From Prediction to Policy in Intelligent Tutoring Systems

Joseph Rollinson (jrollins@andrew.cmu.edu), Advisor: Emma Brunskill (ebrun@cs.cmu.edu)
Carnegie Mellon University, School of Computer Science

Overview

Intelligent tutoring systems (ITS) provide individualized and responsive instruction to students. ITS consist of two parts; a student model and instructional policies that use the student model to make decisions. Despite much research on student modeling, there has been relatively little on automated instructional policies. To address this, we contribute:

- a model agnostic when-to-stop policy.
- a model agnostic skill-choice policy.

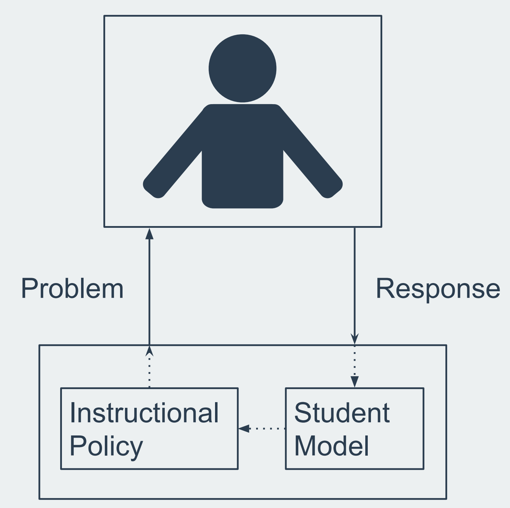

Background: Student Models

Student models can also be used for predicting student performance. Here we provide two frequently used student models.

Bayesian Knowledge Tracing[1]

Bayesian knowledge tracing treats students as a stochastic process. The parameters of BKT are:

- $P(L_0)$: Initial Probability of mastery
- $P(T)$: Probability of transitioning to mastery
- $P(G)$: Probability of correct answer when learning (guess)
- $P(S)$: Probability of incorrect answer when mastered (slip)

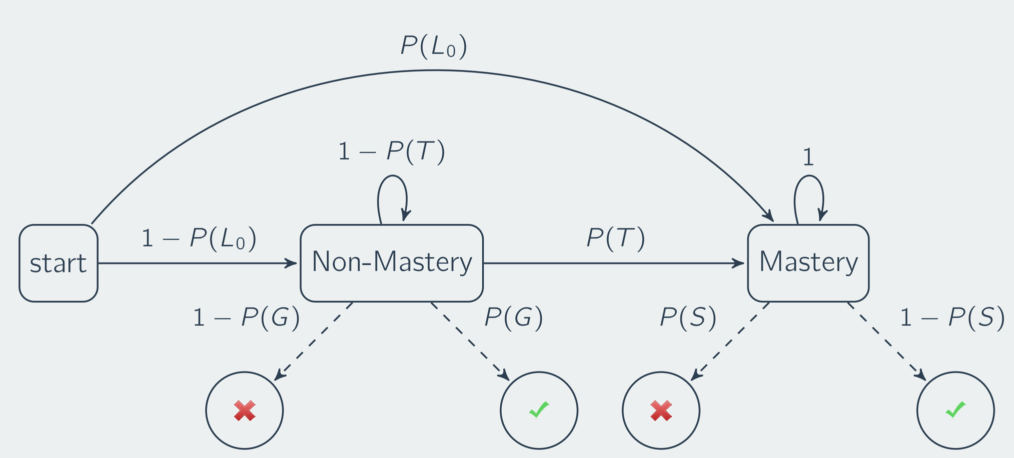

BKT is frequently used with instructional policies, because it allows us to quantify our certainty that a student has mastered a skill.

Logistic Regression

Logistic regression student models combine features to calculate the probability of the student responding correctly to the next question. Performance factors models[2] are a type of logistic regression model that uses the counts of both correct (s) and incorrect (f) responses by the student (i) to the skill (k). It's weights are:

- α_i : Student aptitude
- β_k : Skill difficulty
- μ_k : Correct response count learning rate
- ρ_k : Incorrect response count learning rate

The probability of the student responding correctly to the next question is:

$$P(C) = \frac{1}{1 - e^{-(\alpha_i + \beta_k + \mu_k s + \rho_k f)}}$$

Logistic regression models are popular for student prediction, but are rarely used in decision making.

When-to-stop Problem

When teaching a single skill, intelligent tutoring systems must decide when to stop providing questions to the student. The goal is to ensure that the student masters the skill without needlessly wasting the student's time.

Prior Work: Mastery Threshold

The mastery threshold policy stops when the probability of mastery is above a given threshold Δ. There are two main issues with this policy.

1. It requires a student model that can provide the probability that the student has mastered the skill.
2. It will teach a student indefinitely if they are unable to learn the skill with the given instruction (wheel-spinning)[3].

Predictive Similarity Policy

Our new policy stops when the system is confident that providing another question to the student will not significantly change the probability that the student will answer a question correctly. Mathematically, the policy stops when:

$$P\left(\left|P(C_t) - P(C_{t+1})\right| < \epsilon\right) > \delta$$

This policy will work with any predictive student model, because it only uses the model's predictions. Also, it will stop if the student is unable to learn given the available activities, because the probability of a correct response in these situations does not change.

Comparing Policies

We used ExpOps[4] values (expected number of steps before the given policy stops providing problems to the student) to compare policies. Using skills in the KDD Cup dataset[5], we compared both policies using a BKT and both student models under the predictive similarity policy.

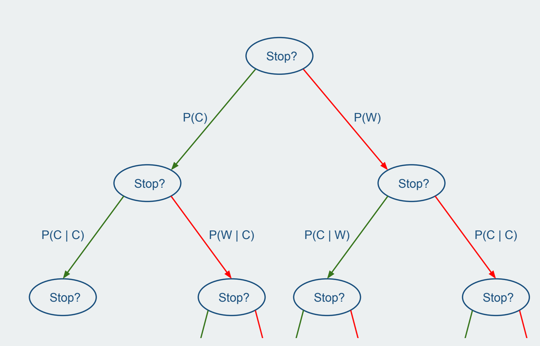

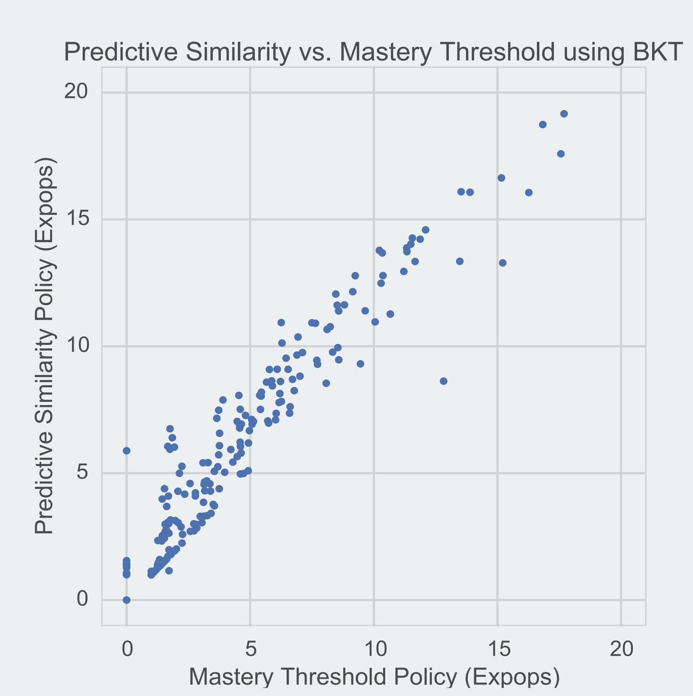

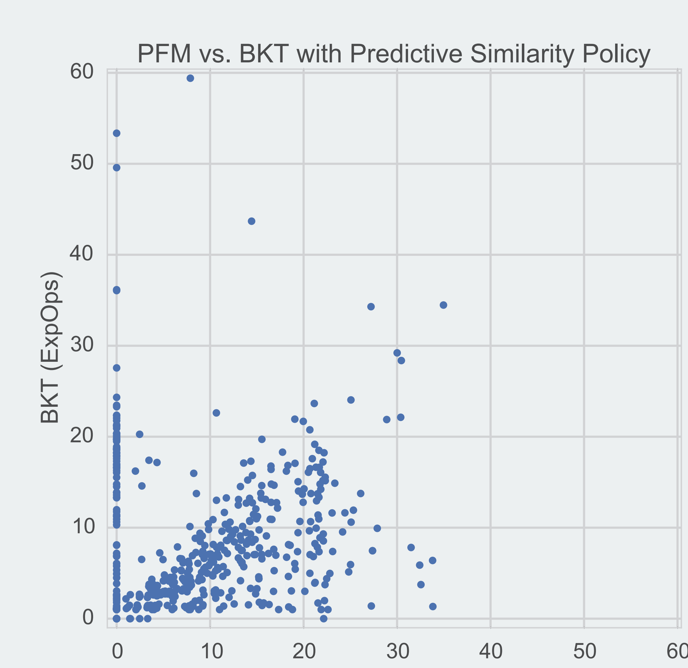

We found that the predictive similarity policy performed similarly to the mastery threshold policy and that the PFM based predictive similarity policy tended to either stop immediately or provide more questions than BKT based predictive similarity policy.

Skill-choice Problem

Often intelligent tutoring systems teach multiple skills within a budget of T problems. If the skills are independent, then this problem is simple. However, skills frequently have a prerequisite structure that may not be immediately obvious.

Horizon-based Policies

An optimal way of deciding which skill to teach is to consider all skills and how the student might respond to each skill repeatedly out to T problems. We then calculate which skill has the best expected score. Although optimal, this full-horizon policy can be computationally prohibitive.

A greedy policy picks the skill that will maximize immediate progress. It sacrifices the promise of optimality for computational efficiency.

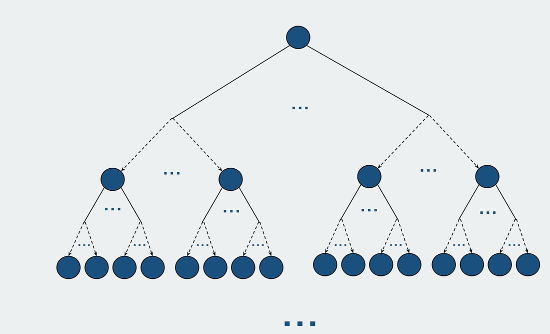

Comparing Policies

Using student observations simulated from different skill hierarchies, we compared the expectation of mastery of skills for both policies with multiple budgets. We developed a hierarchical form of PFM to use as the underlying model.

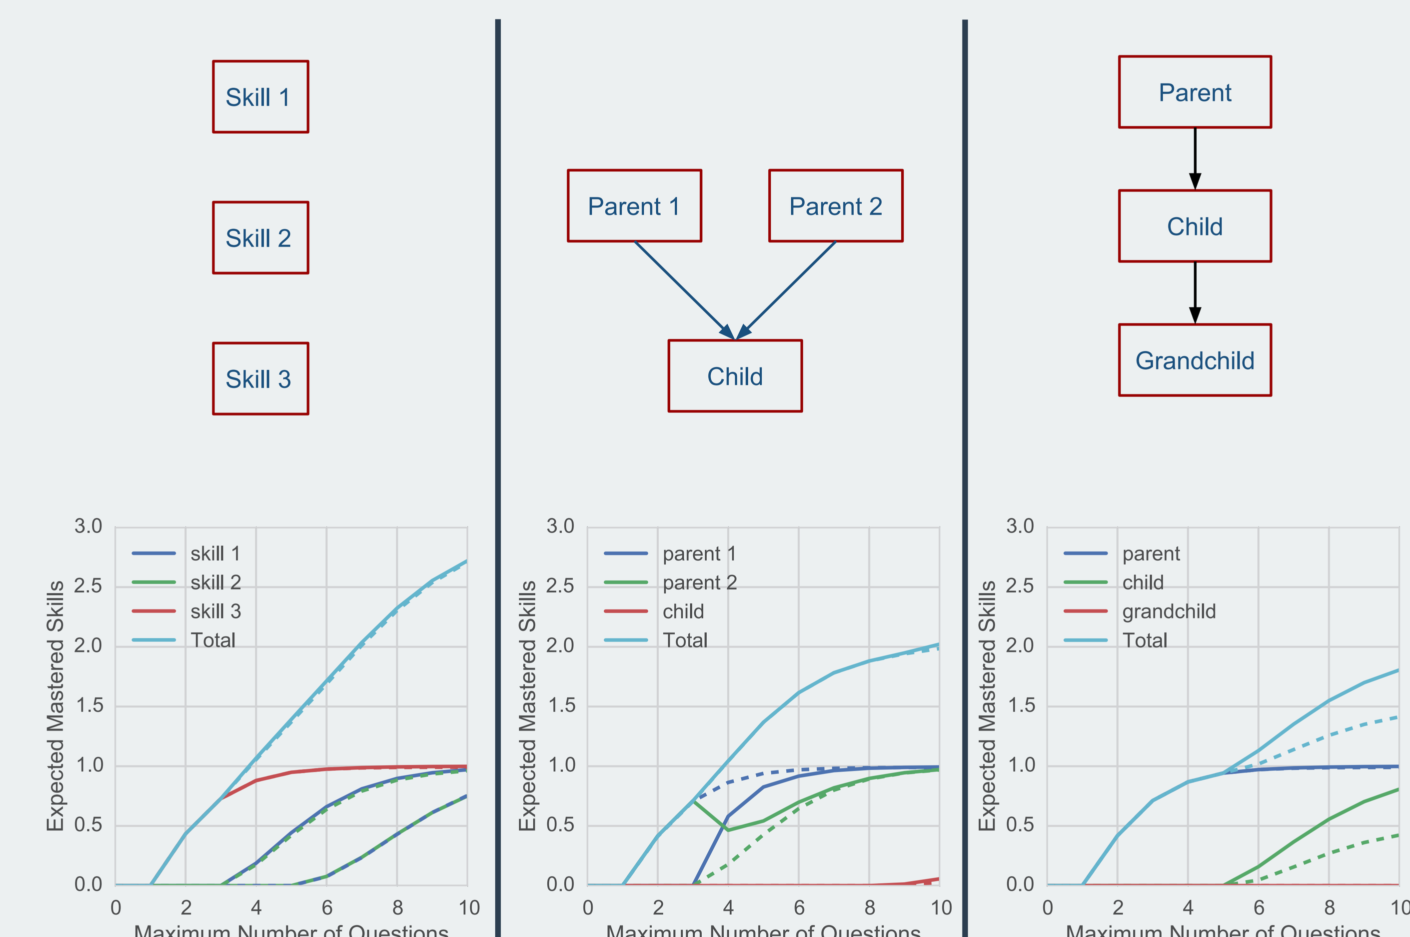

Note: dashed lines are for the greedy policy and full lines are for full-horizon policy

On independent skills the policies acted the same, but on dependent skills the full-horizon policy performed better than the greedy policy.

Future Directions

- Test skill-choice results on data collected from real students.
- Compare greedy and full-horizon policies on more complicated skill hierarchies.
- Investigate how to evaluate instructional policies using student data.

1. Corbett, Albert T., and John R. Anderson. "Knowledge tracing: Modeling the acquisition of procedural knowledge." User modeling and user-adapted interaction 4.4 (1994): 253-278.
2. Pavlik Jr, Philip I., Hao Cen, and Kenneth R. Koedinger. "Performance Factors Analysis--A New Alternative to Knowledge Tracing." Online Submission (2009).
3. Beck, Joseph E., and Yue Gong. "Wheel-spinning: Students who fail to master a skill." Artificial Intelligence in Education. Springer Berlin Heidelberg, 2013.
4. Lee, Jung In, and Emma Brunskill. "The Impact on Individualizing Student Models on Necessary Practice Opportunities." International Educational Data Mining Society (2012).
5. J. Stamper, A. Niculescu-Mizil, S. Ritter, G. Gordon, and K. Koedinger. Algebra i 2008-2009. challenge data set from kdd cup 2010 educational data mining challenge.