

# Reference Materials for Machine Learning and AI

---

## **Deep Learning by Ian Goodfellow, Yoshua Bengio, and Aaron Courville**

This foundational book covers deep learning principles in a rigorous yet accessible manner. It includes theory, architectures, optimization techniques, and practical considerations, making it essential for both researchers and practitioners.

Access: <https://www.deeplearningbook.org/>

## **Pattern Recognition and Machine Learning by Christopher M. Bishop**

Bishop's book is a comprehensive text on probabilistic approaches to machine learning, covering topics such as graphical models, kernel methods, and neural networks with a strong mathematical foundation.

Access: <https://www.microsoft.com/en-us/research/people/cmbishop/prml-book/>

## **The Elements of Statistical Learning by Hastie, Tibshirani, and Friedman**

This classic text provides an in-depth treatment of statistical learning theory, including regression, classification, boosting, and model selection. It is widely used in both academic and applied contexts.

Access: <https://web.stanford.edu/~hastie/ElemStatLearn/>

## **Machine Learning: A Probabilistic Perspective by Kevin P. Murphy**

Murphy's text presents machine learning through the lens of probabilistic modeling. It includes Bayesian inference, graphical models, and deep learning, blending theory with applications and Python code.

Access: <https://probml.github.io/pml-book/>

## **CS231n: Convolutional Neural Networks for Visual Recognition (Stanford University)**

An outstanding course from Stanford covering the foundations and advances in deep learning for computer vision. Includes lecture videos, slides, notes, and assignments.

Access: <https://cs231n.stanford.edu/>

## **CS229: Machine Learning (Stanford University)**

One of the most popular introductory courses in machine learning, taught at Stanford. Covers supervised, unsupervised, and reinforcement learning, with theory and practical insights.

*Access: <https://cs229.stanford.edu/>*

### **MIT 6.S191: Introduction to Deep Learning**

A hands-on MIT course that covers core concepts of deep learning, neural networks, optimization, and real-world applications. Includes labs and tutorials implemented in TensorFlow and PyTorch.

*Access: <http://introtodeeplearning.com/>*

### **Distill.pub: Visual Essays on Machine Learning**

Distill publishes high-quality interactive and visual essays on machine learning topics such as attention, backpropagation, and optimization. It is valued for clarity and pedagogy.

*Access: <https://distill.pub/>*

