

Minimax-optimal and Locally-adaptive Online Nonparametric Regression

We study adversarial online nonparametric regression with general convex losses and propose a parameter-free learning algorithm that achieves minimax optimal rates. Our approach leverages chaining trees to compete against Hölder functions and establishes optimal regret bounds. While competing with nonparametric function classes can be challenging, they often exhibit local patterns - such as local Hölder continuity - that online algorithms can exploit. Without prior knowledge, our method dynamically tracks and adapts to different Hölder profiles by pruning a core chaining tree structure, aligning itself with local smoothness variations. This leads to the first computationally efficient algorithm with locally adaptive optimal rates for online regression in an adversarial setting. Finally, we discuss how these notions could be extended to a boosting framework, offering promising directions for future research.