



Contribution ID: 37

Type: **Invited talk**

Large deviations for linear regression

Friday, 25 September 2026 09:00 (30 minutes)

Linear regression is one of the simplest and most widely used tools to learn patterns from data: it fits a set of coefficients so that a linear combination of predictors best matches observed responses. The quality of the fit is measured by the residual sum of squares, the total squared mismatch between predictions and data, whose minimum defines the training loss. For Gaussian data, the residual sum of squares is a random variable whose statistics is considered for large number P of predictors and N of observations, keeping their ratio $r = P/N$ fixed. Although the typical training loss and the interpolation transition at $r = 1$ are well known, much less is known about rare sample-to-sample fluctuations of the optimal loss. Using the zero-temperature replica method, we compute analytically the full distribution (including large deviation tails) of the minimal regularized training loss for arbitrary penalty function, including Ridge and Lasso. Extensive numerical simulations are in excellent agreement with our theory and clearly show a pronounced deviation from the Gaussian regime of typical fluctuations in the tails.

Primary authors: BARTOLUCCI, Silvia (Department of Computer Science, University College London); Prof. CACCIOLI, Fabio; Dr CARAVELLI, Francesco; Prof. VIVO, Pierpaolo

Presenter: BARTOLUCCI, Silvia (Department of Computer Science, University College London)

Session Classification: Session 7