

Modular experiential learning for secure, safe, and reliable AI

Curricular Initiative to Promote Education in Trustworthy AI

Alvis C. Fong*
Department of Computer Science,
Western Michigan University
alvis.fong@wmich.edu

Ajay K. Gupta
Department of Computer Science,
Western Michigan University
ajay.gupta@wmich.edu

Steve M. Carr
Department of Computer Science,
Western Michigan University
steve.carr@wmich.edu

Shameek Bhattacharjee
Department of Computer Science,
Western Michigan University
shameek.bhattacharjee@wmich.edu

Michael Harnar
Evaluation Center, Western Michigan
University
michael.harnar@wmich.edu

ABSTRACT

Artificial intelligence (AI) is increasingly applied to IT systems. However, AI can be manipulated to perform undesirably, exhibit biases or abusive behaviors. When AI algorithms are parallelized on high-performance computing-based cyberinfrastructure (CI), such misbehaviors and uncertainty can multiply to obscure the root causes. Secure, safe, and reliable computing techniques can mitigate these problems. The project described in this paper aims to inform curriculum and develop materials to educate students who use AI from the outset, so that they will first become aware of the issues and secondly practical considerations will be integrated with theory in classes. Intensive, multi-faceted, modular, experiential learning units are designed to rapidly upgrade the skills of current and future CI users, so they can apply new skills to their tasks. The loosely coupled modules can be taken as standalone self-directed units or integrated into existing classes, starting with CS 1 and CS 2, which are taken by many non-CS STEM students. In a sandpit environment, learners take measured risks when guided on a journey of discovery. The primary purpose of this paper is to present key findings of the research following a 2-year pilot. A secondary purpose of the paper is to disseminate this exciting endeavor broadly, so that likeminded educators and researchers can consider participating in the project.

CCS CONCEPTS

• Information Technology; • Education; • Learning Management Systems;

KEYWORDS

Curricular initiative, artificial intelligence trustworthiness, experiential learning, IT education

*Corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

SIGITE '22, September 21–24, 2022, Chicago, IL, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9391-1/22/09...\$15.00

<https://doi.org/10.1145/3537674.3554756>

ACM Reference Format:

Alvis C. Fong, Ajay K. Gupta, Steve M. Carr, Shameek Bhattacharjee, and Michael Harnar. 2022. Modular experiential learning for secure, safe, and reliable AI: Curricular Initiative to Promote Education in Trustworthy AI. In *The 23rd Annual Conference on Information Technology Education (SIGITE '22)*, September 21–24, 2022, Chicago, IL, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3537674.3554756>

1 INTRODUCTION

Members of the Transformative Interdisciplinary Human+AI Research Group [1] at Western Michigan University (WMU), together with public and private partners at various US locations and beyond, aim to address a critical shortage in STEM IT workforce that understands the anticipated immense technical and societal changes brought about by artificial intelligence (AI). The need for a strong AI-informed workforce is exemplified by the American AI Initiative [2]. Taking a convergent approach, the project described in this paper primarily aims to integrate core literacy and advanced skills at the intersection of Secure, Safe, Reliable (SSR) Computing, High Performance Computing (HPC), and AI into the educational curriculum across STEM disciplines. This will prepare faculty, undergraduate, and graduate students for large-scale data handling and analytics. Our work focuses on institutions that have comparatively lower levels of advanced cyberinfrastructure (CI) adoption, such as second-tiered institutions (Carnegie Classification R2), historically black colleges and universities (HBCU), and community colleges.

The project's secondary aim is to lay the groundwork for future broadening adoption of advanced CI training resources to influence wide segments of CI communities. The course is advanced with outreach activities to establish and maintain a pipeline of talents from pre-college to 2-year and 4-year colleges, and graduate programs. Key project artifacts include:

1. Critical segments of the relevant computer science (CS) and IT curricula will be updated to integrate SSR thinking and practices with AI software running on advanced HPC CI.
2. A series of reproducible, customizable, modular, experiential educational units that can be integrated within existing CS courses and/or taken as standalone self-directed learning activities.

By actively involving regional partners (community colleges, consortia), the project seeks to broaden participation by engaging

segments of diverse underrepresented groups in Michigan and beyond. The target audiences include students (community college students, university undergraduates and graduates) and faculty at participating educational institutions, consortia members, and other IT practitioners. Outreach activities, such as workshops for local high school students, will further diversify the pool of participants. The modular design will ensure compatibility with future initiatives to integrate ubiquitous cloud-based services that provide customized access to training resources by a wide range of CI users both on the job and before, thus contributing to the agenda of the National Strategic Computing Initiative (NSCI).

The rest of this paper is organized as follows. Section 2 explains the importance of SSR AI that runs on HPC CI and highlights key challenges and opportunities that motivated this research. Section 3 presents the project design. Section 4 describes that project deliverables. Evaluation of the project deliverables and key findings are presented in Section 5. Finally, Section 6 concludes the paper and suggests possible future research.

2 CHALLENGES AND OPPORTUNITIES

2.1 What is SSR AI running on HPC CI and Why?

We consider safety, security, and reliability as three vital ingredients of achieving trustworthy and human-friendly AI. AI itself comprises multiple dimensions: perception, reasoning, abstraction, learning, actuation, etc. Machine learning (ML) [3],[4], which grew out of AI, is often considered on its own due to significant recent advances in deep learning (DL). In trying to define SSR AI, we draw inspirations from the thematic pillars of the Partnership on AI [5], the first of which stipulates that AI tools used in safety-critical areas must be safe, trustworthy, and aligned with customary ethical standards. Our working definition is broader, encompassing not just safety-critical applications of AI, but all types of AI deployment that can affect people in both cyberspace and the real world [6]. For example, safe interaction with chatbots is a concern because it can affect users' moods and behaviors (including some that are self-destructive and/or life threatening) both online and offline. This is especially true when it involves children, the elderly, the mentally handicapped, and other vulnerable groups of people. Further, and in alignment with the second thematic pillar [5], SSR AI aims to minimize biases so that disparate user groups across political and socioeconomic spectra will feel fairly treated.

While SSR AI algorithms can be developed to solve real-world problems in principle, these algorithms need to be implemented on architecture to solve problems in practice. Modern algorithms, such as DL neural networks (NNs) that make split-second decisions, often run on advanced HPC CI. It is imperative to ensure that HPC CI is a fully integrated part of any deployed SSR AI solution. Overall, SSR AI that runs on advanced HPC CI is about building trust between humans and AI, such that human users of AI truly feel that all types of human-AI interactions can be safe, secure, and reliable. It is up to current and future CI users and contributors to establish, maintain, magnify, and broadcast this trust as AI continues to permeate every facet of human life and activity moving forward.

2.2 Challenges

An aspect of the current drive toward educating the next-generation STEM IT workforce is a multi-prong / multi-agency response to the projected acute shortage of well-qualified IT professionals for the emerging knowledge-based economy. As articulated by X+Computing and Computing+X, computing permeates numerous aspects of human, social, and scientific endeavors. Disciplines, such as computational biology and evolutionary computation, underscore the intricacies of computing and disparate STEM subjects, and other disciplines further afield, from business analytics and economics to social sciences. CI becomes the critical foundation and catalyst of many convergent human endeavors.

AI and big-data-driven AI that runs on HPC CI are in flux; there are many activities going on in these areas and at the intersection of these areas. With rising popularity of AI in our daily life, there is an urgent need to ensure that the next generation computer scientists, who may be familiar with failsafe computing and related techniques, understand how AI can fail. AI can fail in many ways, e.g., lack of explainability of AI decisions [7], lack of common sense and contextual reasoning, and potential threat of malicious attacks by human or machine agents to negatively influence AI. [7] highlights the dilemma facing users of AI systems. For example, there are end users, such as doctors performing medical diagnosis, who often choose to sacrifice some precision for interpretability by choosing rule-based systems over state-of-the-art DL NNs. Because ML systems trained on data learn facts about the data rather than facts about the real world, these systems can fail miserably when taken out of the operating environment they were originally designed for; how and when they fail can be hard to understand. Conversely, existing methodology in SSR computing can influence future AI algorithms and CIs by making them more robust. Robustness and AI failure models are especially important when uncertainty exists everywhere in daily life. By taking a data-driven approach and leveraging HPC CI, we aim to first establish a two-way connection between SSR computing and AI, and secondly instill cross-discipline active-learning educational materials to enhance the skill sets of SSR computer scientists and IT researchers / practitioners that use AI.

2.3 Opportunities

For all its promises, AI is far from infallible. AI systems have been known to be fooled to perform undesirably, exhibit biases or abusive behaviors [6]. There is still a lack of in-depth understanding of AI failure models. When AI algorithms are parallelized on HPC CI, such misbehaviors and uncertainty can multiply to an extent that the root causes become even more obscured and untraceable. Advanced SSR computing techniques can mitigate these problems.

The project is among the first of its kind that takes a two-prong approach toward addressing the challenges outlined above. The first entails development of innovative experiential modular learning materials to enhance the capability of a broad range of AI users across multiple STEM disciplines. An important aspect of this research direction is to increase the awareness of what AI can and cannot do in terms of problem solving across multiple STEM disciplines. Another key aspect is a fine balance between theory and practice. Learners gain firsthand experiences of the limitations of

using AI, such as AI biases, through practice. They can relate their practical experience with the underlying theory. Once the learners become aware of the issues and limitations, they are then guided through a process of developing a mitigation strategy. In the more advanced learning modules, learners take a further step in implementing and testing an instance of their mitigating strategy. The flexible learning modules are customizable for different domain-specific problems. Examples drawn from different STEM disciplines add relevance to learners' experience. For example, mechanical engineering students develop AI models for fuel consumption predictions; electrical and computer engineering students leverage AI to find optimal load balancing in a power grid scenario or between CPU and GPU usage.

To create an enriching experience for learners, there is an emphasis on practical experiential learning that complements theory. In particular, each learning module begins with the description of a problem that needs theory and knowledge they learn in class, in addition to other available resources, to solve. Learners gain firsthand experience of what AI-related lapses and vulnerabilities look like. Through a guided exploration, learners formulate a goal and develop a problem mitigation strategy (solving in principle). In the more advanced learning modules, learners will be guided toward implementing a solution, which is then evaluated against the set goal (solving in practice). Although modules will be built on top of fundamental knowledge, the modular learning units are designed to be loosely coupled. This means learners can choose in what order they wish to use the modules, and it also means there can be multiple entry points for learners.

The second involves informing relevant curricula, such that AI-thinking can become ingrained in the curricula themselves. The learning modules have been developed to be taken either as self-directed online learning materials or integrated into existing or new classes. It is also possible to use the learning modules in a hybrid mode that combines in-class learning and self-directed online learning. The most fundamental modules can be integrated into foundational CS courses, e.g., CS1 and CS2, which are often taken by non-CS students. For example, artists might find aspects of AI-created digital artwork useful. By exposing learners to what AI can and cannot do early in a STEM curriculum, they can develop a good appreciation – both theory and practice – of what they can meaningfully achieve with AI.

Guided by the principles of Open AI's Charter [8], we aim to promote cooperation among participants in the communities of CI users and CI contributors, system interoperability, and standards pertaining to SSR. As multiple AI research groups, supported by advanced CI, relentlessly race toward the goal of artificial general intelligence (AGI), the Charter reminds us of the importance of balancing technological advances with positive human and societal impacts.

Specifically, it underscores the importance of ensuring broadly distributed benefits across humanity, long-term safety, and the need to establish a cooperative framework toward achieving AGI, e.g., sharing resources and contributing to standards. Even before researchers reach the eventual goal of AGI, incremental advances, e.g., augmented AI aimed at multiplying human cognitive power and capabilities [9], can also have profound human and societal impact. Therefore, it is critically important to seize the opportunity

now and instill SSR knowledge and practices in current and future CI users and CI contributors. It is against this backdrop that we develop our experiential learning modules.

3 PROJECT DESIGN

3.1 Innovative Features

The innovative design features of this project are:

- Holistic view of convergence involving SSR computing, AI, and HPC CI. There is deliberately a two-way connection in which SSR theory and practice influence AI, and is conversely influenced by advances in AI. There is an emphasis on contemporary data-driven AI that runs on HPC CI.
- The design is informed by our in-course trial and will be further refined by our outreach trial. Further inspiration comes from multiple influential sources, e.g., [5],[10],[11], that inform current and future debate.
- Intensive, multi-faceted, modular, experiential learning materials developed to rapidly enhance knowledge and practical skills of a range of CI users and CI contributors. Versatility that is built in to the modules means that they can be integrated with existing courses and/or taken as standalone self-directed learning units. The latter have been found effective [12].
- Groundwork for future community of contributors and consumers of both full versatile experiential learning modules. These are designed to provide users with precise and rapid tools to tackle SSRAI-related problems.
- Use of a sandbox environment to encourage learners to take calculated risks toward solving open-ended problems in their domains. Learners gain confidence in what they can do with AI and see tangible results of their own effort.

As we develop curricular materials, we continue to consult those already developed, e.g., ACM Curricula, JTF on Cybersecurity Education, MOOC such as Udacity [13], Edx [14], and industry efforts, e.g. [15], [16]. Notably, our effort aims to complement, not compete against, existing resources.

3.2 Multidisciplinary Experts

Following the practice of collective impact [17], this project engages a diverse group of expert advisers from the outset. These experts inform initial development, as well as strategies for the design, implementation, and dissemination of the artifacts, plus strategies for maximizing buy-in from a broad spectrum of CI users and CI contributors. Throughout the project, they provide on-going feedback on the efficacy of the artifacts. They also participate in outreach activities. The panel of experts comprises the following:

1. Seven WMU faculty in CS, different branches of engineering (civil, electrical, mechanical), statistics, and business.
2. Eight multidisciplinary faculty from across the US, including those at HBCU minority and 2-year institutions.
3. Four industry experts from computing and computing+X industries, e.g., computational chemistry and pharmaceutical.
4. Four government agency experts, including those from NIST, US Army, etc.

4 PROJECT DELIVERABLES

4.1 Pedagogical Approach

Active learning, which encapsulates the idea of learning by doing, lays the foundation of this project. Furthermore, experiential learning that is informed by domain experts and topical issues, lends realism to a journey of discovery. Learners can relate their learning experiences and education materials to their daily activities and issues related to SSRAI that make headline news. By taking a modular approach, educational content can be flexible in how it is delivered to learners. Segmenting learning content into manageable chunks within each module allows learners to develop a sense of ownership and incremental successes. This helps learners build confidence needed to tackle more difficult problems. Modular design also allows for full integration into existing courses or delivered as standalone mini courses.

The modules are organized in three difficulty levels: Foundational (which can be integrated into CS1 and CS2, which are courses taken by many non-CS STEM students), Intermediate (which can be integrated into relevant upper undergraduate curriculum), and Advanced (which can be integrated into senior undergraduate and graduate courses). Collectively, these learning modules span all levels, from basic literacy to advanced.

To ensure wide adaptability and future proofing, the learning modules are designed to be agnostic in terms of programming languages. Learners can choose their preferred language, e.g., Python, R, but will in any case be encouraged to use open-source libraries and resources. Indeed, to ensure non-CS learners can maximally benefit from the learning experience, the need for programming (in any language) is minimized. For example, instead of getting learners to write their own programs, code snippets that can run in a Jupyter / Colab notebook environment are often provided. To enhance relevance, examples and data are drawn from real-world sources related to learners' disciplines, e.g., social media and news outlets, so learners can relate their study to matters that they care about. For consistency, each module is designed based on the following recommended format:

1. Introduction to background information by course instructor (if integrated) or domain expert.
2. Guided exploration leading to awareness and comprehension of a problem or issue; how it fits in the bigger picture.
3. Implement an instance of the issue to gain insight; it becomes "real" no matter how incredible it seemed initially.
4. Using provided resources, develop a countermeasure strategy based on set goal and problem (solving in principle).
5. Implement and evaluate the countermeasure (solving in practice).
6. Reflect on the outcome: could anything have been done differently / better?

It is important to note that this format is suggestive rather than prescriptive. Indeed, the developed initial modules do not all follow this format strictly. Typically, advanced modules tend to cover all six steps; other modules give more emphasis on some aspects than others. For example, there is more emphasis put on exploration leading to awareness than actually solving a problem and evaluating the solution in the foundational modules. The reason is that early-stage learners are not expected to possess a significant amount of

knowledge and skills to perform substantive problem-solving tasks. They will instead focus on experimenting with strategies toward problem solving based on the theory they have learned.

Although the modules have been developed to run in a series, they are loosely coupled. This means learners can select modules (or sequence of modules) most relevant or interesting, and in whatever order they desire. Furthermore, because in-class modules are fully integrated into existing courses with a balance between theory and practice and between learning and applying, there is no net increase in time to degree completion for degree-seeking students.

4.2 Twelve Initial Learning Modules (freely available from the project website <https://wmich.edu/cs/cybertraining>)

1. Math Toolkit for SSRAI running on HPC CI. (foundational) This module can be integrated with foundational math.

Synopsis: Mathematical preparation for students.

Learning Outcomes: Proficiency in math techniques and tools needed for SSRAI on HPC CI.

2. Algorithmic Exploration and Exploitation of an Intelligent System's weakness. (foundational) This module can be integrated with CS1.

Synopsis: Introduction to computational thinking; guided exploration of weaknesses of intelligent systems.

Learning Outcomes: High-level appreciation of several practical weaknesses of intelligent systems. Focus on one such issue and ability to develop a mitigating strategy using computational thinking.

3. Modular and Structured Software Development for Robust Intelligent Systems that run on HPC CI. (foundational) This module can be integrated with CS2.

Synopsis: Explore and develop strategies for strengthening an intelligent system with advanced software techniques.

Learning Outcomes: Improve robustness of an intelligent system using modular and structured software techniques.

4. Data Structures for SSRAI running on HPC CI. (intermediate)

This module can be integrated with CS 3xxx Big Data.

Synopsis: Application of data and file structures for SSRAI and are amenable to HPC processing.

Learning Outcomes: Understanding of the connections between data and files structures and SSRAI that runs on HPC CIs. Ability to design and implement moderately complex data and file structures that support SSRAI on HPC CI.

5. Deep learning with HPC. (intermediate) This module can be integrated with CS3xx Data and File Structures.

Synopsis: Implementation, training, and evaluation of DL algorithms with open-source resources such as Keras or PyTorch and the use of tensors for data representation. Using the MNIST data set, learners will first discover performance bottlenecks when they run their algorithms on general-purpose CPU architectures. They will then learn how to overcome the bottleneck by running their algorithms on advanced GPU and/or cloud CI.

Learning Outcomes: Understanding of tensors and tensor operations. Ability to develop deep learning algorithms written in Python and with Keras or PyTorch that run on HPC CI.

6. SSRAI Software Development for HPC CI deployment. (advanced, undergraduates and graduates) This module can be integrated with CS 5xxx Secure Software Development.

Synopsis: Application of the robust software techniques for identifying and mitigating threats and vulnerabilities of AI software agent programs that run on HPC CI.

Learning Outcomes: Understanding of the importance of SSRAI software development to minimize threats to AI systems. Ability to design, implement, and evaluate effective SSRAI software.

7. Vulnerabilities of Machine Learning. (advanced, undergraduates and graduates) This module can be integrated with CS 5xxx Machine Learning (ML).

Synopsis: Guided exploration of vulnerabilities of ML that make headlines and mitigating strategies. Learning Outcomes: First-hand account of how easy it is to manipulate ML techniques that are statistically accurate. Ability to design, implement, and evaluate effective countermeasures.

8. Beyond current generation AI and Toward Artificial General Intelligence. (advanced, undergraduates and graduates) This module can be integrated with CS 5xxx Artificial Intelligence (AI) or Cyber-Physical Systems.

Synopsis: An investigative look at harnessing contextual information from disparate sources (e.g., IOT sensors, physical-chemical-biological models). As research groups rush toward AGI, are important SSR considerations overlooked?

Learning Outcomes: In-depth appreciation of how contextual information extracted from disparate data sources can be used to inform SSRAI that runs on HPC CIs. Understanding of SSR technical issues often overlooked in AGI research.

9. Adversarial Machine Learning and Robust Trust Scoring Models. (advanced, undergraduates and graduates) This module can be integrated with ML, Cybersecurity, Networks, Computer Vision course.

Synopsis: Viewed from a cybersecurity perspective, students will gain knowledge in adversarial attacks in different ways and mitigating strategies organized under 4 parts.

Learning Outcomes: Broad view of types of adversarial attacks on ML algorithms and in-depth robust-by-design strategies.

10. Societal Impact of AI. (advanced, undergraduates and graduates) This module can be integrated with CS 5xxx Artificial Intelligence (AI) and/or CS4xxx Software Development II, which has a segment on ethics.

Synopsis: An investigative look at issues with profound consequences, e.g., chatbot misdeeds or future workforce. Learning Outcomes: Broad view of topical issues that affect human societies now and in the foreseeable future. In-depth understanding of one issue (or one group of related issues) that can contribute to meaningful debate.

11. Pitfalls of applying AI to Information Retrieval tasks. (advanced, graduates)

This module can be integrated with CS 6xxx Information Retrieval (IR).

Synopsis: An investigative look at several applications of AI to IR (e.g., learning to rank) in different use cases, e.g., web search, semantic web search, digital library search.

Learning Outcomes: Broad view of pitfalls of AI applications to IR. In-depth knowledge of one such pitfall. Ability to formulate and test an algorithmic approach toward mitigating the pitfall.

12. Real-Time SSRAI with HPC CI. (advanced, undergraduates and graduates) Formulated as a capstone project.

Synopsis: Putting everything together, students will apply knowledge learned toward building SSRAI systems that use HPC CI to repeatedly make good decisions in split seconds.

Learning Outcomes: Insight and practice in development of SSRAI for rapid decision making.

5 EVALUATION

5.1 2-Step Approach

A 2-step approach was taken for evaluation of the developed project deliverables (experiential learning modules) once they were ready in late August 2021. A small-scale trial was launched in the fall 2021 semester (from early September to mid-December 2021). It involved integrating two learning modules into two CS courses: Module 4 for a 3-thousand level course in Big Data (BD) and Module 11 for a 6-thousand level course in Information Retrieval (IR). The BD course made up of mostly lower/middle undergraduates with an equal split between CS and Data Science (DS) students. The IR was exclusively taken by graduate students (both Masters and Doctoral students).

For evaluation purposes, each integrated module + course pair has two separate but interrelated components: an educational component (i.e., completion of learning module) and a research component (i.e., completion of anonymized pre- /post-intervention surveys). Completion of the educational component was mandatory and accounted for 10% of a student's final course grade. However, completion of the research component was entirely voluntary. The voluntary response rate was approximately 75% (N=10) across the two classes. Even though the class compositions were quite different, there was no significant difference in the response rate in the two classes. Qualitatively, the small-scale trial revealed that the experiential learning modules were generally well received. There were no complaints of any kind. However, the small sample size meant no quantitative analysis was performed.

Following the small-scale trial, which led to minor content fine-tuning, a large-scale study was completed in spring 2022 (from early January to end of April 2022). 12 faculty members from different quantitative disciplines (CS, mechanical engineering, civil engineering, statistics, business analytics, etc.) across four US universities were involved in this large-scale launch. All participating faculty collected voluntary pre- and post-intervention surveys for analysis conducted by the independent evaluator in May – July 2022. Approximately 200 undergraduate and graduate students were affected.

5.2 Evaluation Instruments

Dr. Harnar is the independent evaluator responsible for the development of pre- and post-intervention survey instruments that measure efficacy of the developed learning modules. The pre-intervention survey asks participants 5 questions (3 closed, 2 open-ended) related to their preparedness and expectations of using AI in their work. The post-intervention questionnaire asks 20 questions (10

Table 1: Application of Kirkpatrick model to measure efficacy of the learning modules.

Kirkpatrick Stage	Interpretation	Construct
Reaction	The modules are relevant to the students' interest, they foster active engagement, and the ways of delivery of the modules are satisfying to the students.	Satisfaction Engagement Relevance
Learning	The participants gain knowledge or skills or experience positive changes in their attitude, confidence, or commitment to the subject.	Attitudes Knowledge
Behavior	The students see potential uses for the knowledge or skills learned during the module in other contexts or disciplines.	Application

Table 2: Summary of results (N=176 post-intervention)

	Reaction				Learning				Behavior	
	Satisfaction - Content 0..4	Satisfaction - Delivery 0..4	Engaged -2..0..+2	Relevance to Course -2..0..+2	Relevance to Field -2..0..+2	Attitude Worth -2..0..+2	Competence	Practice	Application to Daily Life Course -2..0..+2	
Mean	3.32	3.11	1.28	1.28	1.41	1.71	1.04	1.18	0.67	1.35
Median	3.00	3.00	1.00	2.00	2.00	2.00	1.00	1.00	1.00	2.00
Mode	3.00	4.00	2.00	2.00	2.00	2.00	1.00	2.00	0.00	2.00
Min	0.00	0.00	-2.00	-2.00	-1.00	-1.00	-2.00	-2.00	-2.00	-1.00
Max	4.00	4.00	2.00	2.00	2.00	2.00	2.00	2.00	2.00	2.00

closed, 10 open-ended) about the experiences of the learners. The times needed to complete these pre- and post-intervention surveys are approximately 5-10 minutes and 10-15 minutes, respectively. Evidently, they are not intended to be burdensome but at the same time allow Dr. Harnar to discover insightful information from the anonymized data. The surveys were developed based on the Kirkpatrick model [18] as shown in Table 1.

5.3 Evaluation Procedure

A general procedure was used in this study. This means results across different disciplines and institutions are generally compatible. Specifically, a hybrid mode combining integrated in-class learning and self-directed online learning was used. There was flexibility in how students were rewarded for their efforts. Some faculty members integrated a module into one of their courses and assigned a score that contributed to each student's final course grade. In other words, the learning module became an integral part of a course's assessment structure. Some chose to give students extra credit for completing a learning module. Others adopted a mixed approach combining the two.

5.4 Summary of Key Findings

Qualitatively, the developed experiential learning modules were very well received by the participating students and faculty. Clear guidance and instructions were provided to the learners, and as such very few questions were raised while the learners went through the modules. A few faculty members have furthermore provided valuable suggestions for further improvement of the learning content. Quantitatively, the post-intervention results are summarized in Table 2. Results show that students were mostly satisfied with

both the content and delivery. The lowest scores were competence (ability to explain key concepts) and application to daily life. These are areas of further improvement.

6 CONCLUSION AND FUTURE RESEARCH

Teaching AI and ML to non-CS majors has been found to be difficult [19]. In this research, we take a much broader view toward influencing AI curricula. Our approach to contributing to the Nation's educational curriculum/instructional material fabric is to develop highly innovative, topical, and modular curriculum/instructional material that can be integrated into undergraduate and graduate courses and/or delivered as self-directed online learning. Evaluations conducted in the last two semesters have demonstrated the efficacy of the developed experiential learning modules. Students from diverse backgrounds who participate in the activities will exhibit elevated levels of awareness of SSRAI running on HPC CIs compared to others, and they will possess enhanced abilities to formulate strategies to counter AI vulnerabilities.

Future research directions include evaluating the experiential learning modules in more varied modes of delivery. This will allow us to seek better understanding of whether in-class and/or self-directed online learning can lead to better outcomes. Areas of improvement include making the content more relevant to daily life and strengthening learners' ability to explain key concepts. We are also interested in widening both the breadth of offerings (i.e., creating more experiential modules beyond the initial twelve) and the types of learners that can benefit from the work. Indeed, outreach activities are already underway to customize some of the learning materials to cater to high school students. Another research direction that is in progress is to investigate how best to deeply and

broadly inform future curricular development, not just in CS, but also other relevant disciplines that increasingly apply AI toward solving their domain-specific problems.

ACKNOWLEDGMENTS

This work was supported by the U.S. National Science Foundation under grant no. 2017289.

REFERENCES

- [1] Transformative Interdisciplinary Human+AI Research Group at Western Michigan University. Website: <https://fong.cs.wmich.edu/>
- [2] The American AI Initiative: Accelerating America's Leadership in Artificial Intelligence. Available at <https://www.whitehouse.gov/articles/accelerating-americas-leadership-in-artificial-intelligence/>
- [3] Xiaochun Yang and Jiawei Liang. 2018. Machine Learning Online Education Experience for Non-technical People. In Proceedings of the 49th ACM Technical Symposium on Computer Science Education (SIGCSE'2018). ACM, New York, NY, USA, 1075–1075.
- [4] R.Young and J. Ringenberg. 2019. Machine Learning: An Introductory Unit of Study for Secondary Education. In Proceedings of the 50th ACM Technical Symposium on Computer Science Education (SIGCSE'2019). ACM, New York, NY, USA, 1274–1274.
- [5] Partnership on AI. Thematic Pillars. Available at www.partnershiponai.org/.
- [6] S. Saeedi, A.C. Fong, S. P. Mohanty, A. K. Gupta, and S. M. Carr. 2021. Consumer Artificial Intelligence Mishaps and Mitigation Strategies," IEEE Consumer Electronics Magazine, early access, 2021.
- [7] A.M. Bornstein, Is Artificial Intelligence Permanently Inscrutable? September 2016. Available at <http://nautil.us/issue/40/learning/is-artificial-intelligence-permanently-inscrutable>.
- [8] OpenAI's Charter. Available at <https://blog.openai.com/openai-charter/>.
- [9] A. Fong and G. Hong. 2019. Augmented intelligence with ontology of semantic objects. In Proceedings 4th International Conference on Contemporary Computing and Informatics (iC3I 2019), Singapore, pp. 1-4, December 2019. DOI: 10.1109/IC3I46837.2019.9055577.
- [10] B. Zhang and A. Dafoe "Artificial Intelligence: American Attitudes and Trends." Oxford, UK: Center for the Governance of AI, Future of Humanity Institute, University of Oxford, 2019. Available at https://governanceai.github.io/US-Public-Opinion-Report-Jan-2019/us_public_opinion_report_jan_2019.pdf.
- [11] J.S. Brennen, P.N. Howard, and RK Nielsen, An Industry-Led Debate: How UK Media Cover Artificial Intelligence, Reuters Institute for the Study of Journalism, University of Oxford, December 2018. Available at https://reutersinstitute.politics.ox.ac.uk/sites/default/files/2018-12/Brennen_UK_Media_Coverage_of_AI_FINAL.pdf.
- [12] En-Shiun A. Lee, Karthik Kuber, Hashmat Rohian, and Sean Woodhead. 2021. Pillars of Program Design and Delivery: A Case Study using Self-Directed, Problem-Based, and Supportive Learning. In Proceedings of the 52nd ACM Technical Symposium on Computer Science Education (SIGCSE'2021). ACM, New York, NY, USA, 205–211. DOI:<https://doi.org/10.1145/3408877.3432458>
- [13] Udacity. Available at <https://blog.udacity.com/2016/11/artificial-intelligence-curriculum.html>.
- [14] Edx on AI. Available at <https://www.edx.org/microsoft-professional-program-artificial-intelligence>.
- [15] Google AI education. Available at <https://ai.google/education/>.
- [16] Microsoft AI School. Available at <https://aischool.microsoft.com/en-us/home>.
- [17] J. Kania and M. Kramer. Collective Impact, Stanford Social Innovation Review, Winter, 2011.
- [18] Kirkpatrick evaluation model. Available at <https://www.kirkpatrickpartners.com/the-kirkpatrick-model/>
- [19] E.Sulmont, E. Patitsas, and J. R Cooperstock. 2019. What is hard about teaching machine learning to non-majors? Insights from classifying instructors' learning goals. ACM Transactions on Computing Education (TOCE), Vol. 19, 4 (2019), 1–16.