Journal of Science Innovation for Sustainable Development
https://ph01.tci-thaijo.org/index.php/JSISD/index

BOOK REVIEW

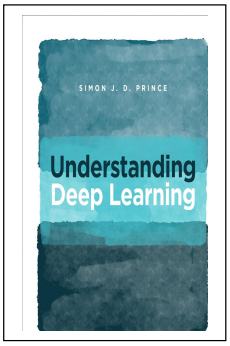

Book Title:	Understanding Deep Learning
Edition Statement:	1st Edition
Author:	Simon J.D. Prince
Format:	544 pages
Publisher:	MIT Press
Language:	English
Bestsellers Rank:	Amazon
ISBN:	978-0-262-04864-4

Published Online: 26 June, 2024

ISSN: **3057-0778**

Dr.Chawalsak Phetchanchai

Faculty of Science and Technology, Suandusit University

(chawalsak@hotmail.com)

Corresponding Author

Received date: 21 June, 2024
Revised date: 23 June, 2024
Accepted date: 25 June, 2024

The book is meticulously organized into chapters that progressively build on each other, beginning with an introduction to fundamental concepts, especially in supervised learning (Chapters 1-2). It starts with 1D linear regression, shallow neural networks with a single hidden layer, and progresses to deep neural networks with multiple hidden layers. Both types of networks use ReLU activation functions to describe piecewise linear mappings from input to output (Chapters 3-4). The importance of loss functions or cost functions is explored to describe the mismatch between model predictions and their corresponding ground-truth outputs (Chapter 5). Iterative optimization algorithms are introduced to help neural networks find the parameters that minimize the loss in mapping inputs to outputs. Performance measurements of training models are also covered (Chapters 6-8). Various types of generalization methods are discussed to examine how well a model can handle new data and make correct predictions after training (Chapter 9). Sophisticated supervised model architectures such as convolutional neural networks, residual networks, transformers, and graph neural networks, which incorporate parameter sharing and allow parallel computational paths, are discussed (Chapters 10-13). Unsupervised learning with generative adversarial networks (GANs) is introduced as a method for passing latent variables through deep networks to create new samples. Normalizing flows, which can learn a probability model by transforming a simple distribution into a more complex one using a deep network, are also discussed (Chapter 16). GANs are further explored for their ability to create samples that are statistically indistinguishable from the training data. Variational autoencoders, based on a solid probabilistic foundation, are introduced (Chapter 17). Diffusion models, which consist of an encoder and a decoder with stochastic mappings, are covered in Chapter 19. The book concludes with discussions on the practical workings and ethical considerations of deep learning (Chapters 20-21). This book provides comprehensive coverage of the topics mentioned, making it suitable for those interested in studying Deep Learning. Additionally, the book offers knowledge in a sequential manner, allowing readers to build their foundation from basic to advanced levels. It is therefore appropriate for use in undergraduate or higher education courses, as well as for self-study.

Each section is thoroughly explained with detailed mathematics, accompanied by illustrations that make understanding easier. However, the book has relatively few real-world application examples. Nonetheless, each chapter includes well-designed and comprehensive questions or problems. Therefore, this book is an excellent choice for anyone interested in studying Deep Learning.

Reference: Prince, S. J. (2023). Understanding Deep Learning. MIT Press.

Reference: Prince, S. J. (2023). Understanding Deep Learning. MIT Press.