Thermodynamic Overfitting: Limits on Complexity in Thermodynamic Learning

Authors

Alexander B. Boyd[1], James P. Crutchfield[2], Mile Gu[3], Felix Binder[1]

[1] Trinity College, Ireland; [2] UC Davis, USA; [3] Nanyang Technological University, Singapore

Abstract

Recent results show that maximizing work production corresponds to thermodynamic learning, guiding information engines toward predictive hidden Markov models. A consequence of this is the thermodynamic Principle of Requisite Complexity: information engines must requisitely match the complexity of their information fuel in order to operate at maximum efficiency. There is a direct parallel between the maximum-work principle in thermodynamic learning and the maximum-likelihood principle that guides machine learning. However, we show that recklessly maximizing work leads to overfitting, which has dire energetic consequences. We demonstrate that irreversible entropy production diverges when engines are unconstrained in their pursuit of energetic advantage during training. We see that the danger of divergent dissipation and overfitting persists for ever longer strings of training data as the complexity of the information engine's model increases. Thus, to counterpose the thermodynamic Principle of Requisite Complexity, we demonstrate an energetic cost to complexity.

References/Recommended reading

Boyd, A. B., Mandal, D., and Crutchfield, J. P.: Thermodynamics of Modularity: Structural Costs Beyond the Landauer Bound, Physical Review X, 8, 031036, 10.1103/PhysRevX.8.031036, 2018.