

Lecture 10:

Sasha Rakhlin

Description

We introduce bagging and boosting algorithms. We discuss the bias-variance tradeoff. The gradient descent view of boosting is introduced and a bound on the performance of Adaboost is proved.

Suggested Reading

- Yoav Freund and Robert E. Schapire. [A short introduction to boosting.](#) Journal of Japanese Society for Artificial Intelligence, 14(5):771-780, September, 1999
 - Other boosting review articles [here](#).
-