


# Diving into deep learning

 Check for updates


## Understanding Deep Learning

By Simon J. D. Prince  
The MIT Press: 2023.  
544 pp. \$90.00

The field of artificial intelligence (AI) has experienced a surge in developments over the past years, propelled by breakthroughs in deep learning with neural networks. This has revolutionized many aspects of society. However, the speed at which AI is advancing highlights the need for textbooks that provide essential resources for educating young researchers and professionals with the latest methodologies and best practices in deep learning, guiding them in exploring uncharted territories and harnessing the full potential of AI.

I am deeply impressed by the new book *Understanding Deep Learning*<sup>1</sup>, written by Simon J.D. Prince and published by the MIT Press (and freely available online in PDF format). There already exist various well-known books of this type, such as *Deep Learning* by Ian Goodfellow, Yoshua Bengio and Aaron Courville<sup>2</sup>, also from MIT Press, which is often considered the ‘bible’ of deep learning. *Hands-on Machine Learning with Scikit-Learn, Keras, and TensorFlow*, by Aurélien Géron<sup>3</sup>, is lauded for its practical approach, implementing deep learning models with popular libraries. *Neural Networks and Deep Learning* by Michael Nielsen<sup>4</sup> is known for its clear, intuitive explanations for beginners. *Deep Learning for Computer Vision* by Rajalingappaa Shanmugamani<sup>5</sup> focuses on computer vision, and thus has a narrower scope compared to the broad coverage of the other books. Most recent is *Deep Learning: Foundations and Concepts* by Christopher Bishop and Hugh Bishop<sup>6</sup>, which is also excellent and is similar in style to Prince’s book. *Understanding Deep Learning* complements the existing literature with its highly accessible explanations—accompanied by 68 Python notebooks immediately usable in the Google Colab environment—while also making a

deep dive into theory. The book’s brilliance lies in its ability to demystify complex ideas without diminishing their sophistication, facilitated by clear explanations and good examples. Also, it goes beyond the basics, challenging readers with advanced topics while providing the foundation for a deep understanding. The book is therefore particularly suitable for an academically oriented audience.

Prince’s book not only covers the latest deep learning models, including diffusion models and large language models, but also incorporates two chapters on topics rarely covered in previous books: a chapter on AI ethics (co-written by Travis LaCroix of Dalhousie University) and one on theoretical issues. The AI ethics chapter focuses on the potential of deep learning for gains and harms. In particular, it raises concerns about misuse, bias, model opaqueness and privacy leakage. This chapter encourages researchers and practitioners to consider the social and ethical dimensions to their work right from the outset. The chapter “Why does deep learning work?” targets the theoretical aspects of deep learning. It raises fundamental questions, such as why deep networks are easy to train, why they generalize well from training to unseen data and why they show a surprising effectiveness, highlighting where further research could lead to significant breakthroughs.

*Understanding Deep Learning* stands out as a comprehensive and contemporary guide to deep learning. In my view, it is an instant classic, and it is a testament to the author’s caliber that he is able to communicate complex ideas in such an engaging and friendly manner. Prince has meticulously selected the most pivotal concepts, ensuring that readers are presented with a rich concentration of vital information, both insightful and straightforward to comprehend. The book begins with the fundamentals of machine learning and delves progressively into increasingly more sophisticated models in 21 chapters. Each idea is initially explained in easy language, and this is followed by precise definitions and complemented by visual illustrations.

The author’s dedication to optimizing the usefulness of this timely book is evident from his establishment of a website related

to it: <http://udlbook.com>. This platform not only demonstrates a commitment to continuous improvement but also offers a dynamic space for readers to find updates, additional resources and interactive content. By creating this website, the author enhances the learning experience and fosters a community of learners and practitioners. This initiative reflects an understanding of the rapidly evolving nature of deep learning and an intention to continuously refresh the content.

Although *Understanding Deep Learning* is, in my opinion, the best textbook on deep learning, it is not without its limitations. One potential weakness lies in its attempt to satisfy a broad audience, which results in a treatment of certain advanced topics that could be perceived as overly simplified. For instance, the chapter on graph networks does not include the spectral approach. Furthermore, despite the inclusion of cutting-edge subjects such as transformers and diffusion models, the rapid pace of innovation in AI might outstrip the book’s updates. For instance, Poisson flow generative models and other advanced alternatives have outperformed diffusion models (for example, see ref. 7). The Python notebooks, although beneficial for hands-on learning, may also pose a barrier for readers less familiar with programming or those using different software environments. They could become out of date when ChatGPT and other models are gradually enhanced to automate programming and bring new type of hands-on experience, perhaps soon. In my own teaching experience, a ChatGPT app can interact with students in new ways, including demo software development. Additionally, the minimal mathematical prerequisites, although they lower the barrier to entry, might not provide sufficient depth for those interested in the theoretical underpinnings of deep learning. This mismatch highlights the challenge of balancing accessibility with the complexity inherent in deep learning, a goal that any single resource must struggle to fully achieve.

Despite the above perceived weaknesses, this book is more than just a textbook; it is a guide to navigating the future of technology with wisdom, insight and foresight.

## Reviewed by Ge Wang

Department of Biomedical Engineering, School of Engineering, Biomedical Imaging Center, Center for Computational Innovations, Center for Biotechnology & Interdisciplinary Studies, Rensselaer Polytechnic Institute, Troy, NY, USA.

✉ e-mail: [wangg6@rpi.edu](mailto:wangg6@rpi.edu)

Published online: 10 May 2024

## References

1. Prince, J. D. S. *Understanding Deep Learning* (MIT Press, 2024); <https://udlbook.github.io/udlbook/>
2. Goodfellow, I., Bengio, Y. & Courville, A. *Deep Learning* (MIT Press, 2016); <https://www.deeplearningbook.org/>
3. Géron, A. *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems* edn 2 (O'Reilly, 2019).
4. Nielsen, M. *Neural Networks and Deep Learning* 2nd edn (Springer, 2013); <https://www.goodreads.com/book/show/24582662-neural-networks-and-deep-learning>
5. Shanmugamani, R. *Deep Learning for Computer Vision: Expert Techniques to Train Advanced Neural Networks Using TensorFlow and Keras* (Packt Publishing, 2018).
6. Bishop, C. M. & Bishop, H. *Deep Learning: Foundations and Concepts* 1st edn (Springer, 2024).
7. Xu, Y. et al. *Adv. Neural Inf. Process. Syst.* **35**, 16782–16795 (2022).

## Competing interests

The author declares no competing interests.