Proposal Example 2

Organizers:

…

Motivation

The shortage of people trained in STEM fields is becoming acute. According to a recent study, for example, there are 2.5 entry-level job postings for each new four-year graduate in STEM [1]. Universities and colleges are straining to satisfy this demand. In the case of Computer Science, for example, the number of US students taking introductory courses has grown three-fold in the past decade. Recently, MOOCs have been promoted as a way to ease this strain, but scaling traditional models of teaching to MOOCs poses many of the same challenges observed in the overflowing classrooms now common in introductory CS courses. Examples of such challenges include assessment of students' knowledge and providing personalized feedback to individual students. This motivates a new agenda for computer science research: formalize tasks such as assessment and feedback as *computational* problems, develop *algorithmic* tools to solve the resulting problems efficiently, and incorporate these tools effectively in *learning environments*.

Emerging Research in Personalized Education

While automated tutoring has been studied at different times in different communities, we are particularly excited about the promise of two recent approaches that bring insights from opposite ends of a spectrum: top-down approaches based on *formal logical reasoning* are an outgrowth of research on formal verification and synthesis, and bottom-up approaches based on *machine learning* that mine data from learning experiences of students.

Logical reasoning: In the last two decades, advances in automated reasoning tools such as model checkers and constraint solvers have led to successful applications of computer-aided verification to industrial-scale problems. A more recent application of automated reasoning is software synthesis, that is, automatic derivation of a program from its high-level specification. Emerging research has shown that reasoning tools developed for verification and synthesis can be effectively used to solve challenging problems in personalized education.

To understand the role of logical reasoning in personalized feedback, consider the task of automatically evaluating a student's submission to a programming problem in an introductory programming course. A commonly used assessment technique is to execute the student's program on a suitably chosen set of test inputs and check whether the resulting outputs match the expected ones. If this is not the case, instead of simply showing an input on which the program did not

[1] http://www.burning-glass.com/research/stem

work correctly, we can try to *synthesize* a variant of the student's program that works correctly. The edits that are needed to obtain such a correction then are used to highlight lines of code that need to be changed or to provide hints. The tool **AutoProf** [3] implements this strategy by relying on contemporary tools for formal verification and synthesis, and its effectiveness has been demonstrated in evaluating students' submissions in introductory programming course at MIT.

Tools rooted in logical reasoning for tasks such as automatic generation of problems, automatic grading, and automatic generation of hints have been developed for problems arising in a diverse set of CS courses such as Algorithms, Automata Theory, Compilers, Databases, Programming, and Embedded Systems (see [2] for a survey). The promise of such tools is indicated by the adoption of the tool **AutomataTutor** [1]: the tool was developed during Summer 2013 and is already used by over a thousand undergraduate students at EPFL, Penn, Reykjavic University, UC San Diego, and UIUC.

Machine learning: Algorithms from machine learning are also beginning to move from industry into education. Current applications range from *learning analytics* tools that help students and instructors keep track of learning progress to *personalized feedback* tools that recommend the next best learning activity for a student based on their activities and progress to date. In contrast to logical reasoning approaches, machine learning analytics typically eschew domain-specific models in favor of statistical models trained from large amounts of student data.

One promising machine learning based tool is the Sparse Factor Analysis (SPARFA) framework [4], which mines student grade book data to learn the latent concepts that underlie a subject. Once these concepts have been identified, SPARFA can assess a student's mastery of the concepts and track it over time to provide useful feedback to both the student and instructor. SPARFA can also autonomously organize the subject's course content (lecture notes, homework problems, feedback hints) by building a graph connecting those items to the latent concepts. Thanks to funding from NSF and the Gates and Arnold Foundations, this new toolset is being integrated directly into the free, open-source OpenStax College textbooks[2] to mine data from 300,000 students at 1000 universities and colleges nationwide.

Workshop Rationale

As discussed above, over the last few years, a number of researchers have demonstrated the benefits of applying tools for logical reasoning and machine learning to problems in personalized education in specific courses. The goal of the proposed workshop is to build on this momentum to develop a long-term research agenda in collaboration with researchers in education, human-computer interaction, and experts in education policy. The workshop can have a transformative impact in the following ways:

- **Shared Research Vision:** The workshop will bring together researchers to share experiences, understand the strengths and limitations of their current approaches, identify common principles, and articulate a shared research agenda for technical challenges and infrastructure.
- **Deep Reasoning:** The success of tools based on logical reasoning can be explained by the observations that for personalized education, (1) deep (or precise) reasoning is essential to generate accurate feedback, and (2) an input instance to the tool consists of a student's solution which is typically small. We will explore the synergies and tradeoffs between such top down approaches and bottom up approaches like machine learning. An issue of particular

[2] http://openstacollege.org

import is how the scalability of machine learning tools can aid techniques based on logical reasoning. The workshop is an opportunity to bring together researchers from these two distinct communities to develop novel solutions to the challenge of deep reasoning.

- **Adaptive Learning:** The success of any educational tool depends on how effectively the tool interacts with a student. The workshop will provide a forum to bring together researchers with expertise in how people learn, in how humans interact with computers, and those developing tools in computer-aided education.
- **From CS Courses to STEM, K-12, and Beyond:** While logical reasoning tools have thus far been focused on computer science courses at the college level, this approach has considerable promise in other STEM subjects as well at lower levels of K-12 education. The majority of emerging tools are focused on concepts in computer science curriculum. However, the techniques are potentially suitable to teach high-school students concepts in algebra, calculus, geometry, and physics. More broadly, computational thinking is increasingly a necessary skill for many people in our society and personalized education tools can help them acquire these skills. The workshop can initiate a conversation aimed at this broader vision.
- **Path to Adoption:** The emergence of platforms for online education has inspired a lot of debate on the education models. The workshop is an ideal forum to invite leaders in education policy from universities, government agencies, and industries to initiate plans for disseminating computer-aided education tools effectively and widely. A key topic will be the role that can be played by freely available, open-licensed content, assessments, and software tools.

Workshop Format

We plan to organize a two-day workshop in Summer 2015. A potential location for the workshop is the campus of University of Pennsylvania in Philadelphia. We will invite about 20 researchers who are currently developing personalized education tools and another 20 researchers from human-computer interaction, education, and education policy. Potential participants are:

1. ….

The workshop itself will be open to all interested participants. The workshop schedule will consist of invited talks, demo sessions, and working group sessions, with ample time for breaks to foster cross-disciplinary discussions.

Anticipated Outcomes

The workshop is expected to articulate a broad research agenda focused on *applications of deep reasoning to personalized learning in STEM subjects*. We expect that the workshop will foster new collaborations among participants from diverse disciplines, suggest new research directions in computer-aided education, inspire other researchers to work on these problems, and ultimately result in technology for effective and personalized learning.

The workshop participants will put together a document that articulates this vision along with research challenges. The workshop talks will be made publicly available and the resulting whitepaper will be widely disseminated. It can be used as a basis for organizing a larger meeting with a call for papers and call for participation.

Another potential outcome of the workshop is an article that summarizes the current techniques and future promise of Computer-aided Personalized Education for STEM subjects. The article will be aimed at a broad audience and can be published in *Communications of the ACM* or magazines on popular science.

Budget

The requested funding is $100,000. All the funds will be used to support the organization of the workshop. The workshop will be open to public and registration for the workshop will be free. Participants will be reimbursed for travel expenses.

Organizers' Biographies

....

References

[1] Rajeev Alur, Loris D'Antoni, Sumit Gulwani, Dileep Kini, and Mahesh Viswanathan. Automated grading of DFA constructions. In *IJCAI*, 2013. Tool accessible at `http://www.automatatutor.com/`.

[2] Sumit Gulwani. Example-Based Learning in Computer-Aided STEM Education. *CACM*, (8):70–80, 2014.

[3] Rishabh Singh, Sumit Gulwani, and Armando Solar-Lezama. Automated feedback generation for introductory programming assignments. In *PLDI*, 2013.

[4] A. Lan, A. E. Waters, C. Studer, and R. G. Baraniuk. Sparse factor analysis for learning and content analytics *Journal of Machine Learning Research*, (15):1959?2008, 2014.