

Personalized Education in the AI Era: What to Expect Next?

Setareh Maghsudi, Andrew Lan, Jie Xu, and Mihaela van der Schaar

Abstract—The objective of personalized learning is to design an effective knowledge acquisition track that matches the learner’s strengths and bypasses her weaknesses to ultimately meet her desired goal. This concept emerged several years ago and is being adopted by a rapidly-growing number of educational institutions around the globe. In recent years, the boost of artificial intelligence (AI) and machine learning (ML), together with the advances in big data analysis, has unfolded novel perspectives to enhance personalized education in numerous dimensions. By taking advantage of AI/ML methods, the educational platform precisely acquires the student’s characteristics. This is done, in part, by observing the past experiences as well as analyzing the available big data through exploring the learners’ features and similarities. It can, for example, recommend the most appropriate content among numerous accessible ones, advise a well-designed long-term curriculum, connect appropriate learners by suggestion, accurate performance evaluation, and the like. Still, several aspects of AI-based personalized education remain unexplored. These include, among others, compensating for the adverse effects of the absence of peers, creating and maintaining motivations for learning, increasing the diversity, removing the biases induced by the data and algorithms, and the like. In this paper, while providing a brief review of state-of-the-art research, we investigate the challenges of AI/ML-based personalized education, and discuss potential solutions.

Keywords: Artificial intelligence, Learning platform, Machine learning, Personalized education

I. INTRODUCTION

The last decade has witnessed an explosion in the number of web-based learning systems due to the increasing demand in higher-level education, the limited number of teaching personnel, advances in information technology and artificial intelligence, and, more recently, COVID-19. In the past few years, to enhance the conventional classrooms, to bridge the constraints of time and distances, and to improve fairness by making high-quality education accessible, most universities have integrated Massive Open Online Course (MOOC) platforms such as the edX consortium in their education systems. Also, several schools have added online labs to their structures, where students, especially those who cannot access physical labs, can

S. Maghsudi is with the Department of Computer Science, University of Tübingen, Germany (email: setareh.maghsudi@uni-tuebingen.de). A. Lan is with the College of Information and Computer Sciences, University of Massachusetts Amherst, MA, USA (email: andrewlan@cs.umass.edu). J. Xu is with the Department of Electrical and Computer Engineering, University of Miami, FL, USA (email: jiexu@miami.edu). M. van der Schaar is with the Faculty of Mathematics, University of Cambridge, UK (email: mv472@damtp.cam.ac.uk).

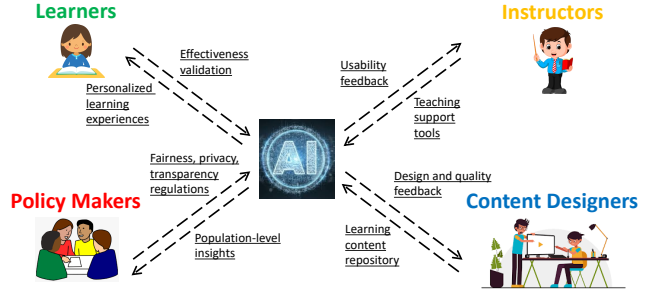


Fig. 1. The baseline ecosystem of AI-empowered personalized education.

perform experiments. Besides, there has been significant growth in the development of other online educational tools that simplify learning. These include, for example, the software for text summarization in different domains, also to produce questions and tests, followed by evaluation, which can be great assistance not only to students but also to teachers. Several advantages of these systems over traditional classroom teaching are: (i) They provide flexibility to the student in choosing what to learn and when to learn; (ii) They do not require the presence of an interactive human teacher; (iii) Often, the capacity in terms of the number of participants is significantly larger than the synchronous presence teaching form. **Figure 1** shows the baseline ecosystem of online personalized education, including all the stakeholders, together with the crucial factors and performance metrics.

However, the currently available online teaching platforms have significant limitations. To a large extent, personalized education has been mainly diminished to a specific type of ‘recommender system’, although its potential goes far beyond advising a series of lectures on an online platform that might be interesting to a specific user. One fundamental difference between existing recommender systems and personalized education is the optimization objective: The former focuses on some form of user engagement to maximize profit, which is system-centric and relatively easy to quantify, whereas the latter focuses on some form of learning outcomes, which is student-centric and hard to define.

ML/AI-enabled education is a response of great potential to the current shortcomings. It creates a new and more flexible learning technology genre that adapts to student learning and allocates resources as obliged. It takes advantage of the strengths of both online tools and individual tutoring. As such, AI-enabled personalized edu-

cation promises to yield many of the benefits of one-on-one instruction at a per-student cost similar to large university lecture classes. The system applies to both online courses and courses that have a hybrid of classroom and online instruction. As shown in **Figure 2**, ML/AI-enabled education comprises of a large set of decision-making strategies that collectively map the available data together with the individual features to a variety of personalized educational materials and recommendations.

Data can be collected on performance in both traditional assignments (problem sets, computer programs, laboratory) as well as online exercises and tests. It includes built-in assessment tools as an essential part of its optimization of lesson sequences. As such, it supports the educational community in developing new teaching modalities in a broad range of disciplines. However, despite intensive research efforts in this decade, a variety of aspects of personalized education remain unexplored, including both dark- and bright sides. In this paper, we discuss six core topics, review existing work, outline their limitations, and propose future research directions; see **Figure 3** for an overview.

When discussing any form of education, *quality* is an inevitable keyword. The quality of education depends largely on the quality of the available learning content and the quality of the personalized recommendations that guide each learner to the most suitable learning content. So far, the researchers have studied the production of learning content, from developing AI-driven smart learning content such as intelligent, interactive textbooks and game-based learning platforms to automatically generating learning content from the wild. Reference [1], for example, develops a sentence deletion method for text simplification. Besides, in [2], the authors investigate the effectiveness of discourse in multimedia to extract the knowledge from the textbooks. Moreover, a large body of the existing work investigates the recommendation of both macro-level and micro-level learning content, including courses in learners' degree plans as well as specific remedial content such as lecture notes, videos, and practice problems. For example, in [3], the authors take advantage of a multi-armed bandit framework to optimize the selection of learning resources and questions to satisfy the needs of each individual student. Moreover, another paper develops an e-learning recommender system framework based on two concepts: peer learning and social learning, which encourage students to cooperate and learn jointly. Despite great efforts, there remain several challenges to address. These include content recommendation at heterogeneous levels, the recommendation of a bundle of connected contents followed by performance evaluation, and the Pareto-optimization of conflicting objectives in the content recommendation. We discuss these progress and the future steps in **Section II**.

Historically, education is tightly coupled with evaluation. In personalized education, assessment and evaluation concern both the learner's performance and the effectiveness of the intelligent learning platform. Early approaches

for learner's assessment such as the 'classical testing theory' (CTT) use summaries of graded standardized tests. Recent approaches include 'item response theory' (IRT) models that enable the estimation of latent knowledge mastery levels and knowledge tracing models that pursue the evolution of a learner's knowledge. In [4], the authors compare the CTT and IRT. Methods such as 'computerized adaptive testing' improve the efficiency of assessments. The current approaches to evaluate the learning platforms use rigorous experiments, often large-scale randomized controlled trials. In this area, open problems include the prediction of learners' future performance, which enables providing better recommendations and more accurate feedback. This is referred to as the knowledge tracing (KT) problem, for which several methods are developed in the past few decades. As an example among many others, some papers discuss a Bayesian framework for KT. Another challenge is to reduce the information loss while grading the arrived input from the learner, by accurate interpretation of knowledge level based on the test design. We elaborate and address such challenges in **Section III**.

The huge advances in science, technology, and health-care have changed the working life of humans. Individuals have way more alternatives to choose a job, they tend to change their job more frequently than before, are more open to mobility, and the career spans a long period of life. As such, continuing education, which aims at advancing one's educational process, as well as lifelong learning, i.e., pursuing additional professional qualifications, are important components of educational policy in the world. Implementing these two concepts successfully has a significant impact on social welfare by developing new skills that enhance personal- and professional life. During the past decade, AI-/ML-based personalized education has been under intensive investigation from several perspectives; nonetheless, the aforementioned aspects are largely neglected. Indeed, personalized education shall accompany the learner throughout her life, which can be difficult and costly to implement. Other challenges include lack of appropriate data, potentially long delay to feedback, high diversity, and fast dynamics in the environment. For example, in [5], the authors design a new genre of educational technology-personal computer systems- that support learning from any location throughout a lifetime. Another research direction is to enable learning system to learn continuously. Reference [6], e.g., investigates the ability of neural networks to enable life-long learning. We elaborate more on this topic in **Section IV**.

Similar to any other task, humans require motivation for learning. Generally, incentives for learning can be defined as an inducement or supplemental reward that serves as a motivational device for intended learning [7]. Presumably, the most conventional models of incentive are the 'grade' and 'certificate', which are implemented as a part of learning platforms to motivate the students. The strength of such motivation depends on the validity and acceptance of such certificates by different authorities such as employers. Nonetheless, employing AI methods enables for incentive




Fig. 2. The basic concept of AI-powered personalized education.

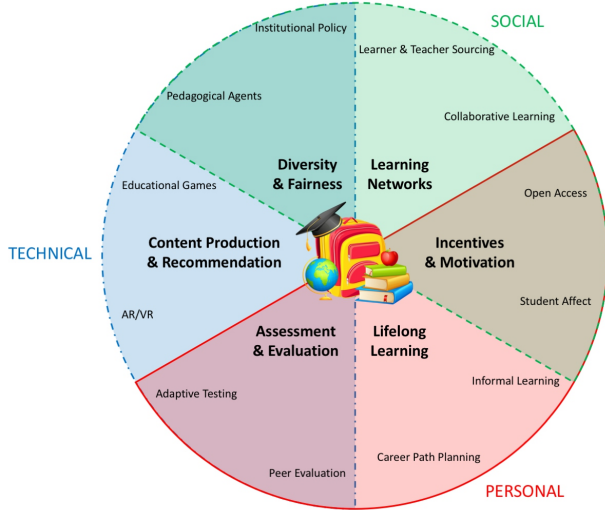


Fig. 3. A list of (some of) the topics in personalized education, organized by three different aspects: technical, personal, and social. We focus on six of them in this paper.

design far beyond handing a certificate. This includes, for example, monetary rewards in the form of bonuses for online learning materials. The incentives can be induced also by soft-methods such as gamification based on the learner's characters to promote continuous learning, or adapting the features of the learning environment based on the learner's traits to engage her in the learning process as far as possible. In [8], for example, the authors investigate the effects of gamification on students' motivation from several perspectives. We discuss such challenges and methods in **Section V**.

Education is social and learners can extremely benefit from their peers. Therefore, it is urgent to develop effective ways to build networks that serve as a conduit of knowledge for learners to interact with each other. In the current form, personalized education suffers from a lack of student-student and student-teacher connections and interactions, which have an unquestionably positive impact on learning through discussions, joint efforts, and brainstorming. For example, Reference [9] investigates and compares the influence of such communities from both students'- and teachers' perspectives. As another example, in state-of-the-art research, the authors study the building and sustaining community in asynchronous learning networks, i.e., when the learners are physically separated. Despite the past research efforts, we believe

that capitalizing on AI and ML methods, online platforms have more to offer, especially, for building the knowledge- and expertise networks that facilitate the assimilation and dissemination of knowledge, and consequently, by enabling close interactions (in terms of mentorship, friendship, coworkers, and the like), creating knowledge. Personalized education platforms can promote autonomous network formation by encouraging learners to interact. Moreover, the platforms can establish links among those learners that satisfy some similarity conditions and hence can be useful to each other for cooperation, inspiration, and motivation. Still, it is vital to notice that online contacts can be lost easily, and the learners, especially at early ages, are more prone to feeling isolated and depressive. We elaborate on these issues in **Section VI**.

In many different ways, education affects the well-being of humans, and thereby the society, both in the short-term and long-term. As such, fairness is a highly important aspect of education, regardless of being in conventional classrooms form or in modern platforms that can personalize the learning experience. Despite this great importance, personalized education, similar to its traditional counterpart, might result in- and strengthen inequality. This arises, for example, due to unequal access to the learning platforms, biases in training data, inaccuracy in algorithm design, and the like. Indeed, existing research shows that some subgroups of students, mainly those privileged also in conventional education forms, would profit from personalized education more than their peers. To address such issue more religiously, there has been intensive effort to develop appropriate fairness models [10]. Moreover, several research works study the fairness of predictive algorithms in educational settings. Another crucial issue is 'diversity'. Today, it is well-established that diversity promotes innovation and efficiency in the working place. Nonetheless, given the social-responsibility of education, only recruiting diverse talents does not suffice. AI-based personalized education platform can be a boost to diversify the education environment, for example, by rewarding collaborative learning in diverse networks. We discuss these topics in **Section VII**.

II. CONTENT PRODUCTION AND RECOMMENDATION

The quality of education ultimately depends on the quality of the learning content. Creating new content requires the wisdom of human content designers and educational experts; to date, AIs have not shown the capability of creating learning content on their own. However,

they still have plenty to offer in content production by automating mundane jobs and helping humans in tasks where human input is necessary. Specifically, the role of AI should be to i) take away repetitive tasks that can be automated and ii) assist humans by providing feedback extracted from data during the process of content production in a human-in-the-loop manner. There are ample future research directions in content production; we list a few below.

- **Content summarization and question generation:**

In many educational domains, knowledge is factual. For example, in History, one often needs to remember specific detailed facts about historical events. Even in scientific domains such as Biology, there is also factual knowledge such as the size and life span of an animal. In this case, there are many natural language processing (NLP)-based tools that can be used for content production. For example, text summarization tools can sort through long, sometimes redundant textbook sections and extract key facts for remedial studies. This is not only helpful but also sometimes crucial to certain learner groups, such as those with learning disabilities. Moreover, automatic question generation can effectively produce high-quality factual assessment questions that have short, textual answers [11]. An example of automated question generation is shown in **Table I**; we reversed a long short-term memory (LSTM) network-driven question answering pipeline trained on common question answering datasets, turned it into a question generation pipeline, and applied it to textbooks. Human experts have indicated that the quality of generated questions is higher than that generated by other methods. See [11] for details.

- **Multi-modal content understanding:** Many educational domains involve multi-modal learning content, such as text, formulas, figures, and diagrams. When a learner fails to answer an assessment question correctly, personalized education systems need to automatically retrieve relevant content to help the learner resolve their confusion (by retrieving examples and explanations) or give the learner more practice opportunities (by retrieving assessment questions). Retrieving content within the same modal is relatively easy; for example, when a learner answers a textual question incorrectly, it is possible to use information retrieval methods to extract relevant textbook chunks or lecture slides. However, when the most helpful content is in another modality, for example, when a Venn diagram is the most effective at helping a learner clear up a misconception in a probability question involving text and mathematical formulas, it is hard to retrieve the diagram. Therefore, more work needs to be done when the domain includes multi-modal content; To understand these content modalities and use them for content production, we need to learn universal representations across all modalities, possibly using

embedding approaches to map multiple modalities into a shared vector space.

- **Human-in-the-loop content design:** Even for humans, learning content is not created in one shot; just like textbooks have different editions, learning content is frequently edited and updated over time. Therefore, during this multi-step process, we can use AIs to act as (possibly even interactive) assistants to content designers. Duties that can be assigned to AIs include i) Analyzing learners' data to identify the areas of priority for new content and assessment questions that need to be improved (see Section III for discussions on how existing learner assessment models can also provide information on question quality); ii) Providing drafts of instructor responses and perform automated checking of human-generated content using NLP tools; iii) Using crowdsourcing to put the learning content together by soliciting on-demand feedback. The last task is especially important in online educational settings, where learning follows during frequent exchanges between learners and human instructors and assistants [12].

Even with high-quality learning content, the presentation, i.e., the personalized recommendation of the right learning content to the right learner at the right time is crucial to optimize the learning outcomes. Fortunately, this is an area where AIs can excel at: By automatically deploying recommendations and analyzing the data of learners' performance, they can quantify the effect of learning content on certain learners in terms of specific learning outcomes to detect the most effective ones. On the contrary, humans, even educational experts in the past, use theoretical models of learning and do not fully take advantage of this data. Among several directions for future research in this area, we discuss a few in the following.

- **Recommendations at the microscopic and macroscopic levels:** Learning content is organized at multiple levels, down to individual paragraphs and assessment questions, and up to courses and textbooks that organize several pieces of learning content together. Therefore, we need to study content recommendations at multiple levels: (i) microscopic level such as individual questions and lecture video suggestions [13]; and (ii) macroscopic level such as course recommendation, especially for learners taking massive open online courses (MOOCs) [14].
- **Efficient experimentation and synthetic learner models:** Traditionally, the fields of learning science and education have relied on rigorous A/B testing to validate the educational impact of learning content, usually in terms of its ability to improve learning outcomes for learners in the experimental group over those in the control group. However, this approach leads to long experimental cycles since i) one can typically validate only one learning content at a time, and ii) metrics such as long-term learning outcomes naturally require long experimental cycles. Therefore,

Context (Biology): On each chromosome, there are thousands of genes that are responsible for determining the genotype and phenotype of the individual. A gene is defined as a sequence of DNA that codes for a functional product. The human haploid genome contains 3 billion base pairs and has between 20,000 and 25,000 functional genes.

3 billion base pairs

between 20,000 and 25,000

How many base pairs are on the human genome?

How many functional genes are on the human haploid genome?

TABLE I

EXAMPLE OF TWO AUTOMATICALLY GENERATED ASSESSMENT QUESTIONS FOR TWO DIFFERENT ANSWERS WITH THE SAME INPUT CONTEXT FROM A TEXTBOOK. THE ANSWERS ARE UNDERLINED AND MARKED WITH DIFFERENT COLORS IN THE INPUT CONTEXT.

it is imperative to search for novel tools that enable rapid experimentation. Possible solutions include using Bayesian optimization to test multiple contents simultaneously, or utilizing reinforcement learning (RL) as more and more learners use a piece of learning content. In the past, using RL to learn instructional policies (content recommendation can be viewed as a form of the instructional policy) has been limited due to the lack of large-scale real learner data; however, recent approaches have looked at using data- or cognitive theory-driven synthetic learner models to simulate learner data.

- **Conflicting objectives:** There is no unified objective in personalized learning since learning outcomes itself is defined at multiple timescales. The optimal action may differ across different objectives. For example, the learning content used in a practice session that maximizes a learner's performance on the midterm exam tomorrow may differ from the one that maximizes their overall course grade, which may differ from the one that maximizes their chance of getting a specific job after graduation. Therefore, We need to develop personalization algorithms that can balance multiple objectives and even resolve potential conflicts among these objectives. We also need to understand how these objectives interact with each other; for example, what skills taught in courses and schools carry over after graduation, which is a key issue in lifelong learning (discussed in detail in Section IV).

Figure 4 shows the interplay between different elements, such as context, prediction, feedback, and the like, to optimize the course recommendation. It is worth noting that the approaches described above are generic in the sense that they have wide applicability to different educational areas including signal processing, possibly with minor domain-dependent adaptations. As an example, in [15], the authors apply several of the aforementioned ideas to develop *eTutor*, a personalized web-based education system that learns the optimal sequence of teaching materials to show based on the student's context and feedback about the previously shown teaching materials. In an experiment, they apply the eTutor system in the following scenario: The students have studied digital signal processing in the past. The role of eTutor is to recommend learning materials to the students with the goal of refreshing their

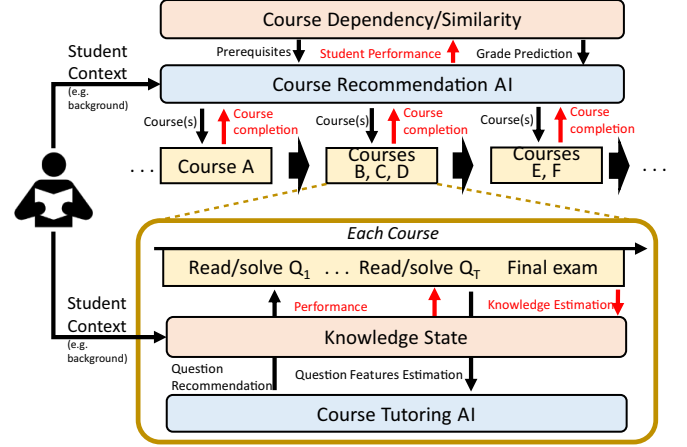


Fig. 4. A detailed framework for course recommendation.

minds about discrete Fourier transform in the minimum amount of time. The e-tutor shows better performance compared to a random- and a fixed-selection rule.

III. ASSESSMENT AND EVALUATION

A key problem in learner assessment is to estimate how well they master each knowledge component/concept/skill from their responses to assessment questions. Related works can be broadly classified into two categories: i) *static* models that analyze the generated data as learners take an assessment and thereby assuming that each learner's knowledge remains constant during the assessment, and ii) *dynamic* models that track learners' progress throughout a (possibly long) period as their knowledge levels evolve. Below we provide a short overview of each category.

- **Static models- Item response theory (IRT):** The basic 1PL IRT model characterizes the probability that a learner answers a question correctly as

$$P(y_{i,j} = 1) = \sigma(a_j - d_i),$$

where $y_{i,j}$ denotes the binary-valued graded response of learner j to question i , where 1 implies a correct response and 0 otherwise. Moreover, $a_j \in \mathbb{R}$ and $d_i \in \mathbb{R}$ are scalars that correspond to the learner's ability and the question's difficulty, respectively. Also, $\sigma(\cdot)$ is a link function that is usually the sigmoid function or the inverse probit link function. Later extensions include 2PL IRT models that add another multiplicative scaling parameter. This parameter cor-

responds to the ability of each question differentiating high-capacity learners from low-capacity ones. Besides, 3PL IRT models add another scalar outside of the link function, which corresponds to the probability that an item can be guessed correctly. Finally, multidimensional IRT models use vectors instead of scalars to parametrize the strengths and difficulties to capture multiple aspects of one’s ability. Using the aforementioned models, one can i) obtain relatively stable estimates of learners’ ability levels by denoising learners’ responses and ii) estimate the quality of each assessment question.

- **Dynamic models- Knowledge tracing (KT):** KT models consist of two parts, learner performance model ($f(\cdot)$) and learner knowledge evolution model ($g(\cdot)$), as

$$y_t \sim f(a_t), \quad a_t \sim g(a_{t-1}),$$

where t denotes a discrete-time index. Early KT models such as Bayesian KT [16] treats knowledge (h_t) as a binary-valued scalar that characterizes whether or not a learner masters the (single) concept covered by a question. The performance and knowledge evolution models are simply noisy binary channels. Later, factor analysis-based KT models use a set of hand-crafted features such as the number of previous attempts, successes, and failures on each concept to represent a learner’s knowledge levels. These models require expert labels to associate questions to concepts, resulting in excellent interpretability since they can effectively estimate the knowledge level of each learner on expert-defined concepts. Recent KT models incorporate deep learning, especially recurrent neural networks into the KT framework, where knowledge is represented as a latent vector \mathbf{a}_t . These models achieve state-of-the-art performance in predicting future learner responses, although in some cases the advantage is not significant despite paying the price of losing some interpretability.

The existing learner assessment models have several bottlenecks. First, there are not many models with both the ability to achieve state-of-the-art performance in data fitting, (i.e., future performance prediction) as well as feedback generation (i.e., providing interpretable feedback to learners and instructors for downstream tasks such as personalization). Therefore, it is imperative to develop new deep learning-based models that not only inherit the flexibility of neural networks to accurately predict learner performance but also build in cognitive theory-inspired structures to promote interpretability and enable the generation of meaningful feedback. As an example, in the recently developed attentive knowledge tracing (AKT) model [17], visualized in **Figure 5**, we combined state-of-the-art attention networks with cognitive theory-inspired modules. We used a monotonic attention mechanism where weights exponentially decay over time and questions embeddings parametrized by the 1PL IRT model

to prevent overfitting. Experimental results show that AKT not only outperforms existing KT models but also exhibits some interpretability; see [17] for details. Moreover, existing learner assessment models almost exclusively operate on *graded* learner responses; however, converting raw learner responses to graded responses leads to considerable information loss. For multiple-choice questions, different distractor options are not created equal; choosing certain incorrect options over others might indicate that a learner exhibits a certain misconception. However, this information is lost when the learner’s option choice is converted to a graded response. Moreover, due to its superior pedagogical value, open-response questions are widely adopted; the specific open-ended response a learner enters contains rich information about her knowledge state. Therefore, it is vital to develop models that go deeper than the graded response level and into the *raw* response level. These models enable personalization at even finer levels, e.g., after each step as a learner solves an open-ended mathematical problem step-by-step, and enable personalized education systems to attend to learner difficulties in a more timely manner.

Another consideration in effective learner evaluation is that assessment and performance prediction models must be tailored to different learning environments and platforms. For example, accurate prediction of students’ future college performance based on their ongoing academic records is crucial to carry out effective pedagogical interventions so that on-time and satisfactory graduation is ensured. However, foretelling student performance in completing degrees (e.g., college programs) is significantly different from that for in-course assessment and intelligent tutoring systems. In what follows, we describe the most important reasons.

- First, students differ tremendously in terms of backgrounds as well as the study domains (majors, specializations), resulting in different selected courses. Even if the courses are similar, the sequences in which the students take the courses might differ significantly. Therefore, a key challenge for training an effective predictor is to handle heterogeneous student data. In contrast, solving problems in intelligent tutoring systems often follow routine steps that are identical for all students. Similarly, predictions of students’ performance in courses are often based on in-course assessments that are identical for all students.
- Second, although the students often take several courses, not all of them are equally informative for predicting the students’ future performance. Utilizing the student’s past performance in all courses that he/she has completed not only increases complexity but also introduces noise in the prediction, thereby degrading the performance. For instance, while it is meaningful to consider a student’s grade in ‘Linear Algebra’ for predicting his/her grade in ‘Linear Optimization’, the student’s grade in ‘Chemistry Lab’ may have much weaker predictive power. However,

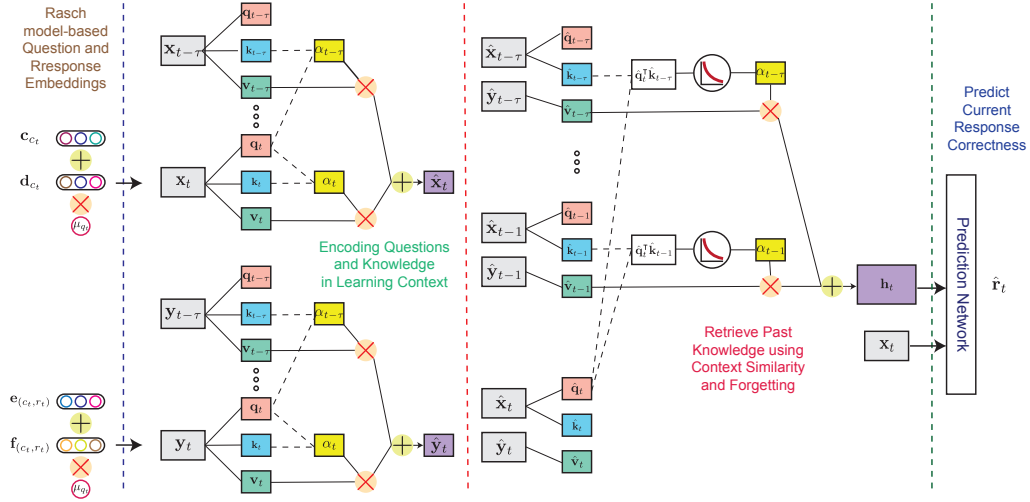


Fig. 5. Overview of the AKT method. We use IRT-based raw embeddings for questions and responses. We compute context-aware representations of questions and responses using two encoders. We then use a knowledge retriever to retrieve past acquired knowledge for each learner using a monotonic attention mechanism, which is used for performance prediction.

the course correlation is not always as obvious as in this example. Therefore, to enhance the accuracy of performance predictions, it is essential to discover the underlying correlation among courses.

- Third, predicting student performance in a degree program is not a one-time task; rather, it requires continuous tracking and updating as the student finishes new courses over time. An important consideration is the following: The prediction shall be made based on not only the most recent snapshot of the student accomplishments but also the evolution of the student progress, which may contain valuable information to improve the prediction's accuracy. However, the complexity can easily explode since even mathematically representing the evolution of student progress itself can be a daunting task. Treating the past progress equally as the current performance when predicting the future may not be a wise choice either since old information tends to be outdated.

Finally, we would like to emphasize the following: Similar to the offline system, in AI-powered personalized education, the assessment does not remain limited to evaluating the performance of individual students in different tests in a single online education portal. Indeed, evaluation might be necessary not only for individuals but also for a collection of students as well as other stakeholders such as educators, policy-makers, and the providers of online education. In particular, fair and precise comparison, analysis, and accreditation of online education portals, as well as the degrees and certificates provided by such portals, are crucial. The reasons include the following: (i) Distance education has grown into a broad industry in the past decade; (ii) The majority of the learners rely on certificates of online classes as approval for obtaining the necessary knowledge and skills; (iii) Online education is inherently international and crosses the boundaries. Similar to improving the evaluation of students, AI and ML methods,

together with big data analysis, can assist in accreditation and comparison of online portals and the issued degrees and certificates; this includes, e.g., comparing the average students' performance with an online degree with that of traditional, yet accredited, degree. A detailed discussion of such topics have several perspectives, and therefore it is out of the scope of this paper.

IV. LIFE-LONG LEARNING

Life-long learning emphasizes holistic education and the fact that learning takes place on an ongoing basis from our daily interactions with others and with the world around us in different contexts. These include not only schools but also homes and workplaces, among several others. Because of its ongoing nature, making foresighted learning plans is crucial for life-long learning to achieve the desired outcome.

In the school context, a specific challenge for developing a learning plan is the course sequence recommendation in degree programs [14]. Recent studies show that the vast majority of college students in the United States do not complete college in the standard time. Moreover, today, compared to a decade ago, fewer college students graduate in a timely manner. While several factors contribute to students taking longer to graduate, such as credit losses in the transfer, uninformed choices due to the low advisor-student ratios, and poor preparation for college, the inability of students to attend the required courses is among the leading causes. If students select the courses myopically without a clear plan, they may end up in a situation where required subsequent courses are offered (much) later, thereby (significantly) prolonging the graduation time. Hence, to accelerate graduation, students shall essentially select courses in a foresighted way while taking the course sequences (shaped by courses being mandatory, elective, pre-requisite) into account. Moreover, it is vital to observe the timing in which the school offers

various courses. More importantly, since the number and variety (in terms of backgrounds, knowledge, and goals) of students is expanding rapidly, the same learning path is unlikely to best serve all students. Therefore, it is crucially important to tailor course sequences to students. To this end, it is necessary to learn from the performance of previous students in various courses/sequences to adaptively recommend course sequences for the current students. Obviously, this depends on the student's background and his/her completion status of the program to maximize any of a variety of objectives including the time to graduation, grades, and the trade-off between the two. To make such plans, AI is a tool of great potential; however, designing AI technologies for personalized, foresighted, and adaptive course planning is challenging in several dimensions, as briefly described below.

- First, course sequence recommendation requires dealing with a large decision space that grows combinatorially with the number of courses.
- Second, there is a great deal of flexibility in course sequence recommendation since multiple courses can be taken simultaneously while it is also subject to many constraints due to prerequisites and availability.
- Third, any static course sequence is sub-optimal since the knowledge, experience, and performance of a student develops and evolves in the process of learning.
- Last but not the least, students vary tremendously in backgrounds, knowledge, and goals.

For example, in [14], we develop an automated course sequence recommendation system to address the aforementioned challenges. To reduce complexity and enabling tractable solutions, we solve the problem in two steps, as illustrated in **Figure 6**: (i) The first step corresponds to offline learning, in which a set of candidate recommendation policies are determined to minimize the expected time to graduation or the on-time graduation probability using an existing dataset of anonymized student records based on dynamic programming; (ii) The second step corresponds to online learning, in which for each new student, a suitable course sequence recommendation policy is selected depending on this student's background using the learned knowledge from the previous students. In other life-long learning contexts (e.g., the workplace), while similar challenges may still be present, new challenges are likely to emerge and hence, foresighted learning plans must be tailored to the specific context.

Recent research shows a significant gap between the lectures offered in schools and job requirements, especially in emerging disciplines like data science. Soft skills such as communication and teamwork are often even more important than typical technical skills. Future research on life-long learning shall bridge this gap. Indeed, there is a systematic demand for the research community to identify and study the skills that significantly contribute to the professional perspectives instead of maximizing achievement in schools. The educators can take advantage of the findings to adjust school curricula and educational

activities to better prepare students for the future. The centerpiece of possible approaches is to fuse a student's school records with future employment outcomes, possibly tracked over a long period, as well as other data sources such as course syllabi and job postings, to identify the crucial skills that extend from the education to profession. There is also a necessity for research in labor studies to conduct interviews with i) employers to understand their requirements, ii) job seekers to identify the skills they are keen to acquire, and iii) training providers to clarify the skills that can be taught in a part-time or on-the-job way rather than through centralized educational programs, given workers' real-life constraints.

V. INCENTIVES AND MOTIVATION

So far, one crucial aspect of personalized education has been largely left aside, namely motivation and incentive-design. This is unfortunate as these factors significantly contribute to the learners' perseverance and engagement, thereby the overall students' achievements. As such, they affect not only the individuals but also the entire society in terms of the efficiency of resource expenditure. In educational sciences, motivation is regarded as a concept that involves several learning-related features such as initiation, goal-orientation, intensity, persistence, and quality of behavior [7]. Therefore, as [18] describes, motivation is an unobservable dynamic process that is difficult to be directly measured but it is inferable from the observations. Similar to other crucial factors of successful education such as talent and interest, motivation originates and is influenced by personal factors such as goals and beliefs. As such, it is reasonable to conclude that intelligent 'personalization' affects the motivation to a large extent. Motivation can be intrinsic or triggered by external factors. As such, various features of personalized education, such as recommending a proper series of content or creating educational networks, might implicitly improve the learner's motivation by increasing the engagement level. Such efforts make the learning experience more pleasant, thereby improving the learner's satisfaction level. This is, however, insufficient. It is imperative to integrate direct motivating methods into personalized education and the learning platform. To this end, in the following, we first describe a few frameworks which can accommodate motivation and its relevant concepts appropriately. More can be found in [18].

- **Behavioral Economics:** Any personalized education platform shall be able to appropriately connect, interact, and interface with humans. Hence, the proper operation significantly depends on various features of the members of the target group that shapes their decision-making behavior. Indeed, a 'utility function' is the most seminal computational model for the interests of the learners. For a rational decision-maker, the utility function is conventionally increasing concave and to be maximized. However, humans often demonstrate unusual patterns in their utility functions and decision-making due to the fol-

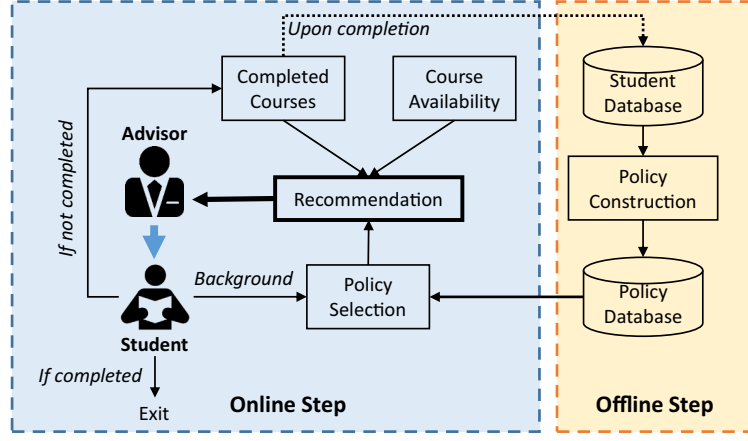


Fig. 6. Illustration of course sequence recommendation.

lowing reasons: (i) Humans make mistakes, often due to inaccurate beliefs and imprecise predictions; (ii) Humans often act irrationally and based on heuristics; (iii) Humans think and act in different manners as a result of their unique background, including personality and experiences [19]. Behavioral economics accommodate and formalizes such aspects; Therefore, one can take advantage of behavioral economics for efficient incentive design and motivation in learning platforms [20].

- **Self-Determination Theory:** This theory asserts that humans have an intrinsic urge to be self-autonomous, competent, and connected, concerning their environment [21]. While behavioral economics is appropriate to investigate the motivation resulted from external rewards, self-determination observes the motivation from an internal perspective. Indeed, any environment, including the learning platforms, that satisfies the aforementioned needs of human, awakes the intrinsic motivation, rendering external triggers rather unnecessary. As such, promoting intrinsic motivation is significantly more effective than extrinsic motivation as it is often associated with lower cost compared to material-rewarding and has a long-lasting effect [18].
- **Self-Efficacy Theory:** This concept corresponds to an individual's confidence in her capability of performing a specific task to be undertaken, for example, learning in an online learning platform or performing at a certain level [22]. Researchers show that humans constantly assess their self-efficacy, mainly based on the observed information from the environment and the past experiences [22]. Similar to the self-determination theory, self-efficacy considers the intrinsic motivation, implying that a feeling of efficacious triggers the internal motivation feeling in learners. Other relevant concepts include 'interest' and 'goal-orientation' [18].

The main challenge is to utilize AI and ML to motivate the learners of a personalized learning platform, based on the

aforementioned theories that formalize and explain human behavior. To clarify this, consider the utility function of a learner in a personalized learning platform as an example [20]. The function quantifies the learner's well-being while using the platform, and, consequently, her (future) engagement. Some learners exhibit hyperbolic preferences, overweighting the present so much that future rewards are largely ignored. Some learners show strong reactions even to non-monetary rewards. Some learners demonstrate reference-dependent preferences, implying that the utility is largely determined by its distance from a reference point, for example, a pre-defined goal or the average performance. By using ML and AI methods, the learning platform can take advantage of the available data and a learner's feedbacks to estimate the utility function of that learner, hence predicting her reaction to potential triggers of incentive and motivation. Consequently, the platform can adjust and allocate the reward among the learners efficiently and fairly. As another example, consider the self-determination theory. Based on this theory, a sophisticated personalized learning platform guarantees choice, connectedness, and the feeling of competence for the learner. To this end, the design of recommendation tools based on AI and ML methods should allow for enough alternatives, both at micro and macro levels, to ensure autonomy. Moreover, the suggested learning content should be based on the learner's feedback and the result's of accurate assessment, to avoid inducing a feeling of incompetence in the learner. Besides, promoting network formation or establishing a link between coherent learners and intensive interaction results in connectedness. This is also in accordance with the self-efficacy theory, in the sense that by providing an appropriate feedback and suitable side-information, the platform increase the positive belief of a learner in her ability to perform well on a learning platform.

VI. BUILDING LEARNING NETWORKS

A potential negative effect of personalized education, especially in an online environment, is a loss of peer interactions and of the sense of community that is usually

present in traditional classrooms. Fortunately, the rise of online social networks seems to facilitate interaction and networking between teachers and learners, also the co-production of content both within and outside the classroom. Learning applications and pedagogy can also be built based on online social networks to bridge formal and informal learning, also to promote peer interactions on both curricular and extra-curricular topics. Moreover, various education-related social networks have been created to facilitate collaboration, post/answer questions, and share resources; however, a formal method to build these learning networks and a deep understanding of their effectiveness are absent.

The core of learning networks is peer interaction, which has important implications for personalized education when teaching resources are limited. For example, peer review serves as an effective and scalable method for assessment and evaluation when the number of students enrolled in a course far exceeds the number of teaching assistants. However, effective peer review in learning networks poses new challenges [23]. On the one hand, the peer reviewers have different intrinsic capabilities, which are often unknown. On the other hand, the peer reviewers can choose to exert different levels of effort (e.g., time and energy spent in reviewing), which is unobservable. Identifying unknown intrinsic capabilities corresponds to the *adverse selection problem* in the game theory. A natural candidate for solving this problem is to use matching mechanisms, i.e., assign reviewers to students. Existing works on matching mechanisms focus on one-shot peer interactions and design one-shot matching rules. However, their assumption does not hold in peer review systems, where the review quality depends crucially on the reviewers' effort. Motivating reviewers to exert high effort corresponds to the *moral hazard problem* in game theory. One way to address this problem is to use social norms, where each peer reviewer is assigned with a rating that summarizes her past behavior and recommended a 'norm' that rewards reviewer with good ratings and punishes those with bad ratings. However, existing works on social norms assume that peer reviewers are homogeneous. This assumption does not hold in peer review systems because different reviewers have different intrinsic capabilities. Because a peer reviewer's ultimate review quality is determined by her intrinsic capabilities and effort, designing effective peer review systems in learning networks becomes significantly more challenging due to the presence of both adverse selection and moral hazard. Therefore, new peer review system designs shall simultaneously solve both problems so that peer reviewers find it in their self-interest to exert high effort and receive ratings that truly reflect their capabilities.

Another primary function of learning networks is to foster learning content co-production and sharing. Building such learning networks is vastly different from building traditional networks such as computer networks and transportation networks, as in learning networks, individual learners create and maintain the links. Because links per-

mit the acquisition and dissemination of learning content, it is theoretically intriguing and practically valuable to have a deeper understanding of the networks that are more likely to be formed by self-interested learners. Game theory is a useful tool to formulate and understand the strategic behavior of learners. The formulation must capture the heterogeneity of learners in terms of goals, capabilities, costs, and self-interest nature [24]; That is, each learner intends to maximize its benefit from content co-production and sharing minus whatever cost it pays to establish links. Our prior work [25] studies the endogenous formation of networks by strategic, self-interested agents who benefit from producing and disseminating information. The results showed that the typical network structure that emerges in equilibrium displays a core-periphery structure, with a smaller number of agents at the core of the network and a larger number of agents at the periphery of the network. Furthermore, we established that the typical networks that emerge are minimally connected and have short network diameters, which are independent of the size of the network. In other words, theoretical results show that small diameters tend to make information dissemination efficient and minimal connectivity leads to minimizing the total cost of constructing the network. These results are consistent with the outcome of numerous empirical investigations. Such theoretical analysis and tools are essential to guide building learning networks. Also, based on this analysis, one can create protocols to motivate selfish learners to take actions that promote the system-wide utility.

Future research in learning networks hinges on understanding the knowledge flow between students via peer interaction. Such an understanding enables educators to effectively modulate peer interactions and to encourage the interactions that promote peer learning. Peer learning is especially valuable as education extends to more diverse settings, such as remote online learning during the COVID-19 pandemic. In such settings, it is hard for instructors to moderate learning activities remotely; hence, peer learning through online course discussion forums becomes essential. Therefore, it is vital to understand the interaction tendencies and students' behaviors in these discussion forums [26], understanding the flow of knowledge by combining discussion forum activities with grades, identifying the factors that enhance knowledge flow, and designing automated strategies to moderate student activities when necessary.

VII. DIVERSITY, FAIRNESS, AND BIASES

Experimental studies show that AI-driven personalization such as student assessment, feedback, and content recommendation improve the overall learning outcomes; nonetheless, certain student subgroups may benefit more than other subgroups due to the biases that exist in the training data [27]. This imbalance jeopardizes students that are already under-served in particular since they often have less access to advanced, digitized educational systems and are less frequently represented in datasets collected by

these systems [28]. Therefore, it is essential to develop AI tools that promote fairness among learners with different backgrounds, thereby making education more inclusive for the next generations.

To mitigate biases and to promote fairness and equity in AIs, currently, the researchers pay significant attention to developing approaches that promote fairness, primarily in the context of predictive algorithms.

- The first major research problem studied is how to properly define fairness. Many definitions of fairness exist, including individual fairness that requires that users with similar feature values be treated similarly, parity in the predicted probability of each outcome across user groups (drawn using sensitive attributes), parity in the predicted probability of each outcome given actual comes and regardless of sensitive attributes, and counterfactual fairness, which requires that the predicted outcome for each user remain mostly unchanged if the sensitive attribute changes.
- The second major research problem is to develop methods that enforce fairness in predictive algorithms. Existing approaches include preprocessing the data to select only the fair features as input to algorithms, and post-processing the output of algorithms to balance across user groups. The most promising approach is to impose regularizers and constraints while training predictive algorithms. These methods result in better fairness at the expense of sacrificing some classification accuracy; however, they are empirically shown to obtain better tradeoffs between fairness and accuracy than other fairness-promoting methods.

Promoting fairness and equity is a crucial necessity of education that requires a comprehensive approach to be fulfilled: We need to not only design fair personalization algorithms but also develop systematic principles and guidelines for their application in practice; In other words, we need a set of tools to regulate the use of AI algorithms. Finally, despite great promises, AI-driven personalization in education can also bring risks that have to be closely monitored and controlled. Recently, there have been calls for a food and drug administration (FDA)-type framework for other AI applications such as facial recognition. It is essential to establish a similar ecosystem in education with a set of legislation and regulations around issues of data ownage, sharing, continuous performance monitoring, and validation, to regulate every step of the process, from ensuring the diversity and quality of the collected data, developing algorithms with performance guarantees across different educational settings, to identifying misuse and implementing a fail-safe mechanism.

VIII. COVID-19 AND AI-ENABLED PERSONALIZED EDUCATION

Among its several other adverse effects, the COVID-19 pandemic has disrupted or interrupted the functionality of conventional education systems around the globe. Not surprisingly, the students perceive the bitterness of this adverse effect at various degrees, depending on several

factors such as country/region, family status, and individual characteristics. The complications vary over a large spectrum and include reduced learning ability, depression, loss of concentration, and a decline in physical fitness. Such issues mainly arise due to spending less- or no time at school, where the students receive educational materials and support in learning, interact with their peers and teachers, develop incentives, and are evaluated. Besides, many students cannot take full advantage of replacements such as online materials, e.g., in the absence of an appropriate technological device or reliable internet connection, or suitable learning ambient at home. As a consequence of its vitality, the impacts of COVID-19 on education has attracted a great deal of attention. For example, in [29], the authors describe the influence of pandemic-triggered growth in online learning on student's performance and equity. Several other papers, such as [30], study domain-specific educational effects of the pandemic, evaluate the available solutions, and provide suggestions for policy-makers to compensate for the pandemic's negative educational consequences.

Personalized- and distance education have already been trending in the past decade; Still, COVID-19 has urged both public and private sectors to rapidly increase the investigations into research and development in this area to earn individual- and/or social profit. For example, the pandemic has boosted the usage of online learning tools for signal processing education, especially at the undergraduate levels. These include, for example, web-based laboratories for digital signal processing [31], and online machine learning education modules [32]. While it is essential to carefully study such a tremendous push towards revolutionizing education from several perspectives, in the scope of our paper, we confine our attention to the role and influence of AI and ML.

As described previously, AI and ML have a great potential to enhance online education in different ways, e.g., through improving the quality of learning materials, enabling fairness and diversity, generating proper tests, and allowing to build knowledge networks. That side is a universal aspect of applying AI and ML methods in distance and asynchronous education regardless of the current pandemic; nevertheless, such methods can additionally assist in accelerating rebuilding the education systems and in mitigating the pandemic's detrimental effects. For example, by using ML methods on the available data, policy-makers can classify the students based on the exposure to the educational effects of a pandemic; using such classification, one can allocate resources efficiently while satisfying fairness constraints. As another example, by taking advantage of ML methods, one can optimize the school closure plan based on different features such as neighborhood, size, grade, and the like.

IX. SUMMARY AND CONCLUSION

Enabling 'personalized education' is one of the most precious merits of AI concerning education. This paradigm significantly improves the quality of education in several

dimensions by adapting to the distinct characteristics and expectations of each learner such as personality, talent, objectives, and background. Besides, online education is of the utmost value under abnormal circumstances such as the COVID-19 outbreak or natural disasters. Indeed, conventional education requires significantly more resources than the online format concerning educational space, scheduling, and human resources, which makes it prone to failure with even a small shift in conditions. As such, emerging alternatives are inevitable. Despite having the potential of a revolutionary transformation from traditional education to modern concepts, personalized education is associated with several challenges. We discussed such challenges, provided a brief overview of the state-of-the-art research, and proposed some solutions. **Table II** summarizes some of the future research directions.

X. ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers and the guest editors for several helpful suggestions.

REFERENCES

- [1] Y. Zhong, C. Jiang, W. Xu, and J. Li, "Discourse level factors for sentence deletion in text simplification," in *Proc. AAAI Conference on Artificial Intelligence*, Feb. 2020.
- [2] M. Sachan, A. Dubey, E. H. Hovy, T. M. Mitchell, D. Roth, and E. P. Xing, "Discourse ineffectiveness of crowd-sourcing on-demand tutoring from teachers in online learning platforms multimedia: A case study in extracting geometry knowledge from textbooks," *Computational Linguistics*, vol. 45, no. 4, pp. 627–665, 2020.
- [3] I. Manickam, A. S. Lan, and R. G. Baraniuk, "Contextual multi-armed bandit algorithms for personalized learning action selection," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 6344–6348.
- [4] C. Magno, "Demonstrating the difference between classical test theory and item response theory using derived test data," *CSN: General Cognitive Social Science (Topic)*, vol. 1, 06 2009.
- [5] M. Sharples, "The design of personal mobile technologies for lifelong learning," *Computers and Education*, vol. 34, no. 3, pp. 177 – 193, 2000.
- [6] G. I. Parisi, R. Kemker, J. L. Part, C. Kanan, and S. Wermter, "Continual lifelong learning with neural networks: A review," *Neural Networks*, vol. 113, pp. 54 – 71, 2019.
- [7] W. Grove and L. Hadsell, *Open Learning Environments*, chapter Incentives and Student Learning, pp. 1511–1517, Springer, 01 2012.
- [8] P. Buckley and E. Doyle, "Gamification and student motivation," *Interactive Learning Environments*, vol. 24, no. 6, pp. 1162–1175, 2016.
- [9] P. Vesely, L. Bloom, and J. Sherlock, "Key elements of building online community: Comparing faculty and student perceptions," *MERLOT Journal of Online Learning and Teaching*, vol. 3, 01 2007.
- [10] J. Gardner, C. Brooks, and R. Baker, "Evaluating the fairness of predictive student models through slicing analysis," in *Proceedings of the 9th International Conference on Learning Analytics & Knowledge*, 2019, pp. 225–234.
- [11] Z. Wang, A. S. Lan, W. Nie, A. E. Waters, P. J. Grimaldi, and R. G. Baraniuk, "QG-net: A data-driven question generation model for educational content," in *Proc. ACM Conference on Learning at Scale*, June 2018, pp. 1–10.
- [12] B. Zylich, A. Viola, B. Toggerson, L. Al-Hariri, and A. S. Lan, "Exploring automated question answering methods for teaching assistance," in *Proc. International Conference on Artificial Intelligence in Education (AIED)*, July 2020.
- [13] A. S. Lan and R. G. Baraniuk, "A contextual bandits framework for personalized learning action selection," in *Proc. International Conference on Educational Data Mining*, June 2016, pp. 424–429.
- [14] J. Xu, T. Xing, and M. Van Der Schaar, "Personalized course sequence recommendations," *IEEE Transactions on Signal Processing*, vol. 64, no. 20, pp. 5340–5352, 2016.
- [15] C. Tekin, J. Braun, and M. van der Schaar, "etutor: Online learning for personalized education," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 5545–5549.
- [16] M. Yudelson, K. Koedinger, and G. Gordon, "Individualized Bayesian knowledge tracing models," in *Proc. International Conference on Artificial Intelligence in Education*, July 2013, pp. 171–180.
- [17] A. Ghosh, T. Heffernan, and A. S. Lan, "Context-aware attentive knowledge tracing," in *Proc. ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Aug. 2020.
- [18] M. Hartnett, *Motivation in Online Education*, chapter The Importance of Motivation in Online Learning, pp. 5–32, Springer, 01 2016.
- [19] S. Maghsudi and M. Davy, "Computational models of human decision-making with application to the internet of everything," *IEEE Wireless Communications*, pp. 1–8, 2020.
- [20] S. D. Levitt, J. A. List, S. Neckermann, and S. Sadoff, "The behavioralist goes to school: Leveraging behavioral economics to improve educational performance," *American Economic Journal: Economic Policy*, vol. 8, no. 4, pp. 183–219, November 2016.
- [21] E. Deci and R. Ryan, "Motivation, personality, and development within embedded social contexts: An overview of self-determination theory," *The Oxford Handbook of Human Motivation*, 01 2012.
- [22] A. Bandura, *Self-efficacy: The exercise of control*, Freeman, 1997.
- [23] Y. Xiao, F. Dörfler, and M. Van Der Schaar, "Incentive design in peer review: Rating and repeated endogenous matching," *IEEE Transactions on Network Science and Engineering*, vol. 6, no. 4, pp. 898–908, 2018.
- [24] S. Maghsudi and M. van der Schaar, "Distributed task management in cyber-physical systems: How to cooperate under uncertainty?," *IEEE Transactions on Cognitive Communications and Networking*, vol. 5, no. 1, pp. 165–180, 2019.
- [25] Y. Zhang and M. van der Schaar, "Strategic networks: Information dissemination and link formation among self-interested agents," *IEEE Journal on Selected Areas in Communications*, vol. 31, no. 6, pp. 1115–1123, 2013.
- [26] A. S. Lan, J. Spencer, Z. Chen, C. Brinton, and M. Chiang, "Personalized thread recommendation for MOOC discussion forums," in *Proc. European Conf. Mach. Learn. and Principle Knowl. Discov. Databases*, Sep. 2018.
- [27] J. Reich and M. Ito, "From good intentions to real outcomes: Equity by design in learning technologies," *Irvine, CA: Digital Media and Learning Research Hub*, 2017.
- [28] S. Doroudi and E. Brunskill, "Fairer but not fair enough on the equitability of knowledge tracing," in *Proceedings of the 9th International Conference on Learning Analytics & Knowledge*, ACM, 2019, pp. 335–339.
- [29] E. Garcia and E. Weiss, "Covid-19 and student performance, equity, and u.s. education policy," 2020.
- [30] G. Di Pietro, F. Biagi, P. D. Mota Da Costa, Z. Karpinski, and J. Mazza, "The likely impact of covid-19 on education: Reflections based on the existing literature and recent international datasets," *Publications Office of the European Union (online)*, 2020.
- [31] A. Dixit, U. S. Shanthamallu, A. Spanias, V. Berisha, and M. Banavar, "Online machine learning experiments in HTML5," in *IEEE Frontiers in Education Conference*, 2018, pp. 1–5.
- [32] U. S. Shanthamallu, S. Rao, A. Dixit, V. S. Narayanaswamy, J. Fan, and A. Spanias, "Introducing machine learning in undergraduate dsp classes," in *International Conference on Acoustics, Speech and Signal Processing*, 2019, pp. 7655–7659.
- [33] N. J. Cepeda, N. Coburn, D. Rohrer, J. T. Wixted, M. C. Mozer, and H. Pashler, "Optimizing distributed practice: Theoretical analysis and practical implications," *Experimental psychology*, vol. 56, no. 4, pp. 236–246, 2009.

TABLE II
SOME RESEARCH DIRECTIONS FOR AI-BASED PERSONALIZED EDUCATION

Challenge	Description	Some References
Content Production/Recommendation	Personalized and profession-oriented production, recommendation, and maintenance of contents	[33], [13]
Evaluation and Assessment	Performance comparison in personalized education, testing without information loss, accreditation	[17], [11]
Lifelong Learning	Continuous education and additional qualification for improvement and pivots in profession	[14]
Incentives	Internal and external motivation for learning, gamification, rewarding, inducing confidence	[18], [20]
Networking and Interaction	Inducing learning networks, forming coalitions for efficient learning, imitating teacher feedback	[23], [26]
Diversity and Fairness	Equal access to quality online education, avoiding biases in platform development	[27]

Setareh Maghsudi is an assistant professor at the University of Tübingen. She is also an adjunct lecturer at the Technical University of Berlin and a senior project manager at the Fraunhofer Heinrich-Hertz-Institute. During 2017-2020, she was an assistant professor at the Technical University of Berlin, where she also received her Ph.D. degree (summa cum laude) in 2015. During 2015-2017, she held post-doctoral positions at the University of Manitoba and Yale University. She has received several competitive fellowships and career awards, including from the German Ministry of Education and Research, the German Research Foundation, and the Japan Society for the Promotion of Science. Her main research interests lie in the intersection of network analysis and optimization, game theory, machine learning, and data science.

Andrew Lan is an assistant professor in the College of Information and Computer Sciences at University of Massachusetts Amherst. He received his Ph.D. degree in Electrical and Computer Engineering from Rice University and has worked as a postdoctoral research associate at Princeton University. His research interests lie at the intersection of artificial intelligence, data science, and education, highlighted by a best student paper award at IEEE Big Data 2020. He has also organized a series of workshops on machine learning for education; see ml4ed.cc for details.

Jie Xu is an assistant professor in the Department of Electrical and Computer Engineering at University of Miami. He received the B.S. and M.S. degrees in Electronic Engineering from Tsinghua University, Beijing, China, in 2008 and 2010, respectively, and the Ph.D. degree in Electrical Engineering from University of California Los Angeles in 2015. His main research interests include machine learning and its applications to education and network systems, mobile edge computing, and network security.

Mihaela van der Schaar is the John Humphrey Plummer Professor of Machine Learning, Artificial Intelligence and Medicine at the University of Cambridge, a Fellow at The Alan Turing Institute in London, and a Chancellor's Professor at UCLA. Mihaela was elected IEEE Fellow in 2009. She has received numerous awards, including the Oon Prize on Preventative Medicine from the University of Cambridge (2018), a National Science Foundation CAREER Award (2004), 3 IBM Faculty Awards, the IBM Exploratory Stream Analytics Innovation Award, the Philips Make a Difference Award and several best paper awards, including the IEEE Darlington Award. Mihaela's work has also led to 35 USA patents and 45+ contributions to international standards. Mihaela's research focus is on machine learning, AI, and operations research for healthcare and medicine.