Issues in Information Systems
Volume 25, Issue 2, pp. 144-153, 2024

DOI: https://doi.org/10.48009/2_iis_2024_111

Racing to build your AI skills: programming RL models with AWS DeepRacer

Natalya Bromall, *Robert Morris University, bromall@rmu.edu*
Steve Mancini, *Robert Morris University, mancinis@rmu.edu*
Jamie Pinchot, *Robert Morris University, pinchot@rmu.edu*
Peter Draus, *Robert Morris University, draus@rmu.edu*
Sushma Mishra, *Robert Morris University, mishra@rmu.edu*

Abstract

Amazon DeepRacer provides an innovative way to integrate complex AI/ML knowledge into a hands-on learning experience for students. Competitions provide a challenging pedagogical way to foster collaborative spirits in students. This study assesses the impact on the knowledge and confidence of students participating in groups in a DeepRacer competition at a small private school in Western Pennsylvania. Students were enrolled in teams and worked with faculty advisors to train their models and for a DeepRacer competition. A pre-posttest assessed their knowledge and confidence level in the AI/ML model training. The results suggest a significant improvement in the knowledge and confidence of the students about reinforcement learning. When used as a pedagogical tool, these competitions motivate students to apply theoretical knowledge to real-world situations and enhance their problem-solving abilities. Implications are drawn, and limitations are noted.

Keywords: artificial intelligence, deep learning, reinforcement learning, autonomous vehicles, AWS DeepRacer

Introduction

Artificial intelligence (AI) and machine learning (ML) advancements have started to impact various industries and can be used to simulate or even surpass human intelligence to accomplish tasks. Machine learning is now used to automate many tasks that are challenging for traditional rule-based programming (Tedre et al., 2021). Computing education efforts need to keep pace with the advancements in this area; therefore, teaching artificial intelligence (AI) and machine learning (ML) concepts has become critically important to train the next generation of science, technology, engineering, and math (STEM) workforce (Fomunyam, 2022).

Amazon Web Services (AWS) DeepRacer is a platform that utilizes a $1/18^{th}$ scale autonomous vehicle equipped with a fisheye lens camera that can be used to navigate a racetrack and can be trained to drive using machine learning. The training models are written in Python in the DeepRacer console (Figure 1, top right); after that, DeepRacer is trained through simulation in an online environment (Figure 1, top left) that allows for experimentation with reinforcement learning (Balaji et al., 2019). Reinforcement Learning (RL) is a type of machine learning that allows a device to interact with an environment and learn through trial and error to discover the most efficient behavior for an assigned task (Navarro et al., 2020). The trained

Issues in Information Systems
Volume 25, Issue 2, pp. 144-153, 2024

model can then be downloaded to the DeepRacer car to allow the car to drive on a physical track (Figure 1, top).

DeepRacer provides an opportunity to teach about autonomous systems, particularly reinforcement learning, through small-scale hardware that can be trained in simulation to provide these hands-on skills. Further, the use of competitions to allow students to compete to create the most efficient ML models for their DeepRacer cars has become a topic worthy of study. Competitions can be used to spark student interest and motivate students to learn about topics in AI and ML (Betz, 2022b).

An AWS DeepRacer Competition was held for students at a private University in Western Pennsylvania in April 2024. The competition format was developed by Amazon's AWS DeepRacer team, and the university planned and funded the event itself. Six faculty and IT staff members volunteered to serve as mentors for student teams. Fifteen student teams, each consisting of 2-3 students majoring in various areas of information systems, signed up to compete. The competition included two brackets, undergraduate and graduate, so there would be a winning and runner-up team from both brackets.

Each team was assigned a faculty mentor and attended a two-hour training workshop on AWS DeepRacer held by the Amazon team. Then, the teams had approximately six weeks to develop and train their models in the AWS online simulation tool to prepare for the competition. Students also worked with faculty mentors to periodically download their models to physical cars and test them on a physical track. Three of the original fifteen teams dropped out of the competition, leaving a total of twelve teams, six undergraduate and six graduate, to compete in the final competition on race day. In summary, the students went through the following activities:

- A 2-hour online training session with the AWS DeepRacer team
- Access to DeepRacer accounts and relevant documentation
- Weekly hands-on sessions with faculty mentors (more frequent sessions were given upon request)

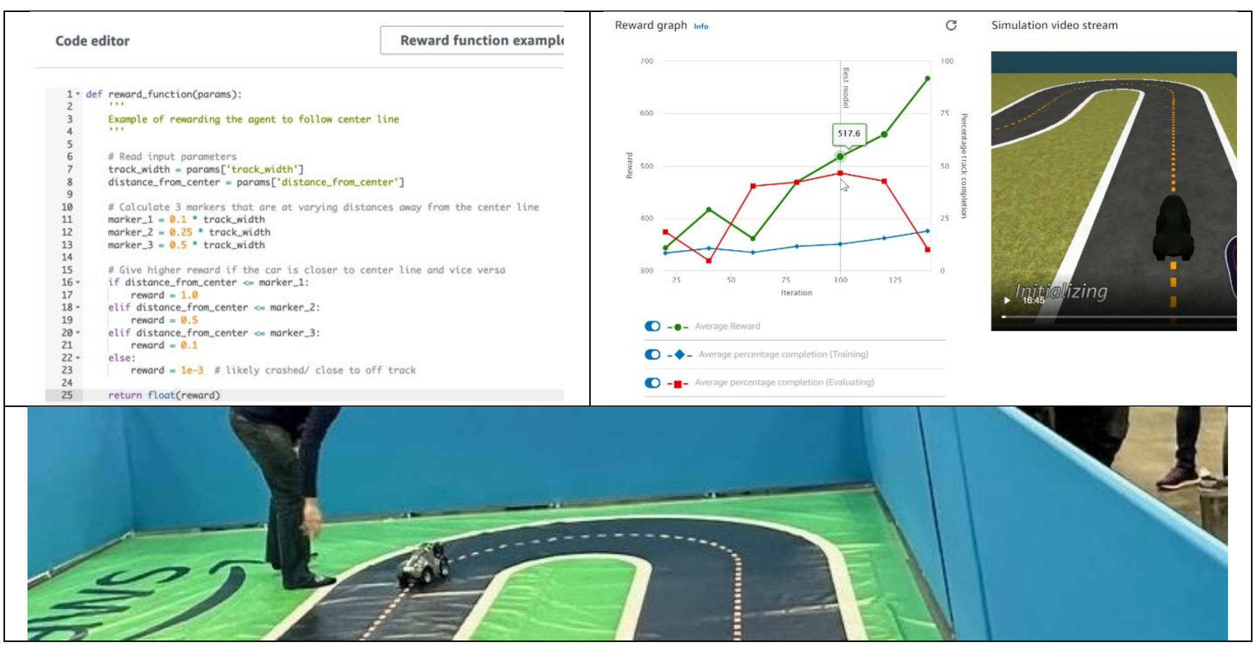

Figure 1: Running AWS DeepRacer in Virtual and Real Environment

Most research surrounding autonomous vehicle competitions focuses on the optimization of the vehicles for racing and challenges related to racing (Li, 2023; Betz, 2022a). Further research is needed to determine if autonomous vehicle racing competitions can enhance students' knowledge of AI and machine learning and their confidence level in approaching AI and machine learning. This research will directly survey students who participated in this AWS DeepRacer competition to investigate the following research questions:

RQ1.*What is the impact of participating in a Deep Racer event on participants' perceived knowledge level in ML and AI and their level of confidence?*

RQ2: *What knowledge of AI/ML is helpful in building RL models for autonomous vehicles?*

Literature Review

In recent years, improving machine learning and autonomous behavior has been necessitated by the race to employ more autonomous devices and systems. Sundberg & Holmström (2024) define ML as a "subfield of artificial intelligence (AI)" that focuses "on the development, application, and analysis of computer systems capable of learning from experience" (p. 56). As such, various research employed the usage of small-scale vehicles, with varying algorithms and approaches, to demonstrate the effectiveness of various machine learning algorithms (Liniger et al., 2014; Petryshyn et al., 2024; Yun et al., 2024).

Liniger et al. (2014) explored various ML optimization strategies for the purpose of racing an autonomous vehicle around a track. The intent was to demonstrate how employing different algorithms and approaches would result in a race car being able to successfully race around a track with speeds that exceeded typical norms for autonomous behavior (Liniger et al., 2014). The requirement to improve required algorithms to enable better and safer operation of autonomous systems has been translated into devices wherein larger numbers of individuals can have access to devices and virtual environments to test various theoretical approaches.

For example, Petryshyn et al. (2024) employed reinforcement learning utilizing Amazon DeepRacer cars and Amazon's virtual environment to demonstrate the viability of reducing collisions among autonomous vehicles. Their approach entailed utilizing Amazon's virtual environment to develop and improve differing algorithms and thus demonstrate marked improvement by modifying their training algorithms (Petryshyn et al., 2024).

Yun et al. (2024) employed a machine-learning approach to ascertain anomalous behavior in vehicle-to-everything (V2X) systems. Like Petryshyn et al. (2024), Yun et al. 's (2024) approach demonstrates the effectiveness of employing various ML algorithms to improve the overall functioning of autonomous systems.

However, the employment of test systems to introduce AI, ML, and reinforcement learning (RL) concepts to students assumes that knowledge among staff and students in said concepts is a given. As Garcia et al. (1993) noted, "Artificial intelligence (AI) is a multidisciplinary subject to which a multitude of interested practitioners in industry, government, and academia are contributing exciting empirical and theoretical developments" (p. 193). They highlighted the challenge of determining which material and how to filter it down to undergraduate level presents challenges (Garcia et al., 1993).

Stadelmann et al. (2021) stated, "we…are faced with the problem of having to prepare experienced students with diverse professional backgrounds for the reality of AI and ML applications and their implications" (p. 1). Coupling the seemingly disjointed nature of curriculum with challenges in technology, various approaches are being employed. Balaji et al. (2019) noted RL learning has been utilized in robotics for decades, but scalability challenges resulted in limitations. To leverage varying capabilities, Fisher (2014), a professor at Vanderbilt University, integrated Stanford University's three massive, open online courses (MOOCs), including a machine learning course, into his courses. In his approach, he required students to view external content, coupled with his course content, to improve engagement with students in his Artificial Intelligence class (Fisher, 2014).

Betz et al. (2022b) developed a hardware and software curriculum utilizing Amazon's DeepRacer vehicles and other devices to teach students about autonomous systems using hands-on learning. Students in the course participated in three competition races. However, race rankings only counted for 30% of the course grade, with 40% based on labs, 5% on a competition document focused on the student's approach to the competition, 5% on peer review, and 20% on a final project.

Four universities utilized the curriculum in spring 2022, including a total of 48 students, and those students were surveyed. More than 70% of the students indicated strong agreement when asked if using the hardware enhanced their understanding of the class subjects. Interestingly, although more than 80% of the students responded that the race competitions motivated them, only 50% indicated strong agreement that the competitions enhanced their learning.

Garcia et al.'s (1993) approach identified three key areas: symbolic and logic programming, introduction to Artificial Intelligence, and knowledge-based and expert systems. Based on these three areas of expertise, a curriculum was built around them which was covered in seventeen modules (Garcia et al., 1993). The seventeenth module, "Module 17: Vision, Robotics, and Learning", was the last module which would then allow students to employ various skills such as programming, software engineering fundamentals, symbolic computation, etc., to employ skills in a more hands-on manner effectively (Garcia et al., 1993).

Stadelmann et al. (2021) proposed the AI-Atlas didactic concept as "a coherent set of best practices for teaching Artificial Intelligence (AI) and Machine Learning (ML)" (p. 1). Recognizing a need to teach AI concepts better, Stadelmann et al. (2021) created this framework in reaction to a recognized need based on the "recent AI surge and corresponding demand for foundational teaching on the subject to a broad and diverse audience" (p. 1). They have three primary didactic principles, "Doing It Yourself," "Intrinsic Motivation," and the principle of "EEE" (Good Explanation, Enthusiasm, and Empathy) (Stadelmann et al. (2021). Stadelmann et al. (2021) approach is about engagement and keeping students motivated.

De Barros et al. (2023) proposed project-based learning (PBL) and agile techniques to teach ML and AI. Utilizing a study consisting of 30 students, De Barros et al. (2023) divided the students into 5 small groups. They intended to show how leveraging agile programming techniques would enable students to utilize a similar approach to employ ML techniques to solve the real-life problem of breast cancer prediction (De Barros et al., 2023). Again, as one can see, varying approaches are being leveraged better to teach concepts of ML and AI to students. In fact, Sundberg & Holmström (2024) highlight the fact that while ML is typically associated with "technical professions, such as computer science and engineering," even non-technical education programs are being impacted by ML and AI and thus require some education to be incorporated into their programs (p. 56).

AlSobeh and Woodward (2023) suggest that AI technologies such as NLP can enhance education effectiveness in various programs. In their approach, they employed no-code AI in a course for students with varying educational backgrounds (Sundberg & Holmström, 2024). Their results showed that "the no-code approach enabled students to engage in collective data work" (Sundberg & Holmström, 2024, p. 64). These key takeaways from the varying research indicate that with the recent surge of AI-driven systems, there is a clear need for different approaches to teaching ML and AI concepts as there is no longer any assurance that students will have the standard, technical background as was the previous case.

Methodology

To answer the research questions, we designed a survey to determine the students' perceived competence (knowledge and experience) in AI and specifically reinforcement learning models, and their confidence that they will successfully complete the competition. In determining the students' knowledge and experience, we outlined four domains needed to succeed in building and training an RL model for an autonomous vehicle:

1. General understanding of AI & ML concepts and techniques and differentiating between the two {gen}
2. General knowledge of ML models {gen mod}
3. Knowledge of RL models {RL mod}
4. Knowledge of cloud platforms – this is necessary to operate with the DeepRacer console, which is a part of AWS platform {cloud}

The first three items are often perceived on different levels by the users and developers, so we split the questions into two groups – {user} and {dev}, to make sure that the respondents distinguished between the two, when answering these questions.

As a result, the survey questions were mapped into the four domains, as shown in Table 1. All questions but one followed the Likert scale (1=Strongly Disagree, 5=Strongly Agree); therefore, higher scores indicated higher levels of knowledge, experience, and confidence, and lower scores indicated the opposite.

Table 1: Mapping between questions and knowledge domains

Domain	Question
{gen}, {user}	- I have a good understanding of AI and ML techniques as a user. - I have a good understanding of the differences between AI and ML techniques.
{gen}, {dev}	- I have good knowledge of AI and ML techniques as a developer.
{gen mod}, {user}	- I have utilized ML models before.
{gen mod}, {dev}	- I have developed ML models before.
{RL mod}, {user}	- I am familiar with various ML training models, such as neural networks and deep neural networks. - I am familiar with the concept of reinforcement learning.
{RL mod}, {dev}	- I am confident that I can write and implement an ML training model that uses reinforcement learning. - I am confident that with a certain amount of research, I can write and implement any ML training model.
{cloud}	- I am familiar with cloud computing ML/AI tools.

In addition, the survey included the overall pre-event confidence rating and a few demographic questions about the students' level (graduate or undergraduate), gender, age, general work experience, and AI/ML-related work experience. We chose the self-reported assessment in this case, assuming that this method will be efficient for collecting most information such as previous experience with the models and familiarity with the concept of RL. In addition, as we expected and as the results showed, most students reported zero or very low knowledge of AI/ML, RI, DeepRacer, and other relevant areas. The post-event survey also included the self-reported assessment. In addition, the formal assessment of the students' learning included the success of the model they created (time of the fastest lap + the ability to stay on track).

In the post-event survey, we removed the pre-event confidence question and the demographic question and added the open-ended questions to learn the students' feedback:

- ✓ What aspects of this event did you find especially helpful?
- ✓ Is there anything about this event that you would like to change, and why?
- ✓ Any other specific comments about the event and your learnings of the topic?

The pre-event survey was distributed to the participants immediately after the competition recruiting was over. After the competition, we asked the remaining participants to complete the post-event survey. In both surveys, respondents were asked their names so that we could match their responses before and after. We then ran a series of t-tests to determine if there were differences in perceived knowledge and confidence between (a) the students who completed the competition and those who dropped out and (b) between the same students before and after the competition.

Discussion

Of the 43 students signed up for the competition (23 undergraduate and 20 graduate), 11 students dropped out on various stages of preparation. Recruiting students to give feedback was the biggest challenge in this project: only 24 enrolled students completed the pre-event survey, giving sufficient answers. Only 11 students from the group who completed the competition agreed to answer the post-event survey, and only 9 students completed both pre-event and post-event surveys.

Table 2: Completion Levels of Students

Total students signed up	**43**	
Completed pre-event survey		**24**
Gave insufficient responses to pre-event survey		**3**
Did not complete pre-event survey		**16**
Signed up and completed the competition	**32**	
Completed post-event survey		**11**
Gave insufficient responses to post-event survey		**0**
Did not complete post-event survey		**21**
Dropped out	**11**	

The respondents reported high confidence in their understanding of Machine learning and Artificial Intelligence, with 83.3% reporting "Strongly Agree" (8.3% and "Agree" (75%) to the statement "I have a good understanding of AI and ML as a user. This drops to 29.2% to the statement "I have good knowledge of AI and ML as a developer, with 45% reporting "Neither agree nor Disagree." On a more practical level, almost half (45.5%) of the participants reported utilizing ML models before participating in the DeepRacer event. A smaller number (26.1%) reported having developed an ML model before the event.

On a theoretical level, the majority (82.6%) either agreed (69.6%) or "Strongly Agreed" (13%) with the statement, "I have a good understanding of the difference between AI and ML techniques. When looking at more specific ML training models, such as Neural networks and deep neural networks, almost half 47.8% reported being familiar with such models. Additionally, 83.3% reported being familiar with reinforcement learning. Slightly more than half (53.3%) reported being familiar with cloud computing ML/AI tools.

Impact on Knowledge and Confidence Level

Our initial plan was to run a paired-sample t-test for a pairwise comparison between the students' perceived knowledge before the race and the same students after the race; however, due to an insufficient number of post-event survey responses, we replaced it with an independent sample t-test with an independent-sample t-test. Figure 2 illustrates the pre and post mean comparison event answers. The simple mean comparison showed improvement in virtually every category except in understanding the difference between AI and ML while the t-Test showed that only a few of these differences were significant.

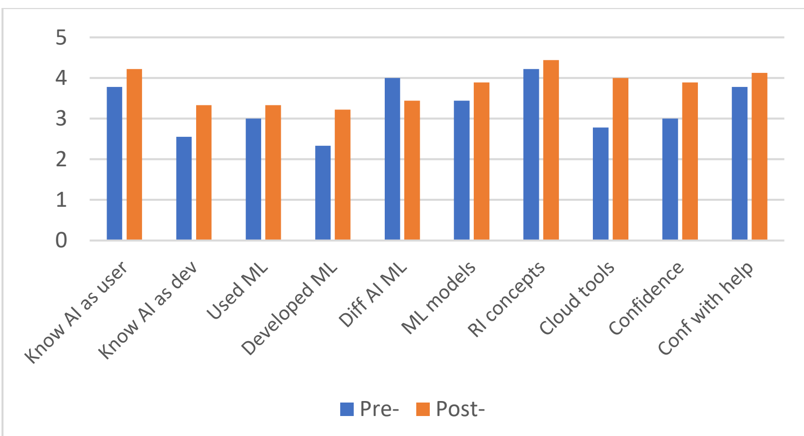

Figure 2: Mean comparison between pre- and post-event answers

The students admitted having a better theoretical understanding of AI and ML after the race and a better understanding of cloud platforms for building and training the models. Evidently, several respondents still did not believe they developed a model in the process. Still, as expected there is a substantial increase in the number of students who believed they did. While the respondents ranked higher in their abilities as developers, this increase was significant only at a .9 confidence level. With the same confidence level, the post-event data showed an increase in the students' belief that they became familiar with the concept of reinforcement learning. Surprisingly, the difference between AI and ML became generally less clear after the completion of the race – this was the only category with an insignificantly lower score on the post-event survey.

As shown in Table 3, both confidence scores significantly improved from pre- to post-event. The students felt more confident about writing and implementing an RL model in a machine learning case (3.91/5.00). Their confidence in writing and implementing an RL model with a certain amount of research was measured even higher (4.20/5.00). The students' confidence level highly correlated (.809**) with the total level of their knowledge and preparedness, which was computed as the sum of their responses for the first set of questions related to their knowledge and experience. In other words, as expected, the students with higher levels of perceived knowledge and experience had higher confidence.

Table 3: Significant findings in students' knowledge and confidence levels

Knowledge and Confidence Category	Pre-	Post-
I have a good understanding of AI and ML techniques as a user.	3.92	4.36**
I have good knowledge of AI and ML techniques as a developer.	3.00	3.64*
I have developed ML models before.	2.61	3.45**
I am familiar with the concept of reinforcement learning (RL).	4.04	4.36*
I am familiar with cloud computing ML/AI tools.	3.35	4.09**
I am confident that I can write and implement an ML training model that uses RL.	3.35	3.91*
I am confident that with a certain amount of research, I can write and implement any ML training model.	3.83	4.20*

** - significant with .9 confidence*
*** - significant with .95 confidence*

Knowledge of AI/ML is helpful in building RL models for autonomous vehicles

We compared the group of students who dropped out of the competition and the ones who went all the way through in an attempt to answer the second research question. This comparison would not, of course, give a complete and conclusive answer. Still, it would help us identify any substantial differences in knowledge domains between the two groups that may have caused the students to drop out of the event.

The comparison did not show any conclusive results. Both groups evaluated their perceived knowledge and experience similarly, with two significant exceptions. The group of students who completed the competition initially admitted a higher level of experience utilizing ML models before (3.39 vs. 2.67) and a lower level of familiarity with cloud ML tools (3.32 vs. 3.83). At the same time, we had a few interesting observations: the group of students who completed the competition admitted a significantly lower confidence level than the group that dropped out.

This study has implications for research and practice. The study shows a clear connection between student learning and enhancing their knowledge and confidence using DeepRacer competition as a pedagogical tool. More studies could be conducted to investigate the impact on knowledge and confidence in various subdomains of AI and ML. Practitioners could benefit from hiring students with hands-on experience in training models using RL/ML techniques.

Limitations and Future Research

First and foremost, the number of participating students who filled in the surveys was a limiting factor. While simple mean comparison was possible, with one of the two groups being so small, t-tests often do not give sufficient accuracy. The AWS DeepRacer team gave a 2-hour workshop to all students enrolled in the competition at the very beginning of the program. This training is aimed at understanding the cloud environment for building DeepRacer models, training them, and transferring them to the actual cars. It does not include coding, RL, or other topics in AI/ML. The second limitation was sending the pre-event test after that initial training, which may have changed their perception. In future events, we intend to give the pre-event survey to the students before the first training takes place.

Conclusion

The study shows that Amazon DeepRacer could be a useful pedagogical tool for teaching AI and ML concepts through student competitions. This experiential teaching method proved effective and engaging for students, boosting their knowledge and confidence in using AI/ML tools and techniques and enhancing their knowledge about the topic. The complexities of reinforcement learning allowed the participants to gain insight into model training, hyperparameter tuning, and understanding a real-world application of machine learning algorithms. This study also shows collaborative teamwork is a good strategy for teaching AI/ML professionals' complex concepts.

References

AlSobeh, A. and Woodward, B. (2023). AI as a partner in learning: a novel student-in-the-loop framework for enhancing student engagement and outcomes in higher education. Conference Proceedings for ACM Special Interest Group for Information Technology Education, Marietta, GA, USA.

Balaji, B., Mallya, S., Genc, S., Gupta, S., Dirac, L., Khare, V., Roy, G., Sun, T., Tao, Y., Townsend, B., Calleja, E., Muralidhara, S., & Karuppasamy, D. (2019). Deepracer: Educational autonomous racing platform for experimentation with sim2real reinforcement learning. *arXiv preprint arXiv:1911.01562*.

Betz, J., Zheng, H., Liniger, A., Rosolia, U., Karle, P., Behl, M., Krovi, V., & Mangharam, R. (2022a). Autonomous vehicles on the edge: A survey on autonomous vehicle racing. *IEEE Open Journal of Intelligent Transportation Systems, 3*, 458-488.

Betz, J., Zheng, H., Zang, Z., Sauerbeck, F., Walas, K., Dimitrov, V., ... & Mangharam, R. (2022b). Teaching autonomous systems hands-on: Leveraging modular small-scale hardware in the robotics classroom. *arXiv preprint arXiv:2209.11181*.

De Barros, Vitor Augusto Menten, Paiva, H. M., & Hayashi, V. T. (2023). Using PBL and agile to teach artificial intelligence to undergraduate computing students. *IEEE Access*, 11, 1-1. https://doi.org/10.1109/ACCESS.2023.3298294

Fisher, D. H. (2014). Leveraging AI Teaching in the Cloud for AI Teaching on Campus. *AI Magazine*, 35(3), 98-100. https://doi.org/10.1609/aimag.v35i3.2546

Fomunyam, K. G. (2022). Machine learning and stem education: challenges and possibilities. *International Journal of Difference Equations; Vol. 17, Issue 2*.

Garcia, O. N., Perez, R. A., Silverman, B. G., Austin, H. S., Baum, R. F., Brady, L. H., Cameron, R. S., Castaneda, S. E., Chen, J., Dey, P. P., DiCristina, G. M., Elmaghraby, A. S., Foster, R., Freeman, C., Kirch, M. R., Lawrence, A. W., Manesh, A., Manickam, S., Ramamoorthy, C. V., . . . Vasquez, R. (1993). On teaching AI and expert systems courses. *IEEE Transactions on Education*, 36(1), 193-197. https://doi.org/10.1109/13.204845

Li, B., Fang, Y., Ma, S., Wang, H., Wang, Y., Li, X.,Zhang, T., Bian, X., & Wang, F. Y. (2023). Toward fair and thrilling autonomous racing: Governance rules and performance metrics for autonomous one. *IEEE Transactions on Intelligent Vehicles*.

Liniger, A., Domahidi, A., & Morari, M. (2015). Optimization-based autonomous racing of 1:43 scale RC cars. *Optimal Control Applications & Methods, 36*(5), 628-647. https://doi.org/10.1002/oca.2123

Navarro, A., Genc, S., Rangarajan, P., Khalil, R., Goberville, N., Rojas, J. F., & Asher, Z. (2020). *Using Reinforcement Learning and Simulation to Develop Autonomous Vehicle Control Strategies* (No. 2020-01-0737). SAE Technical Paper.

Petryshyn, B., Postupaiev, S., Soufiane, B. B., & Ostreika, A. (2024). Deep Reinforcement Learning for Autonomous Driving in Amazon Web Services DeepRacer. *Information, 15*(2), 113. https://doi.org/10.3390/info15020113

Stadelmann, T., Keuzenkamp, J., Grabner, H., & Würsch, C. (2021). The AI-Atlas: Didactics for Teaching AI and Machine Learning On-Site, Online, and Hybrid. *Education Sciences*, 11(7), 318. https://doi.org/10.3390/educsci11070318

Sundberg, L., & Holmström, J. (2024). Teaching tip: Using no-code AI to teach machine learning in higher education. *Journal of Information Systems Education, 35*(1), 56. https://doi.org/10.62273/CYPL2902

Tedre, M., Toivonen, T., Kahila, J., Vartiainen, H., Valtonen, T., Jormanainen, I., & Pears, A. (2021). Teaching machine learning in K–12 classroom: Pedagogical and technological trajectories for artificial intelligence education. *IEEE access*, *9*, 110558-110572.

Yun, K., Yun, H., Lee, S., Oh, J., Kim, M., Lim, M., Lee, J., Kim, C., Seo, J., & Choi, J. (2024). A Study on Machine Learning-Enhanced Roadside Unit-Based Detection of Abnormal Driving in Autonomous Vehicles. *Electronics, 13*(2), 288. https://doi.org/10.3390/electronics13020288