 2013). Despite the conflicting research on their efficacy, they are worth discussing here due to the number of adaptive learning systems that adapt content based on data about students' learning styles.

There are several main theories of learning styles that have been utilized by adaptive learning systems, such as those of the Kolb learning cycle, the Honey-Mumford, and the Felder-Silverman learning styles (Kolb, 1984; Honey and Mumford, 1986; Felder and Silverman, 1988; Essalmi et al., 2010; Graf & Ives, 2010; Klašnja-Milićević et al., 2011; Knauf et al., 2011). The two used most often, Honey-Mumford and Felder-Silverman, share some features. Both situate learners along four dimensions, depending on learners' aptitude for a particular type of learning. In the popular Felder-Silverman learning style theory, students may have preferences across the four dimensions of Active/Reflective, Sensing/Intuiting, Visual/Verbal, or Sequential/Global (Essalmi et al., 2010; Graf & Ives, 2010).

Students learn best, according to some research, when learning material is presented in a way that appeals to their particular style (Knauf et al., 2010). Some adaptive systems incorporate learning style into the learner model and adapt the sequence or form of the content to appeal to those styles for the students' benefit (Knauf et al., 2010). The method of collection for learning style data can be either explicit, such as a learning style survey, or inferred through interactions with the system, such as links clicked and pages visited (Klašnja-Milićević et al., 2011). Many adaptive learning systems use an explicit survey which makes several assumptions, such as the fact that students know the style of learning that fits them, and that such styles are stable over time and across contexts, which is not always true, as discussed earlier (Klašnja-Milićević et al., 2011). The more humanist, student-centered versions of these systems allow for students to select the content type that they feel best fits their learning style, rather than having that style inferred through their performance, and having the content adapted to fit them.

3.3.5 - Data Analysis Algorithms

For the implicit methods of data collection, the adaptive learning system needs some algorithm to be able to translate student interaction data into inferences about their learner profile and model as a learner. A number of different methods are used, such as Bayesian network tracing, hidden Markov models, genetic algorithms, neural networks, and various other data mining techniques. Each has value in particular contexts, and for particular purposes.

In some systems, Bayesian network tracing is used to cluster students into groups of learners based on similarities in their learner models, either similarities in their current knowledge state, or by similarities in their cognitive and non-cognitive characteristics (Brusilovsky, 2001; Magnisalis et al., 2011). In Bayesian statistics, a prior model is updated with new information to reflect a constantly changing statistical model. This is very appropriate to the nature of learning, with a constantly changing model reflecting the students' knowledge state. Typically considered methods for modeling Bayesian probability, Markov chains and hidden

Markov models are often used in adaptive systems for predicting likelihood of student success on particular content elements or assessment, informing the adaptation component to determine which elements should be attempted by which students at which times (Brusilovsky, 2001).

Genetic algorithms are also used to construct optimal learning paths according to patterns of correct and incorrect responses on pre-tests (Chen, 2008). Such algorithms are able to refine their recommendations based on future revisions to the knowledge state of the learner, and by comparing the individual learner model to aggregates of large numbers of learners. Genetic algorithms can be used to circumvent errors of insufficient or inaccurate metadata tags attached to the content elements by a teacher author, though, only after enough students have been recommended a course element that has been seen to be inappropriate for their ability level (Chen, 2008).

Similar to the genetic algorithms, neural networks are used in a “multi-layer feed forward” technique, to conduct pattern recognition on imprecise or incompletely understood data, to generalize and learn from specific examples, such as extrapolating out inferences about students via specific input to their learner model, and to quickly update that model with new input (Chiu et al., 1991; Masters, 1993 and Mullier, 1999; Lo and Shu, 2005; Zatarain, Barron-Estrada, Reyes-Garcfa, and Reyes-Galavia, 2011; Lo et al., 2012). Such neural networks are used most often for inferring student learning styles from interactions and browsing behavior, sometimes employing a genetic algorithm themselves to train the network to understand how to identify learning styles and tailor the learning path appropriately (Lo et al., 2012). Finally, some adaptive systems use other common data mining techniques to discover and group publicly available educational content with techniques such as association rule mining, k-means clustering, and inter-session pattern mining (Zaiane, 2002; Klašnja-Milićević et al., 2011).

3.4 - Adaptation Model

3.4.1 - Adapting the Content

Once the domain model is constructed and the learner model is specified, the system can adapt to the student from a set of options, be it the learning object itself or the sequence of the content presented. The unit of adaptivity, or the content element being adapted, can be modified by either its media form, the topic of the content, the type of content element, or by the difficulty. The media form could be text, video, audio, or graphic, and can be modified based on the particular needs or preferences of the student. Adaptation by topic is used when clustering learning documents based on their relative similarity, determined by either associated metadata tags or semantic search algorithms (Klašnja-Milićević et al., 2011; Chung & Kim, 2012). Adaptation by type refers to the type of content element, such as examples, exercises, assessments, applications, assignments, discussions, or others (Graf & Ives, 2010). Another type of adaptive display is interface based adaptation, where the position, size, and properties of the interface are adapted to fit the learners' needs, typically to address accessibility issues (Burgos 2006). Finally, the learning objects might be adapted based on their difficulty, either through the selection of problems at an appropriate level of difficulty, or through the generation of new problems with sets of parameterized questions or exercises (Mayo & Mitrovich, 2000; Mitrovic & Martin, 2004; Kumar, 2006; Ullrich et al, 2009; Sosnovsky, 2010).

3.4.2 - Adapting the Sequence

Additionally, the sequence of content might be adapted to fit the learner's knowledge state, presenting certain content elements at an appropriate point in the learner's progression through the content, depending on prior mastery (Sosnovsky, 2010; Di Bitonto et al., 2013). This could be done at different levels of scale, from a micro-adaptive approach, which selects content elements directly following the learner's current position in the knowledge space, to a more macro-adaptive approach, which selects course components at a more general scale, typically based off of more static learner profile elements such as learning goals (Burgos et al., 2006). These are also referred to as either outer loop processes, which adapt the progression of course content from item to item, or inner loop processes which recommend next steps for learners within a single problem or task (VanLehn, 2011; Natriello, 2013).

3.4.3 - Adaptive Presentation - Direct or Indirect

The presentation of the adaptations might take the form of more direct or indirect methods, each with their own benefits and drawbacks to the learners. When the adaptation is implemented directly, the students are made aware of an adaptation of the next content element to be accessed (Magnisalis et al., 2011). This could either be achieved through adaptations in the presentation, such as additional explanations given if a student responds to a question incorrectly, or prerequisite explanations, displayed if the student has shown signs of needing extra scaffolding or support before the next content element (Hauger & Köck, 2007). These direct adaptations might also include navigation support, such as a sequential path through the knowledge space presented to the learner with clear indicators as to which elements are the recommended ones for the learner to access. This could also involve link sorting, in which relevant links to recommended content are organized based on their inferred relevance and appropriateness for the student, or link annotation, in which the links are colored, dimmed, or textually annotated in order to indicate degrees of relevance for the learner (Hauger & Köck, 2007; Magnisalis et al., 2011).

Such directly visible adaptations could potentially allow for more student agency, since, if they were able to see what content was recommended for them to learn next, this could prompt conversations with their instructor about the path that is appropriate for them, and potentially lead to more student choice.

On the other hand, the mode of adaptive presentation might be indirect, with the system simply hiding content or links that are deemed irrelevant or inappropriate to the current state of the learner's knowledge (Hauger & Köck, 2007). The indirect presentation of the adaptation might also entail a learning path presented as one among many available options, without clearly indicating which is the most appropriate for the learner. Alternatively, the appropriately adapted learning path might be presented as the only available option, without any indication that there are alternative options available at all (Magnisalis et al., 2011). In this prior case, such as with CogBooks, when a student's performance on an assessment indicates that they need prior knowledge, the system directs students to a set of pre-requisite material, without indicating which may be more useful or relevant than the others. Such an adaptation, along with the invisibility of the adaptation of systems like Knewton, which simply present the single most relevant option for the students, without indicating the possibility for other options, robs students of the chance to make informed choices about the next content they will learn.

3.4.4 - Adaptive Presentation - Mandatory or Optional

Both of these modes of presentation, direct or indirect, are used in some adaptive systems to present adaptations that are mandatory, or are used by other systems to present adaptations as recommendations or suggestions left up to the choice of the individual user. Such mandatory adaptations are incorporated into systems with what is known as "adaptivity", where the system or program creates modifications to the learning content or sequence either unbeknownst to the student, or known to them, but without their ability to control it (Akbulut & Cardak, 2012). However, as Bargel et al (2012) points out, if all of the adaptations are mandatory, then students might have negative feelings towards the paternalism of such a system, in which the algorithm purports to know what is best for the student to learn (Bargel et al, 2012). Whether or not that particular adaptation is actually more effective than the alternatives, if the student perceives that it is not, then they may be less likely to use it or persist with the adaptive system. This paternalism is a prime example of the separation between the claims of autonomy and free choice made by the designers of adaptive learning systems, with the reality of the design of most systems. This prescriptive, mandatory adaptation, while it might be more "optimal" or efficient for students, actually reduces their autonomy over their own learning.

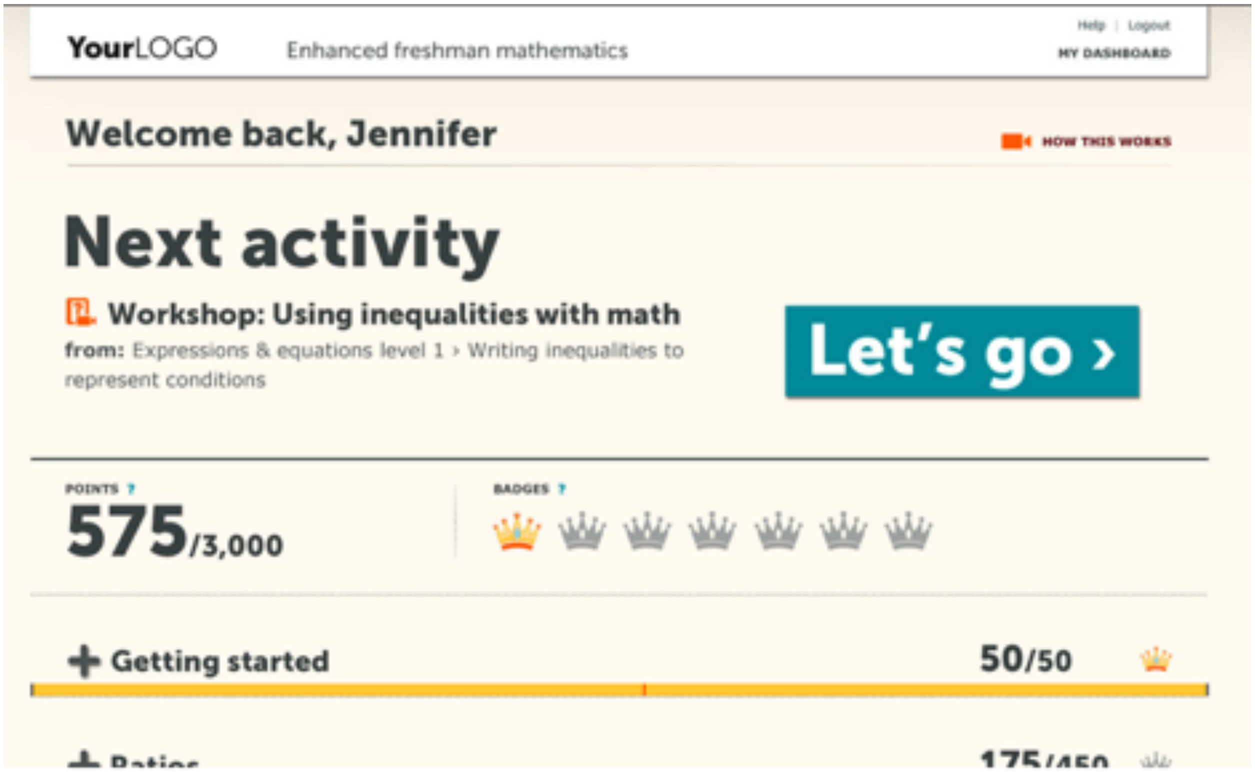

Figure 3.6 - A Mandatory Adaptation Model with Knewton (Knewton, 2015)

On the other hand, "adaptability" is the ability of the system to allow learners to choose certain parameters of the learning experience for themselves (Burgos et al., Tattersall, & Koper, 2007; Akbulut & Cardak, 2012). Often, such systems give the users choice from among a limited range of options, or present the adaptation in a recommendation format without mandating that students access that particular content element. However, in such systems, there are certain assumptions about learners that inform the design, which are not always supported by research in cognitive psychology. For instance, if the student is left entirely alone to choose the next learning object, even from among a limited range of choices, the student must know which one would contribute most optimally to their future learning path, as well as which one they would have the greatest likelihood of success at completing (Burgos et al., 2006). Though the student would not always have that information unaided, this could be indicated with a given percentage of likelihood that that particular choice is the optimal path.

Finally, even if the student knows both of those things, they may still not make the "right decision" about what to learn next (Burgos et al., 2006). In fact, recent research into the psychology of learner choice shows that learners do not always know how to utilize appropriate learning strategies when "left to themselves to manage their learning environment" (Kirschner and Mierrenbeer, 2013). However, it is not clear to what extent the support that the adaptive learning system provides can mitigate that lack of knowledge by providing students with a representation of their own knowledge state in relation to the domain model, and by providing indicators of appropriateness of content and difficulty. This is central to the larger issue of the importance of managing student choice, through recommendations and adaptations, without mitigating student autonomy and agency over their own learning. These tradeoffs and consequences will be discussed in further detail in Chapter 4.

3.5 - Implementation

3.5.1 - Infrastructural Considerations

For the end-users, the students and teachers, the underlying components of the system does not concern them as much as its implementation in the classroom. For students and teachers to use an adaptive learning system, it needs to be embedded in some platform they have access to, be it a learning management system used by their school, or a digital textbook to which they have already obtained access. After this, there are other considerations for the teacher that wants to use an adaptive learning system, involving the clarification of the three components of the system explicated above, the domain model, learner model, and adaptation model, as well as other, potentially more important, infrastructural considerations (Karampiperis & Sampson, 2005). Aside from procuring a platform for the adaptive system and the bandwidth needed for large numbers of students to access the system simultaneously, there are staffing requirements needed, such as content experts and instructional designers to ensure learning objectives and standards are addressed appropriately (Karampiperis & Sampson, 2005).

Additionally, the teacher needs departmental support to give them the freedom to try a new model for student learning, as well as institutional support to allow their class to receive credit on a competency-based model, rather than seat-time. That is to say, the traditional model for receiving credit for a course is dependent on the students attending the class for a given amount of time, during which they are expected to master a given set of skills and knowledge. With a competency-based model, students receive credit not for how long they are in the course, but for how many competencies they have mastered (Nitchot et al., 2010; Newman, 2013). Along with the staffing and policy infrastructure requirements, the teacher should consider the technological mode of delivery, or, whether or not the class is fully online, fully "face-to-face", or a blended mixture of the two (Singh, 2003). This makes a large difference in the way the system will be used, and the way the teacher will be teaching, as, with a wholly online class, students can move through the content at different paces with little cost to their class interactions.

3.5.2 - Pedagogical Considerations

After ensuring the infrastructure is in place, the teacher intending to utilize an adaptive learning system in their class must take the pedagogical implications into consideration. First, they should clearly understand what their pedagogical objectives are for using an adaptive learning system, and how that will influence their intended pedagogical style (Huang et al., 2006; Magnisalis et al., 2011). If the teachers intend to address the discrepancies in students' background knowledge and competencies, then they should choose an adaptive learning system that has a dynamically constructed learner model that modifies data from a statically collected pre-assessment. If teachers are intending to develop students' learning autonomy, then they should select an adaptive learning system that gives students choice over their learning sequence or pace, and create a class culture that supports and cultivates that autonomy (Ku, 2008; Rhode, 2009; Poelhuber et al., 2008).

Regardless of whether the course is online, face-to-face, or blended, the teacher should consider the implications for collaborative learning that a transition to an adaptive learning model would have on the class. In a standard, uniformly paced class, the students have opportunities for spontaneous collaboration that may not be possible when students are at different levels of progress through the course content in an adaptive system (Brusilovsky, 2004).

3.5.3 - Research Considerations

Aside from the pedagogical implications, adaptive learning systems have implications for research due to considerations about data acquisition, use, and privacy. Natriello (2012), in his report on adaptive learning technologies for the National Academy of Education, describes the research resulting from the use of adaptive systems as affecting student users and teacher users, as seen already, as well as informing curriculum development, general education research, and the recursive application of improving the design and development of adaptive systems themselves. I have discussed in great detail the changes to the learning process that research using adaptive systems can offer to students. For teachers, further research in adaptive learning systems can yield immediate insight into student performance at varying levels of granularity, longer range insight into improvements in the design of course content elements and whole course sequences, as well as provide insight into causal factors in student motivation (Natriello, 2012). The general education research community can gain insight into relationships between student cognitive profiles and their performance on various course structures and learning objects, as well as gain a greater understanding of the nature of students' ability to self-regulate, and how that affects, and is affected by, structured adaptive support (Natriello, 2012).

3.5.4 - Data Privacy Considerations

Finally, there are considerations for the data collected via adaptive learning systems that are considerably different than for traditional education research data. The data collected may have a greater volume, with large numbers of students using the learning system, a greater variety, with data captured on both educational and contextual elements of the system, and a greater granularity than traditional data capturing mechanisms. In addition, schools that are using adaptive learning systems consistently with their students have the ability to analyze data in shorter, more tightly iterative cycles (as the systems gather data and adapt the learning process to reflect the inferences from that data), as well as larger, more longitudinal data collection arcs than would be cost-effective with other data collection methods (Natriello, 2013).

However, for all their advantages, the data gathered by adaptive learning systems do have gaps, as they do not adequately take into account the complexity of experiences that students may have had prior to their initial point of access to the system, nor do they account for contextual data or contemporaneous events outside of the system, such as conversations with teachers, other students, or other spontaneous events that may cause the learners' knowledge states to change (Natriello, 2012). Lastly, the educational researcher grappling with data acquired through an adaptive learning system must be able to negotiate the third party of the adaptive platform or provider, a party that is not traditionally involved in the data collection process of typical education research (Natriello, 2012). This addition of the third party adaptive learning provider into the educational research process raises uncomfortable, unresolved questions about students' data privacy. Students generate data on a constant basis through their interactions with the system, much as they do with other modern networked technologies, but it is unclear what protections should be put into place to safeguard students' privacy, without preventing them from taking advantage of the benefits of actionable data on their learning (Natriello 2012).

3.6 - Conclusion

Though the optimal ratio of students to teachers would be, ideally, one to one, this is not feasible in any scalable way, and as such, educators must consider ways to make a 20 or 30 to 1 ratio of students to teachers be more individual and personalized. Each student has various needs, preferences, goals, learning styles, and degrees of background knowledge that they come to a course with, and it is a rare and skilled teacher that can effectively differentiate their instruction and assessment to adequately take those differences into account. Though personalized learning has been a goal of educators, researchers, and policy makers for at least a hundred years, it is only recently that technology has been able to provide effective supports to teachers to allow them to personalize instruction in meaningful, reproducible, and scalable ways.

Adaptive learning systems, if used effectively, have the potential to allow teachers to design a course of study that students can progress through at a pace and sequence appropriate to their current levels of knowledge. However, not all adaptive learning systems allow the same degree of control to the teachers over the design of their course, and not all systems allow the same control to the students over their learning path or pace. As a result, it is important for teachers, administrators, and researchers interested in the possibilities offered by adaptive learning systems to understand the three main components, the domain model, learner model, and adaptation model, as well as the implications for pedagogy and research of such adaptive learning systems. Like other education technologies, adaptive learning systems should be considered a supplemental tool to improve the education process, and not as a wholesale replacement for teachers or existing systems of learning.

Chapter 4 - The Cybernetic Tradition

4.1 - Overview

In earlier chapters, I discussed how values of liberal humanism are supposedly enacted through adaptive learning technologies, according to the rhetoric of the companies that design those systems. In Chapter 2, I analyzed those companies' rhetoric about student autonomy and control, and contrasted their claims with the way those ideas are understood by liberal humanist education philosophers. Then, in Chapter 3, I discussed a general framework to understand the design and functioning of adaptive learning systems, and used that framework to drill down deeper into the specific components and functions of particular systems with a taxonomy.

Though the users of this software, such as students and teachers in particular, may very well have the literacy skills to understand and defend against rhetorical claims written in white papers and position statements (if they even read them), there is another kind of rhetoric at work in the software itself. This "procedural rhetoric" enacts the values of the software's creator through the computational logic at work in the code, manifested through the possible actions that the users can take. Though this is true for any technology, for our purposes, I will be attempting to understand the values enacted through the design of adaptive learning software. It is our argument that the cultural logic embodied in this software is more closely aligned with the tradition of cybernetic command and control systems than that of liberal humanism.

In this analysis of learning technologies, I will first look at the ways in which technologies enact the values of their designers, through what David Golumbia has called the "cultural logic of computation." From there, I will argue that the specific cultural logic of adaptive learning systems is, in fact, in the tradition of cybernetic control systems, which has the effect of reducing the complexity and humanity of students to a set of informational and probabilistic metrics. To understand the unseen influence that cybernetics has had on rhetoric about learning and learning management systems, I will explore the similar rhetorical influence which clocks and self-regulating feedback systems had on political and religious rhetoric in Early Modern Europe.

Finally, I will return to the specifics of adaptive learning technologies, and the particular ways in which they embody cybernetic principles of command and control. The influence of cybernetics on education must also be understood within the context of Taylorist behavioral management ideals of authority, efficiency, and optimization as the ultimate goals of education, and I will close with a discussion of the relationship between these two cultural influences on the design of adaptive learning technologies.

4.2 - Cultural Logic of Computation

The design of technology is never an apolitical act. This is true for both the “moral instrumentalism” of speed bumps and turnstiles, which enforce a particular kind of morality through limiting the actions people are able to take, and the softer, more subtle persuasion of computational logic. In the classroom, educational technology has, since its inception, often been used to reinforce traditional values of centrally located authority in the teacher, as well as the contributing to the reification of established bodies of knowledge. What began with large chalkboards used to direct all students’ attention to a single concept has led to mass produced textbooks used to standardize learning for ever-increasing numbers of students in the schools.

Even before the advent of computers and digital learning technologies in classrooms in the 1980’s, learning technologies had been present in the classroom in the form of blackboards, textbooks, and other media forms adopted and re-appropriated for instruction, such as radio, television, and movies (Collins & Halverson, 2009). As Elizabeth Losh has argued in her book, *The War on Learning*, instructional technology “shapes interaction, mediates communication, participates in social relations, and amplifies the message of the instructor” (Losh, 2014). However, such technologies do not mediate communication between students and teachers neutrally, but always shape and influence the possibilities for interaction. For instance, a televised educational program might amplify a single instructor to allow many students to see and hear them, but it also amplifies a single teacher’s thoughts at the expense of many others’ (Collins & Halverson, 2009). Losh argues that learning technologies do not simply amplify messages, but also “organize information, and thus shape how we make meaning” from the components of our world (Losh, 2014). In this vein, technology is not “simply” an intermediary between its users and the world, acting as a transparent lens to view or magnify the world, but instead reconstructs the world according to its own logic.

Digital tools are no exception. Computationally designed learning systems are equally susceptible to the trend towards reproduction of existing educational models and can, as Losh argues, be used as “a way of conditioning subjects to respond well to the computational model” (Losh, 2014). The term conditioning used here is particularly apt, since many of the ways digital learning technologies have been used have tended towards a sort of behavioral conditioning and reward system. Even Benjamin Bloom, developer of the cognitive taxonomy of learning, has argued that what we think of as “mastery” learning is inherently behaviorist, rewarding preferred outcomes to induce more occurrences of that behavior. Therefore, we must first understand and unpack the underlying assumptions in the design of learning technologies, before we can know how they have been manifested in the design of those tools.

Similar to the aforementioned speed bumps and turnstiles, the design of digital learning technologies influence the behaviors of the students and teachers that use them, in ways that are not always immediately visible. Even with a technology as seemingly apolitical as a learning management system (or, LMS), the design of the interface constrains the choices the students and teachers can take, allowing only certain actions deemed permissible. As mentioned previously, students and teachers may be predisposed to resist rhetorical appeals on the “visual or verbal register,” but the algorithmic rhetoric of the inner workings of a LMS or other digital learning tool is obfuscated on several levels of accessibility (Losh, 2014). First, the students and teachers cannot typically view the source code that powers the system, nor could most of them understand it, even were it to be made accessible, due to its encoding in programming languages many students and teachers cannot parse. This inaccessibility is particularly manipulative, as the students and teachers must either accept the actions permissible by the software, or, in their resistance, must resist blindly, without fully understanding the values expressed by the “complex rule-based systems” of the software algorithms (Galloway, 2012; Losh, 2014).

Just as the design of learning technologies can be persuasive or rhetorical, so too can they be moral in nature. Peter-Paul Verbeek, in his book on the moral implications of technology design, has said that technologies do not only “act as intermediaries between human desires made manifest in the material world, but act as mediators” helping to actively shape that reality (Verbeek). Such a distinction is a crucial one, for it holds the technology (and, implicitly, the designers of the technology) accountable for the nature of the meaning constructed through the users’ interactions. Therefore, moral action can be induced or influenced through the design of technologies that constrain or enable certain desirable actions above others. Such “materially induced behavior” is certainly not limited to education, as one can see in examples from domains such as politics, medicine, and transportation, with devices and tools that have been designed to enforce certain ways of being above others (Verbeek). Unlike moral propositions and beliefs about morality, a tool that is designed to elicit moral behavior of a certain kind cannot be argued against, only resisted or abandoned. As seen in the discussion of distributed cognition in Chapter 2, cognition has never been fully individual, as we have always incorporated tools into our cognitive systems as devices to think with. Humans have thus also never been fully autonomous moral actors, as our moral (or amoral) behavior has always been enabled and constrained by the particular tools we act with.

As we seek to understand the design of current adaptive learning systems, and argue for a more liberal humanist approach to their design, I will be using this idea of ideological design as an entry point into understanding the values which are actually embedded in adaptive learning systems, contrary to those professed by their designers. As Losh has said, the “choices about code, platforms, and infrastructures express particular values,” values not always visible or noticeable to their users, but significant nonetheless (Losh, 2014). Since all teaching is mediated in some way through technologies of inscription and communication, we must understand the ways in which the morals and ideologies of the designers of technology become instantiated through their design. By studying the possibilities for interaction that users have with the technology, and by seeking to understand the political and moral ideologies enacted through the design of the interface and underlying algorithms, we can better understand the actual values and ideological tradition in which such systems exist.

4.3 - Metaphors of Technology

4.3.1 - Automation and Self-Regulation

The influence of technology on people's self-concept and worldview has a long history. In this section, I will examine the ways in which technological metaphors of self-regulating, autonomous systems such as clocks, steam engine governors, and thermodynamic systems began to spread across early modern Europe after their invention, informing many writers' understanding of political and biological systems, among others. In some countries, clocks became a metaphor for a smoothly running government, or, in the Deist religious view, a metaphor for a universe that was created by a clockmaker god whose creation could run independently of his control (Mayr, 1986). However, there was a tension in the interpretation of this metaphor, between continental Europe and England in the early 18th century (Mayr, 1986). Some political philosophers in France and Germany viewed the clock as a benign metaphor for a well-run, orderly, though free state, while others in England saw its pre-built nature and continuous functioning as representative of an authoritarian, conservative regime.

Education, too, has always had a tension between individual, humanist control, and external, authoritarian control. This could be either internal to the classroom, with the individual control desired by students repressed by the authority of the teachers, or it could be the individual desire of teachers to control their curriculum, in tension with the authoritarian curricular control of the school administration or curriculum designers. Adaptive learning systems are one particular learning technology that embodies a tension between the concepts of self-regulating systems and external control and guidance. In Chapter 3, I analyzed where the locus of control was situated in the design and usage of adaptive systems, but now I will look more broadly at understanding how a “self-regulating” system like adaptive systems can be said to offer control at all. This shifting locus of control has consequences, as seen in Chapter 2, for the humanist experience of education, rather than a dehumanizing, mechanized education.

An early use of metaphors of clockwork and automation in education can be seen in the writings of 17th century Czech educational philosopher, John Amos Comenius, who compared a well-run, systematically organized schooling process with clockwork, saying that the art of teaching was “no more than the mastery of time, material, and method” (Mayr, 1986). In such a metaphor, one can see the complex, messy process of education reduced to a simple set of levers and gears that, once mastered and regulated, would be as reliable and automated as a clock. This abstraction, though perhaps pleasant for Comenius to imagine, is of course, unrealistic, as was his prediction that an education carried out by his plan “will be as free from failure as are these mechanical contrivances” (Mayr, 1986). Moreover, such a desire to see schooling automated like clockwork was emblematic of a larger movement among thinkers in early modern Europe to view their world through the metaphor of the machine, viewing the clock as a “symbol of any authority that brings order into human life” (Mayr, 1986).

Because clock metaphors soon proved to be inadequate to explain the functioning of complex systems such as education and political life, they were thus supplemented with newly discovered processes of physics and thermodynamics. Even Isaac Newton is commonly thought to have believed in the notion of a “clockwork universe,” but in fact, his central metaphor was not one of the unmitigated automation of a clock, but that of a “constantly changing dynamic system needing constant attention and periodic adjustment from God” (Mayr, 1986). This dynamic system, as he explains it, is a metaphor drawn from thermodynamic systems, which tended toward entropy and heat-loss, unless regulated by some external actor. When applied to the universe as a whole, such attention and adjustment could come only from some force outside of the system (a deity, in this case), but when applied to the dynamic system of political life or education, there was disagreement between liberal and conservative political philosophers over the exact source and method of the adjustment needed to keep the systems running smoothly.

4.3.2 - Feedback Control

To liberal philosophers, particularly in 17th century England, the ideal dynamic system was one that could could achieve a sustainable equilibrium through its own autonomous, independent actions, without external intervention from a higher authority. Central to this emerging liberal concept of order was the idea of internal self-regulation, and the steam governor became the central metaphor for these self-regulating dynamic systems, much like the clock was for fully automated processes (Mayr, 1986). The steam governor was a dynamic system that, through the process of feedback control, regulated the flow of steam in an engine through raising or lowering the height of two large masses on metal arms. When the speed increased above or below a certain threshold, an aperture was closed or opened, altering the flow of steam, and thus the speed. With a predetermined set of mechanisms and thresholds, this system was able to maintain equilibrium of speed through a process of self-regulation. This provided an extremely useful metaphor for political philosophers attempting to reconcile the idea of multiple branches of government that could regulate their interactions autonomously, without the need for intervention from a higher authority such as a king (Mayr, 1986).

If a liberal political system is one in which all actors were able to regulate their actions according to their own autonomous desires, then the central issue becomes the prevention of conflicts of interest between individuals. In contrast, an authoritarian political system includes a central authority that is able to regulate and control individual behavior before conflicts might arise. The metaphor of a self-regulating engine indicated how systems could govern themselves, given a structure in place to allow for feedback control to maintain equilibrium in a dynamic system. It is not always clear, though, how this feedback control mechanism is established, or by whom. Perhaps the best known example of such a self-regulating system in practice is Adam Smith’s concept of the balancing mechanisms of the free market which guide economic interactions with an “invisible hand” (Mayr, 1986; Hayles, 2008). Unfortunately, as recent economic developments have shown, there continues to be a need for some externally created regulation to ensure the economic system remains in a state of equilibrium. The tension yet remains, then, between how to balance the individual autonomous desires of single actors in the system with an externally regulated control structure. Although self-regulating systems governed by feedback offer one possibility, the central issue lies in who constructs the mechanisms and establishes the thresholds for self-regulation.

Cybernetic systems offered one solution. Such systems were related to the self-regulating feedback-controlled dynamic systems of the steam governor, but with one crucial difference. Human input. The first major application was a human-controlled anti-aircraft gun developed and used in World War II. The human operator would aim at the enemy aircraft to the best of his abilities, and the system would take in his input, incorporate it into its probabilistic aiming algorithm, and correct for the lag time of human responses. As cybernetic control systems began to filter out of military research labs and into the larger culture in the middle of the 20th century, more and more rhetoric in other, non-scientific fields began to be infused with metaphors of self-regulating, cybernetic systems, partially governed by probabilistically determined feedback from internal mechanisms, and partially controlled by humans. Seeing the application of the cybernetic model to other domains where it was not intended, such as political thought, economics, or education, one of the original developers of cybernetics, Norbert Wiener, began to struggle with how cybernetic systems could be designed to protect the the autonomy of the individual user (Wiener, 1954; Hayles, 2008).

Now, in fields like finance, where semi-autonomous machine agents are currently used to purchase and sell microtransactions of stocks at speeds far beyond where humans could intervene, the locus of control is shifting away from humans as the central controllers and operators of the machines (Wiener, 1954; Hayles, 2008). As machines continue to augment and then replace human capacities in a range of disciplines and fields, cybernetic systems must be designed in such a way as to not entirely eliminate the human as the locus of control (Wiener, 1954; Hayles, 2008). For the classroom, and for adaptive learning systems in particular, there has always been a tension between the level of control placed in the hands of individual learners, and that placed in the teacher as the locus of control. The traditional “rhetorical performance of knowledge” as Elizabeth Losh calls it, tends to be a top-down, authoritarian control structure, where the pace and content of the class is dictated to the students as a whole class, from the teacher. One can clearly see the difference between such a model for learning, and the student-centered, constructivist approach as explained in depth in Chapter 2.

In the following sections, I will explicate the functions of cybernetic control systems in order to understand the implications they have for the design of adaptive learning systems, so as to mitigate the reduction of student autonomy in a fully automated, self-regulated, cybernetic model.

4.4 - Cybernetics

4.4.1 - Introduction

Though cybernetics as a concept has existed for hundreds of years, it is commonly thought to have found its modern instantiation in Norbert Wiener and his system for human-operated anti-aircraft guns in the 1940’s. Build with mutually responsive feedback loops between the human operators and machine elements, the cybernetic system of the anti-aircraft gun was used to supplement the limitations of its human operators with the speed and accuracy of machines. The tracking system took in the aiming input from the human operator and modified it by firing at the target the human intended to shoot (the enemy plane), rather than the point where they actually shot.

However, Wiener soon came to regret what he saw as the dehumanization of the human operators of his cybernetic system. As he said in his later work, *The Human Use of Human Beings*, “what is used as an element in a machine, is in fact an element of the machine” (Wiener, 1954). This shift in preposition between the two uses of the term “element” construes his cybernetic systems unfavorably as a factor in the mechanization of humanity. For Wiener, it was important to design and construct cybernetic systems that “reinforce rather than threaten the autonomous self” through clearly determined boundaries between the machine and human elements (Wiener, 1954; Hayles, 2008). However, as discussed in Chapter 2 and in the beginning of this chapter, such boundaries are not always clear or stable, as the design and use of technology will inevitably influence the users’ actions, behavior, and perhaps, self-concept.

4.4.2 - Information and Probability

Cybernetics as a discipline emerged from the mutually interlocking domains of information theory and probability. In order to blend the human and machine’s “experience” of the world effectively, a cybernetic system must have some probabilistically determined measure of accuracy in representing a given state of the world. This work was informed by developments in communication and information theory by Shannon and Weaver in 1949. For humans as well as machines, information was conceived of as a necessary component in the “continuous process by which we observe the outer world, and act effectively upon it” (Wiener, 1954). In this view, humans and machines both take in information about the world, process it in some way, and use that information as feedback to change their future behaviors. However, the mechanisms by which humans and machines “take in” information and use it to construct meaning about the external world are very different, and for Wiener, this was a significant factor in his later rejection of the dehumanizing nature of cybernetics. Though humans use information technologies to observe, record, modify, and share their experience of the world, there of course remain other biological, social, linguistic, and cultural components to the cognitive process of information acquisition and retrieval in humans, which do not exist in quite the same way in information-processing machines.

When the representations of experience, information, and behavior in humans and machines are considered to be relatively equivalent informational patterns, there is a reduction of the unique ways in which humans process information in order to achieve the questionable goal of efficiency. In the cybernetic tradition, where the world is seen as fundamentally probabilistic (rather than deterministic), human and machine actors cannot fully know the state of the world (or the “microstate” of any individual piece of it) with complete detail or accuracy (Hayles, 2008). Therefore, probabilistic values are assigned to events to deal with the uncertainty of fully knowing the current state of the world, and are often used to predicting the likelihood of future states.

Wiener discovered that any pre-determined behaviors for the machine components of his cybernetic systems would not be able to cope with the unexpected developments of life in an uncertain, though probabilistic world. The regulatory control mechanisms for cybernetic systems, then, could not be centrally created or statically pre-determined, but needed to be a “flexible, self-regulating system of control, based on feedback from the system itself” (Mayr, 1986; Hayles, 2008). The success of self-regulating, feedback-enabled cybernetic systems is dependent on consistent levels of probability in their interactions with the world, where “the statistical differences between individuals is [assumed to be] essentially nil” (Wiener, 1954).

The prototypical cybernetic system, that of the anti-aircraft gun that compensates for the deficiencies of its operator, does not create and store a model of its user, and thus, has no way to infer differences between operators. That is to say, in order for the machines to work effectively to, for instance, assist a human operator in targeting an aircraft, they must assume that the statistical difference between humans is low enough to operate effectively with different operators or pilots of the crafts being targeted. Such a probabilistic determination of human identities and behavior is at work in adaptive learning systems as well, particularly in the types of systems that use a stereotype model to cluster students together into related groups and recommend content for them that has been proven to be successful with probabilistically similar students. With such a model, the uniqueness and identity of individual humans and students has been reduced to the abstraction of various sets of informational and probabilistic evaluations of likely behaviors.

4.5 - Cybernetic Influences in Education

4.5.1 - Taylorist Organizational Management

We have seen how metaphors of technology can shape our understanding of complex systems, and I will now use cybernetics as a metaphor by which to understand education systems and adaptive learning systems. For much of the 20th century, there have been political and economic pressures to make the educational process more efficient, driven largely by ever increasing numbers of students in public schools (Collins & Halverson, 2009; Losh, 2014). Despite the best intentions of educational reformers to create a culture of free, autonomous discovery learning in schools, such reform efforts have become subsumed into the standardized model of prepackaged, reproducible curricula. For example, the liberal educational reform efforts of John Dewey and Jean Piaget to allow students to learn by experience and discovery have become disempowered by their adoption in schools, particularly when the students are led to "discover" knowledge by a forced march through a pre-constructed curriculum (Papert). It is not just with adaptive learning systems, then, that despite the best intentions of educational reformers, the desire to empower students has instead tended towards standardized, externally regulated curricula, rather than the organic discovery advocated by Dewey, Piaget, Freire, and others.

One of the major movements in education that has encountered resistance from liberal education philosophers has been the incorporation of the behavioral management science espoused by Frederick Taylor. In contrast to the ideal system desired by liberal education philosophers, ie: a system that is largely self-regulating and driven by the choices of individual members of the system, the public school model as dictated by Taylorist organizational philosophy was a centrally regulated, authoritarian view of order, motivated by probabilistic functions that determined the optimally efficient behavior. According to this model, there would be an optimal set of behaviors for the individuals in an organization that led to the maximum operating efficiency for that organization. Taylor saw the role of the centrally located organizational authority, or in the school's case, the administration, to guide or coerce its employees to that behavior. One can see in this central regulation of employees' time and behavior what became the standardized curricula and regulated student behavior of the current public school system (Collins & Halverson, 2009).

In fact, in 1904, educational theorists at Stanford explicitly called for a redesign of public schools to adopt structures and methods of the "modern bureaucratic organization" (Collins & Halverson, 2009). Such a move, driven by economic and social pressures, would prove to have far-reaching consequences, as the organizational standardization methods later endorsed by Taylorist philosophy resulted in a reduction of the autonomy and agency of teachers and students. In order to properly assess whether or not a given school was being run effectively, a method was developed to regularly measure "production," much like a factory, to continuously assess student learning to see if the teacher and class were performing to specifications (Collins & Halverson, 2009). Such statistical measures were designed to assess student intelligence, amount of content learned, quality of teaching, and whether those parameters were progressing at the specified rate (Collins & Halverson, 2009). In true cybernetic fashion, these probabilistic indicators of school performance were used as feedback to inform the policies and behaviors of those in the school. Though this might seem to be simply good data-driven instruction, the cybernetic consequences arise when the data analytics and feedback are black-boxed into computationally driven learning platforms, as they have become in the last 30 years. This can be seen in a highly evolved form in adaptive learning systems, which have a pre-determined "mastery threshold" and provides recommendations and feedback to students often without the teacher understanding or deciding on the rationale for allowing student progress.

Despite the promise of efficiency, there have been many critics of the infusion of Taylorist management science into education. For instance, these scientific principles of management have led to the creation of modular lessons and curricula from third-party educational content providers that "supposedly deliver information in the most efficient manner," despite a lack of hard evidence to prove those claims (Losh, 2014). Other critics of the influence of Taylorist behavioral management in education have pointed out how the presence of pre-programmed lessons and pedagogy leads to a reduction in teacher agency, and thus, authority (Hayles, 2008; Collins & Halverson, 2009). Perceptions about a lack of teacher agency and authority have also been linked to decreasing numbers and quality of new teachers entering the profession (Hayles, 2008; Collins & Halverson, 2009).

Moreover, with the rise of computerized assessments in “impenetrable black box” proprietary software, some have criticized the fact that such assessment software often focuses on “goals that are most easily measured in quantifiable terms, rather than focusing on the more meaningful results that are difficult for computers to calculate” (Losh, 2014). This is a common objection among teachers and parents who do not believe that an assessment can capture the range and complexity of learning in their child. Others argue against the “false efficiencies of standardization” that value a “one-size fits all” approach to assessment, without taking into account the consequences for individual development that such a model entails (Losh, 2014).

4.5.2 - Adaptive Learning Systems as Cybernetic Systems

No technology is ideologically neutral, and often, the design of the technology itself enacts morals and ideologies that may be contradictory to the stated purpose and ideology of the designers of the tools. As seen in Chapter 2, despite the best efforts of proponents of adaptive learning systems to position their technologies in the tradition of liberal humanism, the nature of the “autonomy” experienced by students and teachers using these systems belies their rhetorical efforts. The cybernetic system, therefore, is a more accurate metaphor for how adaptive learning systems enact control and regulation of their users, both teachers and students.

The promise of adaptive learning systems, as seen in Chapter 2, is that they could address the learning needs and goals of students in an individual way that a more “efficient” standardized assessment could not do. Yet, when teachers and students abdicate their choices to the “optimized” efficient path of the pre-created curricula, they are no longer shaping the direction of learning, but are themselves shaped by the curriculum (Collins & Halverson, 2009; Losh, 2014).

With the presence of computers and the Internet in schools, teachers have already had their authority reduced by virtue of no longer being the sole keepers of knowledge in the classroom. To have their agency reduced as well as their authority threatens to diminish the perceived value of the teacher from an instructor to a facilitator, simply guiding the students along predetermined knowledge paths (Collins & Halverson, 2009). For students, if they were to have complete agency over their learning, they would most likely be unable to decide effectively what the most appropriate or effective next steps for their learning would be (Kirschner). However, adaptive learning systems *could* provide mediated agency for students, recommending appropriate next steps, but allowing students to make informed choices about what they want to learn next.

Elizabeth Losh describes computer-mediated learning in learning management systems as using “continuous feedback loops of error-correction” and progress assessment (Losh, 2014). Her language of feedback loops is directly drawn from the design of cybernetic systems, which use feedback from the outside world (ie: a teacher or grader assessing the students’ performance) to inform the behavior of the system in its current state, or, to recommend which learning materials or assessments it provides for students. According to Norbert Wiener, feedback is “the control of a machine on the basis of its actual performance rather than its expected performance” (Wiener, 1954).

For adaptive systems, the goal therefore would be to use feedback from *actual* student performance to modify the behavior of the system, rather than simply prescribing tutorial videos or assessments based on the *expected* performance of the student. However, even at their best, adaptive learning systems are prescribing future student behaviors based on the expected performance on learning objects of a similar type. Since, by definition, the students will not have had experience with a given learning object or assessment, adaptive recommendations are feedback loops based either on the student’s performance on probabilistically *similar* learning objects, or based on the performance on that learning object of similar *students* and assuming that the outcomes will be the same.

As adaptive learning systems attempt to provide uniquely tailored learning pathways for students that supposedly respond to students' individual needs and desires, they in fact tend to merely reproduce the same centrally determined authoritarian structure they would like to escape from. Except, by replacing the central authority of the teacher with the implicit, cybernetic authority of the predetermined behaviors of the adaptive engine, such systems are less able to "act spontaneously in the face of the unexpected" as a teacher could do (Mayr, 1986).

4.6 - Conclusion

In a cybernetic system, control is shifted away from an individual authority with autonomous decision making power, towards pre-determined patterns of responses based off of probabilistically determined feedback control loops. In educational pedagogy and policy making in the 20th century, such systems of pre-determined behaviors were informed by the Taylorist model of behavioral management and optimization of educational processes. Though efficiency and increased productivity seems like a worthy goal for learning, the exact nature of the optimization process and output has not been thoroughly interrogated. What is being optimized, and what gets lost when certain easily quantified outputs are optimized for at the expense of less easily quantifiable results?

Despite the influence of Taylorist standardization and cybernetic control structures, learning, when done authentically, should be what Elizabeth Losh refers to as a "highly situated activity that resists supposedly rational procedural schemes." (Losh, 2014). That is to say, the particular nature of the rhetorical act between teachers and students should be responsive to the needs and exigencies of a particular class and class culture, rather than attempting to probabilistically predict the optimal outcomes for all situations.

Norbert Wiener was concerned his cybernetic systems would turn people into machinic components in a machine-driven society. Despite a desire for adaptive systems to respond to the learning needs of individual students, the design of those systems always reflects and enacts the implicitly held values of its designers. With the essential role of the curriculum designer in creating the domain model, as well as the faith placed in the authenticity of the assessments of student knowledge, adaptive learning systems fall squarely in the cybernetic conception of order I have laid out. Therefore, the recommendations for student learning they offer should be taken with clear-eyed skepticism and an awareness of the nature of the algorithms that produced those recommendations, an impossibility with the black box design of most adaptive learning software. In order to prevent students and teachers using adaptive learning systems from subsuming their autonomy to the embedded optimization algorithms, the design of adaptive learning systems must be made more easily understandable and modifiable by their users. In the following chapter, I will discuss these issues and other ways to make adaptive learning systems more humanist, in addition to providing a guideline for teachers and administrators attempting to select a humanist adaptive system that reflects their values.

Chapter 5 - Design Guidelines

5.1 - Overview

Because they are no longer the fixed, static curriculum of a printed textbook, adaptive learning systems are lauded for their “personalized” method of responding to dynamic feedback about students’ performance. As I have argued, this responsiveness is ultimately more cybernetic than humanist, due to the design of the self-regulating adaptive system, and as such, often denies control over the learning process to the teachers and students that should have control over their learning. However, this does not have to be the case.

This chapter provides a detailed taxonomy of 16 currently existing systems, using the criteria established in Chapter 3. I then provide case studies of two adaptive systems, which offer differing levels of autonomy and agency for teachers of students. Finally, I provide a set of guidelines for the designers of adaptive learning systems, as well as criteria for administrators, instructional designers, and teachers involved in the selection and implementation of an adaptive learning system, to both create and use more adaptive systems that are more humanist, rather than those that tend towards a cybernetic model.

In Chapter 3, I discussed the components of adaptive learning systems in detail, and then, in Chapter 4, engaged in a cybernetic interpretation of their design at a more broad level. With that in mind, let us now discuss some of the design choices of specific adaptive systems that might make them more or less cybernetic. This explanation will lead towards guidelines for designers of the systems who want to make good on their humanist rhetoric, and criteria for selection and implementation of adaptive systems for educators looking to make informed choices about a system that fits their humanist values.

5.2 - System Taxonomies

In this section, I will be referring to specific components of adaptive systems that I discussed in Chapter 3, and for each, will indicate inflection points for a more humanist design. I have also included a taxonomy at the end of each sub-section that I created in the process of conducting this research, situating 16 adaptive systems within the design framework for each component.

5.2.1 - Provider Taxonomy

For the initial taxonomy (displayed below), I have categorized the systems by whether or not they consider themselves to be a platform or a publisher, by their target users, by their size, and by their ability to have integration with the institution's existing LMS. These criteria are perhaps more important for the selection process than any that follow. First, the platform or publisher distinction is crucial to understanding whether or not institutions will be able to upload and host their own content on the platform, or whether, as a publisher, they will provide content that their own "content experts" have created. Secondly, the target users are fairly straightforward, including K-12 institutions, postsecondary institutions, corporations and job skills retraining, and individual learners. Though they are of course not limited to those users, most of the content and the means of access will be tailored to those audiences. The size of their user base is grouped as either small (<100,000 users), moderate (100,000 - 500,000 users), or large (>500,000 users) (Newman, 2013). Finally, the LMS integration is a binary characteristic, indicating whether they use the LTI interoperability standard to allow for integration.

System Name	Platform	Publisher	Target User	User Base	LMS Integration
Adapt Courseware		✓	Postsecondary	Small	✓
ALEKS		✓	Individual, K-12, Postsecondary	Moderate	
AnewSpring	✓		K-12, Postsecondary, Corporate	Moderate	
Cerego	✓		K-12, Postsecondary, Corporate	Large	✓
Cogbooks	✓		Corporate	Large	✓
Duolingo		✓	Individual, K-12	Large	
DreamBox		✓	K-8	Large	
Enlearn	✓		K-12, Postsecondary	Small	✓
Jones and Bartlett Learning		✓	Postsecondary	Small	✓
Khan Academy	✓		Individual, K-12	Large	
Knewton	✓		K-12, Postsecondary	Large	✓
LoudCloud Systems	✓		K-12, Postsecondary	Small	
LearnSmart Advantage Suite		✓	Corporate, Government	Large	
Open Learning Initiative		✓	Postsecondary	Moderate	✓
Quantum Simulations		✓	Postsecondary, Corporate	Moderate	
Smart Sparrow	✓		Postsecondary	Small	✓

Figure 5.1 - A Taxonomy for Adaptive Learning System Providers

5.2.2 - Domain Model Taxonomy

Next, in a taxonomy of the **domain model**, some systems allow for teachers to construct their own domain model, by creating their own learning objectives, and mapping those onto sets of hierarchical relationships and pre-requisites. Such a design is clearly more humanist and respectful of the teacher's autonomy and judgment than one in which the teacher either selects from a pre-generated list of learning objectives, or, in some cases, simply receives the pre-authored domain model from the publisher on which the adaptive system is embedded.

This design component has multiple elements to it: the creation of the learning objectives, the creation or selection of content elements (lecture videos, text articles, interactive simulations, assessments), and the arrangement of those elements into a hierarchy, either through a content authoring graphic interface, or through tagging them with appropriate metadata. Some systems, such as LoudCloud, use a semantic search to "discover" related content available online, such as Open Educational Resources (OER's), or related videos from Khan Academy, YouTube, or another source. Such a model has obvious problems, as discussed in Chapter 3, such as the reliability of the material presented in those sources, or the relevance of their relation to the current content element at hand.

This method is more obviously of the cybernetic model, as it assumes that the algorithm that evaluates relevancy is more effective or more efficient than allowing a teacher to manually select and input content they feel is related. Finally, some systems allow teacher and student users to view the graphic representation of the domain model, to see which content elements they have completed, and which lay ahead of them. Allowing the students to view the scope of the domain is more humanist than hiding this knowledge from them in a back-end model, since they could be more informed about content they might be about to learn, and be inspired to ask questions about its connection to their current material.

System Name	Instructor-Created Elements	Instructor-Arranged Elements	Algorithmically Discovered Content	Domain Model Visualization
Adapt Courseware	✓	✓	✓	
ALEKS				
AnewSpring	✓	✓		✓
Cerego	✓	✓		✓
Cogbooks	✓	✓		✓
Duolingo				✓
DreamBox				
Enlearn				
Jones and Bartlett Learning		✓	✓	
Khan Academy				✓
Knewton	✓	✓		✓
LoudCloud Systems	✓	✓	✓	✓
LearnSmart Advantage Suite				
Open Learning Initiative				
Quantum Simulations				
Smart Sparrow	✓	✓	✓	✓

Figure 5.2 - A Taxonomy for the Domain Model

In Figure 5.2, 7 out of 16 systems having instructor-created elements, or, allowing instructors to create or upload material that they have already developed for their courses. In the other 9 systems, they use a publisher model, where they publish content developed by so-called domain experts, either hired by their company, or from an external instructional design team. The obvious issue here, aside from the lack of teacher autonomy, is that the quality of the material is only as good as the quality of the instructional designer, and does not allow for institutions to produce their own content. For 8 of 16 systems, the instructors are able to arrange their course elements in a manner of their choosing, establishing hierarchies and relations between content elements. The 4 systems that use semantic searches to algorithmically discover “related” content could potentially reduce the teacher’s autonomy by potentially assigning texts or videos that the system determines are relevant, but which the teacher might not consider to be the case. Some systems, such as LoudCloud, allow teachers to approve the suggested related material before it is displayed to students. Finally, the visualization of the domain model, used by 8 of the 16 systems, allows for students to view the scope of the domain during their interactions with it, which could potentially increase student engagement and confidence, though this needs more research.

5.2.3 - Learner Model Taxonomy

Next, in the **learner model**, the kinds of information the systems collect on students in order to create the learner model situates them as being more or less humanist. All of the systems we analyzed created a model of the learner’s knowledge state, be it a static pre-assessment, or a dynamically constructed continuous assessment of knowledge. However, the other elements of the learner model were less evenly distributed across systems. Some systems attempt to survey or infer the students’ learning style, and adapt content on the basis of how best students are presumed to learn. Such systems are attempting to be more humanist and responsive to students’ individual learning needs, but, due to the questionable reliability or consistency of learning styles over time and context, these systems’ adaptations are not as effective as other, more research-supported methods for constructing the learner model.

Other elements in constructing the learner model are learner preference and learner goals, which some systems collect explicitly through surveys or prompts for student feedback after completion of content elements. If the systems that purport to take learner preference and learner goals into account truly do, then these systems would have considerable regard for student autonomy and choice over their learning, though it is not always clear to what extent even systems that use these components of the learner model base their adaptations off of them. The final category analyzed here is the use of implicitly collected clickstream data to create the learner model, as opposed to explicit prompts for learner responses or feedback via surveys. Though this is not a binary indicator of a humanist or cybernetic approach to adaptive learning, when considered in conjunction with other components, it may be important to consider just how a system is arriving at its inferences about for instance, learning style - whether it is algorithmically determined through the learners' clickstream data with the system, or through a survey that the student is aware of completing.

System Name	Knowledge State	Learning Style	Learning Preference	Learning Goals	Clickstream Data
Adapt Courseware	✓		✓	✓	✓
ALEKS	✓				✓
AnewSpring	✓				✓
Cerego	✓				✓
Cogbooks	✓	✓	✓	✓	✓
Duolingo	✓				✓
DreamBox	✓	✓	✓	✓	✓
Enlearn	✓	✓	✓	✓	✓
Jones and Bartlett Learning	✓				
Khan Academy	✓		✓	✓	✓
Knewton	✓		✓	✓	✓
LoudCloud Systems	✓				✓
LearnSmart Advantage Suite	✓			✓	
Open Learning Initiative	✓				✓
Quantum Simulations	✓				
Smart Sparrow	✓		✓		✓

Figure 5.3 - A Taxonomy for the Learner Model

In Figure 5.3, as mentioned previously, all 16 of the systems use some measure of the knowledge state of the learners to create the learner model. Fewer systems, only 3 of 16, use the students' learning style to inform the learner model, perhaps a reflection of the lack of agreement in the literature on the validity or generalizability of this construct. The next two components, learning preference and learning goals, both have an equal number of systems (7 of 16) using them as elements in the learner model. This is particularly interesting when these supposedly humanist elements occur in systems that do *not* allow for teacher creation of content elements in the domain model. That is to say, some of the above systems, such as DreamBox and Enlearn, notably, do not allow for teachers to have decision over the creation and arrangement of their courses, but they do monitor and track students' preference and goals for their learning. This contradiction is concerning for what it reveals about the values of the designers of the adaptive system. If they do not trust teachers to create their own courses, then why do they allow for student choice over their learning (and how authentic is it, for that matter)? Lastly, only 3 of the 16 adaptive systems analyzed do not include clickstream data in their method for creating the learner model, and instead use a more explicit student feedback and survey method.

5.2.4 - Adaptation Model Taxonomy

Finally, in the **adaptation model**, various systems modify their content and adapt to individual learners on a variety of levels and in a variety of ways. First, the content can be adapted by difficulty or by media type, such as video, text, or interactive simulation. Next, some systems adapt the sequence of instruction, adapting the presentation of content elements for different students at different points in the learning process. All systems surveyed allow for their content to be adapted by pacing, allowing students to move through the content at a pace that fits their level of mastery, and not determined by the teacher or the class cohort. Another level of adaptivity is that of feedback, which some systems provide at a personalized level for individual students depending on their performance on content assessments. Lastly, the method of adaptation can be either direct or indirect, and mandatory or optional. As explained in depth in Chapter 3, a mandatory adaptation is one which the student must accept without choice over the next step in their learning. Additionally, a direct adaptation is one which is made visible to the students, so that they know there are potential options for them to choose, rather than an indirect adaptation which adapts the material without the students seeing the other potential options.

System Name	Difficult y	Media Form	Sequence	Pacing	Feedback	Mandatory
Adapt Courseware		✓	✓	✓		
ALEKS	✓		✓	✓		
AnewSpring				✓		✓
Cerego			✓	✓	✓	✓
Cogbooks			✓	✓	✓	✓
Duolingo			✓	✓		
DreamBox	✓	✓	✓	✓	✓	
Enlearn	✓	✓	✓	✓	✓	
Jones and Bartlett Learning				✓		
Khan Academy	✓		✓	✓	✓	
Knewton	✓		✓	✓	✓	
LoudCloud Systems		✓	✓	✓		✓
LearnSmart Advantage Suite			✓	✓	✓	✓
Open Learning Initiative	✓			✓	✓	✓
Quantum Simulations				✓	✓	✓
Smart Sparrow	✓		✓	✓	✓	✓

Figure 5.4 - A Taxonomy for the Adaptation Model

In Figure 5.4, less than half of the systems (7 of 16) adapting their content based on the difficulty level. This makes sense when one considers the back-end effort required to either create a suite of problems, exercises, or assessments of varying difficulty, or, as is commonly done, create a parameterized exercise generator which could create problems of varying difficulty by modifying parameters. Even fewer systems, 4 of 16, adapt their content based on the media type, which, again is understandable when one considers the cost in time and resources necessary to duplicate or triplicate every content element with a video lecture, a text document (more robust than a transcript), and an interactive simulation. Not only does it take the teacher or instructional designer significant time to invest in re-creating their content in multiple ways, if the student is being led to different media types based off of their learning style (or preference), it has already seen that there are not clear results for this adaptation.

Next, a significant majority (12 of 16) systems adapt the sequence of their content for the learner. This macro-adaptivity (rather than the micro-adaptivity of adapting individual problems), is dependent on a flexibly created domain model, one with robust pre-requisite hierarchies between content elements, for the system to provide recommendations for adaptation for individual learners' knowledge needs or goals. As mentioned previously, all 16 systems allow for individual adaptations of pacing, as these are designed to be used either individually, in an online (and not time-dependent course), or in a blended or flipped learning model where the learner is working at their own pace through the content. More research needs to be done on the implications that a self-paced adaptive model would have for face-to-face student collaboration, peer mentoring and tutoring, and the effect it would have on the overall class culture and community. Next, 10 of 16 systems use adaptive feedback, on a micro-scale, offering feedback to students based off of their performance on questions. This is typically limited to such adaptations as offering hints if students get questions wrong, or supplying pre-requisite or supplemental material if students demonstrate a lack of knowledge over particular elements of the content. Finally, exactly half (8 of 16) systems use a mandatory adaptation, rather than allowing students to choose what content elements they will access, based on a recommendation. This is, as would be expected, a significantly less humanist approach to adaptive learning, as any mandatory adaptation removes the students from the decision process over their learning, and merely replicates the standardized, mandatory curriculum of a print textbook.

5.3 - Case Studies

5.3.1 - Knewton

For this case study, I used the open Knewton Beta platform to go through the course creation process from the perspective of a teacher, and then used the course from the perspective of a student. When I began the course creation process, I selected course goals from a limited set of options, all of which were mathematics domains. If the teacher's course objectives happen to align with the goals delineated by Knewton's instructional designers, then they merely need to select the relevant goals. Otherwise, there is no option for teachers to create their own course goals or objectives, or to use the platform to create a course in a domain other than Mathematics, Algebra, Geometry, Trigonometry, and basic Probability.

Then, once the goals are selected, the course is "built," or, algorithmically generated from the available course materials. The student viewing the first lesson in, for example, a course on recursive sequences will find a lesson taken out of context from a publicly available textbook, in this case, a Creative Commons textbook on Advanced Algebra. Due to this lack of context, the first lesson that I entered began with the statement that "One interesting example is the Fibonacci sequence" without any indication of what this is an example of, or why it is interesting. After reading an example and attempted a practice problem, if the answer was correct, it confirms the correct answer, and if the answer was incorrect, the correct answer was provided without explanation, and the lesson proceeded to video lectures on the topic.

Students can choose to skip the videos and skip the associated questions that are given after the videos, but they will not attain any points towards their proficiency score, which measures how close they are to attaining mastery over the topic. Even after skipping, students still have the option to view and return to previously skipped videos and questions, viewable in a linear History column on the right side of the page. In fact, the videos that are subsequently suggested are the ones the student skips, indicating the system has information about which videos the student has at least allowed to elapse to completion, though they may not have watched or understood it. Also, it was clear that the more items I skipped, the easier the questions became, and the less they counted for this particular topic, as they became increasingly remedial. However, it may not be clear to a student that the questions are getting easier (a motivational concern) or that they are becoming less related to the new material, and perhaps more related to material the student has already mastered. This is a case where a visualization of student progress through the domain would certainly be of assistance.

After using the Knewton Beta for teachers and students, it would seem that the system fails to allow teachers control over the creation and sequencing of course content, other than at a very macro level of “course goals”. Moreover, the adaptations occur in such a black-boxed, indirect way that it is not clear to the student what, if any, adaptations are occurring. In addition, the lack of choice over the learning may be immensely frustrating and disempowering to students and teachers.

5.3.2 - CogBooks

After reaching out to several adaptive providers to speak with a sales representative for a walkthrough, CogBooks agreed to a meeting with an assistant vice president of sales, to discuss options for conducting research with the platform at the Center for 21st Century Universities, at which I am a research assistant. They began the conversation by explaining that CogBooks was a micro-adaptive, algorithm-based, learning sequencing system. That is to say, the system adapts on a content element level, rather than adapting on a whole course level. In addition, it uses their proprietary algorithm to generate learning sequences dynamically for each learner, rather than a more explicit rule-based system of pre-programmed branching learning paths, or an entirely preference based system where the user selects the content they want to learn.

Their system, as they explained, provides support for published textbook content, as well as OER (open educational resources), and, most importantly for our purposes, the ability for teachers to upload their own course content and materials. They made the point that if faculty already have existing video lecture content, perhaps from a MOOC or flipped class, then they would be able to import that content quite easily. They were quite clear that any faculty would be able to create their own courses with this tool, though they might require support from an instructional designer to clearly divide the course into learning objectives in an appropriate sequence. Once the course is divided into a set of learning objectives, and each of those has an associated learning activity (video, text, or interactive simulation) and at least one associated assessment, the teacher or instructional designer must decide what the default path through the domain is, based on those defined learning outcomes. All of the material is then assigned to either the default path, or a pre-requisite path, wherein supplemental material is provided as prerequisites to other content, suggested to students when they need assistance with a particular concept.

From a student perspective, the student first takes a pre-assessment, although the CogBooks representatives advised that this was not the most effective form of assessing prior knowledge - they recommended students be assessed formatively through interactions with the system and subsequent content assessments. Then, the given course module was adapted to our performance on the assessment, indicating that we had already mastered the first several content elements, and should begin on the third.

Students are able to move through a CogBooks course in one of two ways - either a “force-directed” path, or a “self-directed” path, depending on their level of confidence and mastery over the domain. The default option is for students to be self-directed, selecting course elements to view and be assessed on, but when they demonstrate that they need assistance on an element, with an incorrect response, supplementary material is suggested that might be relevant for them. However, this list is generated from teacher-tagged material, and as such has some flaws. First, it is only as good as the quality of the material itself, such as a YouTube video or TED talk on the content topic, which might be relevant, but not rigorous. Additionally, the system provides no actual recommendation or indication of which of those supplementary material might be more relevant, rigorous, or beneficial than the others, which the CogBooks representatives assured me was a feature they were considering, but did not exist yet.

With this in mind, CogBooks allows more control for teachers and students than a system like Knewton, from the course creation process that the teachers are able to use their own material, sequenced in their own way, to the students' learning process where they are able to use the pre-requisite content in ways that are beneficial to them, though it still has issues in the adaptation method.

5.4 - Recommendations

5.4.1 - Design Guidelines

The central issue in the design of humanist adaptive systems is whether they respect the desires, goals, and needs of individual users, and do not elide the individual differences between students in favor of a model that makes mandatory prescriptions about learning. One way to mitigate this issue would be for the designers of adaptive learning systems to engage in practices of participatory design (Muller & Kuhn, 1993). Rather than designing *for* teachers and students, a potential adaptive system designer intending to create a more humanist system would design *with* those teachers and students, or with representative samples of their target population. Understandably, there are barriers to effective and widespread use of participatory design as a scalable design methodology, most notably the demands on teacher time that may preclude their involvement in participatory design workshops. Moreover, it may be difficult for teachers and students to feel involved in the design of the system, without prior training in instructional design principles. However, despite these obstacles, the benefits for teacher and student buy-in in the development process would be enormous for designing systems that are more responsive to the individual needs of their target users.

Another method for a more humanist design would be to develop open source adaptive systems, or, failing a completely open source model, at least provide opportunities for user modification of the system. An open source design for adaptive systems would let teacher users modify the system based on their particular contexts and learning needs. As it stands, the current crop of adaptive systems are designed for a general student audience and decontextualized in such a way that they may potentially encounter resistance due to differences in culture within the class, school, or local community. That is to say, there may be communities in which self-regulation and student autonomy are more strongly valued than others, in which case, the local teacher users should have the ability to modify the adaptive system through an open source or modular approach, to make it more appropriate for their students.

Finally, even without a participatory design approach, or an open source, modular design of the system, adaptive learning system designers can, as some are already doing, make some very simple design choices to make their systems more adaptive. Beginning with the domain model, teachers should have the option to upload and arrange their own content, despite the potential issues in prerequisite tagging of the content. If possible, in the learner model, designers of adaptive systems should include an option for creating a learner model partly informed by student goals and preferences for their own learning. However, this is always at risk for being undermined by the course requirements of the school, district, or state policy. Finally, at the adaptation level, the adaptations should be direct, or, made visible to the students that an adaptation is occurring, and should be optional, not mandatory, with suggestions or recommendations as to what learning path or content should be accessed next, but without the mandatory delivery of the so-called “optimal” content.

5.4.2 - Guiding Questions for Selection

The taxonomies outlined earlier can be a framework which administrators, instructional designers, and teachers can use to select an adaptive system that best fits their values. Because there may not currently be an adaptive system that fully and completely allows for teacher autonomy and meaningfully supports student autonomy, people in a position to select an adaptive provider for their institution should consider the tradeoffs for autonomy explained above. In this section, I provide some selection criteria, in the form of guiding questions, for someone looking to make an informed choice of an adaptive system that supports humanistic educational practices.

i. Should the **teacher** be able to establish the learning objectives for the course?
ii. Should the **teacher** be able to upload their own pre-existing content for the course?
iii. Should the **teacher** be able to establish the relationships between content elements?

Some of those questions are policy-driven questions, and are dictated by decisions made above the level of individual teachers. If particular courses have mandated or standardized learning objectives, then it is quite possible the “teacher” would not be able to establish their own learning objectives, but the instructional designer hired by a school or school district might.

iv. Should the **system** be targeted to K-12, post-secondary, or corporate audiences?
v. Should the **system** be required to have LTI interoperability with your existing LMS?
vi. Should the **system** present a whole course of study or smaller lesson modules?
vii. Should the **system** be able to discover related content that it deems relevant?

These questions speak to the uses to which the system will be put, and may not influence the relative humanism of the system, but are important considerations for implementation nonetheless.

viii. Should the **students’** learning style be considered important to the adaptivity?
ix. Should the **students’** learning goals be considered important to the adaptivity?
x. Should the **students** be able to view the entire scope of the course?

Because all of the adaptive providers I surveyed included a knowledge state as part of their learner model, I did not include that as one of the guiding questions. The other two learner model components - learning style, and learning goals, however, depend on how willing a potential client institution is to have its learning experiences be dependent on student desires and goals. Because the research on learning styles is still inconclusive at best, schools should decide for themselves if they want to include that as a component in the learner model.

xi. Should the difficulty of the course material be **adapted** for students?
xii. Should the media form of the course material be **adapted** for the students?
xiii. Should the sequence of the course material be **adapted** for the students?
xiv. Should the **adaptations** be mandatory for students?

This set of questions speaks to how willing the school is to allow freedom in their courses and learning experiences of the students. If, for a given domain, the sequence of course material needs to be fixed for some policy-related, instructional design, or other non-pedagogical reason, the institution should choose a provider that does not allow for adaptive course sequences. The final question about the mandatory nature of the adaptations again is important for how willing the institutional client is for their students to have a degree of choice in their learning experiences.

Chapter 6 - Conclusion

6.1 - Future Directions

Though there has been a good deal of research into the use of adaptive learning systems in online courses and improving individual learning outcomes, there remain some open questions that still bear researching. The two main constructs assessed in current research on adaptive systems include improvements to learning outcomes and student engagement and retention in the courses. From an autonomy perspective, it would be useful to conduct research on how systems that afford different levels of **student autonomy** over the learning process differently affect the students' engagement and performance on learning outcomes. In addition, from a teacher perspective, it would interesting to research how varying levels of **teacher control** over the course creation process affects the rate of voluntary teacher adoption of the system.

From an adaptation perspective, more research needs to be conducted on the ways in which specific adaptation support mechanisms may affect student perceptions of agency, and thus affecting engagement and retention in the class, and potentially leading to improvements on learning outcome assessments. These adaptation support mechanisms might be such elements as a visualization of the domain space, which many adaptive systems do not include, or the inclusion of a set of available options for subsequent content elements, ranked or rated by their estimated relevance and appropriateness for that student. The visualization of the domain does not exist in many systems, or, typically, it is visible only to the teacher or the instructional designer creating the course, and it would be interesting to see how the presence of this tool affects the students' metacognitive planning about their progress through the course. The ranked suggestions for possible course elements is missing from many systems where, if they provide options for students to select from, such as CogBooks, do not provide any indicator of which of those options would be more relevant or useful for the student.

Finally, the set of research questions that would be most important to pursue are how these adaptive systems affect the nature of student collaboration, peer mentoring, and group projects in face-to-face classes, when students may not be working on the same course element at the same time. Because most of the research and use of adaptive systems occurs in either an individually-paced online course, or in a blended flipped class model, this issue has been circumvented, but there is the potential for an adaptive system to be used to supplement and improve student collaboration, if done effectively. Perhaps the system could be used to pair students together who have asymmetric mastery of various content elements, so that one could tutor the other, and vice versa, or could place students into groups based off heterogeneous mixtures of student mastery levels.

6.2 - Conclusion

The landscape of adaptive learning systems, both at a macro and a micro scale, has typically been full of promises for improvements to student learning, and short on rigorous, research-driven examinations of the methods of adaptation and their actual impact on learning the classroom. This thesis attempts to address the first half of those gaps in the literature, and presents a more organized analysis of the landscape of adaptive systems, so that more research can be conducted in the future, with a greater understanding of the plethora of available systems, and their differences. The major contribution towards that effort that this thesis makes is in presenting a clear and thorough analysis of the components of adaptive systems and their functions, and in categorizing a representative sample of the existing systems according to a taxonomy of the design of those components.

In the process of creating this taxonomy, and in conducting the research that led to it, I discovered a dissonance between the claims that the adaptive system providers were making in their rhetoric and white papers, and the actual design of their systems. While many adaptive providers, and their supporters, claim that these systems empower students, and free them from the traditional constraints of a cohort-based curricular model, in reality, the systems present their own sets of algorithmic constraints, no less constraining than those imposed by the teachers, but simply of a different nature.

Therefore, I undertook an analysis of that rhetoric, through the lens of the rhetorical tradition from which it seemed to poach its language - that of liberal humanist education philosophy. After looking at the language that adaptive providers use to describe the results of their systems, and analyzing them in light of the use of the same language of autonomy, agency, and control by prominent liberal humanist philosophers, I came to realize that the tradition in which they intended to situate their products was not in fact the theoretical tradition of which it was a part. The values enacted through the design of the systems themselves was more accurately situated within the cybernetic tradition, inspired by and modeled after self-regulating systems found in cybernetics (regulated via the algorithm, not by students' choice).

Finally, after arriving at the understanding that this cybernetic model for adaptive systems serves to limit student choice in ways that may not be clearly visible to the administrators, teachers, and students using them, I provided a set of guidelines by which to understand the relative humanism or cyberneticism of the systems. Unfortunately, because many of these systems have proprietary software, teachers and students may not be able to view them before the selection process occurs, and often, the administrators who view the software in action may be unequipped to address the complexity of these systems' functioning, though they may have the best interests of their teachers and students' autonomy at heart. My hope is that this paper helps to elucidate the complex functions of adaptive systems, and may provide a useful guide for administrators and teachers involved in the selection process to find an adaptive learning system that fits their values.

References

Akbulut, Y., & Cardak, C. S. (2012). Adaptive educational hypermedia accommodating learning styles: A content analysis of publications from 2000 to 2011. *Computers & Education*, *58*(2), 835-842.

American Institute for Research (2013). Are Personalized Learning Environments the Next Wave of K-12 Educational Reform? Accessed 12/5/14. http://www.air.org/sites/default/files/AIR_Personalized_Learning_Issue_Paper_2013.pdf

Aspen Institute Task Force on Learning and the Internet (2014). Learner at the Center of a Networked World, Washington, D.C.: The Aspen Institute.

Barab, S., Thomas, M., Dodge, T., Carteaux, R., & Tuzun, H. (2005). Making learning fun: Quest Atlantis, a game without guns. Educational Technology Research and Development, 53(1), 86-107.

Bargel, B. A., Schröck, J., Szentes, D., & Roller, W. (2012). Using Learning Maps for Visualization of Adaptive Learning Path Components. *International Journal of Computer Information Systems and Industrial Management Applications*, *4*(1), 228-235.

Brusilovsky, P. (2001). Adaptive hypermedia. *User modeling and user-adapted interaction*, *11*(1-2), 87-110.

Brusilovsky, P. (2004). Adaptive educational hypermedia: From generation to generation. In *Proceedings of 4th Hellenic Conference on Information and Communication Technologies in Education, Athens, Greece*(pp. 19-33).

Burgos, D., Tattersall, C., & Koper, R. (2006). Representing adaptive eLearning strategies in IMS Learning Design.

Chen, C. M., Lee, H. M., & Chen, Y. H. (2005). Personalized e-learning system using item response theory. *Computers & Education*, *44*(3), 237-255.

Chen, C. M. (2008). Intelligent web-based learning system with personalized learning path guidance. *Computers & Education*, *51*(2), 787-814.

Chung, H. S., & Kim, J. M. (2012). Ontology Design for Creating Adaptive Learning Path in e-Learning Environment. In *Proceedings of the International MultiConference of Engineers and Computer Scientists* (Vol. 1, pp. 585-588).

Collins, A., & Halverson, R. (2009). Rethinking education in the age of technology: The digital revolution and schooling in America. New York: Teachers College Press.

Cuban, L. (1986). *Teachers and machines: The classroom use of technology since 1920*. Teachers College Press.

Dewey, J. (1916). Democracy and education. Champaign, Ill.: Project Gutenberg.

Di Bitonto, P., Roselli, T., Rossano, V., & Sinatra, M. (2013). Adaptive E-Learning Environments: Research Dimensions and Technological Approaches.*International Journal of Distance Education Technologies (IJDET)*, *11*(3), 1-11.

Essalmi, F., Ayed, L. J. B., Jemni, M., & Graf, S. (2010). A fully personalization strategy of E-learning scenarios. *Computers in Human Behavior*, *26*(4), 581-591.

Executive Office of the President (2014). Big Data: Seizing Opportunities, Preserving Values. Accessed 9/22/14. https://www.whitehouse.gov/sites/default/files/docs/big_data_privacy_report_5.1.14_final_print.pdf

Facer, K. (2011). *Learning futures: Education, technology and social change*. Taylor & Francis.

Falmagne, J. C., & Doignon, J. P. (2011). *Learning spaces*. Berlin: Springer.

Franklin, S. & Graesser A. (1996) Is it an Agent, or Just a Program? A Taxonomy for Autonomous Agents Proceedings of the Workshop on Intelligent Agents III, Agent Theories, Architectures, and Languages Lecture Notes In Computer Science; Vol. 1193, pp:21—35.

Freire, P. (2000). Pedagogy of the oppressed (30th anniversary ed.). New York: Continuum.

Galloway, A. (2012). The interface effect. Cambridge, UK: Polity.

Graf, S., & Ives, C. (2010, July). A flexible mechanism for providing adaptivity based on learning styles in learning management systems. In *Advanced Learning Technologies (ICALT), 2010 IEEE 10th International Conference on*(pp. 30-34). IEEE.

Hagel III, J., Seely Brown, J., Mathew, R., Wooll, M., Tsu, W. (2014). The Lifetime Learner: A Journey through the Future of Higher Education. Deloitte University Press.

Hauger, D., & Köck, M. (2007). State of the Art of Adaptivity in E-Learning Platforms. In *LWA* (pp. 355-360).

Hayles, N. K. (2008). *How we became posthuman: Virtual bodies in cybernetics, literature, and informatics*. University of Chicago Press.

Helix Education (2013). Introducing a whole new way of teaching and learning. Accessed 12/18/14. http://www.helixeducation.com/resources/white-papers/

Hense, J., & Mandl, H. (2012). Learning" in" or" with" Games? Quality Criteria for Digital

Learning Games from the Perspectives of Learning, Emotion, and Motivation Theory. *International Association for Development of the Information Society*.

Hill, P. and M. Barber (2014) Preparing for a Renaissance in Assessment, London: Pearson.

Hsiao, A., Brouns, F., Van Bruggen, J., & Sloep, P. (2012). Effects of peer-tutor competences on learner cognitive load and learning performance during knowledge sharing.

Huang, M., Malicky, D., & Lord, S. (2006, October). Choosing an optimal pedagogy: A design approach. In *Frontiers in education conference, 36th annual* (pp. 1-6). IEEE.

Hutchins, E. (1995). Cognition in the wild. Cambridge, Mass.: MIT Press.

Illich, I. (1971). Deschooling society. New York: Harper & Row.

Karampiperis, P., & Sampson, D. (2005). Adaptive learning resources sequencing in educational hypermedia systems. *Educational Technology & Society*, *8*(4), 128-147.

Kirschner, P. A., & van Merriënboer, J. J. (2013). Do learners really know best? Urban legends in education. *Educational Psychologist*, *48*(3), 169-183.

Klašnja-Milićević, A., Vesin, B., Ivanović, M., & Budimac, Z. (2011). E-Learning personalization based on hybrid recommendation strategy and learning style identification. *Computers & Education*, *56*(3), 885-899.

Knauf, R., Sakurai, Y., Takada, K., & Tsuruta, S. (2010, July). Personalizing Learning Processes by Data Mining. In *Advanced Learning Technologies (ICALT), 2010 IEEE 10th International Conference on* (pp. 488-492). IEEE.

Koedinger, K. R., Anderson, J. R., Hadley, W. H., & Mark, M. A. (1997). Intelligent tutoring goes to school in the big city. *International Journal of Artificial Intelligence in Education (IJAIED)*, *8*, 30-43.

Lave, J., & Wenger, E. (1991). *Situated learning: Legitimate peripheral participation*. Cambridge, England: Cambridge University Press.

National Academy of Education. (2012). Adaptive Educational Technologies: Tools for Learning, and for Learning About Learning, G. Natriello (Ed.). Washington, DC.

.

Natriello, G. (2013). Adaptive Educational Technologies and Educational Research: Opportunities, Analyses, and Infrastructure Needs.

Newman, A. (2013). Learning to Adapt: A Case for Accelerating Adaptive Learning in Higher Education. Accessed 2/12/14. http://tytonpartners.com/library/understanding-the-adaptive-learning-supplier-landscape/

Newman, A. (2013b). 2013: The Year of Adaptive Learning. Accessed 8/31/14. http://www.impatientoptimists.org/Posts/2013/04/2013-The-Year-of-Adaptive-Learning

Nitchot, A., Gilbert, L., & Wills, G. B. (2010, July). Towards a Competence based System for Recommending Study Materials (CBSR). In Advanced Learning Technologies (ICALT), 2010 IEEE 10th International Conference on(pp. 629-631). IEEE.

Poellhuber, B., Chomienne, M., & Karsenti, T. (2008). The effect of peer collaboration and collaborative learning on self-efficacy and persistence in a learner-paced continuous intake model. International Journal of E-Learning & Distance Education, 22(3), 41-62

Radich, J. (2013). Technology and interactive media as tools in early childhood programs serving children from birth through age 8.

Singh, H. (2003). Building effective blended learning programs. *EDUCATIONAL TECHNOLOGY-SADDLE BROOK THEN ENGLEWOOD CLIFFS NJ-*, *43*(6), 51-54.

Szilagyi, I., & Roxin, I. (2012, July). Learner ontology for the Active Semantic Learning System. In *Advanced Learning Technologies (ICALT), 2012 IEEE 12th International Conference on* (pp. 393-394). IEEE.

Tomasello, M. (1999). The cultural origins of human cognition. Cambridge, Mass.: Harvard University Press.

U.S. Department of Education (2010). Office of Educational Technology, Transforming American Education: Learning Powered by Technology, Washington, D.C.

U.S. Department of Education (2011). Winning the Education Future: The Role of ARPA-ED. Accessed 12/10/14. http://www.ed.gov/sites/default/files/arpa-ed-background.pdf

U.S. Department of Education (2012). Office of Educational Technology, Enhancing Teaching and Learning Through Educational Data Mining and Learning Analytics: An Issue Brief, Washington, D.C.

U.S. Department of Education (2013). Office of Educational Technology, Expanding Evidence Approaches for Learning in a Digital World, Washington, D.C.

VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, *46*(4), 197-221.

Verbeek, P. P. (2011). *Moralizing technology: Understanding and designing the morality of things*. University of Chicago Press.

Vygotsky, L. (1978). Mind in society: The development of higher psychological processes. Cambridge: Harvard University Press.

Weibelzahl, S., De Bra, P., Paramythis, A., Ertmer, P., & Desjardins, F. (2008, June). Adaptive or Collaborative Learning?. In *World Conference on Educational Multimedia, Hypermedia and Telecommunications* (Vol. 2008, No. 1, pp. 5474-5477).

Wenger, E. (1998). *Communities of practice: Learning, meaning, and identity*. Cambridge, U.K.: Cambridge University Press.

Wiener, N. (1954). *The human use of human beings: Cybernetics and society*. Da Capo Press.

Winograd, T., & Flores, F. (1986). Understanding computers and cognition: A new foundation for design. Norwood, N.J.: Ablex Pub.

Zliobaite, I., Bifet, A., Gaber, M., Gabrys, B., Gama, J., Minku, L., & Musial, K. (2012). Next challenges for adaptive learning systems. *ACM SIGKDD Explorations Newsletter*, *14*(1), 48-55.