The Mathematics Enthusiast

Volume 19
Number 3 *Number 3 (In Progress)*

Article 12

12-2022

Experimental Mathematics – A Course-Based Research Experience in Machine Learning

Mathew Gluck

Recommended Citation

Gluck, Mathew (2022) "Experimental Mathematics – A Course-Based Research Experience in Machine Learning," *The Mathematics Enthusiast*: Vol. 19 : No. 3 , Article 12.
DOI: https://doi.org/10.54870/1551-3440.1580
Available at: https://scholarworks.umt.edu/tme/vol19/iss3/12

Experimental Mathematics–a CURE in Machine Learning

Mathew Gluck

Towson University, Towson, MD

ABSTRACT: Efforts to expand the science, technology, engineering and mathematics (STEM) workforce have been topics of United States policy action for more than 50 years (Hira 2010). Unfortunately, among U.S. undergraduate curricula, STEM has one of the highest attrition rates (Tinto 1993) with less than half of students in the U.S. that enroll in an undergraduate STEM program ultimately receiving a degree in a STEM field (Hayes 2009). Naturally, the high rate of attrition is a topic of persisting concern. Many programs have been designed and implemented to model best practices in retaining students in STEM disciplines. One retention strategy is to engage STEM undergraduates in research experiences, and a number of programs have been implemented to provide such experiences. The Towson University Research Enhancement Program (TU REP) is one such program. This cohort-based program supports faculty in the development of course-based undergraduate research experiences (CUREs). In this note we describe a CURE in machine learning offered by the Towson University Department of Mathematics whose development was supported by TU REP. We categorize this course along the spectrum of traditional, inquiry, CURE and internship in each of the five dimensions characteristic of a CURE.

Keywords: Course-based undergraduate research experience, machine learning

The Mathematics Enthusiast, **ISSN 1551-3440, vol. 19, no. 3**, pp. 822–832

1 Introduction

Recently there has been a nationwide effort to increase retention in undergraduate science, technology, engineering and mathematics (STEM) fields. This is due partly in response the the high attrition rates of undergraduates in STEM majors. As of 2009, fewer than half of undergraduates declaring a STEM major as freshman ultimately receive a degree in a STEM field [Hay]. A variety of intervention strategies have been deployed in response. One such strategy is to engage undergraduate students in research experiences. Research engagement-based approaches show promise in improving retention. For example, it is shown that undergraduate research is strongly correlated to increased rates of graduation [Hat]. Moreover, participation in research is shown to lead to increased persistence toward the undergraduate degree [Nag] as well as positively influencing the selection of a career path in STEM [Zyd].

Traditional undergraduate research experiences, for example those following the internship model or those following the Research Experience for Undergraduates (REU) model tend to be limited in the number of participants that can be accommodated. Reasons for these limitations include constraints on faculty time, constraints on laboratory space and budget constraints. A number of programs across a variety of institutions have been implemented in efforts to increase participation in undergraduate research see, for example [Way, Gat]. One such program, is Towson University's *Towson University Research Enhancement Program: Expanding Inclusive Excellence in Science at Towson University* (TU REP). This program is a Howard Hughes Medical Institute (HHMI)-sponsored program that started in 2018. The goal of TU REP is to retain all TU STEM undergraduates in their declared majors. The general approach to achieving this goal is to support faculty in developing undergraduate courses that incorporate authentic research experiences for students. Specifically, TU REP provides faculty members who agree to develop a CURE with their choice of monetary compensation or teaching reduction-based support. In addition, participating faculty are provided training in STEM pedagogy and inclusive teaching practices.

Based on the premise that involvement in undergraduate research positively impacts student retention, the courses developed in the TU REP program are course-based undergraduate research experiences (CUREs). CUREs are research opportunities that live somewhere between inquiry-based course experiences and an internships. The working definition of a CURE used in the present work is as in [Auc]. Briefly, in a CURE students are afforded the opportunity to engage in research as part of their undergraduate coursework. More specifically, according to [Auc], in order for a course to qualify as a CURE, its participants must engage in each of the following five characteristic dimensions

1. use of scientific practices,
2. discovery,
3. broadly relevant or important work,
4. collaboration, and
5. iteration.

In the 2019–2020 academic year, TU REP supported the development of multiple CUREs in mathematics. One such course, the Spring 2021 installment of Experimental Mathematics at TU, is the subject of this note.

This note will discuss the Spring 2021 installment of Experimental Mathematics, a CURE in mathematics whose topic was machine learning (ML). In Section 2, a description of the course structure, its outcomes and its assessments will be provided. In Section 3, for each of the five characteristic dimensions of a CURE, the course will be categorized along the spectrum of traditional learning experience, inquiry-based learning experience, CURE and internship. Section 4 overviews an example of a student project from the course and Section 5 provides some concluding remarks.

2 Description of the Course

The TU catalog description for Experimental Mathematics allows each installment of the course to have features that are unique to that particular installment. Thus, there is no version of Experimental Mathematics that is representative of all possible versions of the course. The subject matter in a particular installment of Experimental Mathematics depends on the preferences of the instructor of record and may vary from semester to semester. The official catalog description only requires that students in the course be introduced to programming, computational methods, algorithms, and software environments used by research mathematicians. It also requires that students be afforded opportunities to apply these tools to explore patterns, make conjectures, and explore the role of computation in formal mathematical proofs. This description is largely parallel to many accepted notions of what qualifies as experimental mathematics, see for example [Bai]. Example topics provided in the course description include combinatorics, number theory, numerical analysis, modeling and visualization, fractals, computer-assisted proofs and graph theory. In what follows, the course structure for the Spring 2021 installment will be described.

There were two issues that created the need for last minute, ad-hoc adjustments in the course format. First, the total enrollment for the course was 5 students, well short of the target enrollment of 18 students. The targeted enrollment is consistent with the TU REP goal of increasing student access to research opportunities. Factors that contributed to the low enrollment included availability of alternative undergraduate research opportunities and the fact that Experimental Mathematics fulfilled no graduation requirements for mathematics majors (other than allowing students to accumulate additional credit hours). In fact, in the Spring 2021 semester, an alternative undergraduate mathematics research opportunity was available at TU that offered students monetary compensation instead of course credit. The availability of this opportunity was the result of last-minute logistical changes to a TU-hosted mathematics REU in response to the COVID-19 pandemic. It is not anticipated that future installments of Experimental Mathematics will need to compete with such attractive research opportunities for enrolling participants. The second issue that impacted the course format was that the social distancing measures taken in response to the COVID-19 pandemic forced the course to be delivered in a fully online format.

2.1 The Three Phases of the Course

The focus of the course was to give students an opportunity to work on a research project on the machine learning topic of their choosing. In order to accommodate this, the course was divided into the following three phases, each of which will be described further in the subsections that follow.

Phase 1. General knowledge building and project selection (4 weeks)

Phase 2. Development of theory and implementation of algorithm (8 weeks)

Phase 3. Tie up loose ends and report findings (3 weeks)

Phase 1. General knowledge building and project selection

Upon entry to the course, the majority of the students lacked the base-level knowledge necessary to select a project that could be both substantial enough to occupy them for the majority of the semester yet modest enough that meaningful progress could be made in one semester. The primary goal of Phase 1 was to rectify this issue so that students would be in a position to choose a project that would be the focus of Phases 2 and 3 of the course. More concretely, by the end of Phase 1, students should have developed both a well-rounded base-level ML content knowledge and their tastes regarding which ML topics interested them. Moreover, since it was anticipated that most students would want to implement their ideas in code, and since much of the ML community chooses to code in Python, students should have increased their familiarity with Python-based programming techniques by the end of Phase 1.

A secondary goal of Phase 1 was to provide students with opportunities to gauge the interests of their classmates so that they could form groups consisting of members having similar interests. This was accomplished by assigning a variety of in-class activities and presentations. Had the enrollment been more substantial, students would have been required to form groups. However, due to the low enrollment, the ad-hoc decision was made to allow students to choose either to team up with classmates or to pursue their ideas as individuals. All students chose to pursue individual projects. The online format of the course may have created a barrier to group formation that was too substantial for the community building exercises and activities to overcome. The students' decisions to pursue individual projects diminished a large component of the initially intended peer-to-peer collaborative nature of the course.

During Phase 1, the class meetings occurred as a class; the students and the instructor were all in the same online room. In-class content delivery was primarily delivered in lecture format and the topics covered were introductory and general in nature. Examples of such topics included the distinction between supervised and unsupervised learning, multivariate optimization and basic theory of standard fully-connected neural networks. Content during this phase of the course was rarely individualized. Toward the end of Phase 1, class meetings began a transition into more individualized meetings. Had enrollment in the course been sufficiently large or had student groups been formed, each of these meetings would have been between the instructor and all students in a particular group. Instead, these meetings were one-on-one meetings between the instructor and each student. During scheduled class meeting times, students not in one-one-one meetings typically stayed in the virtual classroom to work on their projects.

Phase 2. Development of theory and implementation of algorithm

As was anticipated, by the end of Phase 1, every student wanted to implement their ideas in code. The primary goal of Phase 2 was for students to gain a deeper understanding of the topic they had chosen to pursue and then to develop an algorithm that accomplishes the task they had chosen. In every case, successful development of the algorithm relied on a thorough understanding of the mathematics on which the algorithm is based.

At the start of Phase 2, the structure of class meetings had fully transitioned into individual-based meetings (between the instructor and one student at a time). This format was chosen in order to allow students to discuss aspects of their projects that were specific to them but not relevant to other students. On occasion, similar theoretical or programming-based concepts were encountered across multiple projects. In these cases, the class would meet as a whole for either a discussion or a lecture.

Phase 3. Report findings

The primary goal of Phase 3 was for students to communicate their findings. Communication was required to be delivered both via a written technical report and via oral presentation. The written report was required to be written in the same tone as that of a publication-quality, peer-reviewed academic article. The oral presentations were required to be accessible to those having elementary-level understanding of ML, but little to no understanding of the particular topic being presented (for example other students in the course). During Phase 3, the class meeting format stayed as it was in Phase 2 with the majority of class time being spent in one-on-one discussions between the instructor and particular students. In some class meetings, a substantial portion of the class time was spent discussing topics that were relevant to all students. Compared to the analogous meetings that occurred during Phase 2 of the course, the topics in the class-wide meetings that occurred in Phase 3 were less technical in nature. For example, one such meeting focused on illustrating to the students the writing style and the components of some typical peer-reviewed published articles.

In addition to the mandatory in-class presentations, there was an optional presentation to be delivered at a locally-hosted undergraduate mathematics conference. There was no direct benefit toward the students' grades (e.g. extra credit) for delivering the optional presentation. One of the five enrolled students chose to give the optional presentation.

2.2 Assessment and Outcomes

Compared to a traditional mathematics course, assessment in the Spring 2021 installment of Experimental Mathematics was not typical. For example, there were no quizzes or exams. The three major components assessed were a research journal, multiple oral presentations and a multi-draft written report. In addition to these major components, there were a number of minor assessments given throughout the course. The minor assessments given during Phase 1 were given to all students and were intended both to assist students in their explorations of ML topics and to build a community-like environment amongst classmates. Examples of minor assessments given in Phase 1 include (1) Python-based assignments that help familiarize students with array manipulation and (2) exploratory assignments that require students to find concrete applications of machine learning and to share these findings with their classmates. The minor assessments given in Phases 2 and 3 were individualized assignments given as needed to help students progress with their projects. Because of the unpredictability associated to a typical mathematics research project, the amount of progress students made toward completing their projects was not considered in the assigning of grades.

The research journal was a venue where students were to record their thoughts on their projects. It differed from what one might typically describe as a journal in the sense that it was a collaborative journal where communications between the instructor and the student occurred. The journals were kept on Overleaf, an online collaborative setting for documents produced in LaTeX. This assessment spanned the entire semester. On most weeks students were required to submit three journal entries. Much of the interaction and feedback regarding the projects was provided by the instructor directly in the students' research journals. This format set the stage for productive in-class meetings where time was limited. Based on the journal entries and the content of the one-on-one meetings, small informal individualized homework assignments were given to students, usually written by the instructor directly in the research journal. In most cases there were no grades associated to these types of homework assignments, but their completion was necessary in order to make progress on the project.

Presentations were given in every phase of the course. Due to the online format of the course, all presentations were given virtually via video conferencing software. Students were encouraged to actively participate in other students' presentations. Phase 1 had two presentations. The first presentation was preliminary in nature and its main purpose was to allow students to advertise their ML-related interests to their classmates. The intent of these advertisements was to assist in the formation of groups with similar interests. The second presentation occurred at the end of Phase 1 and its purpose was to force students to solidify their project ideas before moving into the next phase of the course. Phases 2 and 3 had one presentation each.

The written report was a two-draft assignment where students discussed their projects. The report was required to be written in the same tone as a typical peer-reviewed, published article in an academic journal or conference proceedings. Generically, the written report was to include an abstract, an introduction containing historical and contextualizing remarks, a section on mathematical background, a discussion of the results of code when implemented on data sets, and discussion or future work section. Students were given the liberty to deviate from this format if they felt it made sense to do so.

As is consistent with TU REP goals, the primary intended outcome of the course was to promote retention in the students' STEM majors by providing widespread access to research experiences. Additionally, because the majority of the enrolled students had expressed interest in finding work in government or in the private sector upon graduation, a secondary outcome was to develop students abilities to speak fluently and casually about a long-term research project on which they had worked. This was intended to provide the students with talking points that could be used, for example, in a job interview situation.

3 This course is a CURE

From Auchincloss et.al [Auc], the five criteria for a CURE are use of scientific practices, discovery, broadly relevant or important work, collaboration and iteration. In this section we categorize the Spring 2021 installment of Experimental Mathematics along the spectrum of traditional, inquiry, CURE and internship in each of these dimensions, see Table 2 of [Auc]. For convenience this table has been reproduced in Table 1 below.

Use of scientific practices

In [Auc], the use of scientific practices is described as including the following features: (1) asking questions, (2) building and designing studies, (3) selecting methods, (4) using the tools of science, (5) navigating the messiness of real-world data, (6) developing and critiquing arguments and (7) communicating findings. Students enrolled in the Spring 2021 installment of Experimental Mathematics engaged in multiple scientific practices including asking questions, selecting methods, navigating the messiness of real-world data and communicating findings. Thus, along this dimension, this course falls firmly into the CURE/ internship category.

During Phase 1 of the course, students developed the questions whose answers they would pursue in the subsequent phases of the course. Examples of some such questions include "which sequence prediction methods should be used if one is willing to sacrifice accuracy in prediction in return for faster training times?" or "What is the role of the reward function in the use of reinforcement learning to train a computer-based agent to play Tetris?". Many smaller scale and initially unforeseeable questions were confronted during the development of the theory and the algorithm in Phase 2 of the course. Generally, the large-scale questions whose answers the students chose to pursue were entirely their own, but the instructor played a role in helping students narrow the scope of their questions so that meaningful progress toward an answer could be made in one semester. For example, one student's initial goal was to

Dimension		Traditional	Inquiry	CURE	Internship
Use of science practices	Students engage in ...	Few scientific practices	Multiple scientific practices	Multiple scientific practices	Multiple scientific practices
	Study design and methods are...	Instructor driven	Student driven	Student or instructor driven	Student or instructor driven
Discovery	Purpose of the investigation is...	Instructor defined	Student defined	Student or instructor defined	Student or instructor defined
	Outcome is...	Known to students and instructors	Varied	Unknown	Unknown
	Findings are ...	Previously established	May be novel	Novel	Novel
Broader relevance or importance	Relevance of students' work...	Is limited to the course	Is limited to the course	Extends beyond the course	Extends beyond the course
	Students' work presents opportunities for action...	Rarely	Rarely	Often	Often
Collaboration	Collaboration Occurs...	Among students in a course	Among students in a course	Among students, teaching assistants, instructor in a course	Between student and mentor in a research group
	Instructor's role is...	Instruction	Facilitation	Guidance and mentorship	Guidance and mentorship
Iteration	Risk of generating "messy" data are...	Minimized	Significant	Inherent	Inherent
	Iteration is built into the process...	Not typically	Occasionally	Often	Often

Table 1: (Reproduction of Table 2 of [Auc]) Dimensions of different laboratory learning contexts

use sequence prediction methods to quickly predict stock prices. Based on discussions with the instructor, the student ultimately decided to pursue the development of an algorithm that can quickly predict the next entry in a sequence of digits. This pursuit, while far less ambitious than the student's original goal, can be seen as a building block toward achieving the student's original goal.

Navigating the messiness of real-world data is a prevalent theme in ML. In fact, the primary goal of many ML algorithms is to meaningfully interpret vast amounts of messy data. Naturally, this theme was prominent in the Spring 2021 Experimental Mathematics course. As a particular example, one student had to decide how to effectively extract the most important features from screen shots of a video game in order to implement their algorithm, see Section 4 for a description. In other projects, students first developed and tested their algorithms on simulated data; data where the behavior of a correctly working algorithm is known. This allowed students to gauge whether their algorithms were working as expected before applying their algorithms to real-world data.

Communication of findings was emphasized throughout the course. These communications were delivered in both written and oral formats. The course was designed so that the formality in the written communications varied. For example, the research journal was an informal and colloquial venue, while the final written report was expected to be in publication-quality format. Oral communication of the findings from students to the class (as presentations) were emphasized at every phase of the course.

Discovery

According to [Auc], the main components of the discovery dimension are first, the outcome of the investigation should be unknown to both the students and the instructor and second, students should be addressing novel scientific questions aimed at generating and testing new hypotheses.

Students in this iteration of Experimental Mathematics engaged in discovery in a manner consistent with a CURE or internship. As described in Subsection 2.1, students built some familiarity in the topics and ultimately chose the topics of their projects based entirely on their own interests. Because the topics were student-chosen, the outcomes in most of the projects were unknown both to the student and to the instructor. As is consistent with science-based discovery, there was risk for unanticipated outcomes. See Section 4 for a description of one such outcome.

The requirement of student work to be novel was met provided one accepts an appropriate interpretation of the word novel. In most of the student projects, the novelty of the questions students addressed was localized to those with a beginner-level understanding of ML. Because of the technical nature of mathematics and machine learning, although not impossible, it is unlikely that students at such an early stage of development would be able to address questions at the forefront of the field.

Broadly relevant or important work

According to [Auc], CUREs should involve students in work that fits into a broader scientific endeavor that has meaning beyond participation in the course. There are three sub-dimensions of the broader relevance or importance dimension; *novelty*, *relevance of students' work* and *opportunities for action*. Experimental Mathematics qualifies as a CURE/internship in two of the three sub-dimensions and it qualifies as an inquiry in one of the sub-dimensions. Specifically, the findings of the students' work may be novel, but the level of technical skill required to consistently produce novel results in a one-semester timeframe is generally beyond what the target demographic possesses. Thus, in the *novelty* sub-dimension, this course is classified into the inquiry category. In the *relevance of students' work* and the *opportunities for action* sub-dimensions, the course falls firmly into the CURE/internship end of the spectrum. In particular, students' initial project goals were generally too ambitious to be accomplished in one semester. However, these grand ambitions define the scientific endeavor that has meaning beyond the course. The instructor plays an essential role in ensuring that the students have a clear vision of how the outcome of the semester's work will contribute to the overall ambition of the student. In these cases, the students projects can be considered as meaningful progress toward the larger endeavor. From this larger endeavor, the relevance of the work extends beyond the course and presents opportunities for further action.

Collaboration

According to [Auc], collaboration should occur among the students and between students and instructor. Additionally, the instructor's role in a CURE is more accurately described as providing guidance than instruction. The Spring 2021 installment of Experimental Mathematics was initially designed to conform to the CURE model in this dimension. Generically, the vision of the TU Mathematics Department is that the course enrollment should be roughly 18 students. This enrollment would allow for student groups of size 3—5, and thus provide ample opportunities for peer-to-peer collaboration. However, due to the low enrollment and the instructor's decision to prioritize allowing students to pursue their own ideas over student-to-student collaboration, the Spring 2021 installment of Experimental Mathematics is more accurately described as an internship than it is a CURE. The course supported extensive collaboration between the instructor and the students, but there were too few opportunities for student-to-student collaboration for the course to qualify as a CURE along the collaboration dimension.

Iteration

Because the course had a heavy algorithm development component and because iterative approaches are standard in algorithm and software development [Bas], the course naturally involved an iterative component. A typical iteration in the algorithm development cycle included a planning phase, an implementation phase, a testing phase, and an evaluation phase. In most cases, the first iteration of students' projects were bare-bones versions of their algorithms. With each iteration, refinements and improvements to the existing algorithms were incorporated. Although this iterative development approach was the approach recommended by the instructor to the students, students were not required to follow this approach.

4 An Example of Student Work

The student projects covered a variety of topics including sequence prediction and source identification for fake news in social networks. This section describes a student project whose goal was to use reinforcement learning to teach a computer-based agent to play Tetris. The primary inspiration for the project was [Mni13, Mni15, Ste], each of which also used reinforcement learning to train a computer-based agent to play video games. The premise in the student's project (and the premises in the cited works) was that the agent's learning process should mimic that of a human. In particular, the agent does not have access to the internal state of the game. Rather, the agent must interact with the game to accumulate experiences and, based on the accumulated experiences, develop a game-playing policy. The policy is a (probabilistic) rule that indicates which action the agent should take based on the agent's observation of the state of the game. In short, the policy-dictated action the agent chooses to take after observing the game's state is chosen so that the value of a long-term reward function will be maximized. The value of the long-term reward function is constructed by accumulating many short-term rewards. In some cases, a particular action may result in a large short-term reward while eliminating the possibility that the agent receive significant rewards later in the game. Thus, an effective policy should have some long-term strategic capabilities that would dictate that the agent avoid actions that yield rewards that are immediately beneficial but restrict possibilities for attaining future rewards. The student's primary interest was to investigate the role of the reward function in the agent's Tetris-playing policy.

Since there are only finitely many Tetris board configurations, if a long-term reward function is specified, one could in principle represent the reward function exactly (for example in a lookup table). With such a representation, after an agent observes a game state, it simply looks up the corresponding optimal action and executes that action. In practice, the number of states is too large for such a representation to be implemented. Thus, the goal of the training step is to produce an accurate approximation for the reward function. In the training step, the agent plays the game many times and catalogues its experiences along the way. In this context, an experience is a tuple of the form (state, action, reward, next state). At a given state of the game, the agent estimates, based on its experiences, the action that would result in maximizing the long-term reward function. This action is followed most of the time, with an occasional random action taken in order to prevent the policy from converging to local maximizer.

The student used MaTris [Mat], a Tetris emulator written in Pygame in combination with Tensorflow [Tfl] to implement and train a convolutional neural network. In particular, the convolutional layers were used to extract the game states from screenshots of the game board. Ultimately, the agent trained by the student was unable to settle upon a satisfactory Tetris-playing policy. The agent's policy dictated that the active block should be moved left in all cases where the policy was followed, see Figure 1 for a typical game board that occurs with this policy. This outcome exemplifies the possibility that unforeseen outcomes may be realized in practice. The student hypothesized that the poor policy was a result of too few rewards being encountered in training. More specifically, early on in the training process, the agent has very few experiences and so its actions are mostly random. Due to the low probability that a randomly acting agent clear any lines in a particular game of Tetris, it is unlikely that the agent saw enough examples of good game play to develop a successful policy. The student also expressed a disadvantage of their approach of capturing the game's states by using screenshots. Namely, it forced the student to only train the agent at times when they did not need to use their computer for other unrelated tasks. This cut down on the amount of time the student could train the agent.

The student continued working on their project after the conclusion of the Spring 2021 semester and, in the Fall of 2021, provided an update on their progress. In order to verify that the training algorithms were working as intended, the student coded a custom Tetris emulator that included only blocks of shape 4×1 and blocks of shape 2×2. Additionally, as a temporary measure, the student included code in their custom emulator that indicates the state of the game (but does not allow the agent access to the internal state of the game). This eliminated the need to capture screenshots during training and thus allowed for ample training time for the agent. The student expressed intentions of reincorporating the use of screenshots for capturing the state of the game later, once he is confident that his algorithm is working to his satisfaction. The student trained an agent on the simplified version of Tetris and verified that his algorithms were working as intended. The student then added the remaining block shapes to his Tetris emulator so that he had a fully functioning Tetris game. After training the agent on the fully functioning Tetris emulator, the student noticed remarkable improvements in the agent's performance compared to the agent he trained in the Spring of 2021. The student was generous enough to provide example game boards of his improved agent playing Tetris. These are included in Figure 2.

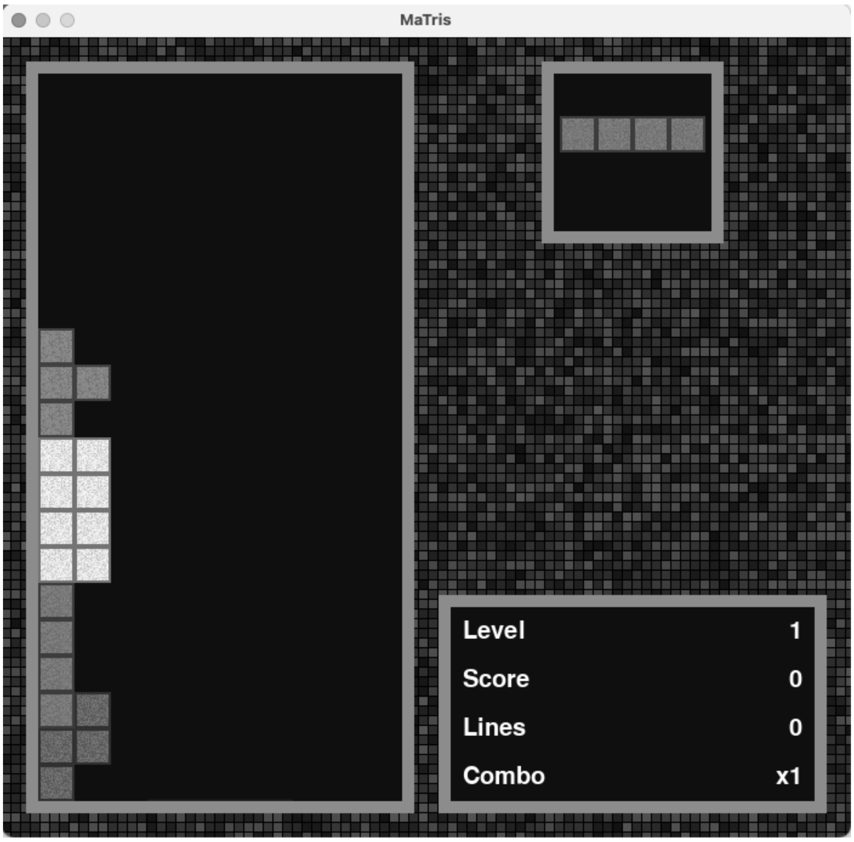

Figure 1: Screenshot of MaTris Tetris emulator. Images of this form were used to capture the state of the game during gameplay. This a a typical game board that results from the trained agent's policy to always move the active block left.

5 Conclusions

Based on the premise that undergraduate participation in research projects positively impacts retention of students in STEM majors, Towson University's TU REP program provided support for the the development of multiple CUREs in mathematics. The Spring 2021 installment of Experimental Mathematics was one such course. It provided students with a course-based undergraduate research experience in machine learning. Students were afforded the opportunity to do research in the machine learning topic of their choosing. In all but one of the five dimensions along which a course can be classified as a CURE, this course was indeed a CURE.

Acknowledgment. *This work was supported by an Inclusive Excellence grant from the Howard Hughes Medical Institute for the Towson University Research Enhancement Program.*

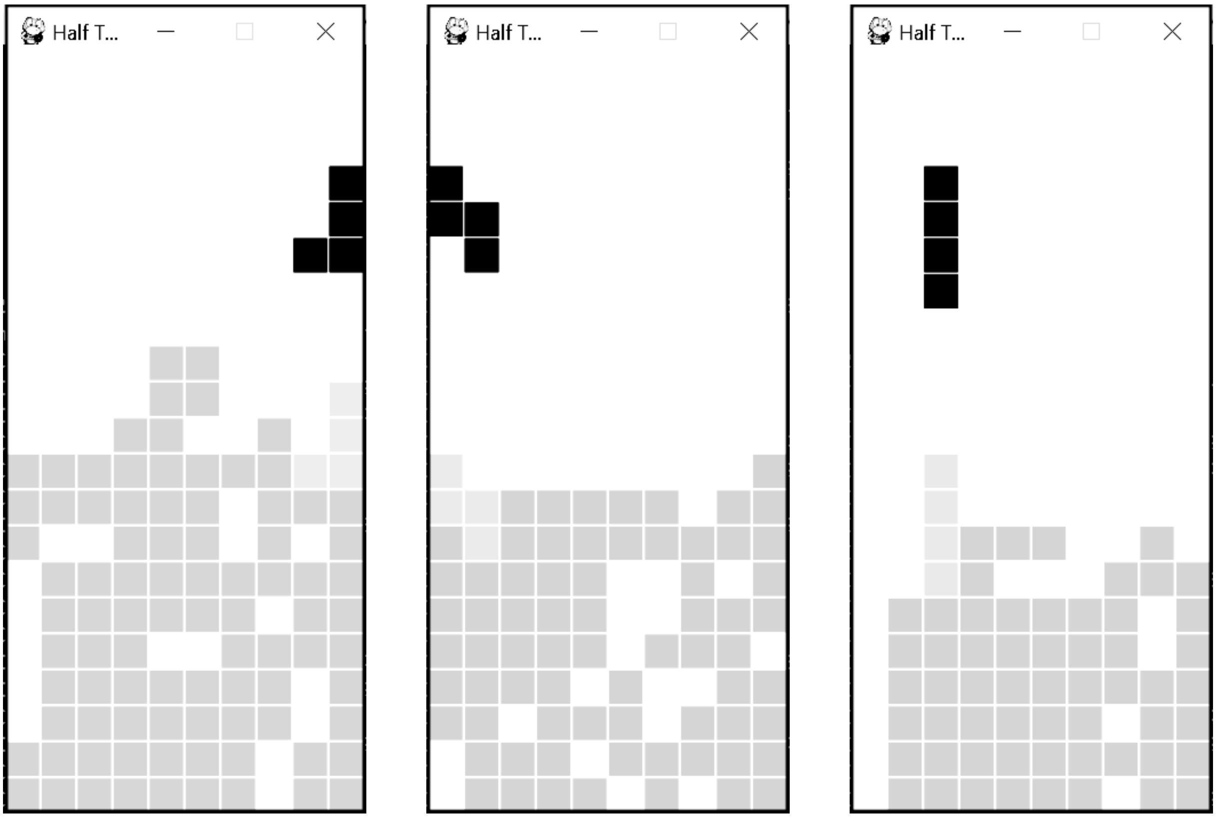

Figure 2: Mid-game states from three different games played by the improved Tetris-playing agent.

References

[Auc] Lisa Corwin Auchincloss, Sandra L. Laursen, Janet L. Branchaw, Kevin Eagan, Mark Graham, David I. Hanauer, Gwendolyn Lawrie, Colleen M. McLinn, Nancy Pelaez, Susan Rowland, Marcy Towns, Nancy M. Trautmann, Pratibha Varma-Nelson, Timothy J. Weston & Erin L. Dolan (2014). Assessment of Course-Based Undergraduate Research Experiences: A Meeting Report *CBE–Life Sciences Education*, 13, 29–40.

[Bai] David Bailey and Jonathan Borwein (2008) Mathematics by Experiment: Plausible Reasoning in the 21st Century (2nd ed.). *A K Peters/CRC Press.*

[Bas] Basil, V. R. and Turner, A. J. (1975), Iterative enhancement: A practical technique for software development. *IEEE Transactions on Software Engineering*, vol. SE-1, no. 4, 390-396.

[Gat] Gates, A.Q., Teller, P.J., Bernat, A., Delgado, N. and Della-Piana, C.K. (1999), Expanding Participation in Undergraduate Research Using the Affinity Group Model. Journal of Engineering Education, 88: 409-414.

[Hat] Russel S. Hathaway, Biren A. Nagda & Sandra R. Gregerman (2002). The Relationship of Undergraduate Research Participation to Graduate and Professional Education Pursuit: An Empirical Study *Journal of College Student Development*, 43

[Hay] Hayes R.Q., Whalen S.K. & Cannon B. (2009). 2008–2009 csrde stem retention report. Center for Institutional Data Exchange and Analysis, University of Oklahoma, Norman

[Mat] GitHub. Matris - a clone of Tetris made using Pygame. `https://github.com/Uglemat/MaTris`

[Mni13] Volodymyr Mnih and Koray Kavukcuoglu and David Silver and Alex Graves and Ioannis Antonoglou and Daan Wierstra and Martin Riedmiller (2013). Playing Atari with Deep Reinforcement Learning *arXiv*

[Mni15] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, et al. (2015). Human-level control through deep reinforcement learning. *Nature* 518, 529–533.

[Nag] Biren A. Nagda, Sandra R. Gregerman, John Jonides, William von Hippel, & Jennifer S. Lerner (1998). Undergraduate Student-Faculty Research Partnerships Affect Student Retention *The Review of Higher Education*, 22.1, 55–72.

[Pud] Lara Pudwell (2017). Teaching the Inquiry Process Through Experimental Mathematics *PRIMUS*, 27 (2), 281–292.

[Ste] Matt Stevens, Sabeek Pradhan (2016). Playing Tetris with Deep Reinforcement Learning `http://cs231n.stanford.edu/reports/2016/pdfs/121_Report.pdf`, Retreived May 25, 2021.

[Sut] Richard S. Sutton, Andrew G. Barto (2018). Reinforcement learning: An introduction *MIT press.*

[Tfl] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Rafal Jozefowicz, Yangqing Jia, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Mike Schuster, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. *TensorFlow: Large-scale machine learning on heterogeneous systems*, 2015. Software available from `tensorflow.org`.

[Way] Heidi A. Wayment and K. Laurie Dickson (2008). Increasing Student Participation in Undergraduate Research Benefits Students, Faculty, and Department. Teaching of Psychology, 35: 194–197.

[Wil] Zakiya S. Wilson, Lakenya Holmes, Karin deGravelles, Monica R. Sylvain, Lisa Batiste, Misty Johnson, Saundra Y. McGuire, Su Seng Pang and Isiah M. Warner (2012). Hierarchical Mentoring: A Transformative Strategy for Improving Diversity and Retention in Undergraduate STEM Disciplines *J Sci Educ Technol*, 21, 148–156.

[Zyd] Andrew L. Zydney, Joan S. Bennett, Abdus Shahid, & Karen W. Bauer (2002). Impact of undergraduate research experience in engineering. *Journal of Engineering Education* 91, no. 2, 151-157.

Department of Mathematics, Towson University, Towson, MD 21252
Email address: `mgluck@towson.edu`