

# The SMART Approach to Instance-Optimal Online Learning

**Siddhartha Banerjee**

*Cornell University*

SBANERJEE@CORNELL.EDU

**Alankrita Bhatt**

*California Institute of Technology*

ABHATT@CALTECH.EDU

**Christina Lee Yu**

*Cornell University*

CLEEYU@CORNELL.EDU

**Editors:** Shipra Agrawal and Aaron Roth

## Abstract

We devise an online learning algorithm – titled *Switching via Monotone Adapted Regret Traces* (SMART) – that adapts to the data and achieves regret that is *instance optimal*, i.e., simultaneously competitive on every input sequence compared to the performance of the follow-the-leader (FTL) policy and the worst case guarantee of any other input policy  $\text{ALG}_{\text{WC}}$ . We show that the regret of the SMART policy on *any* input sequence is within a multiplicative factor  $e/(e-1) \approx 1.58$  of the smaller of: 1) the regret obtained by FTL on the sequence, and 2) the upper bound on regret guaranteed by the given worst-case policy. This implies a strictly stronger guarantee than typical ‘best-of-both-worlds’ bounds as the guarantee holds for every input sequence regardless of how it is generated. SMART is simple to implement as it begins by playing FTL and switches at most once during the time horizon to  $\text{ALG}_{\text{WC}}$ . Our approach and results follow from a reduction of instance optimal online learning to competitive analysis for the ski-rental problem. We complement our competitive ratio upper bounds with a fundamental lower bound showing that over all input sequences, no algorithm can get better than a 1.43-fraction of the minimum regret achieved by FTL and the minimax-optimal policy. We present a modification of SMART that combines FTL with a “small-loss” algorithm to achieve instance optimality between the regret of FTL and the small loss regret bound.

Extended abstract. Full version appears at [arXiv:2402.17720](https://arxiv.org/abs/2402.17720), v2.