Artificial Intelligence Powered MOOCs: A Brief Survey

Simon Fauvel*, Han Yu†*, Chunyan Miao†*, Lizhen Cui‡, Hengjie Song§, Xiaoming Li¶, and Cyril Leung‖
*Joint NTU-UBC Research Centre of Excellence in Active Living for the Elderly (LILY), Nanyang Technological University, Singapore
†School of Computer Science and Engineering, Nanyang Technological University, Singapore
‡School of Computer Science and Technology, Shandong University, Jinan, China
§School of Software Engineering, South China University of Technology, Guangzhou, China
¶Institute of Network Computing and Information Systems, Peking University, Beijing, China
‖Department of Electrical and Computer Engineering, the University of British Columbia, Vancouver, Canada
Corresponding Authors: {sfauvel, han.yu}@ntu.edu.sg, clz@sdu.edu.cn, sehjsong@scut.edu.cn

***Abstract*—Massive Open Online Courses (MOOCs) have gained tremendous popularity in the last few years. Thanks to MOOCs, millions of learners from all over the world have taken thousands of high-quality courses for free. Artificial intelligence (AI) has played an important role in making MOOCs what they are today. By exploiting the vast amount of data generated by learners engaging in MOOCs, AI techniques have been proposed to improve our understanding of MOOC participants and enable MOOC practitioners to deliver better courses. These approaches have also greatly improved student experience and learning outcomes through constructing intelligent and personalized learning trajectories. In this paper, we first review the state-of-the-art AI research making an impact on MOOCs education, emphasizing on works which aim to enhance our understanding of student learning behaviours, improve student engagement, and improve learning outcomes. We then offer an overview of important future research to carry out in sub-fields of AI to enable MOOCs to reach their full potential.**

***Keywords*-MOOCs; artificial intelligence; data science**

I. Introduction

Humans are lifelong learners. We are constantly learning about the world around us and about each other by gathering new knowledge. Massive Open Online Courses (MOOCs) are one of the latest trends in education, making it easier for learners to access free, high-quality learning opportunities. Given a computer and an Internet connection, students from around the world have open access to high-quality courses from the best schools and organisations. Targeting learners who do not have access to formal higher education as well as lifelong students looking to gain new skills and knowledge, MOOCs connect learners from all over the world. Fuelled by advances in cloud computing and widespread access to Internet, MOOCs have gained tremendous popularity in recent years. For example, spearheaded by institutes of higher learning such as Peking University, over 3,200 MOOCs have been created in China as of 2018 which have been taken by over 55 million students.

Much research about MOOCs has already taken place. One research area that has gained significant popularity in the last few years is how artificial intelligence (AI) techniques can contribute to better understanding the MOOC ecosystem, and how they can contribute to improving it. In a nutshell, AI studies how to design intelligent machines and systems that analyse their environment and take actions that maximise their chances of success. An AI-augmented MOOC platform can enable a better understanding of how learning happens in MOOCs as well as providing a more engaging learning experience for learners, thus improving their learning outcomes. MOOCs offer a very rich environment for AI given the large amount of data generated by learners using the platform, but making sense of such data is not a trivial task. In some ways, all AI research in MOOC relies on data to operate - the goal is not to come up with new theoretical developments in AI, but rather to use AI tools to solve a problem in MOOC that generally involves big data.

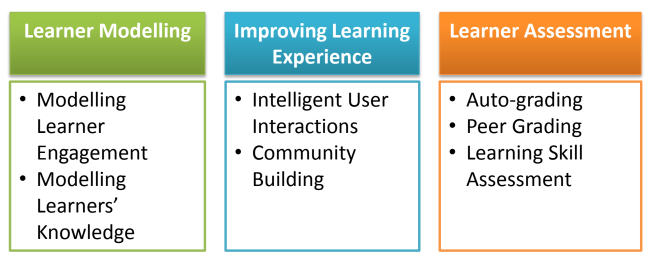

Figure 1. A taxonomy of AI research for MOOCs.

Research in AI applied to MOOCs has contributed to our current understanding of MOOCs, and plays a significant part in the success they have enjoyed so far. More importantly, they have a crucial role to play in solving the challenges facing MOOCs today. In this paper, we review the state-of-the-art AI research for MOOCs. Based on the objectives of AI research for MOOCs, we propose a taxonomy which divides current literature into three main categories as shown in Figure 1: 1) Learner Modelling; 2) Improving Learning Experience; and 3) Learner Assessment. This survey paper will be useful for anyone interested in contributing to AI for MOOCs, allowing them to quickly find their bearing in current research and identify existing gaps. We also hope that this paper will be valuable to AI and MOOC researchers and practitioners, through highlighting the key general findings as well as promising future trends.

II. Learner Modelling

If one wishes to improve MOOC learning outcomes, it is first necessary to better understand how learners engage with MOOC resources and tools. By virtue of being open, MOOCs attract a large range of learners with widely different backgrounds and motivations for taking a MOOC. As such, many different learning styles emerge. Much research has already taken place to understand learning styles.

A. Modelling Learner Engagement

Many MOOCs ask learners to fill in a pre-course survey that covers various elements such as background, expectations, and intentions. When such data is collected, it is possible to link learners' self-reported intentions with their learning behaviour trajectories in the MOOC. In [1], authors use this survey data to map it to learner categories. They identified four categories: no-shows, observers, casual learners, and completers. The clustering is based on simple mappings between intentions and styles.

MOOC platforms also collect a very comprehensive set of data regarding learners' usage behaviours, down to the mouse click level. This offers abundant data that can be used to extract knowledge about their learning styles. In [2], authors focus on two activities for this purpose: 1) viewing a lecture and 2) handing in an assignment for credit. They then use a histogram approach to cluster learners. They developed a taxonomy of learner behaviour comprising five engagement styles: viewers, solvers, all-rounders, collectors, and bystanders.

As every learner is different, taking into account learners' background may provide some insight into achievement. However, [3] showed that looking at learners' demographics alone is not sufficient, as this provides very limited explanatory and predictive information about achievement in MOOCs. Instead, learners' background should be coupled with the activities learners engage in on the platform. In some ways, background is used to augment the learners' behaviour on the platform. While it appears that demographics do not show any clear correlations with achievement, other background information (such as prior competence in the subject matter) has been shown to help explain and predict learning outcomes.

Learners tend to express their sentiments towards a MOOC in forums. This offers a rich data source that can be mined for patterns. In [4], qualitative content analysis of forum posts is first conducted, based on three dimensions: 1) learning, 2) dialogue acts, and 3) topic. The coding is done manually. It is then followed by Bayesian Non-negative Matrix Factorisation to extract communities of learners based on the codes. Different learner communities are identified for two distinct sub-forums. In one sub-forum, four communities are identified: 1) committed crowd engagers, 2) discussion initiators, 3) strategists, and 4) individualists. In the other, five communities are identified: 1) instrumental help seekers, 2) careful assessors, 3) community builders, 4) focused achievers, and 5) support seekers.

B. Modelling Learner Knowledge

An engaged MOOC learner does not necessarily acquire knowledge effectively. Thus, models on how well a learner is acquiring knowledge are needed. An obvious strategy to model learners' knowledge is to look at their performance in standard assessments (homework, quizzes, exams, etc.) Most courses require learners to complete such activities in order to complete the course. In [5], knowledge is measured using discrete knowledge components (KCs) defined by a subject matter expert. Each assessed problem is a KC, and a Bayesian Knowledge Tracing based method is used to model learner knowledge. They show that they can predict with good accuracy whether a learner will successfully answer an exam problem (i.e., whether she acquired the related knowledge during the course). This method requires expert input, which may not always be available. In [6], the modelling of domain-specific knowledge (namely, programming) is investigated. They use a variant of the Additive Factors Model to automatically extract a domain model by exploiting the structure in the Java programming language and then model student knowledge. They show that their models outperform a state-of-the-art model that does not assume learning is taking place. Compared to the previous method, this method is fully automated which minimizes human intervention. Nevertheless, its applicability is limited to the domain of computer programming. In [7], performance on homework is used as a proxy to measure knowledge. Models based on item response theory are designed to overcome the issue that data is missing for a large number of homework problems and that multiple attempts are allowed for each problem. Evaluation is done the same way as for [5], and their results show that their models offer improved estimates when compared to purely correlational models.

Another approach is to jointly consider all activities learners engage in on the MOOC platform. Data obtained from pre- and post-course surveys can also be used. In [8], knowledge is deemed to be related to the scores on assessed components due to the way these assessments were designed. The authors then look at the number of activities learners engage in as well as learners' feedback collected from a post-course survey. Regression analysis and classification are then used to find variables that correlate with knowledge. A higher number of online activities is linked with more knowledge gained. Learners' perceived usefulness (i.e. whether the course appears valuable to the learner) is also linked with better learning outcomes. Similarly, in

[9], assessment grades are used as a proxy to knowledge. Through exploratory data analysis, it is found that students learning through performing tasks achieve better learning outcomes than those only viewing informational assets (e.g. videos). [10] proposes to measure knowledge using the Precise Effectiveness Strategy, a methodology that uses metrics to calculate the effectiveness of learners when interacting with MOOC materials. These metrics must first be defined by experts before they can be computed automatically, based on learners' activities. At this point, no validation experiment has been carried out to verify whether this approach actually correlates well with knowledge.

III. Improving Learning Experience

MOOC learning experience mainly revolves around course contents and community interactions. By virtue of being online, learning experience can benefit from AI.

A. Intelligent User Interactions

Most courses rely on videos to deliver content at scale. In [11], authors study video production features and correlate them with measures of engagement (i.e. how long learners watch a given video, and whether they attempt the post-video exercise). Through the use of simple statistical tools, they make a number of useful observations. These observations mention that shorter videos lead to more engagement, and specify production styles that can lead to more engagement. This can help instructors create MOOC videos that maximise learner engagement. [12] also finds a strong correlation between video length and engagement: the longer the video, the higher the in-video dropout rate (i.e. a learner does not finish watching the video). They find a difference between videos watched for the first time compared to replays, in which higher in-video dropout rates have been observed. Using binning and kernel-based smoothing, they then show second-by-second plots of video interaction (play, pause and skip events) peaks, which can show the typical learner behaviour patterns within a video. They manually categorise each peak into 5 categories that explain the underlying cause of the peak. Their analysis produces an improved understanding of how students interact and learn, which could lead to more effective video interfaces for learning.

Understanding video viewing patterns is useful, but it does not immediately provide value to learners. One way to achieve this is by adding AI tools to facilitate video navigation and watching. [13] aims to facilitate non-sequential video navigation by augmenting the video interface with different tools. Some examples of such tools include a customised dynamic time-aware word-cloud and a 2-D timeline (to provide an overview of the concepts discussed in the video and allow quick navigation to these concepts), video pages (to enable visually searching for information by using automatically extracted visual slides) and a video summarisation method (to provide a summary of the video). [14] automatically identifies and links related topics in MOOC videos. It then creates an interface that enables learners to navigate to topics of interest. Their approach is based on fuzzy formal concept analysis and semantic technologies. Based on analysing learners' interaction trajectory data with MOOC learning contents, [15] is able to predict the likelihood of individual learners dropping out from a given MOOC to help courses determine what measures to use to retain 'at-risk' learners.

B. Community Building

Forums are the main tool used for community building in MOOCs. With better understanding of forum usage patterns in MOOCs, it is possible to improve forum welfare with AI. The most popular strategy to improve forum welfare is through identifying and suggesting relevant forum contents to learners. The high volume of forum posts often prevents learners from keeping up-to-date, so mechanisms are needed to ensure they can view the most relevant contents.

In [16], authors use linguistic modelling to identify content-related threads. They then perform classification on manually labelled threads to demonstrate the effectiveness of their approach. In [17], authors propose a classification scheme of forum posts using non-text features that is language independent and that is reasonably accurate. It can be used to ensure that the labels selected by learners are accurate, thus improving navigation.

In [18], authors build a generative model that enables them to efficiently classify threads as well as assign them a relevance ranking. This can facilitate user navigation by only showing learners the most relevant threads. [19] specifically targets question answering in forums. They develop a constrained question recommendation algorithm based on a context-aware matrix factorisation model that predicts students' preferences over questions, as well as a model to optimise community benefit under multiple constraints. Their results outperform other baseline approaches for question recommendation.

IV. Assessing MOOC Learners

While better understanding learner engagement and achievement and improving learning experience are both important, it is insufficient for determining whether a MOOC is successful in its primary goal: equipping learners with new skills and knowledge. Indeed, a learner may be well engaged in and complete a MOOC without having learned much. As such, assessing learners' knowledge and skills is a crucial aspect to determine a MOOC's success.

A. Auto-grading

For auto-grading assessment, research is mostly targeted at domain-specific problems since each domain has specific requirements. On top of checking for answer correctness, feedback may be offered, either along the way (as hints) or at the end (as suggested improvements).

Programming is by far the domain that has received the most attention in this area of research. While testing the validity of a computer program is a trivial problem that can be solved through the use of test cases, objectively evaluating the quality of a program is much more difficult. [20] does this by considering a program to be constituted of multiple code fragments and measuring the distance between fragments of different students through appropriate distance metrics. By relying on previously graded programs, a quality score is assigned (i.e., programs that have a small neighbouring distance should be of similar quality).

Most works related to programming focus on how to automatically provide feedback to learners. [21] tries to predict how an instructor would encourage a learner to progress towards the correct solution to automatically generate hints. They use Desirable Path algorithms to model the best paths that can lead to the correct solution and steer learners in that direction. In [22], authors develop a technique that uses an error model describing the potential corrections and constraint-based synthesis to compute minimal corrections to students' incorrect solutions, enabling them to automatically provide feedback to learners. Their results show that relatively simple error models can correct a good number of incorrect solutions.

B. Peer Grading

When automated grading is not feasible, peer grading is usually the chosen alternative. In peer grading, a learner's coursework is reviewed and graded by other learners. Usually, on top of the grade, detailed feedback is also offered. A common challenge in peer grading is how to perform grade aggregation so that the outcome is as accurate and fair as possible. A number of works have tackled this task.

In [23], authors use a trust graph to automatically combine grades from peers and tutors. The trust model is used to compute how much weight to give to each grader when computing the final grade. [24] develop methods to improve both cardinal (absolute judgment) and ordinal (relative judgment) peer grading. On one hand, they extend existing probabilistic graphical models to improve cardinal grading. On the other hand, they use cardinal prediction priors to augment ordinal models. [25] seeks to lower the grading burden on peers while maintaining quality. It does so by using a machine learning algorithm to automatically grade an entry. Peers then identify key features of the answer using a rubric, and other peers verify whether these labels are reasonable. Depending on the confidence of the results computed by the algorithm and peer agreement, a different set of peer graders is assigned for each entry. Finally, [26] develops a series of probabilistic models that estimate and correct graders' biases and reliabilities. All of these methods have been shown to improve peer grading accuracy. However, since they do not compare themselves against each other, it is difficult to judge which one performs better.

C. Learning Skill Assessment

MOOC platforms often track interaction data between course instructors and students as well as among students of the same course. Such interactions provide opportunities to generate complex behaviour data which can be used to analyse students' learning skills (e.g. collaboration with others, time management, and critical thinking). The skill assessment results can complement quizzes/exams in MOOCs in order to support personalized interventions for improving learning outcomes [27], [28].

A few works have tackled the problem of assessing MOOC learners' learning skills based on their behaviour data. In [29], the authors leverage topic models to analyze learners' online forum postings in order to infer their learning skills. Measurement theory from education and psychology is also incorporated into the proposed approach to quantify a person's attainment of intangible attributes such as attitudes, abilities or intelligence, thereby enabling the inference of latent skills based on individuals' observed responses on a series of items such as quiz questions. Topic modeling is applied to automatically extract discussion topics from online forum in MOOCs and are used as items on the Guttman scale through matrix factorization. The approach has been shown to perform well through experiments on three Coursera MOOCs and a survey by domain experts.

V. Future Research Directions

Much work remains to be done before AI can enable MOOCs to achieve their full potential. We now summarize a list of key areas where new research and development in AI could significantly improve MOOCs based on our review.

A. Redefining Openness in MOOCs

Current MOOCs can be considered open in the sense that they are freely accessible to all learners (although it should be noted that some MOOC platforms require payment to access courses and/or specific materials). However, we would like to highlight other dimensions of openness that would be highly valuable as well.

Creating MOOCs is a resource-intensive process for instructors. In that sense, MOOCs are not fully open to all instructors, since they must have the resources before they can embark on the MOOC journey. We believe that the biggest contribution in this area would come from improving content reusability and interoperability. There already exist many open materials available online (and more is being created every day), but finding appropriate ones and integrating them into MOOCs is largely impractical at the moment. Standardising knowledge representations specifically for MOOC contents would be a good first step.

Openness for researchers is also a critical issue if we want to improve MOOCs. Specifically, access to data is essential. Most works use data from a very small subset of MOOCs since this is the only data they have access to. Each dataset uses different data sources, and the underlying characteristics of the courses covered by each dataset vary widely from dataset to dataset. What this means is that comparison between studies is very difficult, and the validity and generalisability of results is unclear since the data used to obtain them might not be representative. Open benchmark datasets, as is common in many other fields, could certainly help overcome these issues. On top of the technical challenges associated with the tasks described above, it also requires concerted efforts by MOOC providers and institutions, as well as policy changes. We are aware of legal and ethical issues that currently hinder these changes from taking place, but feel it is necessary for the research community to discuss how to overcome these barriers so that MOOCs can reach their full potential.

B. Complementing AI with Human Effort

AI has improved significantly recently, but there are still a number of challenges that it is unable to solve. For such challenges, combining the power of AI with humans (a concept known as algorithmic crowdsourcing [30], [31]) could prove game-changing.

One important problem in MOOC research is the fact that most data is unlabelled. This makes it much harder to draw causation links, as opposed to simple correlations, from the data. At the same time, supervised algorithms must thus rely on the small amount of labelled data (either provided by student surveys or researchers themselves both approaches are expensive). Examples of data labelling through crowdsourcing include allowing students to rate different portions of a video to assist curriculum redesign, gathering insights about navigation behaviour through short pop-up questions, collaborative annotation to label contents, and more.

Crowdsourcing could also be leveraged to reduce instructor workload. Learners could be organized through crowdsourcing to help MOOC instructors create hints. This approach could also be expanded to other aspects of MOOCs such as exercise creation. Game theory (through incentive mechanisms), trust models and rating mechanisms can be leveraged to ensure the quality of the contributions.

C. From Engagement to Knowledge

Currently, the majority of AI research in MOOCs looking at better understanding learners is concerned with engagement. While this information can be valuable for a number of reasons, it is not sufficient to ensure that MOOCs are in fact able to impart knowledge to students. A student can be engaged in a course without actually learning much (e.g. if the material is too simple). Automatically assessing the instructional quality of MOOCs and identifying the criteria that lead to quality courses could be valuable. Currently, such assessments are done manually, which is not suitable at MOOC scales. Although some research has taken place in modelling learner knowledge, much of it is preliminary and domain-specific. More needs to be done so that methods can be deployed easily across domains. The current research also overlooks important elements related to learner knowledge. Indeed, all works focus on knowledge gain pertaining to the course materials, usually measured through purposely-built assessment tools. However, few research works attempt to identify new knowledge and skills developed at a higher level (beyond technical competency in the subject matter, e.g. 21st Century skills and informal learning events) as a result of learners engaging in MOOCs. It could also prove valuable to look into which variables can predict that a learner will apply what she has learned as opposed to that knowledge remaining theoretical in nature.

VI. Conclusion

In this paper, we reviewed the state-of-the-art AI research applied to MOOCs, emphasising how AI tools and techniques can improve learner modelling, learning experience, and learner assessment. We showed how AI can be seamlessly embedded in virtually every aspect of the MOOC ecosystem as we know it today, and gave an overview of important future research in AI for MOOCs so that they can reach their full potential.

We believe that MOOCs are one part of the puzzle to making universal, lifelong education a reality. By enabling learners from around the world to access free, high-quality courses, we are giving people more opportunities. By connecting learners from all over the world and from all kinds of backgrounds, by enabling them to interact with one another and share ideas and knowledge, we are creating an enabling environment with the help of AI in which boundaries do not matter anymore, an environment in which everyone can satisfy one of the primary urges of being human: to quench our curiosity through learning.

Acknowledgements

This research is supported by the National Research Foundation, Prime Minister's Office, Singapore under its IDM Futures Funding Initiative; Nanyang Technological University, Nanyang Assistant Professorship (NAP); and the NTU-PKU Joint Research Institute, a collaboration between Nanyang Technological University and Peking University that is sponsored by a donation from the Ng Teng Fong Charitable Foundation.

References

[1] J. Wilkowski, A. Deutsch, and D. M. Russell, "Student skill and goal achievement in the mapping with Google MOOC," in *L@S*, 2014, pp. 3–10.

[2] A. Anderson, D. Huttenlocher, J. Kleinberg, and J. Leskovec, "Engaging with massive online courses," in *WWW*, 2014, pp. 687–698.

[3] C. Brooks, C. Thompson, and S. Teasley, "Who you are or what you do: Comparing the predictive power of demographics vs. activity patterns in massive open online courses (moocs)," in *L@S*, 2015, pp. 245–248.

[4] R. Eynon, I. Hjorth, N. Gillani, and T. Yasseri, "'Vote me up if you like my ideas!' experiences of learning in a MOOC," in *ICA Pre-Conference Working Paper: Innovation in Higher Education: Building a Better Future*, 2014, p. http://dx.doi.org/10.2139/ssrn.2439571.

[5] Z. Pardos, Y. Bergner, D. T. Seaton, and D. E. Pritchard, "Adapting Bayesian knowledge tracing to a Massive Open Online Course in edX," in *EDM*, 2013, pp. 137–144.

[6] M. Yudelson, R. Hosseini, A. Vihavainen, and P. Brusilovsky, "Investigating automated student modeling in a Java MOOC," in *EDM*, 2014, pp. 261–264.

[7] Y. Bergner, K. Colvin, and D. E. Pritchard, "Estimation of ability from homework items when there are missing and/or multiple attempts," in *LAK*, 2015, pp. 118–125.

[8] D. Liang, J. Jia, X. Wu, J. Miao, and A. Wang, "Analysis of learners' behaviors and learning outcomes in a massive open online course," *Knowledge Management & E-Learning*, vol. 6, no. 3, pp. 281–298, 2014.

[9] K. R. Koedinger, J. Kim, J. Z. Jia, E. A. McLaughlin, and N. L. Bier, "Learning is not a spectator sport: Doing is better than watching for learning from a MOOC," in *L@S*, 2015, pp. 111–120.

[10] P. J. Muñoz Merino, J. A. Ruipérez-Valiente, C. Alario-Hoyos, M. Pérez-Sanagustín, and C. Delgado Kloos, "Precise effectiveness strategy for analyzing the effectiveness of students with educational resources and activities in moocs," *Computers in Human Behavior*, vol. 47, no. C, pp. 108–118, 2015.

[11] P. J. Guo, J. Kim, and R. Rubin, "How video production affects student engagement: An empirical study of MOOC videos," in *L@S*, 2014, pp. 41–50.

[12] J. Kim, P. J. Guo, D. T. Seaton, P. Mitros, K. Z. Gajos, and R. C. Miller, "Understanding in-video dropouts and interaction peaks in online lecture videos," in *L@S*, 2014, pp. 31–40.

[13] K. Yadav, K. Shrivastava, S. Mohana Prasad, H. Arsikere, S. Patil, R. Kumar, and O. Deshmukh, "Content-driven multimodal techniques for non-linear video navigation," in *IUI*, 2015, pp. 333–344.

[14] C. Maio, V. Loia, G. R. Mangione, and F. Orciuoli, "Automatic generation of SKOS taxonomies for generating topic-based user interfaces in MOOCs," in *EC-TEL*, 2014, pp. 398–403.

[15] J. He, J. Bailey, B. I. P. Rubinstein, and R. Zhang, "Identifying at-risk students in massive open online courses," in *AAAI*, 2015, pp. 1749–1755.

[16] Y. Cui and A. F. Wise, "Identifying content-related threads in MOOC discussion forums," in *L@S*, 2015, pp. 299–303.

[17] L. A. Rossi and O. Gnawali, "Language independent analysis and classification of discussion threads in Coursera MOOC forums," in *IRI*, 2014, p. doi:10.1109/IRI.2014.7051952.

[18] C. G. Brinton, M. Chiang, S. Jain, H. Lam, Z. Liu, and F. M. F. Wong, "Learning about social learning in MOOCs: From statistical analysis to generative model," *IEEE Transactions on Learning Technologies*, vol. 7, no. 4, pp. 346–359, 2014.

[19] D. Yang, D. Adamson, and C. P. Rosé, "Question recommendation with constraints for massive open online courses," in *RecSys*, 2014, pp. 49–56.

[20] A. Drummond, Y. Lu, S. Chaudhuri, C. Jermaine, J. Warren, and S. Rixner, "Learning to grade student programs in a Massive Open Online Course," in *ICDM*, 2014, pp. 785–790.

[21] C. Piech, M. Sahami, J. Huang, and L. Guibas, "Autonomously generating hints by inferring problem solving policies," in *L@S*, 2015, pp. 195–204.

[22] R. Singh, S. Gulwani, and A. Solar-Lezama, "Automated feedback generation for introductory programming assignments," in *PLDI*, 2013, pp. 15–26.

[23] P. Gutierrez, N. Osman, and C. Sierra, "Collaborative assessment," in *CCIA*, 2014, p. http://hdl.handle.net/10261/131373.

[24] F. Mi and D.-Y. Yeung, "Probabilistic graphical models for boosting cardinal and ordinal peer grading in MOOCs," in *AAAI*, 2015, pp. 454–460.

[25] C. E. Kulkarni, R. Socher, M. S. Bernstein, and S. R. Klemmer, "Scaling short-answer grading by combining peer assessment with algorithmic scoring," in *L@S*, 2014, pp. 99–108.

[26] C. Piech, J. Huang, and Z. Chen, "Tuned models of peer assessment in MOOCs," in *EDM*, 2013, pp. 153–160.

[27] R. Luckin, "Towards artificial intelligence-based assessment systems," *Nature Human Behaviour*, vol. 1, no. 0028, pp. doi:10.1038/s41 562–016–0028, 2017.

[28] H. Yu, C. Miao, C. Leung, and T. J. White, "Towards AI-powered personalization in MOOC learning," *npj Science of Learning*, vol. 2, no. 15, pp. doi:10.1038/s41 539–017–0016–3, 2017.

[29] J. He, B. I. P. Rubinstein, J. Bailey, R. Zhang, S. Milligan, and J. Chan, "MOOCs meet Measurement Theory: A topic-modelling approach," in *AAAI*, 2016, pp. 1195–1201.

[30] A. Doan, R. Ramakrishnan, and A. Y. Halevy, "Crowdsourcing systems on the world-wide web," *Communications of the ACM*, vol. 54, no. 4, pp. 86–96, 2011.

[31] H. Yu, C. Miao, Y. Chen, S. Fauvel, X. Li, and V. R. Lesser, "Algorithmic management for improving collective productivity in crowdsourcing," *Scientific Reports*, vol. 7, no. 12541, pp. doi:10.1038/s41 598–017–12 757–x, 2017.