

Diagnosing Algebra Understanding via Inverse Planning

Anna N. Rafferty
Computer Science Division
University of California, Berkeley, CA 94720 USA
rafferty@cs.berkeley.edu

Thomas L. Griffiths
Department of Psychology
University of California, Berkeley, CA 94720 USA
tom_griffiths@berkeley.edu

ABSTRACT

Students' solution processes can offer significant insight into their misunderstandings. However, freeform solutions can be difficult to interpret, leading many educational technologies to examine only students' final answers or to structure problem solving to make it more interpretable. We develop a new approach using Bayesian inverse planning for diagnosing algebra skills that interprets students' step-by-step problem solving, placing no restrictions on how students transform equations to reach solutions. We formalize understanding as several distinct skills, allowing us to identify the causes of errors; for instance, arithmetic skills are separated from systematic misapplications of algebra rules. In simulation, the algorithm recovers the true parameters of simulated learners relatively accurately. Using human data, we show that the algorithm can interpret over 98% of people's actions and its inferences about arithmetic skills are consistent with an assessment of arithmetic ability in isolation. Our work demonstrates that Bayesian inverse planning can successfully scale to the space of algebra and provides new technical solutions that may be relevant in other complex educational domains.

1. INTRODUCTION

The way that students approach and solve problems can provide significant insights into their understanding. Classroom teachers encourage students to "show [their] work," allowing them to gain insight into the students' difficulties. In principle, computers should be able to make such inferences automatically, drawing fine-grained inferences about students' skills based on their choices about problem-solving and the types of errors they make. However, many automated tutoring systems cannot interpret students' worked solutions. We focus on the case of teaching and remediating algebraic equation solving. Existing computer-based systems for helping students learn algebra typically take one of two approaches to assessing and modeling students' algebra understanding. They may structure the problems such that students enter their work in discrete parts, with each part corresponding to a different algebra skill; students gen-

erally must enter the current part correctly prior to moving on. Alternatively, these systems may use only final answers to infer understanding, with many systems solely checking whether an answer is correct.

While it may be pedagogically useful in some cases to structure students' behavior or interrupt their work to point out errors, it should not be necessary to do these things to infer students' understanding from their worked solutions. We develop a Bayesian inverse planning model that can diagnose a student's understanding from observing how she solves linear equations. This approach allows data from multiple problems to be incorporated into the diagnosis, and naturally accounts for inconsistent behavior across problems. The model extends existing work on algebra understanding by using freeform problem solving behavior to diagnose what a student understands and in what ways she misunderstands without requiring that individual steps be correct before the student continues.

2. MODELING ALGEBRA SKILLS USING BAYESIAN INVERSE PLANNING

Bayesian inverse planning uses Markov decision processes (MDPs) to model how people plan their actions in order to achieve their goals, and infers a diagnosis of their understanding as a distribution over possible *hypotheses* [5]. For linear equation solving, the hypotheses represent the possible misunderstandings that people might have about solving algebraic equations and carrying out mathematical operations. By detecting patterns in a person's equation transformations, the algorithm can determine which hypothesis is most likely to represent the person's knowledge. To develop a Bayesian inverse planning algorithm for linear equation solving, we must define how to model equation solving as an MDP and specify the space of possible hypotheses.

MDPs provide a decision-theoretic method for modeling sequential action planning (for an overview, see [7]). MDPs are defined by the set of states that characterize the environment in which the agent is acting and the actions the agent may take. The transition model in an MDP defines the conditional probabilities distributions $p(s'|s, a)$ that the next state will be s' given that the current state is s and the chosen action is a . The reward model then encodes the goals and incentive structure for the agent; in this case, the goal of solving an equation, with fewer actions favored over more actions. We represent the state in linear equation solving as the list of terms on each side of the equation, ignoring the

ordering of these terms. The actions are the possible transformations that a student might apply to an equation. We include six types of actions: moving a term, dividing by a coefficient, multiplying by a constant, combining terms, distributing over a parenthesized term, and terminating solving. These actions are sufficient to represent typical problem solving behavior, with the final action occurring each time a student completes a problem or gives up.

To use Bayesian inverse planning to diagnose students' understanding, we must define the hypothesis space of possible knowledge states. In this case, the knowledge states correspond to possible transition models: how does a student believe the state should be transformed when she chooses a particular action? We represent each knowledge state as a vector θ of six parameter values. Four of these values relate to error tendencies in applying specific actions, based on mal-rules discovered in prior work [4, 6]. For example, one parameter represents the probability a student will make a *sign error* in which she moves a term from one side of the equation to the other without changing the sign: $2x + 3 = 6$ becomes $2x = 6 + 3$. The other two parameters relate to equation solving behaviors not tied to a specific action. The *arithmetic error* parameter is the probability that a student will make an arithmetic error in each operation in a transformation. Separating out this error term differentiates students who get problems wrong due to misunderstandings about the rules of algebra from students who have difficulties with arithmetic. The final parameter relates to the efficiency with which the student solves equations. Following prior work modeling human action planning [1, 5], we assume students choose their actions noisily optimally: $p(a|s) \propto \exp(\beta Q(s, a))$, where $Q(s, a)$ represents the long-term expected value of taking action a in state s . β controls the noisiness of the policy, with increasingly large β corresponding to more optimal action selection. We infer the value of β that best models an individual student's action planning, allowing us to detect how well a student is choosing her actions. Very low inferred values of β might also indicate that the student's data is not well fit by our model.

We represent the diagnosis of a student's algebra understanding as the posterior distribution $p(\theta|d_1, \dots, d_N)$ over the possible parameters given N observed problem solutions. This distribution can be calculated using Bayes' rule, making use of the fact that each θ corresponds to a particular MDP. Because θ contains continuous parameters, we approximate the posterior via Markov chain Monte Carlo sampling [3]. For each sample, we must calculate a Q -function of long-term expected values given the current θ . Since the state and action spaces are infinite, this function can also only be approximated. We use discretization to aggregate both the state and action spaces, a common strategy in large or continuous MDPs (e.g., [2, 7]).

3. DIAGNOSING UNDERSTANDING

To evaluate the effectiveness of Bayesian inverse planning for diagnosing algebra understanding, we first tested the framework in simulation. Simulations allowed us to assess whether there was sufficient information in the problem traces (i.e., the series of equations representing the transformations from the initial equation to the final solution) to recover the true parameters of a learner. We found that recovery of these

values was relatively accurate, with the median difference between actual and inferred values for five of the six parameters less than 0.1. The planning parameter had a median difference of less than 0.25; the larger difference is likely due to its increased range. This suggests that while the diagnosis computed by Bayesian inverse planning is approximate, it still provides accurate information about the learner and might be used to guide remediation.

We then evaluated our model's performance on human data by recruiting participants on Amazon Mechanical Turk to complete an online worksheet and solve twenty problems on the Berkeley Algebra Tutor website, which we designed to collect step-by-step equation solving data. Over 98% of equation transformations could be interpreted by our model, and through manual annotation of a subset of the equations, we found that the model's interpretation of the transformations were generally consistent with human observers' interpretations. By comparing the model's inferred parameter values for individual participants with these participant's worksheet performance, we found that the inferred arithmetic error parameter was correlated with scores on the arithmetic portion of the worksheet. These results demonstrated that the model can interpret real equation-solving data and compute diagnoses of individual algebra skills.

4. CONCLUSION

Developing a Bayesian inverse planning algorithm for algebra offers promise for both extending the scalability of the inverse planning framework and helping to remediate learners' algebra skills. Our work demonstrates that inverse planning can be applied to a complex educational domain, and provides solutions for common technical problems that may arise, such as infinite state and action spaces. Our simulation and experimental results suggest that the algorithm can effectively recover the parameters of simulated learners and interpret people's equation solving. We plan to further evaluate the model using data from current algebra learners and to test the effectiveness of personalized guidance for learners based on our diagnosis of their understanding.

Acknowledgements. This work was funded by an NaEd/Spencer Dissertation Fellowship to ANR and ONR MURI grant number N00014-13-1-0341 and NSF grant number IIS-0845410 to TLG.

5. REFERENCES

- [1] C. L. Baker, R. R. Saxe, and J. B. Tenenbaum. Action understanding as inverse planning. *Cognition*, 113(3):329–349, 2009.
- [2] L. Busoniu, R. Babuska, B. De Schutter, and D. Ernst. *Reinforcement learning and dynamic programming using function approximators*. CRC Press, Boca Raton, FL, 2010.
- [3] W. Gilks, S. Richardson, and D. J. Spiegelhalter, editors. *Markov chain Monte Carlo in practice*. Chapman and Hall, Suffolk, UK, 1996.
- [4] S. Payne and H. Squibb. Algebra mal-rules and cognitive accounts of error. *Cognitive Science*, 14(3):445–481, 1990.
- [5] A. N. Rafferty, M. M. LaMar, and T. L. Griffiths. Inferring learners' knowledge from their actions. *Cognitive Science*, in press.
- [6] D. Sleeman. An attempt to understand students' understanding of basic algebra. *Cognitive Science*, 8(4):387–412, 1984.
- [7] R. S. Sutton and A. G. Barto. *Reinforcement learning*. MIT Press, 1998.