Equity and Fairness of Bayesian Knowledge Tracing

Sebastian Tschiatschek
University of Vienna
Faculty of Computer Science
Research Network Data
Science
Vienna, Austria

Maria Knobelsdorf
University of Vienna
Faculty of Computer Science
Computer Science Education
Vienna, Austria

Adish Singla
Max-Planck Institute for
Software Systems
Saarbrücken, Germany

ABSTRACT

We consider the equity and fairness of curricula derived from Knowledge Tracing models. We begin by defining a unifying notion of an equitable tutoring system as a system that achieves maximum possible knowledge in minimal time for each student interacting with it. Realizing perfect equity requires tutoring systems that can provide individualized curricula per student. In particular, we investigate the design of equitable tutoring systems that derive their curricula from Knowledge Tracing models. We first show that many existing models, including classical Bayesian Knowledge Tracing (BKT) and Deep Knowledge Tracing (DKT), and their derived curricula can fall short of achieving equitable tutoring. To overcome this issue, we then propose a novel model, Bayesian-Bayesian Knowledge Tracing (B^2KT), that naturally enables online individualization and, thereby, more equitable tutoring. We demonstrate that curricula derived from our model are more effective and equitable than those derived from classical BKT models. Furthermore, we highlight that improving models with a focus on the fairness of next-step predictions might be insufficient to develop equitable tutoring systems.

Keywords

equity & fairness, Bayesian Knowledge Tracing, intelligent tutoring systems

1. INTRODUCTION

In recent years Massive Open Online Courses (MOOCs) and online educational platforms have gained significant importance. They hold the opportunity of providing education at scale and making education accessible to a larger part of the world's population. To facilitate learning in online education and enable customized learning paths for all students, intelligent tutoring systems can be employed while limiting the amount of manual work necessary for each student [11].

In that context, moving education from an offline setting to an online setting has the potential to promote *Inclusion, Diversity, Equity, and Accessibility (IDEA)*. In particular, by reducing personnel efforts for tutoring, there is the opportunity to include students with diverse backgrounds and skills, and, importantly, to support their learning equitably. To achieve this, an intelligent tutoring system must be able to adapt to the specific characteristics of each student.

While individualized tutoring has been studied in the community for many years, we consider individualization with a focus on equitable and fair tutoring in this paper. We start by providing a unifying definition of an equitable tutoring system. Our definition is based on the ethical principles of *beneficence* ("do the best") and *non-maleficence* ("do not harm") which are commonly adopted in bioethics and medical applications [1] and strongly related to recent work in the educational data mining field [3]. Concretely, the principle of beneficence dictates that we should provide tutoring which maximizes the achieved knowledge. The principle of non-maleficence dictates that this maximal knowledge should be achieved while minimizing a student's efforts.

In this paper, we particularly focus on how Bayesian Knowledge Tracing (BKT) [2] can be modified to better realize these ethical principles. To this end, we propose the Bayesian-Bayesian Knowledge Tracing (B^2KT) model and demonstrate its advantages for equitable tutoring in several experiments. Furthermore, we investigate the relation of the commonly considered AUC score concerning the derived tutoring policies, finding that even if a BKT model appears fair in terms of the AUC score, the derived tutoring policies can be inequitable. This observation suggests that evaluating models concerning their AUC score can be insufficient to ensure equitable tutoring and that improving models with the sole focus on AUC might not result in more equitable tutoring.

In summary, we make the following contributions:

- We propose a unifying definition of equitable tutoring motivated by the ethical principles of beneficence and non-maleficence.

- We propose the (B^2KT) which allows for effective individualization and demonstrate its benefits concerning equitable tutoring.

- We highlight that focusing on equity in terms of AUC can be insufficient to ensure equitable tutoring in terms of our definition.

2. RELATED WORK

Fairness in online education and BKT. Several works have considered fairness in data-driven educational systems and intelligent tutoring, e.g., [7, 3, 17, 8]. In [7], the authors discussed implications of using data-driven predictive models for supporting education on fairness. They identified sources of bias and discrimination in "the process of developing and deploying these systems". Relevant to promoting IDEA, they discussed several high-level possibilities to improve fairness of systems in the "action step", e.g., through modifying predictions of models for groups of students with different characteristics. It remains open how to operationalize such an approach in cases in which the required knowledge of the group membership is not available in the first place. In [17], it was investigated whether different data sources can provide helpful information to predict students' success in education. They found that different data sources (institutional data and learning management system data) can help to make better predictions but have different characteristics in whether they over- or underestimate students' success for historically disadvantaged subpopulations of students. The authors of [8] investigated fairness trade-offs when students' successes are predicted from university administrative records and identified gender and racial bias in some of the considered fairness measures. Through post-hoc adjustments they investigated how this problem can be alleviated and which trade-offs arise. [3] studied fairness in the context of BKT. In particular, they showed that tutoring policies basing on BKT models can be inequitable when these models are inaccurate and that such inaccurate models can stem from fitting BKT models to populations of students or using models based on incorrect assumptions. They consider the equity-gap as a measure of unfairness, which is the difference between the percentage of fast and slow students mastering a skill.

Individualization in BKT. Several papers have studied individualization of BKT models per student, e.g., [9, 10, 18]. In [10] the *prior per student* model was introduced in which a student-specific parameter characterizing the students' individual knowledge is used. The parameters are fitted using an EM procedure. Importantly, the parameters for each student are fixed during inference by providing the student' identities while in our B^2KT model they are "re-weighted" according to the evidence from interaction. Another example is [18], which considered individualization through defining student and skill specific parameters which are fitted through gradient descent. The models show tangible improvements in prediction performance of unseen students.

Instructional policies. A key necessity for achieving equity according to our definition is using instructional policies which stop practicing a skill at the right time. This problem has for instance been considered in different contexts in [6, 12]. In [6], a stopping condition based on predictive stability, i.e., whether predictions about a student solving some task correctly are stable, was proposed. The resulting policies are particularly useful for student models which allow forgetting and could be combined with our proposed B^2KT model. Further related work has investigated approaches for leveraging deep generative models for creating policies to quickly assess students' knowledge [16] and using reinforcement learning for optimizing tutoring policies [15, 5].

3. BACKGROUND & NOTATION

In this section, we introduce the necessary background and notation for the remainder of this paper.

3.1 Bayesian Knowledge Tracing

Bayesian knowledge tracing (BKT) [2] is a model characterizing the skill acquisition process of students. For a single skill, it can be understood as a standard hidden Markov model (HMM) in which the binary latent state encodes the mastery of the skill, and the observations correspond to binary random variables indicating whether a practicing opportunity of the skill at a specific time was solved correctly or not. Upon practicing the skill, the student acquires the skill with probability $p(T)$ if they have not already mastered the skill. Once a skill is mastered, it remains mastered, i.e., there is no forgetting. If a student has mastered the skill practiced by an exercise, they solve this exercise correctly with probability $1 - p(S)$, where $p(S)$ is the probability of making a mistake. If a student has not mastered the skill, it guesses the correct answer with probability $p(G)$. Before any interaction with the tutoring system, a student has already mastered the skill with probability $p(L_0)$.

In settings with multiple skills, these skills are often assumed to be acquired independently. That is, if there are K skills, the skill acquisition process is modeled by K independent HMMs.

3.2 Deep Knowledge Tracing

In deep knowledge tracing [14] the skill acquisition process is not modeled by an HMM as above but rather implicitly modeled through a recurrent neural network (RNN). In particular, an RNN is trained to predict whether a student will answer the next exercise correctly based on input consisting of which exercises were practiced in the past and whether these practicing opportunities were solved successfully. There is typically no explicit representation of the knowledge state of students in DKT. In the context of DKT previous work has mainly looked into so-called EXPECTIMAX policies which select exercises to practice in order to maximize the average probability that the next exercise will be solved correctly. These policies can be computed for different lookaheads, e.g., for a one-step look-ahead in a greedy and myopic way or for longer-lookaheads using techniques for approximate planning or using exhaustive search.

3.3 Notation

We consider the interaction of students $s \in \mathcal{S}$ with an intelligent tutoring system. The interaction history up to time t is collected in a data set $\mathcal{D}_t^s = \{(z_1, c_1), (z_2, c_2), \ldots, (z_t, c_t)\}$, where $z_{t'} \in \mathcal{Z}$ is the skill practiced through an exercise at time t', $c_{t'} \in \{0, 1\}$ is an indicator of whether the exercise was solved correctly, and $\mathcal{Z}$ is the set of skills. In the context of BKT, we refer to the random variables (RVs) indicating whether skill $i \in \mathcal{Z}$ is mastered at time t as Z_t^i and to the RVs indicating whether an exercise practicing that skill would be solved correctly as C_t^i. Sometimes we use an additional superscript s to indicate the student the random variables correspond to. We use upper-case terms like Z_t^i to denote RVs and their lower-case counterparts like z_t^i to denote particular instantiations of these RVs.

4. OPERATIONALIZING EQUITY

In this section, we provide a definition of equity in intelligent tutoring and discuss its operationalization for promoting IDEA.

4.1 Setting

We consider a tutoring setting in which a total of K sills ought to be taught. These skills are taught by an intelligent tutoring system that employs a tutoring policy $\pi \colon \mathcal{H} \to \mathcal{I} \cup \{\top\}$. This policy maps histories $h \in \mathcal{H}$ consisting of (previous) observations of a student's learning process to an exercise $e \in \mathcal{I}$ to be practiced next or to a stop-action $\top$, which ends the teaching process. We collect all tutoring policies π of the above form into the set Π.

We assume interaction with a set of students $\mathcal{S}$. The students can have different characteristics in terms of their learning behavior. That is, if the learning behavior of the students can be characterized by BKT models, the corresponding parameters $p(L_0)$, $p(S)$, $p(G)$ and $P(T)$ can differ for different students. Every tutoring policy π results in an expected stopping time $T^s(\pi)$, i.e., the expected time of executing the stop action, and an expected knowledge $L^s(\pi)$ acquired by the end of the teaching process, i.e., $L^s(\pi)$ is the expected number of mastered skills when the stop action is executed.

For defining a notion of equity in this setting, we turn to the ethical principles of *beneficence* and *non-maleficence*. We understand them to translate into the objective of maximizing a student's knowledge ("do the best") using as little of the student's resources as possible ("do not harm"), i.e., performing a minimal number of exercises:

Definition 1. *Consider a tutoring system employing a tutoring policy $\pi \in \Pi$. The tutoring policy π is equitable for student s if and only if*

$$T^s(\pi) = \min_{\pi' \in \Pi, L^s(\pi') = K} T^s(\pi') \quad \text{and} \quad L^s(\pi) = K,$$

where K is the total number of skills. A tutoring system is equitable if its tutoring policy is equitable for all students $s \in \mathcal{S}$.

Thus, informally, a tutoring system is equitable if it can teach all K skills in the minimal possible amount of time to any student. Please note that our notion of equity is strongly related to that introduced in [3] (cf. discussion below).

In the above definition, we assume that all students can master all K skills.[1] Importantly, a tutoring system can only be equitable if it is adaptive to the students which are interacting with it. In particular, it has to individualize the assignment of exercises and needs to carefully select the "stop action", in order to achieve equity. The above definition describes an idealized notion of equity which in general cannot be achieved as the tutoring policy would have to be

---

[1]Otherwise, if a student existed that was not able to acquire all skills, no tutoring system would be equitable according to Definition 1. Our definition can be generalized in a straightforward manner to account for an individual student's maximal achievable knowledge.

informed by a kind of "oracle" in order to teach using the optimal policy within Π right from the beginning.

Nevertheless, we can compare tutoring policies π in the spirit of the above definition. In particular, given two tutoring policies π and π' which both teach the same number of skills, we consider the policy π to be *more equitable* as compared to π' if for all students $s \in \mathcal{S}$ it holds that $T^s(\pi) \leq T^s(\pi')$.

We note that our notion of equity is strongly related to that introduced in [3]. In [3], the authors "assume that an equitable outcome is when students from different demographics reach the same level of knowledge after receiving instruction". As such, there is no need to reach this level of knowledge fast; hence this basic notion of equity is in principle trivial to achieve by providing vast amounts of exercises (as also noted by the authors of [3]). The desideratum of achieving knowledge fast is later added on top of their notion of equity whereas in our case it is a fundamental constituent in the first place. Furthermore, our interest extends to downstream implications of such a definition of equity, namely the individualization of knowledge tracing.

4.2 Theoretical Implications

Our definition of equity immediately leads to the (probably obvious) but important observation, that equity of a tutoring system can in general only be achieved by individualization:

Observation 1. *A tutoring system for a population of students with different learning characteristics can only be equitable if its tutoring policy is adaptive to the students.*

Thus, we note that if the tutoring policy is derived deterministically from a non-adpative model of the students, the tutoring system will in general not be equitable.

Furthermore, as already mentioned, our definition of equity is in general impossible to achieve. In particular, achieving equity would require basing a policy on rich side information in order to employ an optimal tutoring policy for any particular student right from the beginning. Such side information might not be available or insufficient to determine the optimal policy.

4.3 Operationalization

In practice, tutoring policies are often either simple fixed strategies, e.g., repeating exercising a skill until a certain number of consecutive exercises related to that skill were solved successfully, or derived from a model, e.g., a BKT model, such that each knowledge component is repeatedly exercised until it is mastered with a certain probability.

Tutoring policies based on incorrect or non-adaptive models can result in a student not acquiring all skills or suggest to perform too many practicing opportunities. Thus the following two general directions are important for building equitable tutoring systems:

1. **Using side information.** To start the interaction with a student as effectively as possible, any available side information about a student should be used to individualize the underlying models. For instance, in the
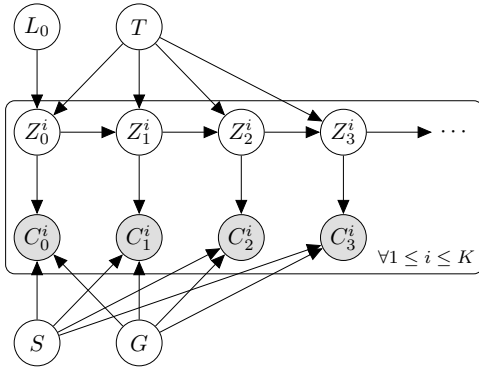
Figure 1: Representation of Bayesian-Bayesian Knowledge Tracing as a probabilistic graphical model. The acquisition and application of K skills is controlled by global parameters $p(L_0)$, $p(S)$, $p(G)$, and $p(T)$, i.e., they apply to all K skills.

context of RNN based DKT models, the side information could be used to predict the initial hidden states of the RNNs before the interaction with the student starts. In the context of classical BKT models, the side information could be used to make an initial guess about the key parameters of the model (initial knowledge, transition, slip and guess probabilities). Unfortunately, we are not aware of publicly available datasets which come with such side information and therefore it is difficult to develop such approaches.

2. **Online adaptation.** Even when using side information, a model is likely not perfectly individualized to all students it encounters. To further adjust the models in such cases, online adaption of the models during interaction seems promising. This can for instance be operationalized by training models for different subpopulations of students and during usage decide on which model is best fitted for a particular student. An alternative approach, which we propose in this paper, is a Bayesian BKT model, which performs posterior parameter inference about students' properties and leverages the results for further predictions.

5. BAYESIAN-BAYESIAN KNOWLEDGE TRACING

As outlined in the previous section, an equitable tutoring system must employ policies that are adaptive to students' properties. As these properties might not be known before actually interacting with a student, online adaption during interaction is important. To this end, we propose a Bayesian variant of the classical BKT model from which individualized policies can be derived, cf. Figure 1.

In particular, we assume that each student s has its own learning dynamics, described by student-specific parameters θ^s. If the learning dynamics can be described using a BKT model, which we assume in the following, $\theta^s = (p(L_0^s), p(T^s), p(S^s), p(G^s))$. We assume these learning dynamics to apply for the acquisition of *all* skills, i.e., they are *global* parameters of a student. If we knew these parameters for a student, we could use them in a classical BKT model and would likely observe good tutoring performance. How-

ever, in practice, we don't know these parameters and need to infer them. To this end, we take a Bayesian approach. Therefore, we assume a set of possible parameters Θ such that $\theta^s \in \Theta$ and a prior distribution $p_0(\theta^s)$. Based on t observations of a student practicing exercises collected in data set $\mathcal{D}_t$, we can compute the probability that a student has mastered a specific skill and base tutoring policies thereon. Note that as we don't know θ^s, this requires marginalizing out the (unknown) parameters θ^s. In this way the different possible parameters and their influence for predicting the knowledge state get re-weighted according to the available data. In particular, we need to compute

$$p(Z_t^{s,i} \mid \mathcal{D}_t) = \int_{\theta \in \Theta} \underbrace{p(Z_t^{s,i} \mid \theta, \mathcal{D}_t)}_{=:(\#1)} \underbrace{p(\theta \mid \mathcal{D}_t)}_{=:(\#2)} \, \mathrm{d}\theta, \qquad (1)$$

where $Z_t^{s,i}$ is a random variable indicating whether skill i is mastered at time t by student s. For only a few possible parameters θ, the above equation can be solved exactly by enumeration and by observing that both terms (#1) and (#2) can be computed efficiently by leveraging the forward algorithm through the following recursion:

$$\alpha_0^\theta(1) = p(Z_0^{s,i} = 1 \mid \theta) = p(L_0)$$
$$\alpha_0^\theta(0) = p(Z_0^{s,i} = 0 \mid \theta) = 1 - p(L_0)$$
$$\alpha_{t+1}^\theta(l) = p(Z_{t+1}^{s,i} = l, \mathbf{c}_{t+1}^i) = \sum_{z_t^{s,i}} p(Z_{t+1}^{s,i} = l, Z_t^{s,i} = z_t^{s,i}, \mathbf{c}_{t+1}^i)$$
$$= \sum_{z_t^{s,i}} p(c_{t+1}^i \mid Z_{t+1}^{s,i} = l) p(Z_{t+1}^{s,i} = l \mid Z_t^{s,i} = z_t^{s,i}) \alpha_t(z_t^{s,i})$$

where $\mathbf{c}_t^i$ collects all observations with respect to practicing the ith skill up to time t, and $c_{t'}^i$ is the t'th entry of $\mathbf{c}_t^i$. Then

$$(\#1) = p(Z_t^{s,i} = 1 \mid \theta, \mathcal{D}_t) = \frac{\alpha_t^\theta(1)}{\alpha_t^\theta(0) + \alpha_t^\theta(1)}, \text{ and}$$

$$(\#2) = p(\Theta = \theta \mid \mathcal{D}_t) = \frac{p_0(\theta) \cdot (\alpha_t^\theta(0) + \alpha_t^\theta(1))}{\sum_{\theta' \in \Theta} p_0(\theta') \cdot (\alpha_t^{\theta'}(0) + \alpha_t^{\theta'}(1))}.$$

If Θ is large/continuous MCMC-techniques can be used to approximate Equation (1).

In our experiments, we mainly consider a concrete instantiation of our proposed model in which we consider the learning rate to be unknown. This is motivated by previous work which has identified the learning rate as a key parameter for improving BKT based models in terms of AUC [18].

Advantages of B^2KT. Through posterior parameter inference based on observed evidence, the proposed model can often yield (approximately) equitable tutoring policies. In contrast, classical BKT models cannot adjust to the students' properties, possibly resulting in inequitable policies.

Furthermore, Bayesian treatment of the parameters can enable us to leverage ideas from active learning to actively seek information about students with the goal of teaching more effectively later. This could be a clear advantage over other approaches for individualization, e.g., those using EM algorithms — in such a case it is not obvious how to actively seek information about a students' parameters. We did not explore this idea in this paper but consider this as a promising direction for future work.

6. EXPERIMENTS

In this section, we compare tutoring policies derived from B^{2}KT, BKT, and DKT models in terms of equity.

6.1 Experimental Setups

In all presented results we denote the average stopping time of a policy for a population of students by T_{stop} and the average number of acquired skills by % skills.

Curricula. We consider the following curricula:

- THRESHOLD(τ): These curricula are based on knowledge tracing algorithms which allow computing the posterior probability of skill mastery. In particular, these curricula repeatedly exercise a skill until it is mastered with a probability of at least τ.

- EXPECTIMAX(1): These curricula are derived from DKT models which do not enable us to compute a posterior probability of skill mastery. Skills that maximize the average probability that a practicing opportunity for a randomly selected skill at the next step would be solved correctly are practiced. The stop action is invoked when the predicted probabilities change by less than a threshold from one step to the next.

Models. We consider the following models:

- BKT: the classical BKT models with fixed parameters.

- B^{2}KT: the proposed Bayesian BKT model.

- DKT: deep knowledge tracing models [14] based on the implementation provided by [4].

6.2 Experimental Results

Students with different learning behaviors. We study the equity of tutoring policies when the students are sampled uniformly from two groups, each containing students with learning dynamics described by a ground truth BKT model. In particular, we build on the experimental setup from [3] where there is a group of slow learners (BKT slow) and fast learners (BKT fast). In [3], the authors also fitted a BKT model to interaction data from students from both groups; we refer to the corresponding BKT model as BKT mixed. The parameters of these models are as follows:

	$p(L_0^s)$	$p(S^s)$	$p(G^s)$	$p(T^s)$
BKT slow	0.0	0.2	0.2	0.05
BKT fast	0.0	0.2	0.2	0.3
BKT mixed	0.071	0.203	0.209	0.096

We considered the interaction with 400 students, 200 from each group, and we compared the performance of THRESHOLD(0.95) tutoring policies based on the different BKT models for different numbers of skills that ought to be taught in Table 1. We observe that in the case of mismatch of the student properties and the BKT models used for the threshold policy, either only a small fraction of the skills (clearly below 95 %) is acquired or that more than necessary time is spent exercising. Furthermore, in order to teach the desired 95 % of the skills to slow students when basing the policy on the BKT model for fast students, the threshold τ would have to be increased significantly, thereby providing
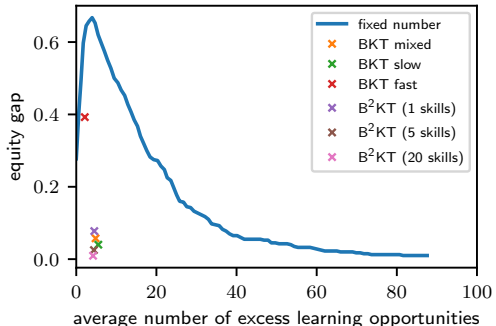

Figure 2: Equity gap versus number of excess learning opportunities for different policies. We use a Threshold(0.95) tutoring policy derived from the knowledge tracing models. B^{2}KT becomes more equitable as more skills are taught.

unnecessary extra exercises to fast students. For example, by setting the threshold τ to 0.998, the fraction of skills learned by the slow learners increased to 0.95 and the stopping time to $T_{\text{stop}} = 464.08$ (much worse as compared to the THRESHOLD(0.95) policy for the correct model). At the same time, the stopping time for the fast learners increased to $T_{\text{stop}} = 174.72$. The mismatch issue is alleviated in the case of the B^{2}KT model (assuming a uniform prior over both types of students), in particular for a larger number of skills. Intuitively this is because, in the case of multiple skills, the model has more opportunities to learn about the students' properties and use these insights in later tutoring. This is an important characteristic of a model with respect to equity as the model becomes "more equitable" as it learns about the students. This fact is also illustrated in Figure 2 in which we reproduce and extend an experiment from [3] in which they compare the "equity gap" (the difference in the percentage of skills mastered by fast students to the percentage of skills mastered by slow students) to the number of excess learning opportunities. In particular, we extended the figure to the case of teaching multiple skills and show the performance of B^{2}KT. Importantly, B^{2}KT becomes more equitable (the corresponding point moves to lower left) as more skills are taught.

Out-of-distribution generalization. We evaluate whether B^{2}KT can help with aspects relevant to *inclusion and diversity*, the first two principles of IDEA. In particular, we consider a stylized mismatch setting in which a tutoring system interacts with students who have a learning behavior that was not encountered in the data that was used for constructing the BKT models. In particular, in addition to the previous two types of students, we assume a third type of learner (BKT med) with the following parameters: $p(L_0^s) = 0.0, p(S^s) = 0.2, p(G^s) = 0.2, p(T^s) = 0.18$. We considered THRESHOLD(0.95) policies based on BKT models of slow and fast learners and the B^{2}KT model with a uniform prior over slow and fast learners. Our results are presented in Table 2. We observe that the performance of the policies derived from the B^{2}KT model have comparable performance to those derived from the true model whereas other models yield policies clearly worse in terms of stopping at the right time or teaching the right amount of skills. Although the correct model has 0 prior probability (and thus also zero posterior probability) in the B^{2}KT model,

Table 1: Equity trade-offs of curricula derived from different models/parameterizations.

| | 1 skill | | | | 5 skills | | | | 20 skills | | | |
| | slow learners | | fast learners | | slow learners | | fast learners | | slow learners | | fast learners | |
THRESHOLD(0.95)	% skills	T_{stop}	% skills	T_{stop}	% skills	T_{stop}	% skills	T_{stop}	% skills	T_{stop}	% skills	T_{stop}
BKT slow	97.00	24.14	99.50	9.49	97.20	122.80	99.90	66.00	97.55	492.64	99.90	183.84
BKT fast	61.00	13.85	97.50	5.96	62.60	71.98	96.10	29.81	64.20	288.76	97.23	120.59
BKT mixed	95.00	23.51	100.00	8.33	95.40	113.67	99.90	40.93	94.53	466.55	99.68	169.86
B^2KT	94.50	24.04	100.00	7.88	97.70	120.87	98.40	32.61	96.68	493.00	96.66	120.05

the model performs well. Such a characteristic can be of key importance for promoting inclusion, e.g., when interacting with students who were not represented in a data set that was used for building an intelligent tutoring system or students who were underrepresented in that data.

Fair next step predictions do not necessarily imply equitable tutoring. We show empirically that models which might appear fair when looking at their AUCs for different groups of students do not necessarily yield equitable tutoring policies. In particular, we again focus on a student population consisting of two groups of learners:

	$p(L_0^s)$	$p(S^s)$	$p(G^s)$	$p(T^s)$
Group 1	0.0	0.1	0.4	0.1
Group 2	0.0	0.1	0.2	0.3

We generated data for 400 students in a setting with 20 skills and 1000 random exercises from BKT models, where 50% of the students are from group 1 and group 2, respectively. The true model of group 1's students achieved an AUC of 0.7393 for group 1's students, while the true model of group 2's students achieved an AUC of 0.6710 for group 2's students.

Looking only at the AUC, the two models appear rather inequitable (there is no group parity). Thus it might appear sensible to aim to use a BKT model for tutoring which has comparable AUCs for both groups in order to promote equity. For instance, a BKT model using parameters $p(L_0) = 0$, $p(S) = 0.4$, $p(G) = 0.1$, $p(T) = 0.65$ achieves an AUC of 0.6719 on group 1's students and of 0.6733 on group 2's students, respectively. That is, the AUCs on the two groups are approximately equal.

However, when looking at the different models with respect to the tutoring performance when using a THRESHOLD(0.95)-policy, we observe a very different picture. We summarize our findings in Table 3. In particular, the fraction of skills taught differs significantly between the two groups: In group 1 only 28.68% of the skills are acquired by the students on average while in group 2 74.70% of the skills are acquired on average. This finding is closely related to the observation that models with greatly different characteristics can have similar AUCs [13].

In this context, we also evaluate EXPECTIMAX(1) policies based on DKT models (trained using the implementation provided by [4]; using default settings for hidden sizes, etc.). The AUCs of the DKT model are 0.6757 and 0.6551 for slow and fast students respectively and so roughly comparable to that of the BKT models appearing to be equitable in terms of their AUC. However, the performance of EXPECTIMAX(1) policies we observed was very low (although the average

predicted probability for answering questions correctly increased monotonically in most cases). In particular, for a threshold $\tau = 0.001$ it achieved a surprisingly low performance of 8.3 % skills / $T_{\text{stop}} = 41.01$ (slow learners) and 10 % skills / $T_{\text{stop}} = 29.80$ (fast learners). Interestingly, a difference in stopping time for slow and fast learners could be observed despite the model making overall bad recommendations. When looking into this issue in more detail we observed that the policies repeatedly taught the same skill, i.e., DKT did not correctly model the learning process— it appeared to only have picked up that the probability of answering exercises correctly increases over time.

Table 2: Out-of-distribution generalization.

| | 1 skill | | 5 skills | | 20 skills | |
| | BKT med | | BKT med | | BKT med | |
THRESHOLD(0.95)	% skills	T_{stop}	% skills	T_{stop}	% skills	T_{stop}
BKT slow	99.50	11.28	99.55	56.45	99.59	225.25
BKT fast	90.75	7.61	91.55	37.59	91.70	151.36
BKT mixed	99.50	10.46	99.00	51.71	99.21	211.29
BKT med	98.25	8.82	97.35	45.59	97.84	184.20
B^2KT	98.75	10.33	97.50	48.80	94.19	168.36

Table 3: Fairness in terms of similar AUCs on different groups does not imply fairness in terms of the derived teaching curricula. See the main text for details.

| | group fair wrt AUC | | | true model wrt group | | |
group	AUC	% skills	T_{stop}	AUC	% skills	T_{stop}
group 1	0.6719	28.68	61	0.7393	96.13	308
group 2	0.6733	74.70	64	0.6710	96.35	105

7. CONCLUSIONS & FUTURE WORK

We considered the equity and fairness of curricula derived from knowledge tracing models. We provided a unifying definition of an equitable tutoring system motivated by the ethical principles of beneficence and non-maleficence. Our definition is, in many practical settings, not realizable but suggests that the individualization of tutoring policies to students is key for realizing equity. We proposed the B^2KT model, a Bayesian variant of the classical BKT model, and demonstrated in various experiments that it can be beneficial for realizing equitable tutoring systems and promoting IDEA more generally. Furthermore, we highlighted that developing models with the sole focus on (fair) next-step predictions might be insufficient to develop equitable tutoring systems. In future work we aim to build novel models and policies which allow individualization through leveraging side information and improve scalability of B^2KT by approximate inference techniques.

8. REFERENCES

[1] T. L. Beauchamp, J. F. Childress, et al. *Principles of biomedical ethics*. Oxford University Press, USA, 2001.

[2] A. T. Corbett and J. R. Anderson. Knowledge tracing: Modeling the acquisition of procedural knowledge. *User modeling and user-adapted interaction*, 4(4):253–278, 1994.

[3] 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*, pages 335–339, 2019.

[4] T. Gervet, K. Koedinger, J. Schneider, T. Mitchell, et al. When is deep learning the best approach to knowledge tracing? *Journal of Educational Data Mining*, 12(3):31–54, 2020.

[5] J. He-Yueya and A. Singla. Quizzing policy using reinforcement learning for inferring the student knowledge state. *International Educational Data Mining Society*, 2021.

[6] T. Käser, S. Klingler, and M. Gross. When to stop? towards universal instructional policies. In *Proceedings of the sixth international conference on learning analytics & knowledge*, pages 289–298, 2016.

[7] R. F. Kizilcec and H. Lee. Algorithmic fairness in education. *arXiv preprint arXiv:2007.05443*, 2020.

[8] H. Lee and R. F. Kizilcec. Evaluation of fairness trade-offs in predicting student success. *arXiv preprint arXiv:2007.00088*, 2020.

[9] J. I. Lee and E. Brunskill. The impact on individualizing student models on necessary practice opportunities. *International Educational Data Mining Society*, 2012.

[10] Z. A. Pardos and N. T. Heffernan. Modeling individualization in a bayesian networks implementation of knowledge tracing. In *International conference on user modeling, adaptation, and personalization*, pages 255–266. Springer, 2010.

[11] G. Paviotti, P. G. Rossi, and D. Zarka. Intelligent tutoring systems: an overview. *Pensa Multimedia*, pages 1–176, 2012.

[12] R. Pelánek. Conceptual issues in mastery criteria: Differentiating uncertainty and degrees of knowledge. In *International Conference on Artificial Intelligence in Education*, pages 450–461. Springer, 2018.

[13] R. Pelánek. The details matter: methodological nuances in the evaluation of student models. *User Modeling and User-Adapted Interaction*, 28(3):207–235, 2018.

[14] C. Piech, J. Bassen, J. Huang, S. Ganguli, M. Sahami, L. J. Guibas, and J. Sohl-Dickstein. Deep knowledge tracing. *Advances in Neural Information Processing Systems (NeurIPS)*, 28, 2015.

[15] A. Singla, A. N. Rafferty, G. Radanovic, and N. T. Heffernan. Reinforcement learning for education: Opportunities and challenges. *arXiv preprint arXiv:2107.08828*, 2021.

[16] Z. Wang, S. Tschiatschek, S. Woodhead, J. M. Hernández-Lobato, S. Peyton Jones, R. G. Baraniuk, and C. Zhang. Educational question mining at scale: Prediction, analysis and personalization. *AAAI Conference on Artificial Intelligence*, 35(17):15669–15677, 2021.

[17] R. Yu, Q. Li, C. Fischer, S. Doroudi, and D. Xu. Towards accurate and fair prediction of college success: Evaluating different sources of student data. *International Educational Data Mining Society*, 2020.

[18] M. V. Yudelson, K. R. Koedinger, and G. J. Gordon. Individualized bayesian knowledge tracing models. In *International Conference on Artificial Intelligence in Education*, pages 171–180. Springer, 2013.