
Enabling End-To-End Machine Learning Replicability: A Case Study in Educational Data Mining

Josh Gardner
School of Information
The University of Michigan
jggard@umich.edu

Yuming Yang
Department of Statistics
The University of Michigan
yangym@umich.edu

Ryan S. Baker
Graduate School of Education
The University of Pennsylvania
rybaker@upenn.edu

Christopher Brooks
School of Information
The University of Michigan
brooksch@umich.edu

Abstract

The use of machine learning techniques has expanded in education research, driven by the rich data from digital learning environments and institutional data warehouses. However, replication of machine learned models in the domain of the learning sciences is particularly challenging due to a confluence of experimental, methodological, and data barriers. We discuss the challenges of end-to-end machine learning replication in this context, and present an open-source software toolkit, the MOOC Replication Framework (MORF), to address them. We demonstrate the use of MORF by conducting a replication at scale, and provide a complete executable container, with unique DOIs documenting the configurations of each individual trial, for replication or future extension at <https://github.com/educational-technology-collective/fy2015-replication>. This work demonstrates an approach to end-to-end machine learning replication which is relevant to any domain with large, complex or multi-format, privacy-protected data with a consistent schema.

1 Introduction

The repeated verification of scientific findings is central to the construction of robust scientific knowledge, particularly in a fast-growing field such as machine learning. This can take the form of (a) reproduction (reproducibility), using the original methods applied to the original data to reproduce the original results, and (b) replication (replicability), applying the original methods to *new* data to assess the robustness and generalizability of the original findings. Since reproducibility is a necessary condition for replicability (an experimental procedure cannot be applied to *new* data if the procedure cannot even be reproduced), any replication procedure also requires solving the problem of reproducibility.

In this work, we discuss the reproducibility crisis in machine learning, noting specific challenges faced by applied researchers in the learning sciences, particularly in the sub-fields of educational data mining and learning analytics. We argue that existing frameworks for reproducible machine learning, such as open code-sharing platforms and public code notebooks (such as those solicited by this workshop) are valuable steps, but are insufficient to fully address the challenges both within our subfield of interest and the broader machine learning community. In particular, we argue that

code-sharing does not address the breadth of challenges – experimental, methodological, and data – we face as practitioners, as Section 3 details.

Instead, we propose a paradigm of *end-to-end* reproducibility for machine learning: fully reproducing (or replicating) the pipeline from raw data to model evaluation. End-to-end reproducibility is possible with current freely-available computing technologies, namely containerization, which can encapsulate an experiment as well as operating system and software dependencies in a single package. We describe a novel open-source platform for conducting reproducible end-to-end machine learning experiments on large-scale educational data, the MOOC Replication Framework (MORF), and describe additional benefits, beyond reproducibility, afforded by this platform in Section 6.

2 Prior Work

2.1 The Reproducibility Crisis in Machine Learning

Much has been written about the reproducibility crisis in science, particularly in fields which conduct human subjects research such as social psychology. Recent empirical evidence has shown that issues with reproducibility are also widespread in the field of machine learning. In some ways, this realizes concerns that have been voiced in the field for nearly twenty years concerning the complex technical steps and dependencies involved in conducting a machine learning experiment which are rarely described or shared in adequate detail with existing publication methods. A recent survey of 400 research papers from leading artificial intelligence venues shows that none of the works surveyed document all aspects necessary to fully reproduce the work; only 20-30% of the factors evaluated were adequately reported in the works surveyed [15]. In a case study of replication of deep reinforcement learning algorithms, [17] show that the variance inherent to statistical algorithms, the use of different hyperparameter settings, and even different random number generation seeds contribute to a lack of reproducibility in machine learning research and have a direct impact on whether experimental results and baseline model implementations replicate. Olorisade et al. [23] survey 30 machine learning studies in text mining, and identify poor reproducibility due to lack of access to data, software environment, randomization control, and implementation methods. None of the 30 works surveyed provided source code, and only one of 16 applicable studies provided an executable, to reproduce their experiments.

These reproducibility issues are partly attributable to culture and convention. A survey of authors published in the *Journal of Machine Learning Research* found that roughly one third intentionally did not make their implementations available, for reasons including a lack of professional incentives, a reluctance to publish messy code, and the convention that doing so is optional [27]. Peng [24] observes that only five of 125 published articles in the journal *Biostatistics* have passed the (voluntary) reproducibility review since its inception two years prior, yet considers this effort “successful” compared to the reproducibility of previous work.

As big data and machine learning permeate disciplines, this crisis in replication has also affected other fields of study, such as the learning sciences. This is especially relevant in cases of very large datasets where the majority of learning is computer-mediated, such as in Massive Open Online Courses (MOOCs). For example, [14] showed that a large-scale replication of machine learning models led to substantially different conclusions about the optimal modeling techniques, with several findings replicating significantly in the *opposite* direction of the original study (which was conducted on only a single MOOC). Khajah et al. [19] attempted to replicate the findings of the “deep knowledge tracing” method introduced in [25], and demonstrated that much simpler baseline methods could achieve equivalent performance, and that the initial performance gains demonstrated in the original work were at least partially due to data leakage.

2.2 Existing Tools for Reproducible Machine Learning

An exhaustive survey of tools and platforms to support reproducible machine learning research is beyond the scope of this work. However, we include a survey of those tools which most closely align with our interests in building machine learned models for predictive analytics in education.

OpenML is “an open, organized, online ecosystem for machine learning” and allows users to create data science pipelines to address specific “tasks”, such as classification and clustering [29]. OpenML requires users to run their analyses locally and upload the results. OpenML’s workflow may




Figure 1: End-to-end reproducibility requires addressing data, technical, and methodological issues with reproducibility. Replication of the computational environment, in particular, is key to replicating this complete pipeline from raw data to results.

contribute to “openness,” but does not ensure reproducibility, as it relies on user-submitted code to fully reproduce the results which are generated and submitted by users. OpenML functions as a code-, data-, and results- sharing platform, but these analyses are subject to key limitations (dependency management, code rot) described in more detail in Section 3. Finally, OpenML presumes data can be shared freely for replication, an often incorrect presumption in the field of the learning sciences, where institutional and governmental regulations and policies manage how data can be shared.

The OpenAI Gym is an open-source interface for developing and comparing reinforcement learning algorithms [3]. Its wide use for both teaching and research serve as an example of how a field can create and adopt shared tools to support researchers’ needs and support reproducibility. For example, [17] uses this platform for its reproducible baseline models. However, OpenAI Gym supports only reinforcement learning (not other machine learning tasks such as supervised learning).

Several publishing platforms dedicated to reproducible computational research also exist, such as ReScience ¹, CodaLab ², and WholeTail [2]. These platforms unify code, data, computation, and presentation in a single location. CodaLab and WholeTail also use Docker containerization to ensure reproducibility. However, these platforms generally provide no support or restricted-access data; for example, WholeTail specifically limits its scope to public (non-restricted) datasets [2].

Each of these platforms is an important step toward reproducible machine learning research, and many of them address key barriers. However, these tools are insufficient for many types of machine learning tasks, including supervised learning with large-scale behavioral data from MOOCs. In particular, none of these platforms supports replication where the underlying data sets are privacy-restricted and cannot be publicly shared. In educational data, many of the types of unanonymizable data that are privacy-restricted are also necessary for analysis (such as the text of discussion forum postings, IP addresses, or student names). Such restrictions are also likely to drive away machine learning researchers from working with this data, as gaining access to unprocessed raw educational data can be difficult or impossible without close collaborators and strong institutional support. Even with institutional support, government regulations such as the Family Educational Rights and Privacy Act (FERPA) may restrict or complicate data sharing.

3 Barriers to End-to-End Replication in Machine Learning

The replication crisis is the result of a confluence of forces which must be collectively addressed in order to achieve end-to-end reproducibility. We group these challenges into experimental, methodological, and data challenges. No existing solution discussed in Section 2.2 currently addresses all three barriers.

Experimental challenges with reproducibility relate to reproducing the exact experimental protocol.³ Many have advocated for the open sharing of code as a potential solution to address technical issues with reproducibility [e.g. 28, 24]. However, while necessary, code-sharing is insufficient to guarantee reproducibility in machine learning research. For example, Collberg et al. [7] showed that the published code accompanying 20% of their large sample of 613 published computer systems papers failed to build or run, and in total, it was not possible to verify or reproduce 75.1% of studies surveyed using the artifacts provided in publication.

¹ <http://rescience.github.io/about/>

² <http://codalab.org/>

³The term “experimental reproducibility” is adopted from [15].

Even when code is available, other technical issues can prevent reproducibility in computational research workflows [10, 20]. Buckheit and Donoho [4] noted over 20 years ago that the complete software environment is a necessary condition for reproducing computational results; however, the open sharing of such environments along with the published results of scientific work remains rare. These include *code rot*, in which code becomes non-functional or its functionality changes as the underlying dependencies change over time (for example, an update to a data processing library which breaks backwards compatibility, or a modified implementation of an algorithm which changes experimental results), as well as *dependency hell*, in which configuring the software dependencies necessary to install or run code prevents successful execution [1]. This complex web of interdependencies is rarely described or documented in published machine learning and computational science work [10, 15, 23].

Fortunately, modern containerization tools were developed to resolve such technical issues in software development contexts [22]. Docker containers, for instance, are frequently used in both industrial software applications as well as computational and computer systems research [1, 6, 18]. A major advantage of Docker over simple code-sharing is that Docker containers fully reproduce the entire execution environment of the experiment, including code, software dependencies, and operating system libraries. These containers are much more lightweight than a full virtual machine, but achieve the same level of reproducibility [22, 18]. While some of the platforms discussed in Section 2.2 utilize containerization, these platforms “hide” this functionality from the user, assembling containers from user-submitted code. This severely limits users’ ability to fully leverage containerization by building the complex, customized environments many machine learning experiments may require.

Methodological challenges to reproducibility reflect the methods of the study, such as its procedure for model tuning or statistical evaluation. Existing work on reproducibility largely focuses on strictly technical challenges, but (as our experiment in Section 4 shows), methodological issues are at least as important. Methodological challenges include the use of biased model evaluation procedures [5, 30], the use of improperly-calibrated statistical tests for classifier comparison [9], or “large-scale hypothesis testing” where thousands of hypotheses or models are tested at once, despite the fact that most multiple testing corrections are not appropriate for such tasks [11]. A machine learning version of these errors is seen in massive unreported searches of the hyperparameter space, and in “random seed hacking” wherein the random number generator itself is systematically searched in order to make a target model’s performance appear best or a baseline model worse [17]. Replication platforms can address methodological issues – working toward what [15] terms *inferential reproducibility*– by architecting platforms which support effective methodologies and adopt them by default, effectively nudging researchers to make sound choices.

Data reproducibility concerns the availability of data itself. In many domains, making raw data available is more an issue of convention than a true barrier to reproducibility. However, in the case of educational data mining, data are governed by strict privacy regulations which protect the privacy of student education records. Similar restrictions affect many other fields, from the health sciences to computational nuclear physics [20]. As a result, researchers are often legally prohibited from making their data available. Efforts such as the Pittsburgh Science of Learning Center DataShop [21] and the HarvardX MOOC data sets [16] have attempted to address this problem in education by only releasing non-identifiable data, but many analyses require the original, unprocessed data for a full replication. Indeed, restricted data sharing is one of the main factors (in our experience) hindering generalizability analysis in educational data mining: investigators are generally limited to one or two courses worth of data (e.g. the courses they instruct), and models are often overfit to these datasets.

4 The MOOC Replication Framework

MOOC Replication Framework (MORF) is a Python toolkit, accompanied by a platform-as-a-service (the “MORF Platform”), which collectively address the challenges faced by researchers studying large-scale online learning data noted above.⁴ MORF addresses (i) experimental barriers to replication via the use of containerization; (ii) methodological issues by providing support for true holdout sets and reliable model evaluation; (iii) data issues by providing a large-scale dataset while preventing the download of sensitive information (currently, this includes the complete raw data exports from over 270 MOOCs offered by the University of Michigan and the University of Pennsylvania). Additionally,

⁴The MORF website, which includes documentation and short tutorials, is at <https://educational-technology-collective.github.io/morf/>

MORF eases the computational expense of conducting such research at scale by providing nearly an order of magnitude greater computational infrastructure than any of the platforms discussed in Section 2.2, and out-of-the-box parallelization to utilize it.

Users submit jobs to the MORF Platform using a simple Python API. A short, four-line “controller” script is used to guide the execution of each stage of the end-to-end supervised learning pipeline (extract, train, test, and evaluate). The use of controller scripts is a best practice in computational research [20], and MORF’s combination of containerization and controller scripts allow the user to manage low-level experimental details (operating system and software dependencies, feature engineering methods, and statistical modeling) by constructing the Docker container which is submitted to MORF for execution, while MORF manages high-level implementation details (parallelization, data wrangling, caching of results). The containers used to execute each job run on MORF are persisted in MORF’s public Docker Cloud repository, and the configuration file and controller scripts are persisted in Zenodo and assigned a unique Digital Object Identifier (DOI). This yields a reproducible end-to-end pipeline that is highly flexible, easy to use, and computationally efficient.

4.1 Experimental Reproducibility via Containerization

MORF uses containerization to resolve many of the experimental challenges described in Section 3. The Docker containers submitted to MORF fully encapsulate the code, software dependencies, and execution environment required of an end-to-end machine learning experiment in a single file, ensuring end-to-end reproducibility and enabling sharing of the containerized experiment. Building Docker containers requires only a single Dockerfile (akin to a makefile) which contains instructions for building the environment. This imposes minimal additional burden on researchers relative to configuring, programming, and executing an experiment, but achieves a considerable increase in reproducibility. While other existing machine learning research platforms sometimes utilize docker “under the hood” for reproducibility, this limits users’ ability to configure or share these environments. We are not aware of any which allow users to build and submit Docker images directly for execution.

As part of MORF, we are assembling an open-source library of pre-built Docker containers to replicate experiments conducted on MORF to serve as shared baseline implementations. These containers can be loaded with a single line of code, allowing the research community to replicate, fork, interrogate, modify, and extend the results presented here⁵.

4.2 Methodological Reproducibility via Platform Architecture

MORF provides sensible default procedures for many machine learning tasks, such as model evaluation. For example, MORF avoids the use of cross-validation for model evaluation: The prediction tasks to which most MOOC models aspire are prediction of *future* student performance (i.e., in an ongoing course where the true labels – such as whether a student will drop out – are unknown at the time of prediction). As such, using cross-validation within a MOOC session, when the outcome of interest is accuracy on a *future* MOOC session, provides an unrealistic and potentially misleading estimate of model performance. Prior work has demonstrated that cross-validation provides biased estimates of independent generalization performance [30], and in the MOOC domain, that cross-validation can produce biased estimates of classification performance on a future (unseen) course [31]. Adopting more effective model evaluation techniques by default requires no additional work for MORF users, and ensures that work produced on the MORF platform follows effective model evaluation procedures.

4.3 Data Reproducibility via Execute-Against Access

MORF achieves data reproducibility while also meeting data privacy restrictions by providing strictly “execute-against” access to underlying data. Most MOOCs are generated by a small number of platforms (e.g. Coursera, edX), and all courses from a given platform use publicly-documented data schemas [e.g. 8]. Thus, users can develop experiments using their own data from a given platform – or even the public documentation – and then submit these experiments for MORF to execute against *any* other course from that platform. This enables MORF to currently provide an interface to over

⁵The experiment presented below can be loaded by running `docker pull themorf/morf-public:fy2015-replication` in the terminal of any computer with Docker installed.




Figure 2: Original results from [12] (left) and replication results using the MOOC Replication Framework evaluated using a held-out future course session (center) and cross-validation (right). 95% confidence intervals shown. IOHMM models not replicated due to lack of an open-source implementation which supported prediction.

200 unique sessions of more than 70 unique courses offered by two different institutions on the Coursera platform, and to execute containerized experiments against this data in a secure, sandboxed environment, by utilizing the shared, public schema of MOOC datasets. These shared public data schemas also ensure that existing experiments in MORF can be replicated against new data (from the same MOOC platform) as it becomes available.

5 Case Study in Replication with MORF

We demonstrate the use of MORF through a case study. We replicate a comparison conducted in [12], which compares several machine learning algorithms using a set of seven activity features (e.g. number of lecture videos viewed, quizzes attempted, and discussion forum posts for each student) over each week in a MOOC in order to predict a binary dropout label indicating whether a user showed activity in the final week of a course.

This study is an ideal candidate for replication because it compares several models (effectively testing many pairwise hypotheses) using cross-validation on only a single dataset. This testing of many hypotheses/comparisons, with only a single observation for each, can lead to poor methodological reproducibility and provides no information about the variability of the estimates, relative to their magnitude. Particularly because this experiment was concerned with empirical performance (in order to inform future “early warning” dropout prediction systems), obtaining an accurate estimate of models’ expected performance on *future* course sessions across a large, representative dataset can provide insight into the generalizability of these findings. This work is indicative of prediction tasks in the fields of learning analytics and educational data mining (see [13] for a more thorough discussion).

We replicated the original experiment across 45 unique MOOCs using MORF. Results are shown in Figure 2.⁶

The original work argued that a Long Short-Term Memory (LSTM) neural network model “beats the ... other proposed methods by a large margin” [12, pp. 1]. Our results show, however, that (1) LSTM is actually one of the worst-performing models, with the lowest average performance of any model tested in weeks 4 and 5; (2) in most cases, the 95% confidence intervals for algorithm performance overlap (indicating no significant difference); and (3) observed performance of all models is lower than observed in [12], particularly in later weeks of the course.

We hypothesize that (1) may be due to overfitting on the original training data. Particularly when using cross-validation on a single dataset with a highly flexible model such as LSTM, the original work was quite susceptible to overfitting. Overfitting seems particularly likely because no procedure for selecting hyperparameters was reported in the original work, and some relevant hyperparameter settings for the final LSTM model (e.g. batch size) were not reported at all. These hyperparameters were not available even after correspondence with the authors, who did not record them and no longer had the original code available (which itself points to the need for reliable long-term reproducibility solutions such as MORF). Point (2) shows the advantage of using MORF’s large data repository, which allows us to observe variability in each algorithms’ performance across many MOOCs. This result suggests that while differences in average performance may exist, these are too small to be

⁶We would like to acknowledge the helpful assistance of the original authors in conducting this replication.

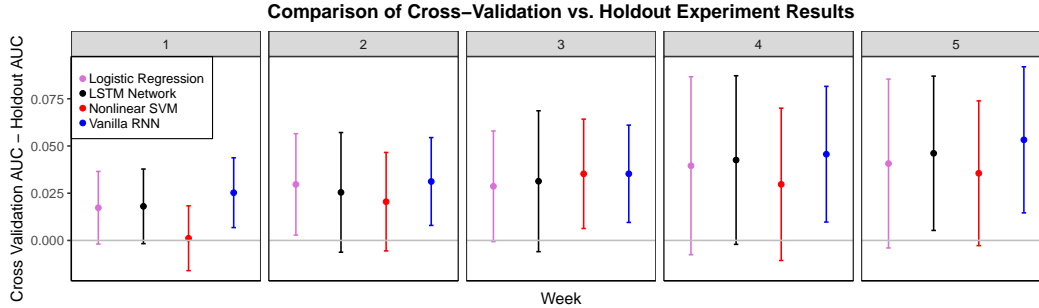


Figure 3: Comparison of AUC estimates from identical experiments using a holdout vs. cross-validation evaluation architecture. These results show a persistent positive bias when cross-validation is used to estimate predictive performance on a future, held-out session of a MOOC.

interpreted as genuine and not spurious – particularly in light of the results shown in Figure 3, which shows that the differences due to cross-validation bias are, in many cases, larger than the observed differences between algorithms. Finally, (3) is likely due to the cross-validation used in the original experiment (as opposed to the holdout architecture used in MORF). The original experiment allowed models to train and test on subsamples from the same overarching population, which the holdout architecture used in MORF requires the (more challenging) task of predicting student performance on a future course session – which is the true target of prediction in this work.

The right panel of Figure 2 demonstrates the results of an identical replication experiment, but evaluated using five-fold cross-validation (as in [12]) instead of the future held-out course session used in the center panel. The contrast between the center (holdout) and right (cross-validation) panels demonstrates the optimistic bias which can be introduced by evaluating generalization performance via cross-validation without the use of an independent hold-out set. This comports with previous results demonstrating that the bias of performance estimates when models are optimized over cross-validation folds can often exceed the difference between learning algorithms [5]. These results are further demonstrated by Figure 3, which shows a small but persistent positive bias for model evaluation performed by cross-validation versus the “true” performance on a future course session. A two-sided Wilcoxon signed rank test of a null hypothesis of equivalence between the holdout and cross-validated experimental results rejected with $p < 2.2 \times 10^{-16}$.

The results of this experiment demonstrate the importance of using large, diverse datasets for machine learning experiments, and of performing multiple experimental replications. Additionally, these results demonstrate that simpler models – such as RNN, radial SVM, and logistic regression – achieve equivalent or better performance relative to LSTM on our specific task of MOOC dropout prediction. Our results, which are contrary to the findings of the original study, also suggest that further replication is necessary to identify the most effective algorithms for MOOC dropout prediction, as we perform no hyperparameter tuning and only replicate the original models and features.

6 Additional Advantages of Reproducible Machine Learning Frameworks

Much prior work on reproducibility has focused on verification – ensuring that published results are true and can be reproduced. However, end-to-end reproducible machine learning frameworks, such as MORF, provide benefits beyond mere verification, including:

“Gold standard” benchmarking: open replication platforms allow for the comparison of results which were previously not comparable, having been conducted on different data. The use of such benchmarking datasets has contributed to the rapid advance of fields such as computer vision (e.g. MNIST, IMAGENET), natural language processing (Penn Tree Bank, Brown corpus), and computational neuroscience (openfMRI). These datasets have been particularly impactful in fields where it is difficult or expensive to collect, label, or share data (as is the case with MOOC data, due to legal restrictions on sharing and access). These help to evaluate the “state of the art” by providing a common performance reference which is currently missing in many fields.

Shared baseline implementations: We noted above that variability in so-called “baseline” or reference implementations of prior work has contributed to concerns about reproducibility in the field [17]. By providing fully-executable versions of existing experiments, MORF ameliorates these issues, allowing for all future work to properly compare to the exact previous implementation of a baseline method.

Forkability: containerization produces a ready-made executable which fully encompasses the code and execution environment of an experiment. These can be readily shared and "forked" across the scientific community, much in the same way code is "forked" from a git repository. This allows machine learning scientists to build off of others’ work by modifying part or all of an end-to-end pipeline (for example, by experimenting with different statistical algorithms but using the same feature set as a previous experiment) within the same software ecosystem.

Generalizability analysis: Each successive replication of an experiment provides information about its generalizability. Evaluating the generalizability of experiments has been a challenge in MOOC research to date, where studies conducted on small, restricted, and often homogenous datasets tend to dominate the literature. These challenges are shared by many other fields where data is scarce and sharing is restricted, such as nuclear physics or medical research [20]. When containerized end-to-end implementations are available, replicating these analyses on new data – even data which are not publicly available but share the schema of the original data – becomes as straightforward as running the containerized experiment against new data.

Sensitivity Analysis: This technique, used widely in Bayesian analysis, evaluates how changes to the underlying assumptions or hyperparameters affect model fit and performance. Such an evaluation can provide useful information about a model’s robustness and potential to generalize to new data. Without being able to fully reproduce a model on the original data, sensitivity analyses are not possible. In MORF, such analyses can be conducted by simply forking and modifying the containerized version of the original experiment, then re-executing it against the same data. These analyses can also include so-called *ablation analyses*, wherein individual components are removed from a model to observe their contribution to the results, as well as *slicing analyses*, where fine-grained analysis of performance across different subgroups (e.g. demographic groups) is explored [26].

Full Pipeline Evaluation: Each stage of an end-to-end machine learning experiment (feature extraction, algorithm selection, model training, model evaluation) can be done in many different ways. Each stage also affects the others (for example, some algorithms might perform best with large feature spaces; others might perform poorly with many correlated features). However, current research usually evaluates only one or two components of this pipeline (e.g. training several algorithms and tuning their hyperparameters on a fixed feature set). Not only are the remaining stages often described in poor detail or not at all [15]; such work also leaves future researchers unable to evaluate the synergy between different aspects of the end-to-end pipeline in a published experiment (for example, exploring whether an algorithm’s performance improves with a different feature set). MORF fully encapsulates this end-to-end pipeline for a given experiment and it makes it available for modification to any other researcher.

Meta-Analysis: While meta-analyses are common in fields with robust empirical research bases, such analyses have been less common in the field of machine learning, which has an emphasis on novelty. The open availability of executable machine learning experiments affords detailed meta-analyses by providing complete results of all modeling stages for meta-analysis.

7 Conclusion

While MORF currently supports MOOC data, the broader workflow it uses to enable end-to-end machine learning reproducibility is domain-agnostic, and the problems it seeks to address are faced by many fields which use machine learning. We believe that the paradigm of end-to-end reproducibility can be adopted by any domain which uses large, complex, multi-format data. A platform such as MORF is fully capable of being adapted to such domains, particularly those which use restricted-access data. We hope that future work utilizes and extends our open-source implementation, and that MORF contributes to the construction of robust machine learning research.

References

- [1] Carl Boettiger. An introduction to docker for reproducible research. *Oper. Syst. Rev.*, 49(1):71–79, January 2015.
- [2] Adam Brinckman, Kyle Chard, Niall Gaffney, Mihael Hategan, Matthew B Jones, Kacper Kowalik, Sivakumar Kulasekaran, Bertram Ludäscher, Bryce D Mecum, Jarek Nabrzyski, Victoria Stodden, Ian J Taylor, Matthew J Turk, and Kandace Turner. Computing environments for reproducibility: Capturing the “whole tale”. *Future Gener. Comput. Syst.*, February 2018.
- [3] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. OpenAI gym. *arXiv preprint*, June 2016.
- [4] Jonathan B Buckheit and David L Donoho. WaveLab and reproducible research. In Anestis Antoniadis and Georges Oppenheim, editors, *Wavelets and Statistics*, pages 55–81. Springer New York, New York, NY, 1995.
- [5] Gavin C Cawley and Nicola L C Talbot. On over-fitting in model selection and subsequent selection bias in performance evaluation. *J. Mach. Learn. Res.*, 11(Jul):2079–2107, 2010.
- [6] Jürgen Cito, Vincenzo Ferme, and Harald C Gall. Using docker containers to improve reproducibility in software and web engineering research. In *Web Engineering*, Lecture Notes in Computer Science, pages 609–612. Springer, Cham, June 2016.
- [7] Christian Collberg, Todd Proebsting, Gina Moraila, Akash Shankaran, Zuoming Shi, and Alex M Warren. Measuring reproducibility in computer systems research. Technical report, Univ. Arizona Dept. of Comp. Sci., 2014.
- [8] Coursera. *Coursera Data Export Procedures*. Coursera, June 2013.
- [9] T G Dietterich. Approximate statistical tests for comparing supervised classification learning algorithms. *Neural Comput.*, 10(7):1895–1923, September 1998.
- [10] David Donoho. 50 years of data science. In *Princeton NJ, Tukey Centennial Workshop*, pages 1–41, 2015.
- [11] Bradley Efron and Trevor Hastie. *Computer Age Statistical Inference*. Cambridge University Press, 2016.
- [12] M Fei and D Y Yeung. Temporal models for predicting student dropout in massive open online courses. In *Intl. Conf. on Data Mining Workshop (ICDMW)*, pages 256–263, 2015.
- [13] Josh Gardner and Christopher Brooks. Student success prediction in MOOCs. *User Modeling and User-Adapted Interaction*, 2018.
- [14] Josh Gardner, Christopher Brooks, Juan Miguel Andres, and Ryan Baker. Replicating MOOC predictive models at scale. In *Proceedings of the Fifth Annual Meeting of the ACM Conference on Learning@Scale*, 2018.
- [15] Odd Erik Gundersen and Sigbjørn Kjensmo. State of the art: Reproducibility in artificial intelligence. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence and the Twenty-Eighth Innovative Applications of Artificial Intelligence Conference*, 2017.
- [16] Mitx And HarvardX. HarvardX-MITx Person-Course academic year 2013 De-Identified dataset, version 2.0, May 2014. Title of the publication associated with this dataset: HarvardX-MITx Person-Course Academic Year 2013 De-Identified dataset, version 2.0.
- [17] Peter Henderson, Riashat Islam, Philip Bachman, Joelle Pineau, Doina Precup, and David Meger. Deep reinforcement learning that matters. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18)*, pages 3207–3214, 2018.
- [18] Douglas M Jacobsen and Richard Shane Canon. Contain this, unleashing docker for hpc. *Proceedings of the Cray User Group*, 2015.
- [19] Mohammad Khajah, Robert V Lindsey, and Michael C Mozer. How deep is knowledge tracing? In *Proceedings of the 9th International Conference on Educational Data Mining*, pages 94–101, 2016.
- [20] Justin Kitzes, Daniel Turek, and Fatma Deniz. *The Practice of Reproducible Research: Case Studies and Lessons from the Data-Intensive Sciences*. Univ of California Press, October 2017.

- [21] Kenneth R Koedinger, Ryan Sjd Baker, Kyle Cunningham, Alida Skogsholm, Brett Leber, and John Stamper. A data repository for the EDM community: The PSLC DataShop. *Handbook of educational data mining*, 43, 2010.
- [22] Dirk Merkel. Docker: Lightweight linux containers for consistent development and deployment. *Linux J.*, 2014(239), March 2014.
- [23] Babatunde K Olorisade, Pearl Brereton, and Peter Andras. Reproducibility in machine Learning-Based studies: An example of text mining. In *Reproducibility in ML Workshop at the 34th International Conference on Machine Learning*, 2017.
- [24] Roger D Peng. Reproducible research in computational science. *Science*, 334(6060):1226–1227, December 2011.
- [25] Chris Piech, Jonathan Bassen, Jonathan Huang, Surya Ganguli, Mehran Sahami, Leonidas J Guibas, and Jascha Sohl-Dickstein. Deep knowledge tracing. In C Cortes, N D Lawrence, D D Lee, M Sugiyama, and R Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 505–513, 2015.
- [26] D Sculley, Jasper Snoek, Alex Wiltschko, and Ali Rahimi. Winner’s curse? on pace, progress, and empirical rigor. In *ICLR 2018 Workshops*, February 2018.
- [27] Sören Sonnenburg, Mikio L Braun, Cheng Soon Ong, Samy Bengio, Leon Bottou, Geoffrey Holmes, Yann LeCun, Klaus-Robert Müller, Fernando Pereira, Carl Edward Rasmussen, Gunnar Rätsch, Bernhard Schölkopf, Alexander Smola, Pascal Vincent, Jason Weston, and Robert Williamson. The need for open source software in machine learning. *J. Mach. Learn. Res.*, 8(Oct):2443–2466, 2007.
- [28] Victoria Stodden and Sheila Miguez. Best practices for computational science: Software infrastructure and environments for reproducible and extensible research. *Journal of Open Research Software*, 2(1):1–6, 2013.
- [29] Jan N van Rijn, Bernd Bischl, Luis Torgo, Bo Gao, Venkatesh Umaashankar, Simon Fischer, Patrick Winter, Bernd Wiswedel, Michael R Berthold, and Joaquin Vanschoren. OpenML: A collaborative science platform. In *Machine Learning and Knowledge Discovery in Databases*, pages 645–649. Springer Berlin Heidelberg, 2013.
- [30] Sudhir Varma and Richard Simon. Bias in error estimation when using cross-validation for model selection. *BMC Bioinformatics*, 7:91, February 2006.
- [31] Jacob Whitehill, Kiran Mohan, Daniel Seaton, Yigal Rosen, and Dustin Tingley. Delving deeper into MOOC student dropout prediction. *arXiv preprint*, February 2017.