Invited Talks (abstracts)

Can AI help MOOCs?

Jie Tang
Department of Computer
Science and Technology at
Tsinghua University
jietang@tsinghua.edu.cn

ABSTRACT

Massive open online courses (MOOCs) boomed in recent years and have attracted millions of users worldwide. It is not only transforming higher education but also provides fodder for scientific research. In this talk, I am going to first introduce the major MOOC platforms in China, for example, XuetangX.com, a similar platform to Coursear and edX, is offering thousands of courses to more than 7,000,000 registered users. I will also introduce how we leverage AI technologies to help enhance student engagement on MOOCs.

The evolution of virtual tutors, clinician, and companions: A 20-year perspective on conversational agents in real-world applications

Ronald Cole
Boulder Learning Inc.
rcole@boulderlearning.com

ABSTRACT

The talk will present an overview of research projects initiated in 1997 and continue today in 2017, in which 3-D computer characters interact with children and adults with the aim of improving their language communication skills, educational achievement, and/or personal well-being. The talk examines how advances in human language and character animation technologies, and research leading to a deeper understanding of how to apply these technologies to optimize engagement and learning, led to positive experiences and learning outcomes similar to experienced teachers and clinicians,individuals from 5 to 80 years of age, The talk concludes with a consideration of how recent advances in machine learning algorithms, coupled with cloud-based delivery of automated assessment and instruction, delivered by virtual agents, can save teachers millions of hours of time annually, and provide EDM researchers with vast amounts of speech and language data that can be mined to improve students' learning experiences and outcomes.

JEDM Track Journal Papers (abstracts)

Identifiability of the Bayesian Knowledge Tracing Model

Junchen Feng
17zuoye.com
Greenland Center Tower B 16th Floor
Beijing China
junchen.feng@17zuoye.com

ABSTRACT

The three "unidentified" model specifications proposed by Beck and Chang (2007) are identified by the Bayesian Knowledge Tracing model with a non-informative Dirichlet prior distribution and an observed sequence that is longer than three periods. Although these specifications have the same observed learning curve, they generate different likelihood given the same data. The paper further shows that the observed learning curve is not the sufficient statistics of the data generating process stipulated by the Bayesian Knowledge Tracing model. Therefore, it cannot be used in parameter inference of the Bayesian Knowledge Tracing model.

RiPLE: Recommendation in Peer-Learning Environments Based on Knowledge Gaps and Interests

Hassan Khosravi
University of Queensland
h.khosravi@uq.edu.au

Kendra Cooper
Independent Scholar
kendra.m.cooper@gmail.com

Kirsty Kitto
University of Technology Sydney
kirsty.kitto@uts.edu.au

ABSTRACT

Various forms of Peer-Learning Environments are increasingly being used in post-secondary education, often to help build repositories of student generated learning objects. However, large classes can result in an extensive repository, which can make it more challenging for students to search for suitable objects that both reflect their interests and address their knowledge gaps. Recommender Systems for Technology Enhanced Learning (RecSysTEL) offer a potential solution to this problem by providing sophisticated filtering techniques to help students to find the resources that they need in a timely manner. Here, a new RecSysTEL for Recommendation in Peer-Learning Environments (RiPLE) is presented. The approach uses a collaborative filtering algorithm based upon matrix factorization to create personalized recommendations for individual students that address their interests and their current knowledge gaps. The approach is validated using both synthetic and real data sets. The results are promising, indicating RiPLE is able to provide sensible personalized recommendations for both regular and cold-start users under reasonable assumptions about parameters and user behavior.

Keywords

Peer-Learning Environments, Recommender Systems, Knowledge Gaps

Modeling Wheel-spinning and Productive Persistence in Skill Builders

Shimin Kai
Teachers College Columbia University,
525 W125th Street, New York, NY 10027
+1 212-678-3000
smk2184@tc.columbia.edu

Ma. Victoria Almeda
Teachers College Columbia University,
525 W125th Street, New York, NY 10027
+1 212-678-3000
mqa2000@tc.columbia.edu

Ryan S. Baker
Graduate School of Education, University of Pennsylvania
3700 Walnut St., Philadelphia, PA 19104
+1 877-736-6473
ryanshaunbaker@gmail. com

Nicole Shechtman
Center for Technology in Learning, SRI International
333 Ravenswood Avenue
Menlo Park, CA 94025
+1 650-859-2000
nicoleshechtman@sri.com

Cristina Heffernan
Worcester Polytechnic Institute
100 Institute Rd, Worcester, MA 01609
+1 508-831-5000
cristina.heffernan@gmail.com

Neil Heffernan
Worcester Polytechnic Institute
100 Institute Rd, MA 01609
+1 508-831-5000
nth@wpi.com

ABSTRACT

Research on non-cognitive factors has shown that persistence in the face of challenges plays an important role in learning. However, recent work on wheel-spinning, a type of unproductive persistence where students spend too much time struggling without achieving mastery of skills, show that not all persistence is uniformly beneficial for learning. For this reason, it becomes increasingly pertinent to identify the key differences between unproductive and productive persistence toward informing interventions in computer-based learning environments. In this study, we attempt to address this by using classification models to distinguish between productive persistence and wheel-spinning in ASSISTments, an online math learning platform. Our results indicate that wheel-spinning is associated with shorter delays between solving problems of the same skill, more attempts to answer problems, and the heavy use of bottom out hints except for the first problem. These findings suggest that encouraging students to engage in spaced practice and avoid over-using bottom-out hints is likely helpful to reduce their wheel-spinning and improve learning. These findings also provide insight on which students are struggling and how to make students' persistence more productive.

Modeling MOOC Student Behavior With Two-Layer Hidden Markov Models

Chase Geigle
Department of Computer Science
University of Illinois at Urbana-Champaign
Urbana, Illinois, USA
geigle1@illinois.edu

ChengXiang Zhai
Department of Computer Science
University of Illinois at Urbana-Champaign
Urbana, Illinois, USA
czhai@illinois.edu

ABSTRACT

Massive open online courses (MOOCs) provide educators with an abundance of data describing how students interact with the platform, but this data is highly underutilized today. This is in part due to the lack of sophisticated tools to provide interpretable and actionable summaries of huge amounts of MOOC activity present in log data. To address this problem, we propose a student behavior representation method alongside a method for automatically discovering those student behavior patterns by leveraging the click log data that can be obtained from the MOOC platform itself. Specifically, we propose the use of a two-layer hidden Markov model (2L-HMM) to extract our desired behavior representation, and show that patterns extracted by such a 2L-HMM are interpretable, meaningful, and unique. We demonstrate that features extracted from a trained 2L-HMM can be shown to correlate with educational outcomes.

Acknowledgments

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant Number DGE-1144245.

Closing the loop: Automated data-driven cognitive model discoveries lead to improved instruction and learning gains

Ran Liu
Human-Computer Interaction Institute
Carnegie Mellon University
ranliu@cmu.edu

Kenneth R. Koedinger
Human-Computer Interaction Institute
Carnegie Mellon University
koedinger@cmu.edu

ABSTRACT

As the use of educational technology becomes more ubiquitous, an enormous amount of learning process data is being produced. Educational data mining seeks to analyze and model these data, with the ultimate goal of improving learning outcomes. The most firmly grounded and rigorous evaluation of an educational data mining discovery is whether it yields better student learning when applied. Such an evaluation has been referred to as "closing the loop", as it completes cycle of system design, deployment, data analysis, and discovery leading back to design. Here, we present an instance of “closing the loop” on an automated cognitive modeling improvement discovered by Learning Factors Analysis (Cen, Koedinger, & Junker, 2006). We discuss our findings from a process in which we interpret the automated improvements yielded by the best-fitting cognitive model, validate the interpretation on novel data, use it to make changes to classroom-deployed educational technology, and show that the changes lead to significant learning gains relative to a control condition.