Human-
Computer
Interaction
Institute

SimStudent: Improving Tutor Quality and Reducing Authoring Costs

Christopher J. MacLellan, Eliane Stampfer Wiese, Noboru Matsuda, and Kenneth R. Koedinger

Intelligent Tutors are Effective

- Multiple studies have confirmed their effectiveness (Koedinger et al. 1997; Ritter et al., 2007; Vanlehn et al., 2005; Murray, T., 2004)
  - Double the learning gains over traditional instruction (Pane et al. 2013)
  - More learning in half the time (Lovett et al., 2008)
  - Good (or better) learning in 25% of the time (Bowen et al., 2012)

Intelligent Tutors are Widely Used

- For example, more than 500k students per year complete a Carnegie Learning tutor course

But, they are not generally adopted…

- Online learning platforms such as Khan Academy, Coursera, etc. do not use them
- Why?
  - Cost of tutor development too high?
  - Learning gains not enough to outweigh costs?

Our Approach: SimStudent Authoring

- Interested in improving both sides of this cost-benefit equation
  - Building higher quality tutors that lead to more robust learning by building Cognitive Science theory into the authoring tool
  - While simultaneously decreasing authoring costs

Model Tracing Paradigm

- The expert model defines the actions that can be taken on a given state (usually in the form of if-then rules)
  - Whenever an action is taken, if the model suggests this action, then it is marked as correct
  - Otherwise, it is marked as incorrect

Example Tutor Rule

- If one side of an equation has a single variable term with a coefficient, then divide both sides by that coefficient
  - 3x = 12 -> divide 3 — CORRECT
  - 3x = 12 -> subtract 3 — INCORRECT

Cognitive Tutor Authoring Tools

- Tools for building interfaces via a drag-and-drop editor
- Tools for authoring cognitive models (if-then rules) by demonstration
  - Example-Tracing
  - SimStudent
- Tools for deploying tutors (offline or online)
  - See: https://mathtutor.web.cmu.edu/

Example-Tracing

- A domain expert solves problems directly in the tutor interface
  - The system learns state to action mappings from author demonstrations
  - Produced model only works for the specific problems authored
    - There exist some limited generalization techniques, such as mass production
  - Easy to produce (demo to come)

SimStudent

- Uses machine learning to create general production rules from demonstrations and feedback

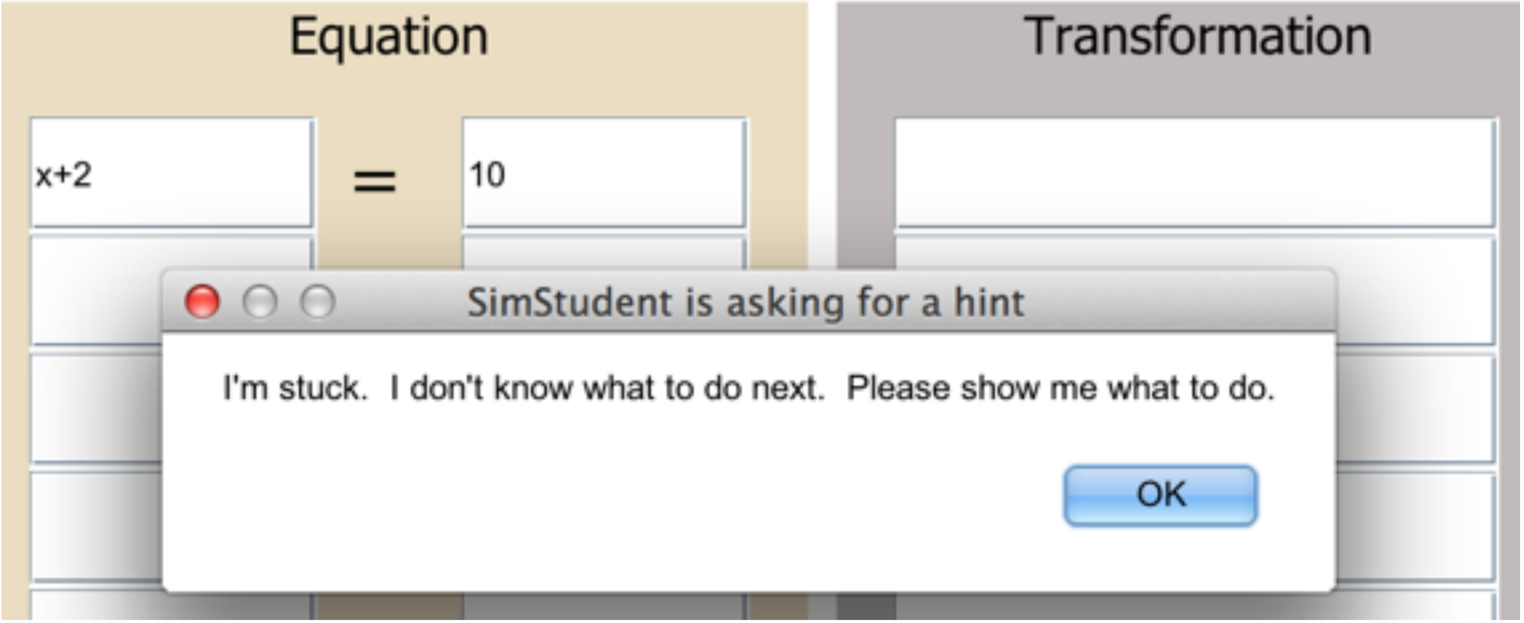

SimStudent Architecture

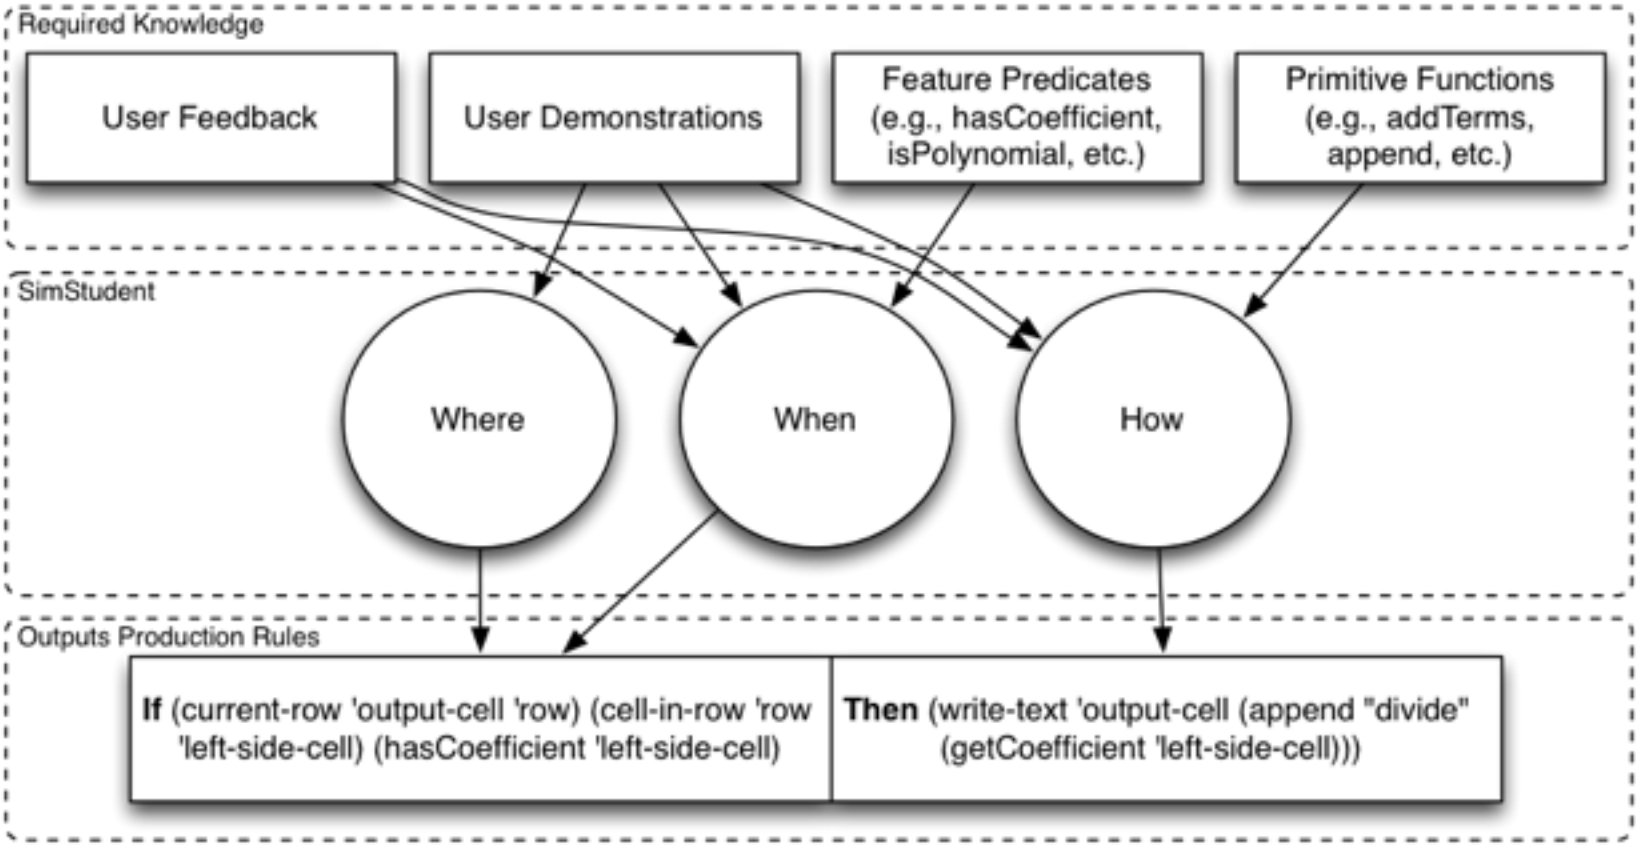

SimStudent Theory of Learning

- Able to identify knowledge that human experts miss
  - For example, -1x = 3 vs. -x =3
- Can make predictions about generalization errors that students are likely to make
  - Can be used to identify “buggy” rules

Demo

- Using SimStudent to author a basic Algebra tutor
  - Show interface builder
  - Look at background knowledge
  - Load interface and SimStudent
  - Author expert model

Discussion - Better Learning

- SimStudent has been shown to produce better cognitive models than experts (Li et al., 2013)
- SimStudent can learn incorrect productions due to incorrect induction from background knowledge (Matsuda et al., 2010)
  - These incorrect rules are plausible "bug" rules that might be missed by experts
- Better cognitive models have been shown to improve student learning (Koedinger et al., 2013)

Human-Computer Interaction Institute

Discussion - Lower Costs

- Example-Tracing tutors reduce authoring time by as much as 50% over hand-authoring cognitive tutors (Aleven et al., 2009)
- SimStudent has been shown to decrease authoring time compared to Example-Tracing tutoring, once background knowledge has been authored (MacLellan et al., 2014)

Human-
Computer
Interaction
Institute

Future Work

- Automating the construction of background knowledge (Li et al., 2012)
- Using SimStudent to learn cognitive models for open-ended tasks, such as educational games (Harpstead et al., 2013)

Thank you!