Research Using the MOOC Replication Framework and E-TRIALS

Ryan Baker
Graduate School of Education
University of Pennsylvania
Philadelphia, USA
rybaker@upenn.edu

Stephen Hutt
Graduate School of Education
University of Pennsylvania
Philadelphia, USA
hutts@upenn.edu

Michael Mogessie
Human-Computer Interaction Institute
Carnegie Mellon University
Pittsburgh, USA
michaelmogessie@cmu.edu

Haripriya Valayaputtar
Graduate School of Education
University of Pennsylvania
Philadelphia, USA
haripriyavalayaputtur@gmail.com

Abstract—MOOCs have emerged as an important venue for educational research, but the community of researchers who can conduct research in MOOCs is limited. With a few exceptions, running a study in a MOOC is limited to researchers at universities that host MOOCs, and the MOOC data sets available to the broad community of research often redact key information such as discussion forum posts. In this paper, we discuss efforts at our university to make our MOOCs and MOOC data available for research to the broader community, using the E-TRIALS platform for experimentation and the MORF platform for secondary data analysis. We provide technical details on both of these tools, discuss the researcher capacity needed to use these tools, and review the research opportunities these tools provide.

Keywords— MOOC, A/B Test, Data Enclave

I. INTRODUCTION

While there has been a large amount of research involving MOOCs over the last decade, researchers have faced limitations in data availability and platform capabilities. Aside from a small number of MOOC researchers who have insider access to data sets, most have had to work with public data sets such as the edX RDX data package, which redact key information such as discussion forum data and demographic information. This has led to a substantial proportion of MOOC research focusing on a somewhat narrow range of topics and features. Correspondingly, although the major platforms offer support for A/B testing, the range of possible tests is somewhat limited by the platforms' functionality.

Learning in any context is a deeply personal experience. Aspects of a student's identity can influence how and when they succeed, as well as what supports most benefit them. Identity has many dimensions, many of them Personally Identifying Information (PII). Studies have consistently shown that the efficacy of online methods and interventions varies by student identity, including race, gender, and national origin [13, 15]. As we strive for MOOCs that are effective for the full range of learners who use them, we must factor in student identity as a moderator in our A/B testing. For this, we need an infrastructure that allows us both to test at a large scale (in order to provide appropriate representation) and with access to PII for moderation effects.

Similarly, as we strive to improve equity across education and education research, PII becomes critical to the evaluation of algorithmic biases. As Baker & Hawn (2021) argue [13], we cannot achieve fairness without knowing whether an algorithm is biased, and we cannot determine that without demographic data, data currently redacted from public releases of MOOC data. As such, redacted datasets do not always facilitate these goals, and additional research tools are needed to address questions of equity and bias while still protecting individual students.

Furthermore, the current redaction of data sets available to researchers limits their ability to conduct linguistic analyses of student work within MOOCs. Some of the most interesting and important activity within MOOCs occurs within discussion forums, in the textual details of students' interactions with each other and with learning material [11]. However, it is highly difficult to fully redact textual data for full public release – another type of tool is needed [12].

To address these issues, this paper discusses two tools we have adopted at our university, the MOOC Replication Framework [3] and E-TRIALS [7]. These tools support researchers (both internal and external to our university) in conducting research on MOOCs. The MOOC Replication Framework (MORF) is a system that enables researchers to conduct analyses on a complete and unredacted repository of MOOC data. MORF is a data enclave that enables the analysis of a wide variety of research questions while maintaining student privacy by requiring analyses to be conducted on a controlled server with limited output functionality. MORF is under continual development to expand functionality and support additional data sources. The codebase is open source and is used by multiple universities. This paper also discusses our university's partnership with the ASSISTments platform to use their Ed-Tech Research Infrastructure to Advance Learning Science (E-TRIALS) experimental infrastructure. In this paper, we will discuss the ways that MORF and E-TRIALS have been used at our university, focusing on the research that has been enabled. This paper serves as an invitation to the broader LWMOOCS community to join our research network and leverage the discussed infrastructure for their own research.

II. MORF FRAMEWORK

A. MORF Design

MORF, the MOOC Replication Framework, has been developed in a partnership between researchers at multiple universities in the United States [3]. MORF allows researchers to conduct analyses on a sizeable repository while preventing direct viewing (or export) of the data. By limiting access in this way, MORF provides a complete and unredacted data set for analysis while protecting student privacy.

In order to achieve this goal, we have implemented a workflow that leverages Amazon Web Services (AWS) and the subsequent cloud services that they provide. Jobs are submitted to MORF via a public URL using an API key that is unique to each MORF user. This API facilitates job submission and links each job to the submitting researcher. The key is provided to a researcher once a basic data agreement has been signed and Institution Review Board research approval (or equivalent) is verified.

A MORF job consists of two key components: (1) specifications for the job environment and (2) analysis code to be run in that environment.

Environment specifications are provided using Docker, and give information of any pre-requisites for the proposed analysis (e.g., programming languages or libraries). Using Docker makes it possible for researchers to use almost any programming environment so long as the software pre-requisites (i.e., libraries) are readily available via a package manager or similar tool. When the job is started, this environment is built in an isolated container that can access internal MORF resources but no external resources (e.g., the internet). This provides the job with full access to the data repository without risking the security of the data. This structure also preserves the exact running environment of a job for later replication purposes.

The provided code is then run within the environment specified. All intermediate outputs used in these processes are outputted and stored privately on the MORF servers but are not immediately directly available to the researcher (they can be used later in coordination with MORF administrators if any debugging is necessary). Currently, the primary type of analysis run within MORF is predictive modeling; as such, once the job is complete, the model is evaluated using a predefined set of functions and evaluation metrics are sent to the user's email (linked through the API key). By preventing viewing of personally identifying information and offering access to a restricted set of evaluation metrics, MORF allows extensive analysis but protects student privacy.

B. Current Data

Our university's instance of MORF currently provides access to data from 52 MOOCs (45 taught in English), with additional data being continuously added. This data consists of interactions from around two million learners drawn from over 100 countries. The following data is available:

- Clickstream: all the learner clicks within their course
- Discussion forum posts: all threads, posts, and comments made within course discussion forums (including learners and instructors).
- Course data: data on every other part of the MOOC experience, including learner viewing of lecture videos, quiz-taking, and assignment submission.

A complete list of courses and the number of learners enrolled in each course (accurate at the time of this writing) can be seen in Table 1. This count is continually changing as students take courses. Learners may also be removed from the system at their request under data privacy agreements and legislation (e.g., GDPR)

Table 1. List of courses included in the current installation of MORF including number of learners (#Learner), number of learners that completed the course and received a certificate (#Comp), and number of times offered (Num)

Course	**#Learner**	**#Comp**	Num
ADHD Through the Lifespan	19,591	6,248	1
American Education Reform	3,224	961	1
Analyzing Global Trends for Business and Society	52,671	17,507	3
Applying to U.S. Universities	38,230	16,768	2
Arts Culture & Strategy	7,908	1,635	1
Better Leader, Richer Life	86,942	15,210	3
Calculus: Single Variable	122,412	27,739	6
Cardiac Arrest, Resuscitation Science, Hypothermia	18,897	7,195	1
Design: Creation of Artifacts in Society	115,661	25,866	6
Designing Citie	36,393	10,434	2
Experimental Genome Science	31,106	14,675	2
Business Foundations: Accounting	87,760	39,172	3
Business Foundations: Accounting*	514	223	1
Business Foundations: Accounting*	4,001	1,809	2
Business Foundations: Corporate Finance	42,656	17,893	2
Business Foundations: Corporate Finance*	639	147	1
Business Foundations: Marketing	106,827	46,681	4
Business Foundations: Marketing*	840	384	1
Business Foundations: Marketing*	4,003	1,576	2
Business Foundations: Operations	66,055	30,856	4
Business Foundations: Operations*	446	256	1
Business Foundations: Operations*	2,788	1,251	2
Fundamentals of Pharmacology	33,998	11,015	1
Gamification	204,476	74,917	4
Going Out on A Limb	15,127	6,170	1

Greek and Roman Mythology	106,819	25,059	4
Growing Old Around the Globe	12,693	3,532	2
Health Policy and the Affordable Care Act	40,598	12,932	2
History of the Slave South	18,213	4,257	2
Intro to American Law	15,349	7,221	1
Intro to Corporate Finance	61,280	26,635	1
Intro to Dental Medicine	5,537	2,325	1
Intro to Financial Accounting	184,440	48,407	2
Intro to Key Constitutional Concepts	15,100	4,772	2
Intro to Marketing	206,786	76,626	3
intro to Operations Management	259,651	74,227	5
Listening to World Music	31,969	6,816	1
Modern and Contemporary American Poetry	84,403	19,158	4
More Financial Accounting	12,697	5,678	2
Networked Life	57,049	15,819	3
Neuroethics	10,775	3,232	1
New Health Policy	2,627	1,231	1
New Health Policy II	1,607	652	1
Principles of Microeconomics	20,931	4,852	1
Probability	33,664	8,370	2
Rationing and Allocating Scarce Medical Resources	5,054	1,457	1
Revolutionary Ideas: An Intro to Legal	16,315	3,628	1
Social Entrepreneurship	22,731	6,902	2
Sustainability in Practice	13,012	2,939	1
The Global Business of Sports	5,856	2,526	1
Vaccines	36,898	15,154	3
Vital Signs	64,368	21,374	3

Note. * indicates identical course content taught in a language other than English.

Each of the courses in Table 1 consists of a variety of data actions completed by students. Though these may vary by context or course (i.e., some courses use different pedagogical features than others), there are consistencies across courses that allow for cross-course analysis without extensive data pre-processing. Table 2 details the number of actions held in MORF for common action categories such as forum actions or video actions. As with the data in Table 1, this is correct at the time of writing but is subject to change as the repository continues to grow.

C. Research Supported To Date

MORF has supported several analyses by researchers both at our university and other universities. To give a few examples, [4] used MORF to establish and study a new metric for measuring model performance differences between demographic groups of students. [5] used MORF to study the impact of different feature sets on MOOC dropout prediction, attempting (and failing) to replicate a previously-published set of results using a much larger data set. [1] used MORF to study the degree to which predictive models transfer between countries. [9] used MORF to study the longitudinal positive career impacts of taking MOOCs, linking MOOC participation data to a survey data set on career impacts. In general, MORF enables a wide range of machine learning and statistical analyses for researchers.

Table 2. Action counts for data contained in the current installation of MORF

Feature	Total Number
Total number of clicks related to any forum activity (e.g., viewing, posting, commenting)	20,883,493
Total number of clicks related to any quiz activity (e.g., viewing, answering, submitting)	19,308,718
Total number of clicks to any peer-assessment-related activity	19,022,016
Total number of clicks related to any video lecture activity (e.g., playing, pausing, increasing video speed, etc.)	3,542,058
Total number of forum threads started	360,433
Total number of responses to others' forum posts	38,691
Total number of others' responses on one's own forum posts	48,277

MORF has also been used internally to examine the effects of studies on various student outcomes within our courses. Some examples of this involving the E-TRIALS architecture are given in the next section, but MORF has also been used to analyze studies conducted solely using standard MOOC functionality. For example, researchers in one department of our university conducted a study on an intervention (originally published in [6]) that asked students to write about their values prior to starting the course. Students were encouraged to consider how those values connected to their goals for taking the MOOC. This study (conducted in one MOOC) found increased completion and higher assignment grades, a finding the researchers are now replicating in other MOOCs. Researchers in another department of our university conducted a study on messages that give students information on when successful students start each assignment. The study investigated whether to give these messages for every assignment or more rarely, and found that students who received more messages achieved marginally significantly higher performance in the second half of the course.

D. Opportunities to use MORF

There are two main opportunities for researchers looking to use MORF in their own MOOC research. The first opportunity is for researchers with MOOC data that they wish to share via this framework. All MORF code is open source, with detailed information regarding how the individual elements connect and interact. This code and documentation can then be used to implement additional instances of MORF, separate from the instance at our university and the other universities currently using MORF. This allows institutions to provide the same privacy-preserving data sharing as described above whilst keeping all their own data directly under their own control (i.e., it is not necessary to send data to our university to share data via MORF). Researchers wishing to set up their own instance of MORF will require detailed knowledge of Python, Databases, relational calculus to set up database queries, and Amazon Web Services. In theory, MORF could be run on an alternate cloud computing service (other than AWS); however, this would likely involve making substantial edits to the publicly available source code.

The second opportunity for researchers is to use the instance of MORF available at our institution (or other universities using MORF) for their own research questions. In order to use our university's instance, researchers must sign a data agreement to receive an API key. The key then allows researchers to submit jobs to MORF and receive the results. There are currently no restrictions on who can receive a MORF API key, researchers simply need to apply[1]. As noted above, we currently have a predefined selection of evaluation metrics. Users may define their own evaluation metrics and submit these to the MORF administrators for inclusion. Processes are checked for any potential security risk and then included on the MORF production server.

We recommend that researchers wishing to use the installation of MORF currently running at our institution (or at other institutions) have the following technical skills: (1) Elementary knowledge of JSON and Python to use the MORF job submission API. (2) Basic knowledge of Docker to make sure their job runs as expected. (3) Basic knowledge of SQL to pull the appropriate data for their analysis. (4) Sufficient knowledge of a programming language to conduct the chosen analysis. Note that, because of the use of Docker containers, researchers can run analysis in the language of their choice, providing it can be installed in a Docker container. By providing this flexibility, we reduce the limitations placed upon researchers and facilitate the investigation of a wider variety of research questions.

We provide a number of resources to further aid researchers as they set up their experiments using MORF. The MORF repository contains documentation and a set of minimum working examples (MWEs) that researchers can build upon. The MWEs detail two parts of MORF analysis: (1) designing an analysis, and (2) submitting the job to MORF. Researchers can use an MWE as a proof of concept, confirm that job submission works, and that results are as expected. We recommend that all researchers start by submitting an MWE to inform any later troubleshooting. In addition to the online resources for researchers using the MORF installation at our university, our team is also available to provide continual support in accessing the data, designing analysis, and troubleshooting any issues that arise.

Regardless of the option used to pursue research with MORF (i.e., using our installation or setting up an institution-specific installation), we encourage researchers to use this tool to also support the replicability of work. Using the Docker architecture facilitates researchers in providing exact details of their analysis for future researchers, including specific library versions and idiosyncrasies of feature engineering and algorithm implementation. This, in turn, means that all analyses performed in MORF can be reproduced without compromising student privacy. This approach facilitates a more standardized analysis of MOOC data and makes it easier to replicate/verify results, strengthening the field's ability to adhere to open research principles.

III. E-TRIALS FUNCTIONALITY FOR MOOCS

E-TRIALS, the Ed-Tech Research Infrastructure to Advance Learning Science (previously called the ASSISTments Testbed), was developed by researchers at Worcester Polytechnic Institute and the ASSISTments Foundation [7]. E-TRIALS has been used in research by dozens of researchers to conduct randomized controlled trials in middle school mathematics. Our university integrated E-TRIALS with our MOOC platforms using LTI integration. Doing so enabled researchers to conduct complex experiments and deliver assignments and activities that are more complex than the assignments and activities natively available in our MOOC platforms. After conducting an experiment, the data from the experiment is then made available for analysis within MORF, and data from ASSISTments activities can be linked to other MOOC data (such as forum data, video watching, and course completion).

A. Experimentation Options

To support experimentation, E-TRIALS (and ASSISTments) enables random assignment within either individual assignments or across a student's course experience (e.g., across multiple assignments within ASSISTments). This partnership currently offers researchers two types of experiments that are outlined below.

1) Item Varying Assignments

The first option for experimentation is item varying assignments. In this category, the experimental design relies upon variations in the design of individual questions that students are being asked in the assignment. The E-Trials system randomly assigns students to a given condition; this condition assignment can then be preserved across multiple assignments if necessary.

Condition differences may include variations in wording or presentation. For example, in a two-condition experiment, one condition might receive a question as plain text, whereas students in the second condition would receive a question with images to support the text. Such an experiment could support instructional design goals as well as provide insights into student preferences and effective communication with a given population.

Another option for item varying assignments would be to totally change the items themselves. For example, students in

[1] API requests can be sent to gse-pcla-morf@gse.upenn.edu

one condition may receive problem set A, and another may receive problem set B, for classic A/B testing. Such an experiment could assess a wide variety of research questions, including those surrounding assessment and auto-grading policies.

Further, E-TRIALS presents the opportunity to vary the number of items in an assignment or present items in different situations or orders, allowing researchers to examine potential ordering effects or counterbalance their study design.

2) Feedback Varying Assignments

The second option for experimentation is feedback varying assignments. In this category of experiments, students are randomly assigned to receive different feedback on their correct or incorrect answers. Researchers can vary how and when students receive feedback.

Considering first the "when", researchers can change if a student receives feedback only after submitting an incorrect answer, or if a student can request feedback/assistance prior to submitting an answer. If researchers choose to provide feedback after an incorrect answer, they can also choose if the student must resubmit an answer, work through multi-step scaffolding, or can proceed without making any changes.

Next considering variations in "how" feedback is delivered, researchers have additional options for experimentation. Problem feedback can take many forms. For example, it could simply be a hint, highlighting an important part of the problem text, or prompting a student to consider a certain fact, rule, or proof from the course content. Feedback could also be more detailed and provide answer scaffolding or split the problem into multiple parts to make the problem easier for the student. These variations have been shown to have different effects for different age groups [see 10] and to be highly context-specific (e.g., scaffolding may be more beneficial depending on the abilities and needs that are common in a given course [14]), making it useful to determine which form of feedback is most appropriate for a given MOOC.

As noted above, the E-TRIALS framework facilitates the creation of these experiments, providing support for experimental designs such as counterbalancing techniques and control for ordering effects.

B. Data provided by E-TRIALS

The E-TRIALS system provides a variety of data regarding each student's behavior and performance within a given assignment. Researchers receive both the raw answers from the students as well as the answers coded for correctness, along with semantic information such as the skill each problem is tagged with.

Interaction data provided includes when a student started and completed the assignment and start and end times for each individual item within the assignment. From this information, the total time for the assignment and time per item is also calculated. Data is also provided regarding the number of attempts made and actions taken within each item. All actions are timestamped for synchronization with other data streams (e.g. discussion forum data).

In order to match E-TRIALS data to the external MOOC, E-TRIALS records two identifiers per student. The first is an internal identifier for matching across multiple E-TRIALS documents. The second is the LTI identifier. This is a unique ID generated by the MOOC and passed through the LTI integration between the two applications. This identifier is what is used to match a student's E-TRIALS/ASSISTments data to their MOOC data. Note that the LTI ID is not the same as their ID within the MOOC; this ID is used for the pairing of data between the two systems.

C. Advantages for MOOC Research

Through integrating the ASSISTments platform into MOOCs, E-TRIALS also enables richer assignments with scaffolding (assigning a specific next problem based on a specific next error), on-demand hints (with video embedding), mathematical formula input, and a range of other features. It also outputs rich description of student interaction within the activity, including automated detection of student disengaged behaviors and affect [2].

From a research perspective, this rich interaction behavior can be aligned (via timestamps) with MOOC behavior and clickstream data for a more in-depth analysis of student patterns. Researchers can relate individual assignment behaviors to behaviors such as broader help-seeking strategies. For example, a researcher may track students who move from the assignment to a MOOC resources (e.g., text files or video lectures) mid-assignment or examine how many attempts were made before the student posted on the discussion forum requesting assistance from their peers/the instructor. By combining the two data streams, researchers can access a complete picture of student behavior and how they are approaching their work, providing more detail than either datastream alone.

E-TRIALS has also been used in partnerships between course designers and instructors (providing a venue for experiments) and researchers (providing study ideas) both at our university and externally. For example, a multi-institution collection of researchers proposed a study to replicate a past finding that scaffolding activities are more effective than hints at promoting learning in middle school students [8]. A replication conducted in a MOOC at our university found the opposite result from this earlier work [10], suggesting a difference in how feedback should be designed for adult learners versus school-age learners. An upcoming study brings together researchers from two countries to study the impacts of different dosages of in-video quizzes in one MOOC.

D. Opportunities for Future Research

We encourage any and all researchers interested in running a study of this kind to contact us, whether they have a suitable MOOC available or not. We can then act as an intermediary, connecting researchers with similar interests and maximizing resources. Similarly, take a researcher with a scientifically interesting research question and/or experiment design but no MOOC course in which to run it. In that case, we can leverage our existing network of MOOC instructors to find a potential course and foster a research partnership.

With regard to the technical skills required for conducting research in E-Trials, both E-Trials and MOOC platforms provide a detailed GUI for setting up an experiment and the broader MOOC that will contain the experiment. As such, researchers need no detailed technical knowledge to set up their experiments. E-Trials automatically provides basic reports on differences between conditions. If a researcher wishes to analyse the data generated in the experiment in more complex fashions, they will need to use MORF and will need the technical skills outlined in Section II.

IV. CONCLUSIONS

The use of MORF and E-TRIALS in our university's MOOCs provides us with a robust infrastructure for conducting two key types of modern educational research: A/B testing and secondary data analysis. This infrastructure can be used both by researchers at our university and external researchers. Credentials for our university's installation of MORF are now available upon request, and studies can be conducted in our MOOCs using E-TRIALS by researchers anywhere in the world, again by request. We are aware that not all researchers in the world have had access to research infrastructure of this nature, and we are committed to making it accessible to the wider scientific community. This project aims to increase access for scholars who previously have not had the opportunity to conduct research with large-scale data sets and MOOC courses of this nature.

V. ACKNOWLEDGEMENTS

This research was supported by the National Science Foundation (NSF) (NSF-OAC#1931419). Any opinions, findings, and conclusions, or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the NSF.

REFERENCES

[1] J.M.L. Andres-Bray, "Replication in Massive Online Open Course Research Using the MOOC Replication Framework". Doctoral Dissertation, University of Pennsylvania, 2021.

[2] A.F. Botelho, R. Baker, N. Heffernan, "Improving Sensor-Free Affect Detection Using Deep Learning". In Proceedings of the 18th International Conference on Artificial Intelligence in Education, pp. 40-51, 2017

[3] J. Gardner, M. Andres-Bray, C. Brooks, R. Baker, "MORF: A Framework for Predictive Modeling and Replication at Scale With Privacy-Restricted MOOC Data". In Proceedings of the 3rd Workshop on Open Science in Big Data, 2018.

[4] J. Gardner, C., Brooks, R. Baker, "Evaluating the Fairness of Predictive Student Models Through Slicing Analysis." In Proceedings of the 9th International Learning Analytics and Knowledge Conference, pp. 225-234, 2019.

[5] J. Gardner, Y. Yang, R. Baker, C. Brooks, "Modeling and Experimental Design for MOOC Dropout Prediction: A Replication Perspective". In Proceedings of the 12th International Conference on Educational Data Mining, pp. 49-58, 2019.

[6] R.F. Kizilcec, G.M. Davis, G. L.,Cohen, G. L. "Towards equal opportunities in MOOCs: affirmation reduces gender & social-class achievement gaps in China". In Proceedings of the fourth (2017) ACM conference on learning@ scale, pp. 121-130, 2017.

[7] K.S. Ostrow, N.T. Heffernan, "Studying learning at scale with the ASSISTments TestBed". In Proceedings of the Third (2016) ACM Conference on Learning@ Scale, pp. 333-334, 2016.

[8] L. Razzaq, N.T. Heffernan, "Scaffolding vs. hints in the Assistment System". In International Conference on Intelligent Tutoring Systems, pp. 635-644, 2006.

[9] A. Trumbore, "Learner Behavior and Career Benefits in Business Massive Open Online Courses". In Proceedings of the 15th International Conference of the Learning Sciences-ICLS 2021.

[10] Y. Zhou, J.M. Andres-Bray, S. Hutt, K. Ostrow, R.S. Baker, "A Comparison of Hints vs. Scaffolding in a MOOC with Adult Learners". In Proceedings of the International Conference on Artificial Intelligence, 2021.

[11] Dowell, N. M., Skrypnyk, O., Joksimovic, S., Graesser, A. C., Dawson, S., Gaševic, D., Hennis, T., de Vries, P., Kovanovic, V. "Modeling Learners' Social Centrality and Performance through Language and Discourse". In Proceedings of the International Educational Data Mining Society, 2015.

[12] Bosch, N., Crues, R. W., Paquette, L., Shaik, N. ""Hello,[REDACTED]": Protecting Student Privacy in Analyses of Online Discussion Forums". In Proceedings of the International Conference on Educational Data Mining, 2020.

[13] Baker, R.S., Hawn, M.A. "Algorithmic Bias in Education". To appear in International Journal of Artificial Intelligence and Education , in press.

[14] Kim, N. J., Belland, B. R., Axelrod, D. "Scaffolding for optimal challenge in K–12 problem-based learning". In The Interdisciplinary Journal of Problem-based Learning, 13 (1), Article 3, 2018.

[15] Penn Center for Learning Analytics, "Algorithmic bias in Education," Algorithmic Bias in Education - Penn Center for Learning Analytics Wiki. [Online]. Available: https://www.pcla.wiki/index.php/Algorithmic_Bias_in_Education.