



ELLIS talk

Tilted Losses in Machine Learning: Theory and Applications

Virginia Smith

Carnegie Mellon University

Host: Marco Mondelli

Exponential tilting is a technique commonly used to create parametric distribution shifts. Despite its prevalence in related fields, tilting has not seen widespread use in machine learning. In this talk, I discuss a simple extension to ERM---tilted empirical risk minimization (TERM)---which uses tilting to flexibly tune the impact of individual losses. I make connections between TERM and related approaches, such as Value-at-Risk, Conditional Value-at-Risk, and distributionally robust optimization (DRO), and present batch and stochastic first-order optimization methods for solving TERM at scale. Finally, I show that this baseline can be used for a multitude of applications in machine learning, such as enforcing fairness between subgroups, mitigating the effect of outliers, and handling class imbalance---delivering state-of-the-art performance relative to more complex, bespoke solutions for these problems.

Thursday, February 10, 2022 05:00pm - 06:00pm

Zoom Link:

<https://istaustria.zoom.us/j/64653423567?pwd=eEYvUkjWY2VwdlBWMEFZT1BYTGdhUT09>

Meeting ID: 646 5342 3567 Passcode: 469530



This invitation is valid as a ticket for the ISTA Shuttle from and to Heiligenstadt Station.

Please find a schedule of the ISTA Shuttle on our webpage:

<https://ista.ac.at/en/campus/how-to-get-here/> The ISTA Shuttle bus is marked ISTA Shuttle (#142) and has the Institute Logo printed on the side.