Introducing Deep Learning to Undergraduate Engineering Majors

Lin Zhang
University of Central Arkansas

Abstract

Artificial intelligence (AI) is one of the fastest growing technologies. The undergraduate students, especially those majoring in engineering, are demanding more encounters with AI. Deep learning (as one of the most popular and fundamental subcategories of the AI) is introduced to the undergraduates through a 4-course education. The students will acquire intuitions on the concepts, the theoretical/mathematical backgrounds and practical skills and knowledge. This article introduces the curriculum for Engineering Physics majors at the University of Central Arkansas. A deep learning based students' project of autonomous racing vehicle is showcased in this article.

Keywords

Artificial intelligence, deep learning, curriculum design.

Introduction

Artificial intelligence (AI) is reshaping our world especially through the blast of deep learning. Deep learning models are getting more and more powerful to impact our daily lives (such as ChatGPT) and to make novel scientific findings (such as AlphaFold) [1], [2]. The college students in this generation are living in another time of technology revolution powered by AI. As the undergraduates (especially those majoring in engineering) showing increasing interest and the job market showing the increasing preference over AI, offering courses of AI can benefit both parties [3-5].

At University of Central Arkansas (UCA), we have been developing the curriculum for undergraduates majoring in Physics Engineering. We offer a bundle of four courses to gradually introduce AI to our students. Our students will be recommended to take the following courses in order. Generally, our students' AI journey begins at ENGR 4421: Robotics 2. They will integrate a neural network based object detection function into a robotics project. The students can better understand the neural network and deep learning by taking ENGR 3321: Introduction to Deep Learning for Robotics. The fundamental concepts and background mathematics will be introduced. At the same time, our students can practice their knowledge and skills by enrolling in ENGR 4311 Senior Design 1 and ENGR 4312 Senior Design 2. Their goal is to build an autonomous racing vehicle using the AI technologies they've learned.

The contents of this article are organized as follows. In the **Curriculum Development** section, we will introduce more details of the curriculum design. In the **Student Outcomes** section,

demonstrations of student projects will be given. In the **Summary and Future Work** section, the existing curriculum and upcoming upgrades will be summarized and discussed.

Curriculum Development

Our goal is to help students to live their academic and career lives with a handy tool – AI. However, AI is such a huge and complex subject that we are not able to cover every aspect. Our plan is to start from a simple yet widely spread technology: deep learning. As shown in Fig. 1, students will gradually master the power of deep learning through a 4-course journey.

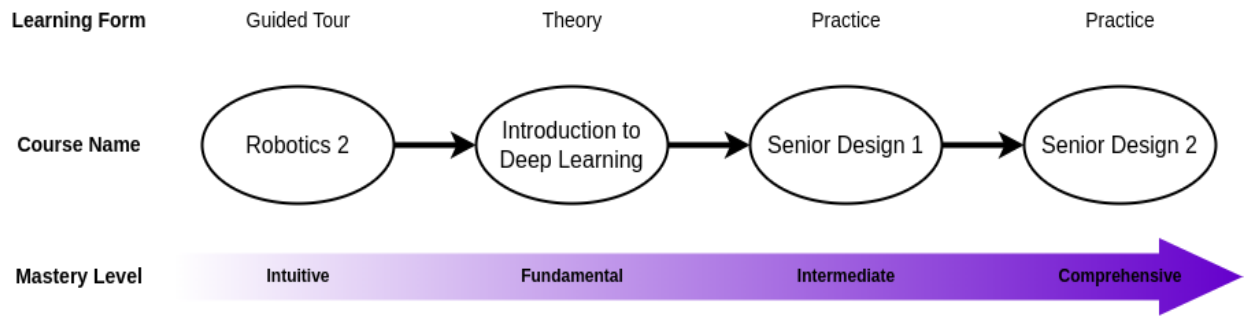

Figure 1 Curriculum of AI education at UCA

Deep Learning Intuition

In the course of *Robotics 2*, the students will spend two weeks to integrate an AI powered function to their own robots. Object detection is a computer vision task for locating instances of objects in images [6]. It is one of the most notable beneficiaries of deep learning, yet, the cost to implement it can be little. The goal is to offer the students the intuition and to stimulate their interests by encouraging them to utilize such deep learning backed AI technology. No advanced theory nor mathematical formulae are needed to be taught. A workshop will be organized to help the students to learn the basic concepts and the usage of an open-sourced application programming interface (API).

Deep Learning Fundamentals

In the course of *Introduction to Deep Learning for Robotics*, the students will spend a whole semester exploring the under-the-hood concepts and theoretical fundamentals of deep learning. This course focuses on vision oriented neural networks. The goal is to encourage the student to both think and put hands on the more fundamental contents of deep learning. The low-level mechanisms to drive and optimize neural networks will be taught. A variety of neural network architectures will be demonstrated. PyTorch, one of the most popular Python libraries to implement neural networks, will be introduced [7]. The students will upgrade their cognition via lectures and coding labs.

Deep Learning Practices

The students will have opportunities to practice their deep learning skills and knowledge in Senior Design 1 and Senior Design 2. Through a two-semester project, they need to prototype and upgrade autonomous driving vehicles with neural networks backbones. The goal is to encourage the students to apply their knowledge and get ready for the post-graduation lives.

Several workshops and brainstorming will be hosted to guide the students with key techniques. Students' hands-on experience will be maximized in this project.

Students Outcomes

We have been offering Senior Design 1 & 2 with deep learning based autonomous driving projects for two academic years. The rest of the AI contents were offered in the past academic year. These educational resources were welcome by our students. Most of them acknowledged that the experience benefited their academic and post-graduation lives. As the following shows, our students also made notable outcomes during their journey of learning AI.

1. *Robotic Object Detection* – In *Robotics 2* class, a team of three students investigated tensorflow's object detection API back in 2022 [8]. Their goal was to enable the in-house built mobile robot to track specific objects (e.g. human, cup, ball, etc.). For interested readers, the open-sourced project can be accessed using the following link: https://github.com/willward20/the-bot-bros. Due to the time limitation and lack of deeper understanding, the product was rudimentary. However, this first-touch buried the seeds of AI into these students' minds. They later developed a remarkable deep learning based software to drive an autonomous racing vehicle.
2. *Autonomous Driving Vehicle* – In *Senior Design 1 and 2,* our students implemented their knowledge of AI and successfully built an autonomous driving software from scratch.
    a. DonkeyCar: In the first year, a team of two students explored an open-source project: DonkeyCar (https://github.com/autorope/donkeycar). They i) found a solid solution to modify a toy RC car to an autonomous driving platform; ii) figured out the usage of the API provided by DonkeyCar; iii) located and documented the neural network architecture that powered the autonomous driving system. This team eventually participated in the National Robotics Challenge and won first place in the section of Autonomous Vehicle Challenge. Although they were relying on the DonkeyCar software, their experience and documentation established a solid foundation for the next year students to build their own software.
    b. WHAM Module: In the second year, a team of four students took the legacy from the previous team and upgraded the autonomous driving platform with i) new hardware configuration, ii) completely student developed software, iii) more detailed documentation and an open-source project repository named as WHAM. The project can be accessed via Github: https://github.com/willward20/WHAM. Interested readers may use this repository to start their own autonomous driving project. The WHAM module is not as universal as the DonkeyCar software, but it is way easier to use or to develop. Refer to Fig. 2 for the difference between using the DonkeyCar software and using/developing WHAM software. The DonkeyCar project provides universal modification solutions towards a variety of ground vehicles. This results in a huge amount of configurations and documentations. Our students found it was annoying to locate the most related information that applied to our platform. Accordingly, manipulating the core management script was cumbersome because of the overwhelming configurations. The WHAM team,

however, discarded the generality and brought up a flatter design. Since our students only consider the RC car modified platform, tons of options that work with other platforms can be ignored. In addition, the core scripts are not wrapped deeply in the software. Users or developers only need to interact with three Python scripts to get the autonomous vehicle fully functional.

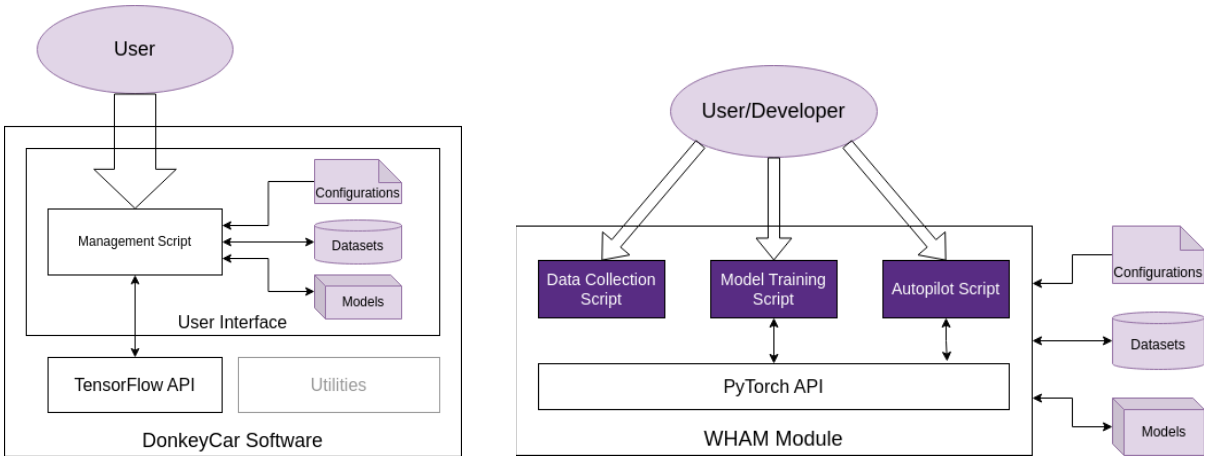

Figure 2 Student developed autonomous driving software

Summary and Future Work

At University of Central Arkansas, we offer a pragmatic AI learning experience. The students will have the intuition via learning the usage of an object detection software. They can dive deeper to learn the foundations of deep learning by taking an introductory course. The senior design project provides a practical opportunity for them to apply and master the skills. Whether a student needs a tryout or a more comprehensive understanding, the curriculum introduced in this article can satisfy his/her demands.

The current curriculum is still under development. The following aspects are considered to be improved in the near future.

1. *PyTorch Ecosystem* – As PyTorch is winning the deep learning programming market, we will use it for all the coding education and practices. For example, the future object detection project will be shifted to the PyTorch based API.
2. *Introduction to Transformer* – Since the "transformer" architecture is dominating the deep learning tasks, it is necessary to introduce more theoretical fundamentals of such a powerful instrument to our students [9].
3. *Discard the Bulky DonkeyCar* – Our students started a much more lightweight autonomous driving software. The more generic DonkeyCar project is no longer a good reference. Instead, the future senior design team will inherit WHAM's heritage and continue the development from there.

We will appreciate it if the readers are inspired or have suggestions.

References

[1] OpenAI, “GPT-4 Technical Report,” 2023, arXiv: 2303.08774.
[2] J. Jumper, R. Evans, A. Pritzel, T. Green, M. Figurnov, O. Ronneberger, K. Tunyasuvunakool, R. Bates, A. Žídek, A. Potapenko, and A. Bridgland, “Highly accurate protein structure prediction with AlphaFold,” Nature, vol. 596, no. 7873, pp. 583–589, Aug. 2021.
[3] X. Gong, L. Zhao, R. Tang, Y. Guo, X. Liu, Y. Tang, J. He, W. Shi, X. Niu, X. Wang, and FY. Wang, “AI educational system for primary and secondary schools,” in 2019 ASEE Annual Conference & Exposition., Jun. 2019.
[4] A. ARIS, A. Uluagac, L.P. Rondon, D. Ortiz, M. Ross, and M. Finlayson, “Integrating Artificial Intelligence into Cybersecurity Curriculum: New Perspectives,” in 2022 ASEE Annual Conference & Exposition., Aug. 2022.
[5] C. Tucker, K.S. Jackson, and J.J. Park, “Exploring the future of engineering education: Perspectives from a workshop on artificial intelligence and the future of stem and societies,” in 2020 ASEE Virtual Annual Conference Content Access, Jun. 2020.
[6] Z. Zou, K. Chen, Z. Shi, Y. Guo and J. Ye, "Object Detection in 20 Years: A Survey," in Proceedings of the IEEE, vol. 111, no. 3, pp. 257-276, Mar. 2023
[7] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, and A. Desmaison, “Pytorch: An imperative style, high-performance deep learning library,” Advances in neural information processing systems, 32. 2019.
[8] J. Huang, V. Rathod, C. Sun, M. Zhu, A. Korattikara, A. Fathi, I. Fischer, Z. Wojna, Y. Song, S. Guadarrama, and K. Murphy, “Speed/accuracy trade-offs for modern convolutional object detectors,” in Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 7310-7311, Jul. 2017.
[9] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A.N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” Advances in neural information processing systems, 30, 2017.

Lin Zhang

Dr. Lin Zhang received the B.S. degree in automation from the Harbin Institute of Technology, Heilongjiang, China, in 2007, and the Ph.D. degree in engineering from New Mexico State University, Las Cruces, NM, USA, in 2016. He is currently an assistant professor with the Department of Physics and Astronomy, University of Central Arkansas. His major research interests are deep reinforcement learning and robotics.