

Abstract Submitted
for the MAR16 Meeting of
The American Physical Society

Statistical Physics of High Dimensional Inference¹ MADHU ADVANI, SURYA GANGULI, Stanford University — To model modern large-scale datasets, we need efficient algorithms to infer a set of P unknown model parameters from N noisy measurements. What are fundamental limits on the accuracy of parameter inference, given limited measurements, signal-to-noise ratios, prior information, and computational tractability requirements? How can we combine prior information with measurements to achieve these limits? Classical statistics gives incisive answers to these questions as the measurement density $\alpha = \frac{N}{P} \rightarrow \infty$. However, modern high-dimensional inference problems, in fields ranging from bio-informatics to economics, occur at finite α . We formulate and analyze high-dimensional inference analytically by applying the replica and cavity methods of statistical physics where data serves as quenched disorder and inferred parameters play the role of thermal degrees of freedom. Our analysis reveals that widely cherished Bayesian inference algorithms such as maximum likelihood and maximum a posteriori are suboptimal in the modern setting, and yields new tractable, optimal algorithms to replace them as well as novel bounds on the achievable accuracy of a large class of high-dimensional inference algorithms.

¹Thanks to Stanford Graduate Fellowship and Mind Brain Computation IGERT grant for support

Madhu Advani
Stanford University

Date submitted: 06 Nov 2015

Electronic form version 1.4