Journal of Machine Learning

A Mathematical Framework for Learning Probability Distributions

Hongkang Yang

DOI: 10.4208/jml.221202, J. Mach. Learn., 1 (2022), pp. 373-431.

Communicated by: Weinan E

Category: Theory

Summary for general readers:

The modeling of probability distributions is an important branch of machine learning. It became popular in recent years thanks to the success of deep generative models in difficult tasks such as image synthesis and text conversation. Nevertheless, we still lack a theoretical understanding of the good performance of distribution learning models. One mystery is the following paradox: it is generally inevitable that the model suffers from memorization (converges to the empirical distribution of the training samples) and thus becomes useless, and yet in practice the trained model can generate new samples or estimate the probability density over unseen samples. Meanwhile, the existing models are so diverse that it has become overwhelming for practitioners and researchers to get a clear picture of this fast-growing subject. This paper provides a mathematical framework that unifies all the well-known models, so that they can be systemically derived based on simple principles. This framework enables our analysis of the theoretical mysteries of distribution learning, in particular, the paradox between memorization and generalization. It is established that the model during training enjoys implicit regularization, so that it approximates the hidden target distribution before eventually turning towards the empirical distribution. With early stopping, the generalization error escapes from the curse of dimensionality and thus the model generalizes well.

Sponsored by the Center for Machine Learning Research, Peking University & AI for Science Institute, Beijing.

Global Science Press https://www.global-sci.org/jml