Deep Learning with Python

Artificial intelligence, machine learning, and deep learning

First, we need to define clearly what we're talking about when we mention AI. What are artificial intelligence, machine learning, and deep learning (see figure 1.1)? How do they relate to each other?

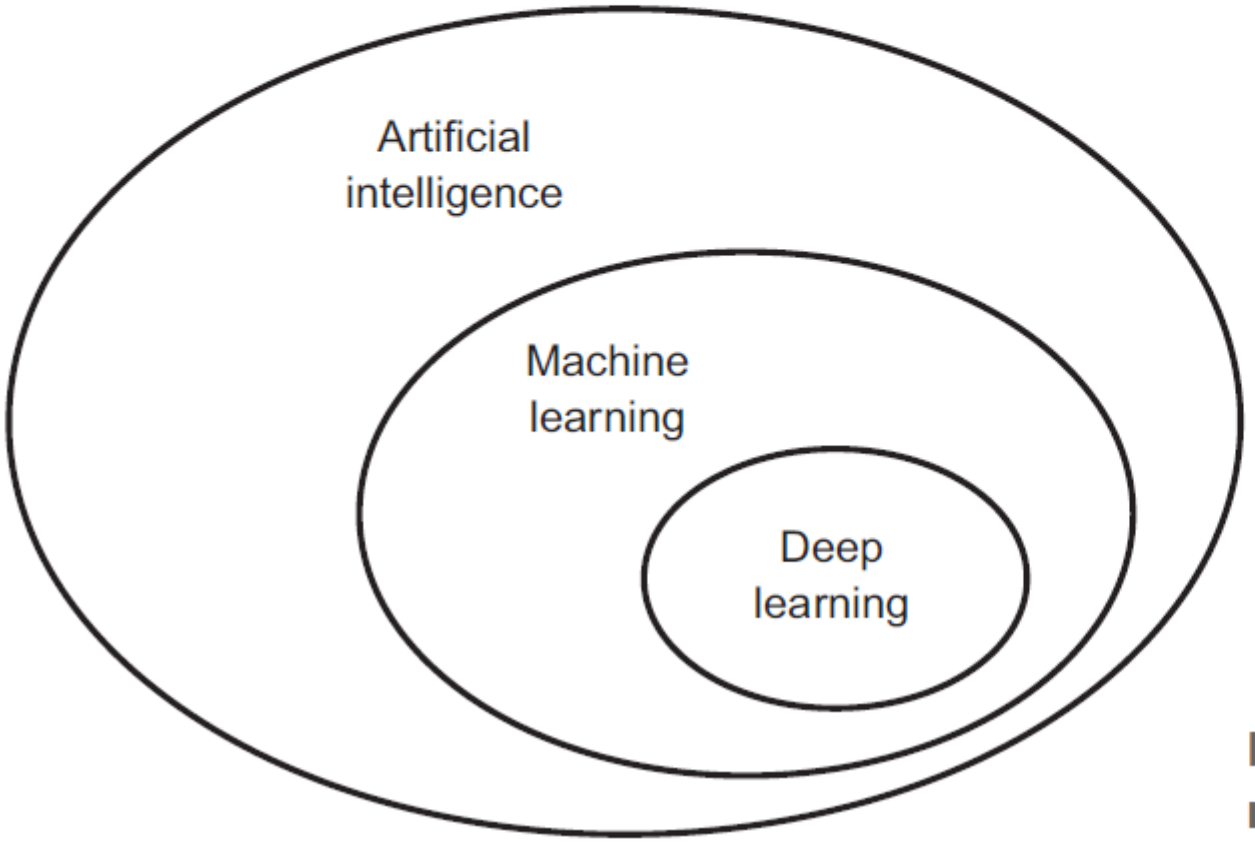

Figure 1.1 Artificial intelligence, machine learning, and deep learning

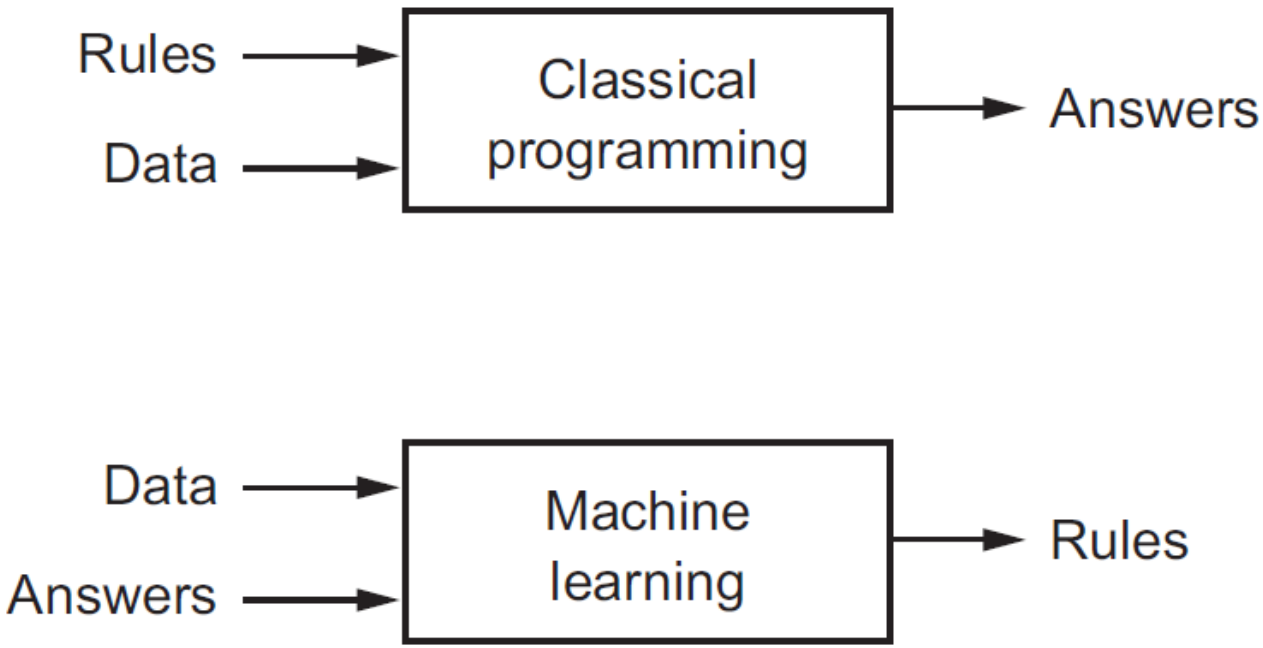

Figure 1.2 Machine learning: a new programming paradigm

Learning representations from data

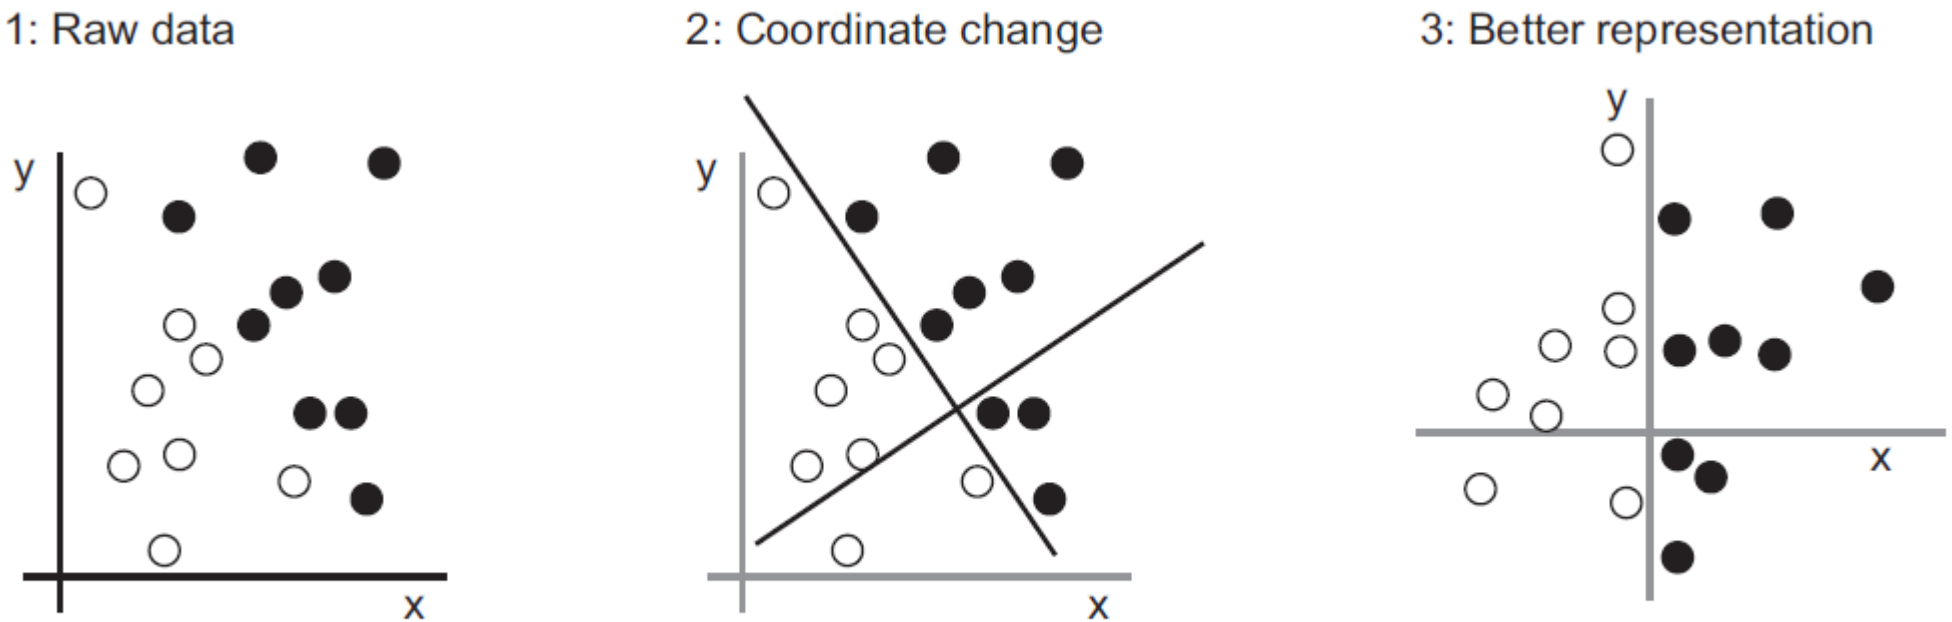

Figure 1.4 Coordinate change

All machine-learning algorithms consist of automatically finding such transformations that turn data into more-useful representations for a given task.

The "deep" in deep learning

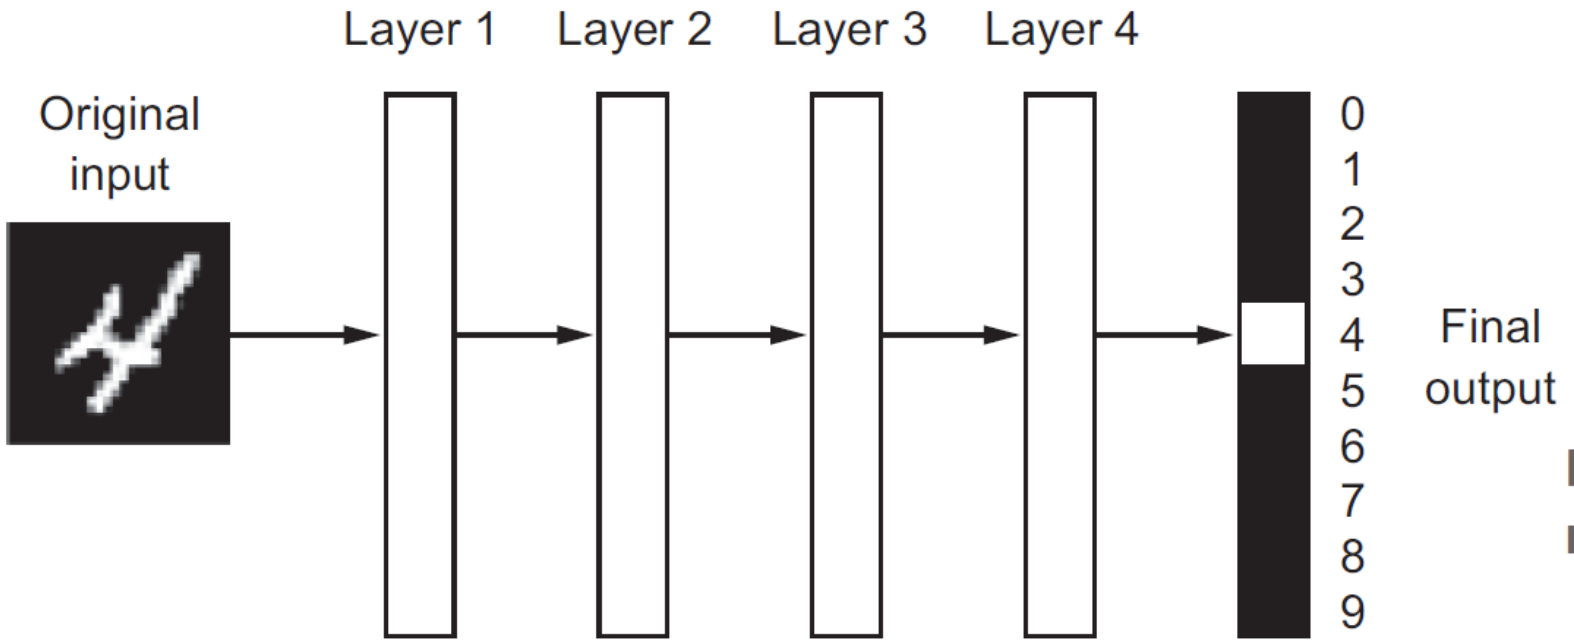

Figure 1.5 A deep neural network for digit classification

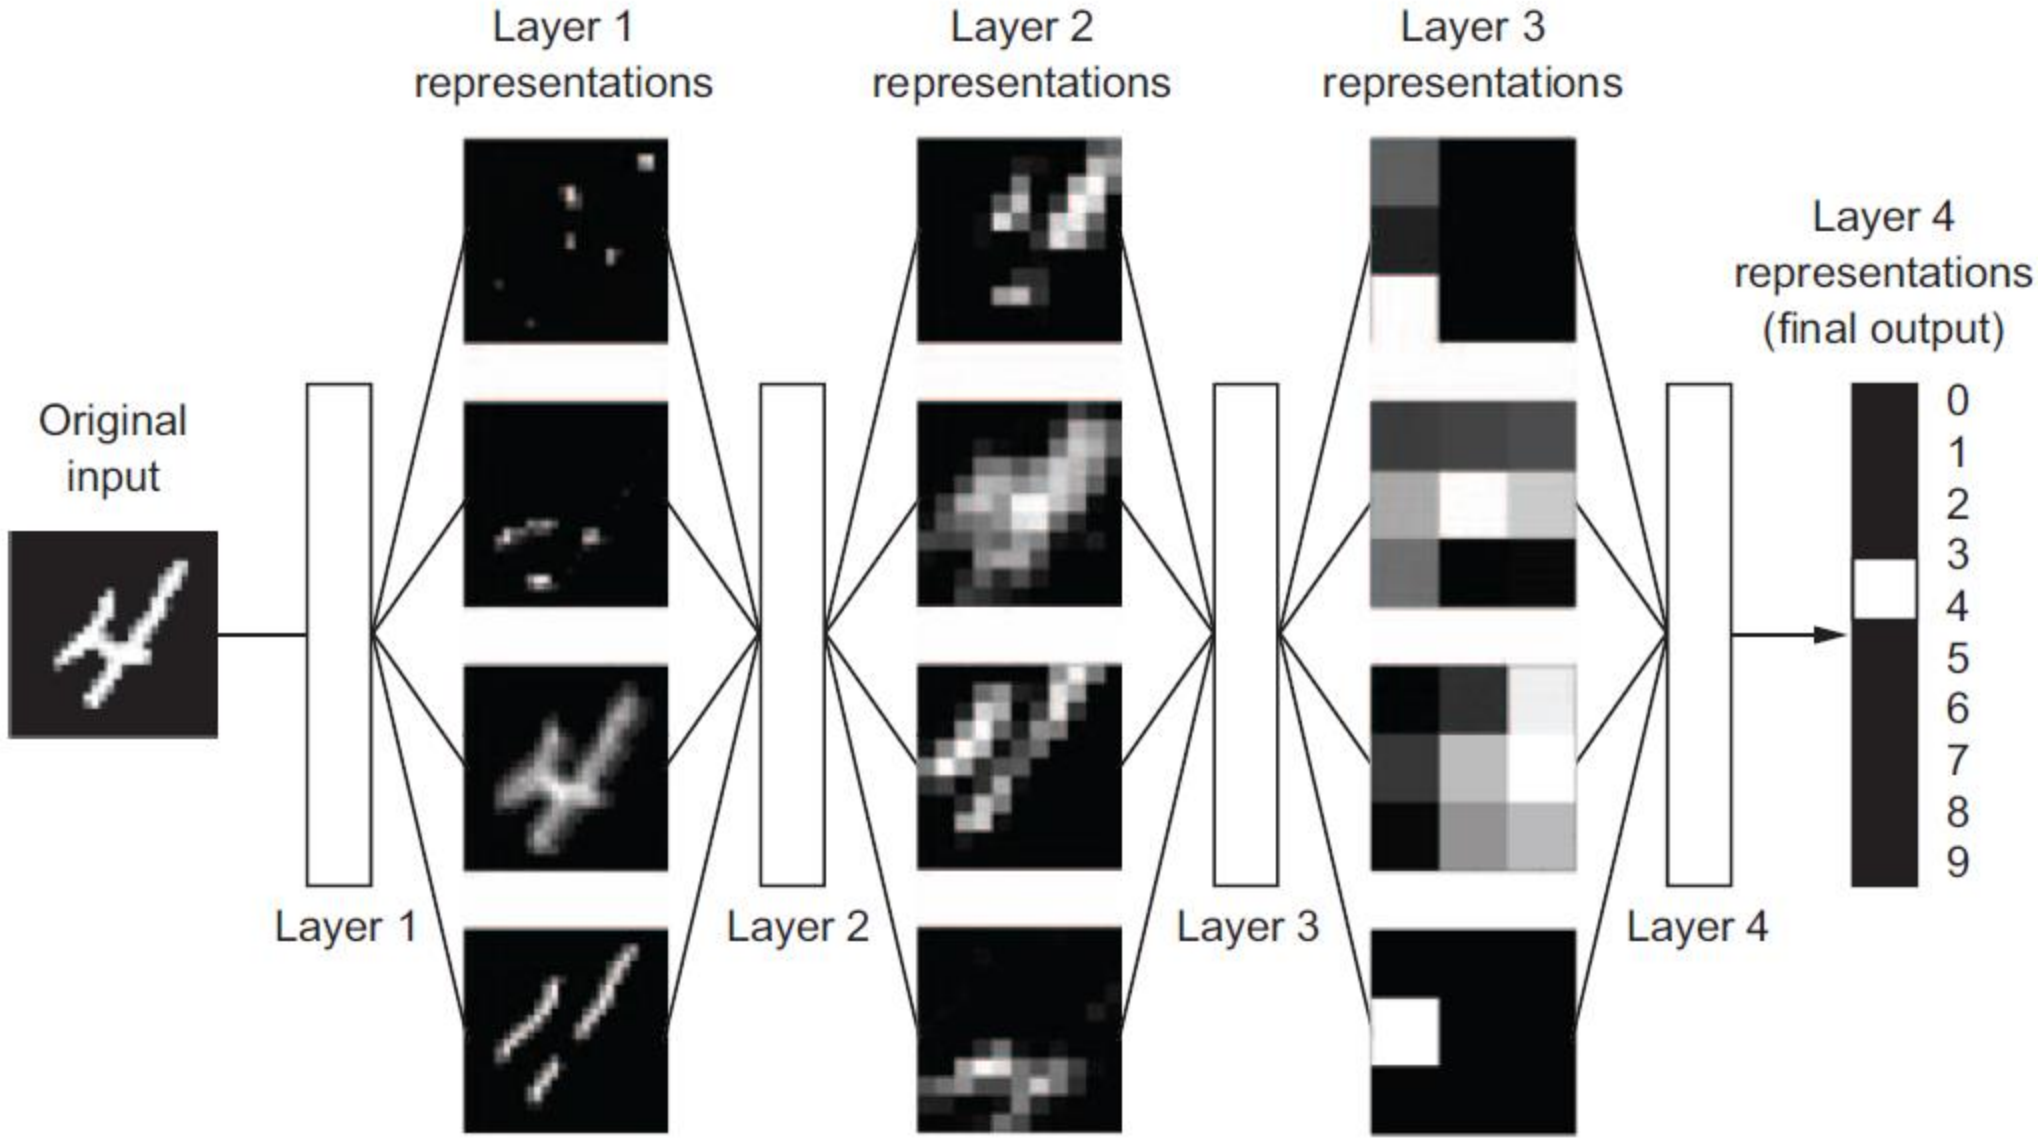

Figure 1.6 Deep representations learned by a digit-classification model

Understanding how deep learning works, in three figures

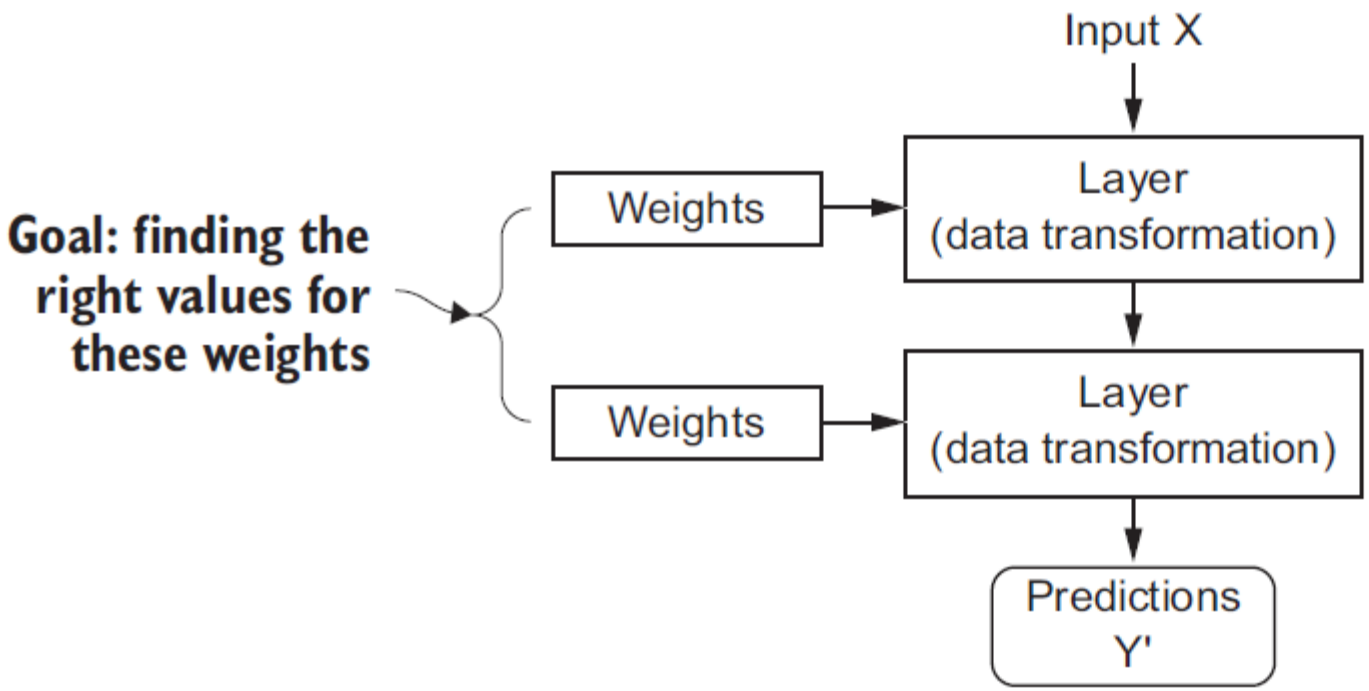

Figure 1.7 A neural network is parameterized by its weights.

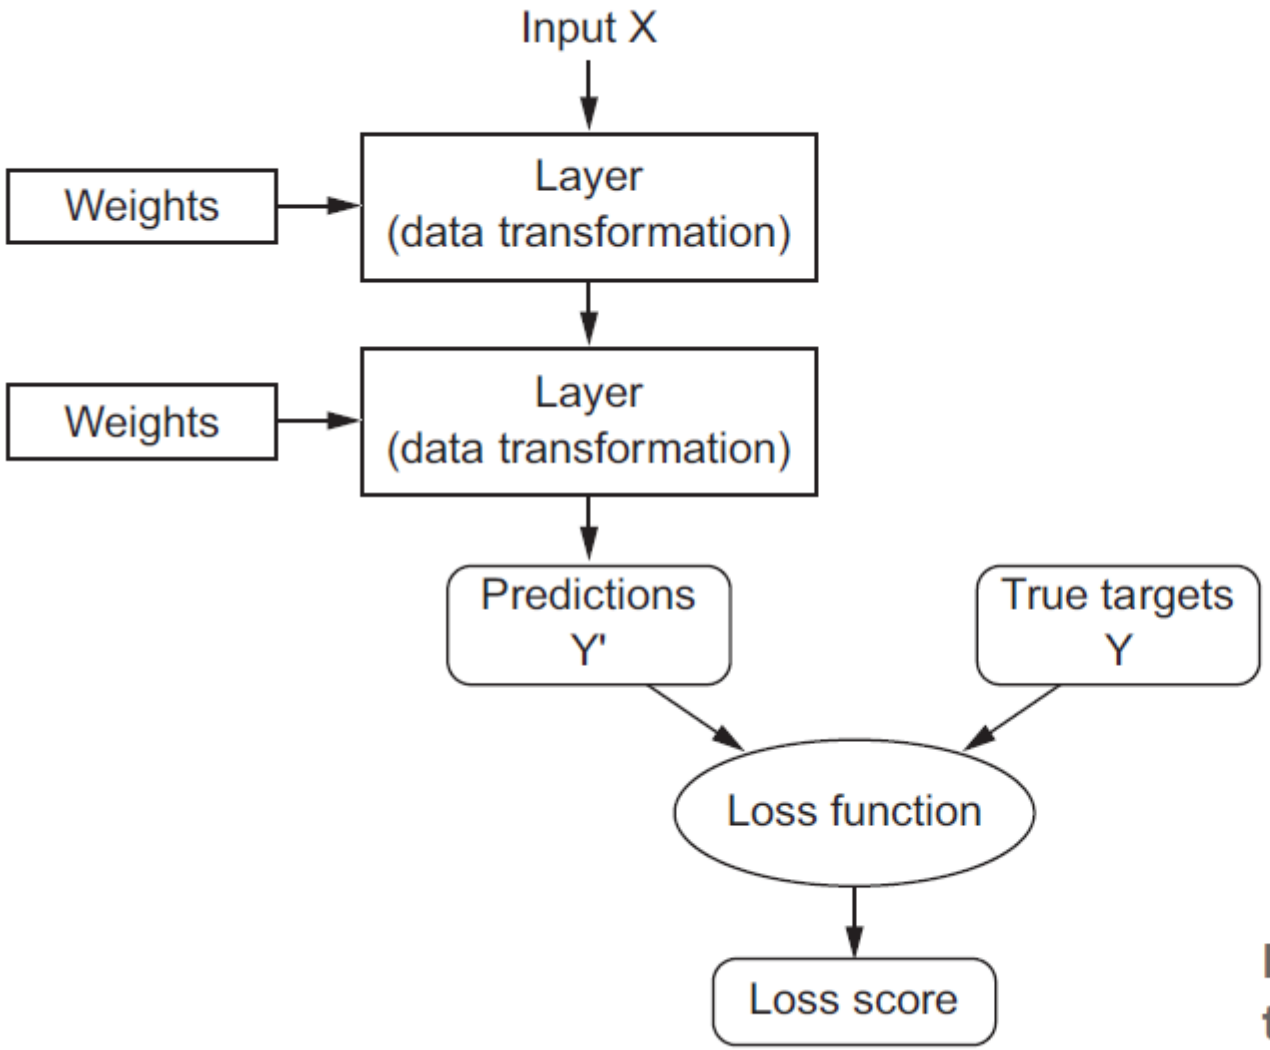

Figure 1.8 A loss function measures the quality of the network's output.

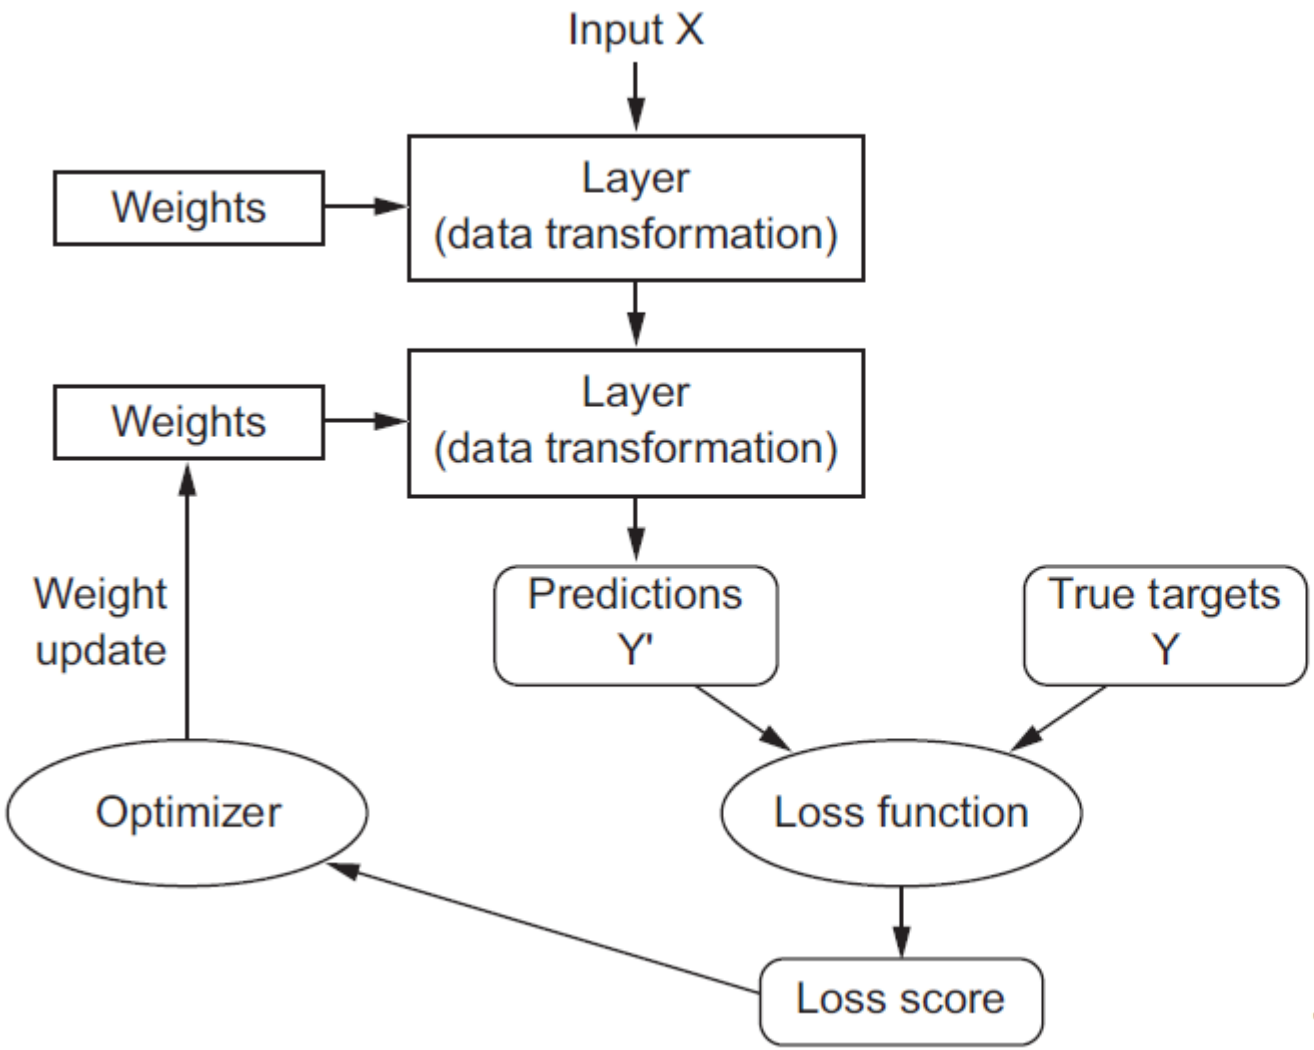

Figure 1.9 The loss score is used as a feedback signal to adjust the weights.

Before deep learning:
a brief history of machine learning

Probabilistic modeling

Early neural networks

Kernel methods

Decision trees, random forests, and gradient boosting machines

Around 2010, although neural networks were almost completely shunned by the scientific community at large, a number of people still working on neural networks started to make important breakthroughs: the groups of Geoffrey Hinton at the University of Toronto, Yoshua Bengio at the University of Montreal, Yann LeCun at New York University, and IDSIA (Dan Ciresan) in Switzerland.

In 2011, 74.3%

In 2012, a team led by Alex Krizhevsky and advised by Geoffrey Hinton was able to achieve a top-five accuracy of 83.6%—a significant breakthrough.

In 2015, 96.4%

ImageNet challenge was notoriously difficult at the time, consisting of classifying high resolution color images into **1,000 different categories** after training **on 1.4 million images.**

A great way to get a sense of the current landscape of machine-learning algorithms and tools is to look at machine-learning competitions on Kaggle.

In 2016 and 2017, Kaggle was dominated by two approaches: gradient boosting machines and deep learning.

XGBoost library

Keras library