Infusion of Machine Learning Concepts and Skills: A Workshop Review*

Sambit Bhattacharya, Bogdan Czejdo, Valentin Milanov
Department of Mathematics and Computer Science
Fayetteville State University
Fayetteville, NC 28301
{sbhattac,bczejdo,vmilanov}@uncfsu.edu

Abstract

With the recent successes of research in Artificial Intelligence (AI), specifically Machine Learning (ML), infusion of AI and ML concepts and use of tools can help increase the responsible adoption of AI and ML in different disciplines. We report on the design of a workshop on ML, results from the evaluation of the workshop, and suggests topics and pedagogical approaches that may be adopted for disseminating the core concepts of ML which are among the most prevalent data-driven modeling approaches in AI. The workshop individually targeted a diverse range of participants including college and university faculty members from different disciplines at 2-and-4 year institutions. The materials of the workshop were made publicly available. Future workshops and educational modules building on this work will be able to infuse knowledge of the frontiers of AI and problems & benefits.

1 Introduction

When AI pioneer Geoff Hinton says "Deep learning is going to be able to do everything," [7] many are likely to take notice, but understanding and agreeing on what exactly he means by everything and what it means to be able to do those things can become a very challenging task. Dr. Hinton is in a group of esteemed individuals who have been making similar statements. Even after these pronouncements there seems to be a significant number of examples where the actions of AI in the real world are of major concern [13].

Disseminating AI knowledge is an important goal and institutions of higher education need to lead in this effort by engaging educators and researchers from multiple disciplines. The overall goal of the workshop development and implementation project reported in this article was to conceptualize, design, and implement a model for individually targeted infusion of Machine Learning (ML) concepts and skills. It describes the experiences of the authors who organized a summer workshop (July 20-24, 2020) on teaching ML techniques at Fayetteville State University (North Carolina). The foundation for launching this effort targeted professors at 2-and-4-year colleges and universities. It had two goals: 1) Increase understanding of the successes of AI and its limitations as described above, 2) Provide an enriching teaching and learning experience using active learning pedagogical techniques.

In addition to existing data trends from national market analyses on current and future jobs in AI and allied fields, the authors collected their own data to better understand and assess the local demand for teaching AI in colleges and universities. This was done by gathering feedback from faculty members from local community colleges and universities using online surveys, personal correspondence, conducting site visits to colleges and universities, doing lectures, and having follow up conversations.

2 The Need for Teaching ML at Colleges and Universities

The authors conducted a series of lectures at various institutions within the University of North Carolina System including: Winston Salem State University, Elizabeth State University, NC A&T State University, and UNC Pembroke. The lectures were attended by both faculty and students from Engineering disciplines and Sciences. The authors also conducted phone conversations with local officials from Cumberland county (where Fayetteville is located) and the adjacent Bladen and Hoke Counties that are primarily rural and underserved counties of the state.

As a follow up to the visits and conversations, an online survey was administered to assess the biggest need area(s). A total of 39 responses were recorded. The most common response (above 70%) was Artificial Intelligence (AI). The respondents indicated that it is a high priority to educate underrepresented students, students from underserved areas such as rural and/or remote communities, and first-generation college-enrolled students, to help prepare them for future education and the STEM workforce and particularly AI within STEM. In line with the directions suggested by the survey respondents, there are conversations happening at international meetings like the AAAI 2020 Workshop on Equity and Diversity in AI [3]. AI researchers and educators are also researching the discrepancy in the quality of education worldwide, including AI education.

Fayetteville State University is also in the proximity to Fort Bragg, one of the largest military installations in the country. The military personnel and their spouses make up a significant portion of the student population and contractors seek talented graduates. The students trained in computing and developing ML applications are in increasing demand. Job searches on Google and Glassdoor using simple terms such as "data scientist," "AI," "Fort Bragg," and "Fayetteville," yield in 50+ hits.

3 Machine Learning Explainability

Numerous barriers prevent students from acquiring AI skills [3]. Underserved communities typically lack resources—equipment such as computers among other tools. While adequate supply of equipment remains a challenge, another gap is limited physical space for conducting hands-on AI activities such as robotics and distributed sensing for decision making. Other challenges are lack of enough trained instructors qualified to teach AI with qualified members leaving underserved communities in search of better opportunities. Because of these reasons, most underserved communities face critical challenges in educating local students in AI and therefore, fail to accommodate the rapidly changing state of art in AI, especially development of new ML models and their applications.

In addition to the above challenges, there are specific challenges related to the limited explainability of AI models [12]. AI models are very heterogenous requiring a substantial effort to construct an integrated teaching approach. Explainability is a growing area of research in ML [4, 11]. Most of the research is oriented towards explainability as an indicator of limitations of ML models in various disciplines [14, 10, 15]. ML explainability is strongly related to educational "explainability." In order to design the individually targeted AI workshop, the authors have extracted from the literature a taxonomy of explainability components [6, 5]that directly relate to the educational applications and outcomes to help construct educationally-oriented ML "explainability." An important characteristic of the workshop plan was to emphasize diversified ML content and teaching methods that are responsive to the needs of learners from both Computer Science and other disciplines.

Several core elements of each ML process and the taxonomy of the "explainability" of ML processes in a broader sense are presented in the table below. According to this classification, "transparency" is the first component and it addresses the fundamental issue of understanding any ML process. In a basic ML process, a model is learned from the given input data and with a specific learning algorithm. Then, in the next phase, the output results are computed using the learned model. Transparency of an ML process can therefore be further classified into the transparency of an overall process structure and function, the transparency of individual process components, the transparency of the learning algorithm, and transparency of how the specific solution is obtained by the algorithm. This document refers to these four transparency categories as shown in Table 1 as: process transparency, component transparency, transparency of learning algorithm, and transparency of testing algorithm.

Informally, full transparency is the opposite of the ”black-box” restricted visibility, where using an ML model to perform predictions is the main goal, but the question of how the model outputs its predictions is left unanswered. Transparency thus indicates the level of understanding on how the model actually works. Transparency was explained further by referring to simulatability (considering the entire model) decomposability (considering individual components) and algorithmic transparency for the learning and testing algorithm. All transparency categories determine how well each component's function and structure can be explained in understandable terms to a human.

The analysis of the role of hyper-parameters is also important in order to consider all ML process complexities. Hyper-parameters can determine general model structure, components, optimization algorithm, learning rate, and the size of stochastic sampling, and are often chosen in a heuristic fashion [12]. Due to the possible presence of several local minima, the solution is usually not easily reproducible; therefore, these decisions are not transparent. They can be considered as part of the transparency of the process or associated with the learning algorithm. Transparency of ML processes is an important part of student learning knowledge since this knowledge is needed to optimize the models with respect to high accuracy. Recently there is also growing need to provide the explanation of output results with respect to domain knowledge [11].

In addition to “transparency,” another core element of each ML process is identified in Table 1. This category is titled “explainability for domain knowledge” and it requires techniques to discover the underlying reasons for the ML to produce specific results. The explanability of a complex ML process can be improved by using simpler models with only less accurate result. Traditionally, decision trees were used for the explanation of results. More recently the model-agnostic module LIME [11] by Ribeiro et al. was proposed. In the latter case, feature importance can be detected by local linear proxy model in the neighborhood of a focused data. The educational content analyzed here, namely ML processes, is diversified in terms of explainability as shown in the Table 1. Such content is significantly different from content in other areas/disciplines and making the educational process effective requires special effort. Frequently, ML processes cannot be easily explained, and the user has to deal with "black-box" models. Mapping the hard-to-explain content with the proper pedagogical methods is critical. There are multiple pedagogical approaches to teaching various algorithms, procedures and rules developed for the "black-box" or "gray box" approach.

4 POGIL as a Teaching Method To Address ML Teaching Challenges

The Process Oriented Guided Inquiry learning (POGIL) teaching technique [9] supported the authors efforts and help them design hands-on activities in AI education. They aligned explainability (Table 1) into the POGIL teaching method. As a result, they were able to choose the best teaching method instead of applying one solution that may not fit all.

Table 1: Taxonomy of AI Processes Explainability

Model/ Algorithm	**Transparency Of Whole Process**	**Transparency Learning Algorithm**	**Transparency Testing Algorithm**	**Transparency Hyper-Parameter**	**Explainability for domain knowledge**
Linear Regression	High	High	High	High	High
Decision Tree	High	High	High	High	High
SVM	Moderate to High	Moderate	Moderate to High	Low	Low
Generic NN	Moderate to High	Moderate	Moderate to High	High	Low
Generic Convolutional Networks	Low to Moderate	Low to Moderate	Low to Moderate	Low to Moderate	Low to Moderate
Generic NL Deep Learning	Low	Low	Low	Low	Low

As a pedagogical approach, POGIL allows students to socially construct knowledge through iterative cycles that include three steps: exploring a model, inventing a concept, and applying the resulting ideas. A growing body of research indicates that relative to lecture-based approaches, POGIL supports student learning more effectively [8]. It allows for a structured approach that requires participants to work in self-managed or regulated teams to explore content, ask questions, solve problems, conduct analysis, record the proceedings, and cooperate to draw valid conclusions. While there are any number of student-centered active learning classroom techniques, POGIL is unique in that it makes students responsible for their own learning, in collaborative teams, so it helps them develop group process skills while they are gaining content knowledge.

Within the POGIL methodology, the authors used an AI platform called Google Colab which hosts notebook service for AI and Data Science [2]. The main intent behind using the POGIL technique was to enable workshop participants to learn and work in teams, collaborate to understand a concept and solve a structured problem or set of questions rather than being given the content via a lecture by a teacher. In summary, the POGIL approach included: 1) Faculty generated model and related content and 2) Specific problem or defined set of questions for small groups to solve/answer with little guidance from the facilitator.

5 Workshop Organization

In summer 2020, the authors of this paper hosted a week-long workshop titled Strategies to Train and Engage Students In Artificial Intelligence. The workshop brought together educators from across North Carolina who learned to use new tools and strategies for teaching AI that would benefit students from diverse backgrounds [1]. The workshop was evaluated by MN Associates, Inc. (MNA). A total of 19 participants engaged in hands-on sessions offered at the AI workshop. Of these, one of the educators was from a community college and 18 were from various universities. The following were goals of the workshops:

- Develop strategies to provide AI knowledge and skills to college students
- Develop techniques for exposing students to the frontier of AI knowledge–problems & benefits emphasizing Machine Learning
- Developing and applying AI computing frameworks including Cloud computing for ML
- Using Python with AI libraries
- Integrate AI computer coding in teaching methods

In the pre-workshop phase the authors and MNA co-developed an online participation interest form (pre-survey) that was sent to various UNC system colleges, universities, and high schools to help recruit workshop participants. Questions related to demographic questions as well as prior/current knowledge of AI topics such as Regression, Decision Tree, and Artificial Neural Networks were posed. Upon completion, the authors extended a formal e-invite to the participants in preparation of the workshop.

Daily post-survey questions were posed as a series of multiple-choice and open-comments questions related to each day's sessions. In both pre- and post-workshop surveys, the participants were asked to rate on a 5-point Likert scale (Low to Very High) the level of their understanding of three topics- Regression, Decision Tree, and Artificial Neural Network before and after the workshop.

6 The Workshop and Results

During the workshop three model building approaches were introduced: Regression, Decision Tree, and Artificial Neural Networks. The materials from this workshop are available on Github [1].

6.1 Regression

The relatively well-known regression model was introduced first. Understandably, the initial level of the understanding of the topic among the participants was already very high compared to other topics. Still, majority of participants reported that their understanding of regression model improved after the exposure to POGIL lessons. During this phase of the workshop it was evident that there were two relatively different groups of participants in respect to previous exposure to ML concepts. The percentage points change in understanding from pre-to-post workshop is shown in Figure 1.

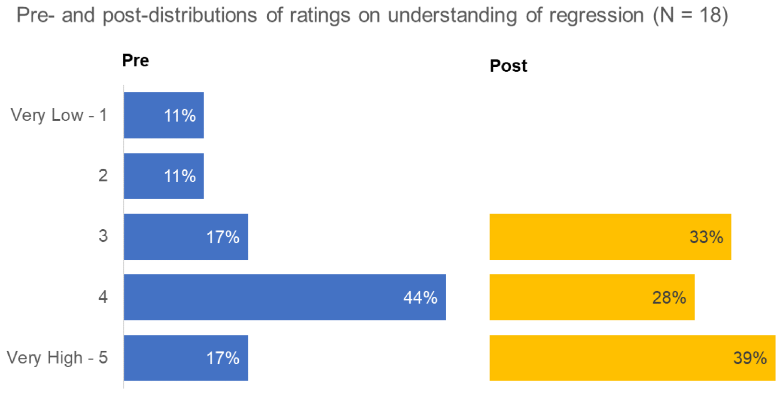

Figure 1: The percentage points change in understanding from pre-to-post workshop for regression model.

6.2 Decision Tree

The less-known (relative to Regression) Decision Tree learning model was introduced next. Surprisingly, the initial level of the understanding of the topic among the participants was already high compared to other topics. Our hypothesis is that the answers were related to using the decision tree and not necessary constructing it in a bottom-up method from labeled data, as done in ML. Based on the survey, majority of participants also reported that their understanding of Decision Trees model improved after the exposure to POGIL lessons. During this phase of the workshop it was observed that the first group moved from level 1 and 2 to level 3 and the second group moved from level 3 to 4 or 4 to 5. There were no big jumps as in the case of regression but smaller and consistent growth of reported knowledge. The percentage points change in understanding from pre-to-post workshop is shown in Figure 2.

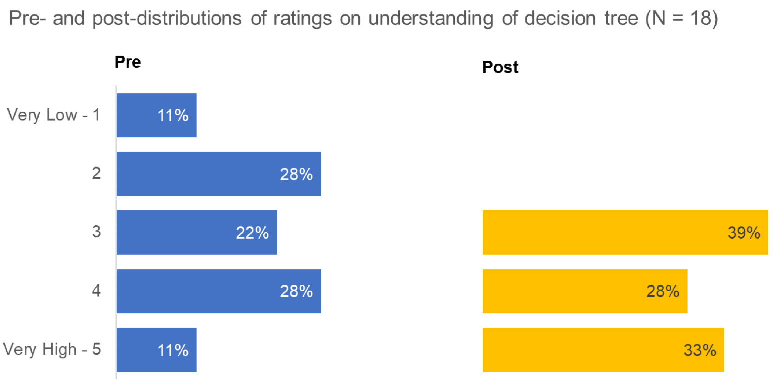

Figure 2: The percentage points change in understanding from pre-to-post workshop for Decision Tree model

6.3 Artificial Neural Network (ANN)

The ANN model was introduced last. The results of workshop here show a considerable impact. First of all, the initial level of the understanding of the topic among the participants was much lower compared to other topics. The majority of participants also reported that their understanding of ANN model improved significantly after the exposure to POGIL lessons. The percentage points change in understanding from pre-to-post workshop is shown in Figure 3.

7 Summary

Based on the assessments conducted by the authors during a summer workshop on teaching AI, results from the evaluation of the workshop, and understanding the burgeoning needs of AI education, this article suggests topics and pedagogical approaches that may be adopted for disarming the core concepts of

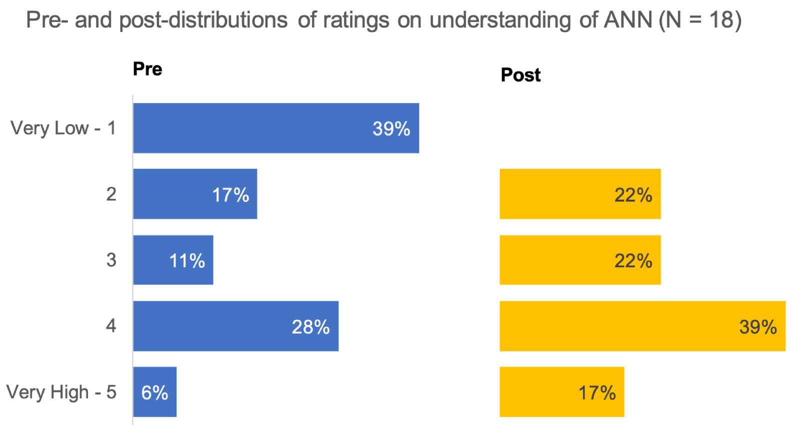

Figure 3: The percentage points change in understanding from pre-to-post workshop for ANN model

ML which are among the most widely used data-driven modeling approaches in AI. These workshops may be individually targeted to a diverse range of participants including both Computer Science and non-Computer Science students, and college and university faculty members from different disciplines at 2-and-4 year institutions. The content of future workshops and educational modules should continue to expose students to the frontier of AI knowledge and problems & benefits.

Acknowledgements

This article is based upon work supported by the National Science Foundation under Grant No. 1818694.

References

[1] AI workshop. `https://github.com/sbhattac/ai-workshop`.

[2] Welcome to colaboratory. `https://colab.research.google.com/`.

[3] AAAI. AAAI-20 call for diversity and inclusion activities. `https://aaai.org/Conferences/AAAI-20/aaai-20-diversity-inclusion/`.

[4] Robert Andrews, Joachim Diederich, and Alan B Tickle. Survey and critique of techniques for extracting rules from trained artificial neural networks. *Knowledge-based systems*, 8(6):373–389, 1995.

[5] Alejandro Barredo Arrieta, Natalia Díaz-Rodríguez, Javier Del Ser, Adrien Bennetot, Siham Tabik, Alberto Barbado, Salvador García, Sergio Gil-López, Daniel Molina, Richard Benjamins, et al. Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion*, 58:82–115, 2020.

[6] Riccardo Guidotti, Anna Monreale, Salvatore Ruggieri, Franco Turini, Fosca Giannotti, and Dino Pedreschi. A survey of methods for explaining black box models. *ACM computing surveys (CSUR)*, 51(5):1–42, 2018.

[7] Karen Hao. AI pioneer Geoff Hinton:"deep learning is going to be able to do everything", 2020. `https://www.technologyreview.com/2020/11/03/1011616/ai-godfather-geoffrey-hinton-deep-learning-will-do-everything/`.

[8] Helen H Hu, Clifton Kussmaul, Brian Knaeble, Chris Mayfield, and Aman Yadav. Results from a survey of faculty adoption of process oriented guided inquiry learning (POGIL) in computer science. In *Proceedings of the 2016 ACM Conference on Innovation and Technology in Computer Science Education*, pages 186–191, 2016.

[9] Clifton Kussmaul, Helen H Hu, and Matthew Lang. Guiding students to discover CS concepts and develop process skills using POGIL. In *Proceedings of the 45th ACM technical symposium on Computer science education*, pages 745–745, 2014.

[10] Jesus Mena. *Machine learning forensics for law enforcement, security, and intelligence*. CRC Press, 2011.

[11] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. " why should I trust you?" explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144, 2016.

[12] Ribana Roscher, Bastian Bohn, Marco F Duarte, and Jochen Garcke. Explainable machine learning for scientific insights and discoveries. *IEEE Access*, 8:42200–42216, 2020.

[13] Cynthia Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215, 2019.

[14] Cynthia Rudin and Berk Ustun. Optimized scoring systems: Toward trust in machine learning for healthcare and criminal justice. *Interfaces*, 48(5):449–466, 2018.

[15] Jacky Visser, John Lawrence, and Chris Reed. Reason-checking fake news. *Communications of the ACM*, 63(11):38–40, 2020.