

Personalized Curiosity Engine (Pique): A Curiosity Inspiring Cognitive System for Student Directed Learning

Safat Siddiqui¹, Mary Lou Maher¹, Nadia Najjar¹, Maryam Mohseni¹ and Kazjon Grace²

¹University of North Carolina at Charlotte, NC, U.S.A.

²University of Sydney, Sydney, Australia

Keywords: Personalized Learning, Curiosity, Recommender Systems for Education, Computational Models of Novelty.

Abstract: Pique is an AI-based system for student directed learning that is inspired by a cognitive model of curiosity. Pique encourages self-directed learning by presenting a sequence of learning materials that are simultaneously novel and personalized to learners' interests. Pique is a web-based application that applies computational models of novelty to encourage curiosity and to inspire learners' intrinsic motivation to explore. We describe the architecture of the Pique system and its implementation in personalizing learning materials. In exploring the use of Pique by students in undergraduate and graduate courses in Computer Science, we have developed and implemented two computational models of novelty using Natural Language Processing techniques and concepts from recommender systems. In this paper, we describe the Pique model, the computational models for measuring novelty in text-based documents, and the computational models for generating sequences of personalized curiosity-eliciting learning materials. We report the response from students in the use of Pique in four courses over two semesters. The contribution of this paper is a unique approach for personalized learning that encourages curiosity.

1 INTRODUCTION

One of the most significant challenges in education at scale is how to personalize learning (Sampson and Karagiannidis, 2002). This problem is particularly prevalent when considering problem- (Wood, 2003), project- (Krajcik and Blumenfeld, 2006) and studio- (Carter and Hundhausen, 2011) based learning, in which content is open ended and students have autonomy in deciding how to focus their learning. How can each student be presented with knowledge and challenges that fit their interests and encourage curiosity? One approach is to devote instructor time to providing personalized learning materials and advice to each student, which is highly effective in small classrooms but does not scale. We present the use of AI methods — specifically Natural Language Processing— for personalized learning, based on eliciting curiosity using computational models of novelty. We contextualize our models of novelty in an interactive system for recommending course relevant publications to University students.

Curiosity can be defined as the desire to learn or know. Curiosity can be both a trait and a state (Berlyne, 1966). Curiosity-as-trait refers to an innate

desire possessed by different people to different degrees, while curiosity-as-state refers to a motivation to seek novel stimuli. This latter definition is the basis for encouraging curiosity in Pique. The curiosity state can arise from exposure to appropriately novel stimuli, with insufficiently novel stimuli being boring and overly novel stimuli being alienating, a model first proposed by early psychologist Wilhelm Wundt (Berlyne, 1966). This creates a region of optimal novelty within the space of possible stimuli, in which curiosity is maximally likely to be stimulated. The parameters of this region are dependent on experiences, context, and personal preference for novelty (Boyle, 1983; Kashdan and Fincham, 2004). The goal of Pique is to stimulate the curiosity state of the student and recommend resources that place them in a state of maximal curiosity by combining their interests and novelty in selecting course materials.

The Pique system is built on the principle theorized by Loewenstein (Loewenstein, 1994) that denotes curiosity as the result of an 'information gap' — the distance between what is known and what is desired to be known. This leads to a model of curiosity as a kind of intellectual hunger, in which small amounts of new knowledge prime further desire to

learn, but larger amounts have a satiating effect (Kang et al., 2009). A related notion from developmental psychology is Vygotsky's 'Zone of Proximal Development' (Vygotsky, 1978): the space of all knowledge that is adjacent to current knowledge and thus is comprehensible given it. These theories of knowledge acquisition all use spatial metaphors for describing curiosity's influence, which suggests approaches for how it might be operationalized. Computational models of curiosity based on these approaches have led to AI systems that value the unexplored and unexplained (Schmidhuber, 2010; Merrick and Maher, 2009; Grace and Maher, 2015b; Grace and Maher, 2015a). Pique contributes to this computational modeling of curiosity by proposing to stimulate the curiosity of each individual student. This is then combined with a representation of that student's preferences and knowledge to produce personalized recommendations.

Throughout our project, we have explored a variety of algorithms for both novelty and the generation of sequences of learning resources. In this paper we present two computational models of novelty and three ways to select learning resources to stimulate students' curiosity. Each of our models are based on the concept of unexpectedness as a cause of novelty and surprise, which consequently leads to curiosity (Grace et al., 2017; Grace et al., 2018). For instance, information that students find interesting but are not expecting creates a surprising response when presented as a stimulus to their learning process. When students are presented with information related to their existing knowledge but contain different perspectives, those learning materials seem novel to them. Recently novelty and surprise have been proposed as components of a new kind of recommender system that attempts to expand its users' preferences (Niu et al., 2018; Adamopoulos and Tuzhilin, 2014). In the Pique system, we use AI-based computational models to identify novel documents from a data set of learning resources and develop algorithms to generate a sequence of learning materials to encourage curiosity personalized to individuals' interests. This approach enables instructors to personalize the learning experience after identifying the corpus of learning materials in open ended project based learning.

The four major components of the Pique system are illustrated in Figure 1: Learning Materials, Artificial Intelligence Methods (AI), Learner Model, and User Experience (UX). These components are described in Section 3. Pique was applied in undergraduate and graduate courses in Human Centered Design (HCD) as well as a Graduate Teaching Seminar for PhD students. The Pique system was used over sev-

eral semesters, throughout which we continually developed the models of novelty and sequence generation based on student and instructor feedback. Platform data from the students' usage of Pique as well as their reflections on the recommended learning content were used to answer the following research questions:

- RQ1: How does the experience of using Pique enable self-directed exploration and personalized learning?
- RQ2: How does the experience of using Pique assist students in expanding their learning interests?

The remainder of this paper describes the related work and theoretical background (Section 2), the components of the Pique system (Section 3), the experiences of the students who used it in the classroom (Section 4), and directions for extending this research (Section 5).

2 BACKGROUND

Pique contributes to the broad field of AI in education. Baker and Smith (Baker et al., 2019) identify three perspectives on educational AI: learner-facing (which focus on assisting students), teacher-facing (which focus on reducing teachers' workloads), and system-facing (which focus on institutions' administrative and management capabilities). Zawacki-Richter et al. (Zawacki-Richter et al., 2019) similarly identify four areas of AI applications in higher education: adaptive systems and personalization, assessment and evaluation, profiling and prediction, and intelligent tutoring systems (ITSs). In this paper, we present Pique - a unique approach to adaptive systems and personalization in educational AI as it aims to include curiosity-inspiring content in the students' learning process.

Pique combines elements of the Educational Recommender Systems (i.e., recommendation of customized course materials) and ITS approaches to providing personalized learning activities. The difference between Pique and existing educational recommender systems is that we adopt a cognitively-inspired model of curiosity as the basis of our recommended sequences of resources, rather than a focus on explicit learning goals.

2.1 Intelligent Tutoring Systems and Personalized Learning Technologies

The broad goal of intelligent tutoring is to leverage AI techniques to provide the instructional capability that adapts to the needs of individual students. To support students in effective ways, models have been de-

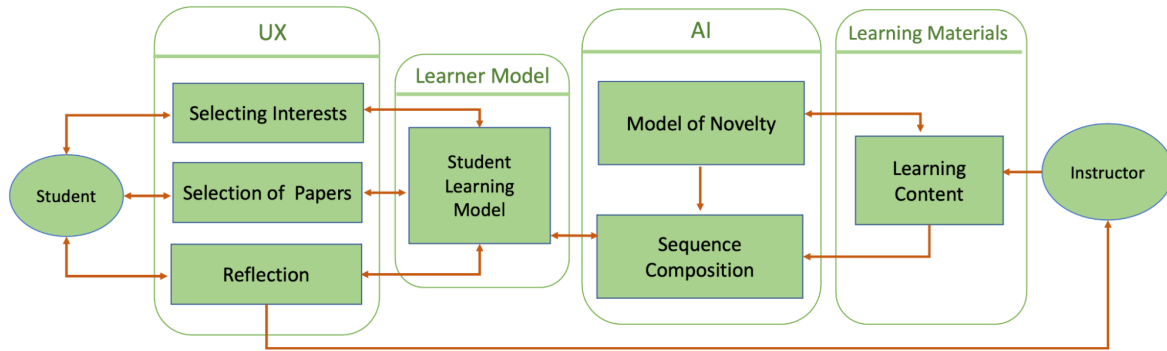


Figure 1: Architecture of the Pique Cognitive System.

veloped to distinguish students' search behaviors (An et al., 2020). Learning Management Systems (LMS) can monitor students' learning patterns and categorize students based on their behavioral patterns (Kuo et al., 2021; Papamitsiou and Economides, 2014). Ai et al. (Ai et al., 2019) have applied deep knowledge models to trace students' knowledge status and have adopted reinforcement learning to recommend exercises. Intelligent tutoring systems have been developed to provide personalized feedback for learning programming courses (Keuning et al., 2018) and to support students' mathematical problems solving process (Pozdniakov et al., 2021). While we do not present Pique as an ITS, it can be thought of as analogous to one: it models student preferences, represents the learning resources to be recommended, and has a strategy for composing resource sequences that focus on students' motivation and maximizing their learning curiosity.

Models of student motivation have previously been used to augment the instructional capabilities of ITSs. Del Soldato & Boulay (Del Solato and Du Boulay, 1995) describe an approach to planning communication with a student performing a series of learning tasks based on a model of the student's motivations. They base the motivational reasoning on Keller's (Keller, 1987) model of motivation as consisting of curiosity, challenge, confidence and control, previously used in computer-supported collaborative learning (Jones and Issroff, 2005). The Pique system focuses on the first of these motivational factors: curiosity, which is stimulated by surprise and novelty. We explore the computational models of novelty and recommend resources calculated to increase student familiarity with concepts, but additionally aim to elicit their surprise and curiosity.

2.2 Educational Recommender Systems (ERS)

The goal of ERSs is to recommend learning resources, and they can be applied in either formal or informal educational contexts, and used by either students or instructors. The Pique system is intended to be applied to open-ended learning tasks, which — while part of a formal educational context — share many of the features of informal learning. Open-ended learning tasks require students to select a scope of focus for their work within the proposed problem space (Hill and Land, 1998), which means that students are self-directed to a degree, an aspect shared with informal educational contexts.

Educational recommender systems have been used to recommend course material for students in Computer Science (Kose and Arslan, 2016) and Business and Administration studies (Hall Jr and Ko, 2008). Cobos et al. (Cobos et al., 2013) have developed a recommendation system for the instructors to prepare course content. Educational recommender systems include explainability to justify the recommendation. Barria-Pineda et al. (Barria-Pineda et al., 2019) have helped students minimize misunderstandings related to solving programming problems and have explained the reasoning behind the recommended learning activities. Barria Pineda and Brusilovsk (Barria Pineda and Brusilovsky, 2019) identify that students spend more time on the exploratory interface and suggest the effectiveness of the transparent recommendation process. In Pique, we show students the calculated novelty scores of the papers to facilitate students' paper selection process.

3 THE PIQUE SYSTEM

3.1 The Learning Materials Component

The Learning Materials component is the source of documents provided by the instructor for a specific course. For Pique, these documents are represented as unstructured text, drawn from articles from relevant conferences, journals, and digital libraries. The preparation of the learning materials for use in Pique includes data collection, dataset preparation, and dataset description to prepare the text documents for applying computational models of novelty.

Pique has been included in two courses in a Computer Science program. The first, ‘Human Centered Design’, has a focus on human computer interaction. The learning materials are papers published in the ACM Digital Library under the classification of Human Centered Computing. The second, ‘Graduate Teaching Seminar’, has a focus on educational research in computer science, and the relevant learning materials are papers published in the ACM SIGCSE (Special Interest Group on Computer Science Education) proceedings.

For the Human Centered Design Course we extracted 9,452 conference, journal and magazine papers with publication dates from 2008-2018. For each publication we collected the following metadata: title, ISSN, location, abstract, publisher, address, ACM ID, journal, URL, volume, issue date, DOI, number, month, year, pages, and tags/keywords. For the Graduate Teaching Seminar we extracted 1172 papers with publication dates from 2008-2018, with the following metadata: title, author, conference, year, DOI, keywords, and abstract.

3.2 The AI Component

The AI component of the Pique system applies NLP and novelty detection algorithms to rank the novelty of items in the learning content module as a basis for generating a personalized sequence of learning materials. The AI component has two subcomponents: the Model of Novelty and the Sequence Composition. The Model of Novelty module receives a set of documents from the Learning content module and generates novelty scores for each of the learning items in the corpus. This subcomponent uses a list of features to represent each learning item based on topic models or keywords associated with each item. These features provide the basis for computing a novelty score for each item. The Sequence Composition subcomponent uses the novelty score, information about the learner’s interests, and the information from their in-

teraction with Pique to generate a sequence of learning materials. These materials are personalized to encourage the learner’s curiosity and their intrinsic motivation to explore the learning resources.

3.2.1 Models of Novelty

In developing Pique, we implemented two computational models of novelty, each based on the probabilities of keywords and topic models associated with each document. We refer to these models as the ‘Keyword co-occurrence model’ and the ‘Topic co-occurrence model’. The keyword co-occurrence model represents each item in the learning materials as a bag of keywords. The topic co-occurrence model represents each item in the learning materials as a vector of topic distributions by applying a topic modeling algorithm (Blei and Lafferty, 2007) to the corpus of learning materials. This section describes how these computational approaches generate novelty scores for the corpus of learning materials provided by the instructors.

Keyword Co-occurrence Model. The keyword co-occurrence model is based on the probability of the co-occurrence of each pair of keywords in the corpus. The challenge in this model is to identify the keywords for the learning materials. In our corpus, each paper has two fields in the metadata that can serve as the keywords for this model: the keywords selected from the ACM’s Computing Classification System (CCS) and the author defined keywords. The ACM Computing Classification System has been developed as a poly-hierarchical ontology that results in common topics relevant to all papers, but they do not specifically represent the content in each paper. Conversely, author-defined keywords are defined specifically for each paper, but do not follow any standard representation.

To calculate the probability of keyword co-occurrence, we synthesized the list of keywords from each paper into a master list of keywords for the corpus. Due to the large number of keywords in the master list of keywords, we manually curated a reduced set that can be used to create a mapping from a user’s interests to the concepts in the learning materials. In reducing the number of keywords we tried to target the largest number of keywords that are reasonable to present to students for selection: too many keywords would be overwhelming, and not enough keywords would not represent the dataset with enough fidelity. In our pilot study of using Pique in the Human-Centered Design course, when using the CCS classification as keywords, we gave students 118 different

keywords to select from. This was perceived by the students as too many, and it resulted in a very sparse user preference vector. As a result, we reduced this set in subsequent semesters. We manually replaced keywords that were not in the reduced list to be the most relevant keyword in the reduced set. Across the semesters, feedback from students indicated that our reduced set of 35-55 keywords was sufficient for students to express their interests.

With the keywords for each paper, we created a bag-of-keywords representation to calculate the co-occurrence of keywords for measuring novelty. First we eliminated papers that had fewer than two keywords. We then measured the probability of each pair of keywords appearing together in the corpus. We used these probabilities (see eq. 1, eq. 2) to calculate the probability of keywords x_1 and x_2 occurring together in the corpus (eq. 3) and took its logarithm as the novelty score for that pair of keywords (as in (Niu et al., 2018), (Bouma, 2009)). We prepared a novelty matrix, NM (eq. 4) by applying this process to all pairs of keywords in the corpus. This matrix serves as the look-up table for identifying the novelty scores among the keyword pairs in the papers. To convert from this keyword-pair novelty score to the score for a paper, we took the highest value of all keyword pairs present in the paper (eq. 5) as surprising combinations stand out (Grace et al., 2017).

$$prob(x_1) = \frac{\# \text{ of papers have } x_1}{\# \text{ of total papers}} \quad (1)$$

$$prob(x_2) = \frac{\# \text{ of papers have } x_2}{\# \text{ of total papers}} \quad (2)$$

$$prob(x_1, x_2) = \frac{\# \text{ of papers have both } x_1 \text{ and } x_2}{\# \text{ of total papers}} \quad (3)$$

$$NM(x_1, x_2) = \frac{\log_2(prob(x_1, x_2))}{prob(x_1) * prob(x_2)} \quad (4)$$

$$NoveltyScore_{P_n} = \max(NM(x_1, x_2), NM(x_1, x_3), \dots) \quad (5)$$

Topic Co-occurrence Model. The Topic Co-occurrence Model calculates the novelty score based on frequency and proportion of topics present in the corpus. We applied the topic modeling approach on the abstract of the papers. Topic modeling produces a set of ‘topics’ each comprising a distribution over all the words in the corpus (Grace et al., 2017). The advantage to using topic modeling compared to author defined keywords is that there is consistency in the identification of features across the entire data set in

topic modeling, where author defined keywords provide features relevant to the author of a single item in the corpus.

We used the R package ‘STM’ (Structural Topic Model) (Roberts et al., 2019), a topic model extension that is equivalent to CTM (Correlated Topic Model) (Blei and Lafferty, 2007). Correlated topic models relax the assumption made by earlier topic modeling algorithms that all the topics in a corpus are independent and therefore no one pair of topics is more likely to occur together in a document than any other. The STM algorithm was run on the dataset to obtain a vector of topic proportions for each paper and a topic correlation matrix. Each paper in the corpus is represented with a 20-dimensional vector containing the prevalence of topics in that document. The correlation matrix is a 20×20 matrix including the correlation coefficient for all topic pairs.

We calculate the novelty of a document as equal to the most novel concept or combination of concepts within that material (Grace et al., 2017). The novelty of a document is the highest negative correlation coefficient among all pairs of topics present in that document (above a certain threshold), weighted by the proportion of the document which contains that pair (Grace et al., 2017). In order to determine whether a topic is significantly present in a document a topic proportion threshold of 0.1 is used (i.e. the document should be at least 10% of that topic). This novelty formula is based on previous work in topic-model approaches to novelty (Grace et al., 2017). Equation 6 shows the formula for a paper p given $p = [t_i, t_j, \dots, t_n]$ as the set of topics significantly present in p . The pair of topics with the highest negative correlation coefficient are denoted by t_i and t_j . This coefficient is normalized against the most novel pair of topics in the whole corpus (here denoted t_a and t_f) and then weighted by the proportions of each topic in p to calculate the novelty score.

$$NoveltyScore_{P_n} = \frac{CovMat_{(t_i, t_j)}}{CovMat_{(t_a, t_f)}} \times 2(\min(prop(d, t_i), prop(d, t_j))) \quad (6)$$

$CovMat$ is the covariance matrix obtained from the STM model. $CovMat_{(t_i, t_j)}$ is the correlation of the document’s most atypical topic combination (t_i, t_j) , and $CovMat_{(t_a, t_f)}$ is the correlation of the most atypical topic combination among the whole corpus $(t_a$ and $t_f)$. $prop(d, t)$ is the proportion of document d that consists of topic t . This corresponds to the novelty of the document’s most novel topic combination, relative to the corpus’s most novel combination, weighted by how much of the document consists of that combination. The reason for using the minimum of the

two topic proportions rather than their average is to prevent favoring documents that just passed the significance threshold with one topic, and were thus not particularly novel in combining it with another, much more weighted topic (Grace et al., 2017). In the Topics Co-occurrence Model, the novelty rating for documents containing many relatively novel topic combinations will be higher than for documents containing only a little of the most novel pair of topics (Grace et al., 2017). Equation 6 assigns a novelty score of 1 to a document that is made up of 50% of each t_i and t_j . The equation assigns 0 as the novelty score to the document when the rarest topic combination in it is independent, and assigns negative scores when the only topic combinations above the threshold in the document are positively correlated.

3.2.2 Sequence Composition

The purpose of the Sequence Composition subcomponent is to take the novelty ratings of each document in the corpus and construct a sequence of learning resources that will maximize the chance of a student experiencing optimal novelty. The Sequence Composition subcomponent generates a sequence of learning resources to support student-directed learning and to stimulate students' curiosity about learning. During the course of our project we explored three sequence generator models, which we call the 'Origin-Destination model', the 'Destination model', and the 'User-Directed model'.

Pique generates a personalized sequence of nine documents in sets of three papers from the corpus of learning resources in the Learning Materials component, based on information from the Learner Model. Students choose one paper from each set of three, read it, and reflect on it before they are presented with the next set. The different sequence generator models are based on different representations of student interests.

The Destination Model uses a single set of student-specified interests as the input to the algorithm, while the Origin-Destination Model uses two student-specified sets of keywords: one that they self-report as already knowing about (the 'origin' set) and one that they want to learn more about (the 'destination') set. The User-Directed Model builds on the Origin-Destination Model to include other keywords from the papers most recently selected by the student. The sequence generator use these keywords to represent student preference, and combine that with the novelty score for each paper to select and sequence learning resources with the goal of inspiring curiosity.

Destination Model. The Destination Model prioritizes what students desire to learn. It recommends a set of nine novel documents containing information related to their stated desires. When used with our keyword co-occurrence novelty model, the student interests can be directly mapped to corpus keywords. When using the topic co-occurrence model, a mapping was manually built between the keyword set we had constructed and the automatically generated topic model topics. Here we refer to 'novel documents' generally, without specifying which of the novelty models labeled them as such.

Students select their learning interests, which become the destination set, D . The destination model then identifies candidate documents from the learning materials corpus for which the top N topics within that document include at least one of the user's selections. After some experimentation we decided on $N=3$, as most documents in our corpus included at least this many topics at reasonable proportions. From this set of candidate documents the nine most novel papers are selected and sorted based on their novelty, with the most novel last. In this way the Destination Model returns nine documents as output containing information that students' want to learn, starting with a document of moderate novelty but then scaling up to highly-novel documents as the student reads more and learns about the topics they are interested in.

Origin-destination Model. The Origin-Destination model intends to inspire students to explore learning materials that contain some information that they already know, combined with some new information that they don't. This is based on ideas from educational psychology like Vygotsky's Zone of Proximal Development (Vygotsky, 1978), in which new material is only learnable if it is at least somewhat connected to topics already known. The model selects a sequence that interpolates from what the student already knows to what they want to know. This recommendation generator is inspired from the surprise walks algorithm (Grace et al., 2018) that similarly tries to interpolate from an unsurprising source to a surprising destination.

The Origin-Destination Model stimulates individuals' curiosity by presenting learning materials in three steps: 'close', 'far', and 'farther'. By recommending the learning materials step by step, the model assists students to learn new materials similar to what they already know and inspire them to explore without recommending materials that are so novel as to be unfamiliar or alien for them (Berlyne, 1966). In the first step, the model recommends papers that are similar to what students already know and labels

those papers as the ‘close’ category of learning materials for that student. In the second step, the model recommends papers that are similar to both what the students already know and what they want to learn, and labels those papers as the ‘far’ category. Finally, in the third step, the model recommends papers that contain materials related only to what students want to learn, and labels those papers as the ‘farther’ category.

In the ‘close’ category, the model identifies candidate papers that contain at least one common keyword (or topic) from the students’ ‘source’ interest set. The model uses the k-means algorithm and clusters the candidate papers based on their novelty scores to distinguish papers with three novelty levels: high, medium, and low. The model also calculates the paper’s familiarity score, which is the number of keywords in common between the paper and the ‘origin’ set of topics/keywords the student already knows. The papers with highest familiarity scores in each novelty level are selected. The algorithm recommends one low, one medium, and one high novelty paper.

In the ‘far’ category step, the model recommends another three papers intended to extend students’ learning from what they are familiar with to the new topics they desire to learn. The candidate papers of this category contain at least one common keyword from the ‘origin’ keywords set (O) and at least one common keyword from the destination keywords set (D). The model uses the same clustering approach to identify low, medium, and high novelty candidate papers, and identifies the candidate paper in each level with the highest number of common keywords.

In the ‘farther’ category step, the model presents papers that contain information that students desire to learn. The candidate papers of this category contain at least one keyword from the destination keywords set (D), and are categorized into three levels of novelty just like the other two sets.

User-directed Model The User-Directed model is an adaptation of the Origin-Destination Model that considers students’ decisions during the recommendation process to recommend materials aligned with their evolving interests. The model recommends papers step by step (close, far, and farther) as in the Origin-Destination model, but additionally keeps track of students’ selections of papers from the previous step. The keywords in the papers from the previous step are used to prioritize similar resources in the recommendations of the next step.

Specifically, the User-Directed model filters the candidate papers for the far step to those that share at least one keyword papers selected in the close step.

Likewise, the model first identifies candidate papers for the farther step that contain at least one keyword match with the keywords of the paper selected in the far step. Other than this additional filtering, the model is identical to the Origin-Destination model: it recommends one low, one medium, and one high novelty paper in each of the close, far, and farther steps.

3.3 The Learner Model Component

The Learner Model in Pique is primarily focussed on collecting information about the learner to support the selection and presentation of learning materials as well as information needed to analyze the use of Pique. The Learner model is not a comprehensive model of the learner. This component stores two kinds of information: information about the students and how they have used Pique to date. Most information about students remains constant: their name, ID, email address, and course. The IDs are automatically generated by the Pique system and serve to de-identify students as required by our IRB approval. The final component of the student profile is the one that can change: their interests, which they select when they start using Pique but are prompted to change each recommendation ‘cycle’. Each time the student uses Pique, their data is updated with a new cycle record, containing timestamps, the papers they selected, the options they chose from, and their reflections. Their reflections comprise responses to three questions: 1) why the student has selected the paper, 2) whether the selected paper matches their interests, and 3) what topics the student expects to learn from the paper when they read it.

3.4 The UX Component of Pique

The User Experience Component of Pique supports students’ interaction with the following three subcomponents: Selecting interests, Selection of papers, and Reflection.

3.4.1 Selecting Interests

The Selecting Interests subcomponent captures students’ interest by prompting them to identify what they want to know. This prompt assists students to formulate their learning goals and provide them more control over their learning choices and enables self-directed learning. Figure 2 shows the user interface with the learning options for the students as they were in the Graduate Teaching Seminar course.

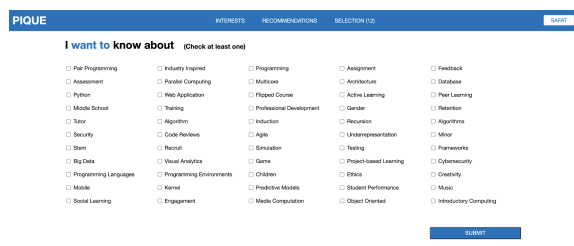


Figure 2: UX for selecting interests in Pique.

3.4.2 Selection of Papers

The Selection of Papers subcomponent of Pique enables students' self-regulated learning, with the intention of stimulating their intrinsic motivation to learn and explore. This module presents the papers that are recommended by the sequence composition module of Pique's AI component. Pique presents the nine papers in sets of three, based on the Sequence Composition subcomponent (see Section 3.2.2). Figure 3 shows an example of papers being recommended in the Graduate Teaching Seminar course based on the Origin-Destination sequence composition model: the top three are closely related to what the student already knows, the middle three are related to both what they know and what they are interested in, and the bottom three what they are interested in only. The Selection of Papers subcomponent informs students about how novel a particular paper is, and allows them to manually choose more or less novel papers by selecting the drop-down menu labeled 'show me papers' in the top right corner of Figure 3.



Figure 3: The UX for selecting learning content based on interests and novelty scores.

3.4.3 Reflection

The third subcomponent of the Pique UX component is Reflection. Cognitive studies of students demonstrate that reflection is key to effective learning [(Schön, 2017), (Kolb, 1999), (Cowan, 2006)]. The Pique system includes two types of reflection: one appears when students select a paper to read (see Figure 4), and one at the end of the semester. The first allows them to reflect on their paper selection and to describe what they expect to learn from it.

The second type of reflection asks students to reflect on their overall learning experience. Students are asked to summarize the papers they read and categorize those papers into groups. Students are asked to identify the paper they found most interesting and justify why. This reflection allows students to organize their newly acquired knowledge where the learning paths are constructed by the students rather than the instructors. It was also critical for evaluating the impact of this educational innovation on the student experience.

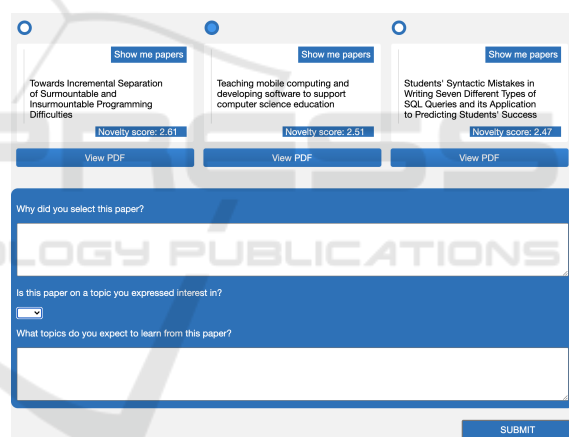


Figure 4: Pique nudges students to reflect on their paper selection and learning expectations.

4 THE COURSE EXPERIENCE OF USING PIQUE

In this section, we present the experiences of students who have used Pique in the classroom to respond to the Research Questions presented in Section 1. We used Pique over four semesters in both undergraduate and graduate courses in Human Centered Design as well as a PhD course, a Graduate Teaching Seminar. In the Human Centered Design course students were asked to use Pique for six weeks, and had to submit weekly and end of semester writing assignments about the papers they had read. Each week they were

asked to submit a summary of the three papers they downloaded and read, and identify the most interesting paper among the three. For the end of semester report, the students were asked to describe their experience of using Pique, what they learnt, and the most interesting paper they found (and why). For the Graduate Teaching Seminar course, students were asked to use the Pique system for the whole semester, but submitted only a final report without any weekly submissions. This was due to the PhD students' greater familiarity with reading published articles, as well as their overall greater autonomy as learners.

In response to our first research question concerning how the use of Pique helped enable self-directed exploration we investigated how the student cohort differed in the resources they explored, as a measure of how self-directed their experiences were. In Table 1 we summarize our results. Though students' options for selecting interests remain constant (39 and 55 interests in Human Centered Design and Teaching Seminar courses, respectively), we found that students were presented with very diverse sequences of learning resources. Pique recommended a total of 621 unique papers for one semester of the Graduate Teaching Seminar course, even though that course only included five students. 55% of those papers were recommended to at least two students, due to overlaps in topics of interest. Those five students selected a total of 66 papers to read, with 86% of the selected papers being selected by only one student. Across all four courses we saw 72% of recommended papers being recommended to at least another student, but the selections made by students were highly diverse, with 70% of the selected papers being unique to that individual student.

Our second research question asked how using Pique assisted students in expanding their learning interests. In response to this we investigated the change in students' interests over time, illustrated in Figure 5. The top two sub-figures are for HCD courses (Spring and Fall) and the bottom two sub-figures are for Graduate Teaching Seminar courses (Spring and Fall), where X-axis represent the number of Pique cycles and Y-axis represents the average cumulative growth of interest selections. We calculated the number of interests selected by each student during each cycle. We aggregated this across students within a cohort to give the average number of interests selected by the students in that cycle. The cumulative number of interests in Figure 5 shows the expansion of stated interests over the semester. When starting to use Pique, total students in the HCD courses selected an average of just four interests. The searching of learning interests increased as students used the Pique

system, and at the end of the semester all students had explored an average of 67 interests. Similarly, total students in the Graduate Teaching Seminar on average started with only two interests, and over the semester their average number of searching learning interests increased to 42.

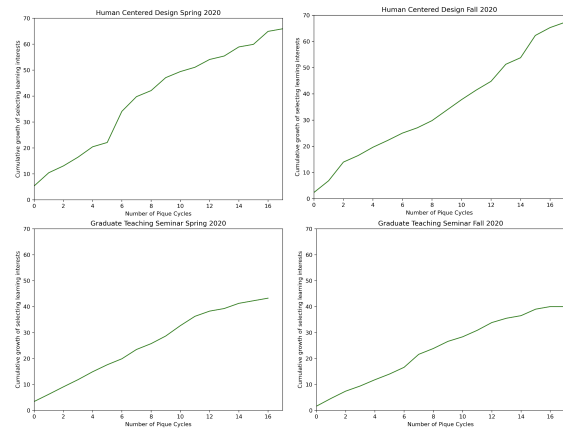


Figure 5: The growth of learning interests over the courses.

We noticed a difference between the students in the HCD courses and those in the Graduate Teaching Seminar, illustrated in Figure 6. The top two sub-figures are for HCD courses (Spring and Fall) and the bottom two sub-figures are for Graduate Teaching Seminar courses (Spring and Fall), where X-axis represent the number of Pique cycles and Y-axis represents the percentage of students searched for new interests that they had not selected in earlier Pique cycles. The students in the HCD courses were undergraduate and graduate students who initially expanded their learning interests and over time they reduced the number of new interests. In contrast, the students in the Graduate Teaching Seminar courses were PhD students who kept exploring new interests. For instance, all the PhD students in the Fall semester of Graduate Teaching Seminar continued to add new interests until the end of the semester. 71% of Graduate Teaching Seminar PhD students in the Spring semester included new interest in their 8th Pique cycle, but only 18% of the undergraduate students in the HCI course continued exploring in the 8th Pique cycle. This result indicates that students use the Pique system differently to expend their learning selections.

After the students had finished using Pique, students were asked to reflect on which paper from the system they found most interesting and why. Two researchers performed a thematic analysis on students' written responses to identify meaningful patterns in the data (Braun and Clarke, 2006). The use of multiple coders provided investigator triangulation to our analysis (Patton, 1999). To establish a broad consen-

Table 1: Distribution of learning materials to personalize learning.

Course name	Graduate Teaching Seminar	Graduate Teaching Seminar	Human Centered Design	Human Centered Design
Semester	Spring 2020	Fall 2020	Spring 2020	Fall 2020
Number of students	24	5	12	12
Number of unique learning sequences generated by students	24	5	12	12
Total papers selected by students over the Pique cycles	221	66	76	77
% of selected papers uniquely picked by individuals	50% (111 papers)	86% (57 papers)	71% (54 papers)	74% (57 papers)
Total papers recommended by Pique	1987	612	729	774
% of papers recommended to at least one other	84% (1669 papers)	54% (333 papers)	75% (548 papers)	72% (558 papers)
% of papers recommended to only one student	16% (318 papers)	46% (279 papers)	25% (181 papers)	28% (216 papers)

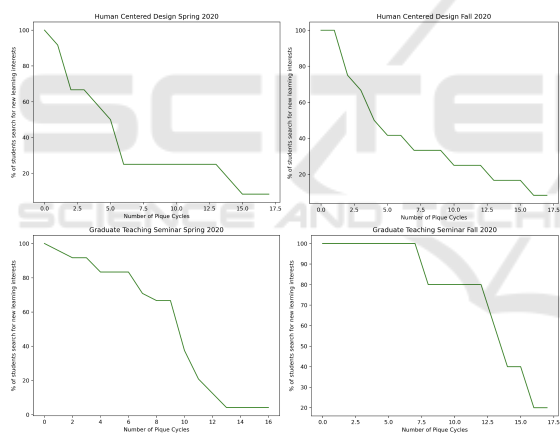


Figure 6: Percentage of students searching for new learning interests over the pique cycles.

sus, the two researchers conducted a parallel coding workshop on the first 10% of the written responses. After identifying this initial set of themes, they each coded the rest of the data separately and then converged on a set of collaboratively authored themes through follow-up workshops.

Through this analysis we identified three major themes underlying why most students found papers interesting: novelty, personal relevance, and curiosity. The first theme captured how students found papers interesting because of the novelty and innovation of the idea presented in it. The second theme captured how students found papers particularly interesting when they could connect its contributions or

implications to their current work or personal experience. For example, one student found a paper describing a VR gaming application named ‘Spider Hero’ interesting because he was a huge Spiderman fan. The third theme captured how the recommended papers made students curious about the research field as a whole (HCI or computing education) and helped to grow their interest in the field. For example, one student expressed that they learned something new from each of the recommended papers, and became so curious that they did extensive further personal research to learn more about those specific topics. We observed that this curiosity theme was related to the idea of students connecting their class lessons with the recommended papers. For example, one student learned the concept of a ‘Wizard-of-Oz’ study through the HCI class lectures and got excited when he found the same concept in a research paper. Taken overall, these written responses show that the recommended papers motivated students to explore and learn more in the domain.

5 CONCLUSION AND FUTURE WORK

We present a cognitively inspired system architecture, Pique, that presents students with personalized sequences of novel learning resources. We have shown that these sequences encourage curiosity and may be

helpful to support self-directed learning. Pique uses computational models of novelty to identify documents from a corpus of learning materials that are both relevant to the student's interest and novel with respect to the corpus. Pique is effectively an educational recommender system with a goal of inspiring individuals' curiosity to learn rather than shepherding them through a specific curriculum.

Through our four-semester evidence-based exploration of how to inspire students' curiosity, we developed two separate computational models of novelty: one based on keyword co-occurrence and one based on the co-occurrence of topics from a topic modeling algorithm. These computational models rely on the same underlying information theoretic approach to novelty or surprise as features that negatively correlate, but differ in the way we generated a keyword or topic representation of the documents. Both had their strengths and their weaknesses, and each was able to identify some surprising-seeming papers that the other missed. In future work we aim to explore a variety of other approaches to representing our corpus, including NLP and machine learning techniques. These new representations would extend our current computational models of learning resource novelty.

Throughout the Pique project we also developed three recommendation models: one based only on the student's stated interests (their 'destination'), one based on taking them on a journey from what they already knew (their 'origin') to their interests (their 'destination' again), and one based on blending that origin-destination effect with similarity to the things they've recently explored. Each of these recommendation models combined student preferences with our computational models of novelty to encourage curiosity in the learning process. We did not compare our sequence recommendation models directly, although we do believe, from the evidence of using them in the classroom, that the latter two models both offer advantages over their predecessors.

This paper presents a proof of concept and deployment of Pique, a personalized curiosity engine for education. A limitation of our study is that we did not collect data on how much time the students spent on reading the papers before reflecting on them. From evaluating the experiences of students who used Pique extensively as part of their courses, we identified three aspects that made recommended learning materials interesting: how novel they were, how personally relevant they were, and the curiosity and further self-directed learning that they evoked. These findings are evidence of how curiosity can be elicited from students as part of a course experience, at least when self-directed and open-ended engagement with learn-

ing resources is desirable. While we cannot claim that student curiosity was entirely due to Pique, we conclude that the approach of encouraging curiosity Pique shows promise for future research on computational novelty in open-ended learning environments.

REFERENCES

- Adamopoulos, P. and Tuzhilin, A. (2014). On unexpectedness in recommender systems: Or how to better expect the unexpected. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 5(4):1–32.
- Ai, F., Chen, Y., Guo, Y., Zhao, Y., Wang, Z., Fu, G., and Wang, G. (2019). Concept-aware deep knowledge tracing and exercise recommendation in an on-line learning system. *International Educational Data Mining Society*.
- An, S., Bates, R., Hammock, J., Rugaber, S., Weigel, E., and Goel, A. (2020). Scientific modeling using large scale knowledge. In *International Conference on Artificial Intelligence in Education*, pages 20–24. Springer.
- Baker, T., Smith, L., and Anissa, N. (2019). Educ-ai-tion rebooted? exploring the future of artificial intelligence in schools and colleges. Retrieved May, 12:2020.
- Barria-Pineda, J., Akhuseyinoglu, K., and Brusilovsky, P. (2019). Explaining need-based educational recommendations using interactive open learner models. In *Adjunct Publication of the 27th Conference on User Modeling, Adaptation and Personalization*, pages 273–277.
- Barria Pineda, J. and Brusilovsky, P. (2019). Making educational recommendations transparent through a fine-grained open learner model. In *Proceedings of Workshop on Intelligent User Interfaces for Algorithmic Transparency in Emerging Technologies at the 24th ACM Conference on Intelligent User Interfaces, IUI 2019, Los Angeles, USA, March 20, 2019*, volume 2327.
- Berlyne, D. E. (1966). Curiosity and exploration. *Science*, 153(3731):25–33.
- Blei, D. M. and Lafferty, J. D. (2007). A correlated topic model of science. *The annals of applied statistics*, 1(1):17–35.
- Bouma, G. (2009). Normalized (pointwise) mutual information in collocation extraction. *Proceedings of GSCL*, 30:31–40.
- Boyle, G. J. (1983). Critical review of state-trait curiosity test development. *Motivation and Emotion*, 7(4):377–397.
- Braun, V. and Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2):77–101.
- Carter, A. S. and Hundhausen, C. D. (2011). A review of studio-based learning in computer science. *Journal of Computing Sciences in Colleges*, 27(1):105–111.
- Cobos, C., Rodriguez, O., Rivera, J., Betancourt, J., Mendoza, M., León, E., and Herrera-Viedma, E. (2013).

- A hybrid system of pedagogical pattern recommendations based on singular value decomposition and variable data attributes. *Information Processing & Management*, 49(3):607–625.
- Cowan, J. (2006). *On becoming an innovative university teacher: Reflection in action: Reflection in action*. McGraw-Hill Education (UK).
- Del Solato, T. and Du Boulay, B. (1995). Implementation of motivational tactics in tutoring systems. *Journal of Artificial Intelligence in Education*, 6:337–378.
- Grace, K. and Maher, M. L. (2015a). Specific curiosity as a cause and consequence of transformational creativity. In *ICCC*, pages 260–267.
- Grace, K. and Maher, M. L. (2015b). Surprise and reformulation as meta-cognitive processes in creative design. In *Proceedings of the third annual conference on advances in cognitive systems ACS*, page 8.
- Grace, K., Maher, M. L., Davis, N., and Eltayeb, O. (2018). Surprise walks: Encouraging users towards novel concepts with sequential suggestions. In *Proceedings of the 9th International Conference on Computational Creativity (ICCC 2018)*. Association of Computational Creativity.
- Grace, K., Maher, M. L., Mohseni, M., and PEREZ Y PEREZ, R. (2017). Encouraging p-creative behaviour with computational curiosity. *ICCC*.
- Hall Jr, O. P. and Ko, K. (2008). Customized content delivery for graduate management education: Application to business statistics. *Journal of Statistics Education*, 16(3).
- Hill, J. R. and Land, S. M. (1998). Open-ended learning environments: A theoretical framework and model for design.
- Jones, A. and Issroff, K. (2005). Learning technologies: Affective and social issues in computer-supported collaborative learning. *Computers & Education*, 44(4):395–408.
- Kang, M. J., Hsu, M., Krajbich, I. M., Loewenstein, G., McClure, S. M., Wang, J. T.-y., and Camerer, C. F. (2009). The wick in the candle of learning: Epistemic curiosity activates reward circuitry and enhances memory. *Psychological science*, 20(8):963–973.
- Kashdan, T. B. and Fincham, F. D. (2004). Facilitating curiosity: A social and self-regulatory perspective for scientifically based interventions.
- Keller, J. M. (1987). Strategies for stimulating the motivation to learn. *Performance and instruction*, 26(8):1–7.
- Keuning, H., Jeuring, J., and Heeren, B. (2018). A systematic literature review of automated feedback generation for programming exercises. *ACM Transactions on Computing Education (TOCE)*, 19(1):1–43.
- Kolb, D. A. (1999). *Learning style inventory*. McBer and Company Boston, MA.
- Kose, U. and Arslan, A. (2016). Intelligent e-learning system for improving students’ academic achievements in computer programming courses. *The International journal of engineering education*, 32(1):185–198.
- Krajcik, J. and Blumenfeld, P. C. (2006). The cambridge handbook of the learning sciences. r. keith sawyer.
- Kuo, R., Krahn, T., and Chang, M. (2021). Behaviour analytics-a moodle plug-in to visualize students’ learning patterns. In *International Conference on Intelligent Tutoring Systems*, pages 232–238. Springer.
- Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological bulletin*, 116(1):75.
- Merrick, K. E. and Maher, M. L. (2009). *Motivated reinforcement learning: curious characters for multiuser games*. Springer Science & Business Media, .
- Niu, X., Abbas, F., Maher, M. L., and Grace, K. (2018). Surprise me if you can: Serendipity in health information. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–12.
- Papamitsiou, Z. and Economides, A. A. (2014). Learning analytics and educational data mining in practice: A systematic literature review of empirical evidence. *Journal of Educational Technology & Society*, 17(4):49–64.
- Patton, M. (1999). Enhancing the quality and credibility of qualitative analysis. *Health services research*, 34:1189–1209.
- Pozdniakov, S., Posov, I., and Anton, C. (2021). Interaction of human cognitive mechanisms and “computational intelligence” in systems that support teaching mathematics. In *International Conference on Intelligent Tutoring Systems*, pages 259–266. Springer.
- Roberts, M. E., Stewart, B. M., and Tingley, D. (2019). Stm: An r package for structural topic models. *Journal of Statistical Software*, 91(1):1–40.
- Sampson, D. and Karagiannidis, C. (2002). Personalised learning: educational, technological and standardisation perspective. *Digital Education Review*, (4):24–39.
- Schmidhuber, J. (2010). Formal theory of creativity, fun, and intrinsic motivation (1990–2010). *IEEE Transactions on Autonomous Mental Development*, 2(3):230–247.
- Schön, D. A. (2017). *The reflective practitioner: How professionals think in action*. Routledge.
- Vygotsky, L. (1978). Interaction between learning and development. *Readings on the development of children*, 23(3):34–41.
- Wood, D. F. (2003). Problem based learning. *Bmj*, 326(7384):328–330.
- Zawacki-Richter, O., Marín, V. I., Bond, M., and Gouverneur, F. (2019). Systematic review of research on artificial intelligence applications in higher education—where are the educators? *International Journal of Educational Technology in Higher Education*, 16(1):1–27.