

Abstract Submitted
for the MAR16 Meeting of
The American Physical Society

Compression and regularization with the information bottleneck

DJ STROUSE, Princeton University, DAVID SCHWAB, Northwestern University
— Compression fundamentally involves a decision about what is relevant and what is not. The information bottleneck (IB) by Tishby, Pereira, and Bialek formalized this notion as an information-theoretic optimization problem and proposed an optimal tradeoff between throwing away as many bits as possible, and selectively keeping those that are most important. The IB has also recently been proposed as a theory of sensory gating and predictive computation in the retina by Palmer et al. Here, we introduce an alternative formulation of the IB, the deterministic information bottleneck (DIB), that we argue better captures the notion of compression, including that done by the brain. As suggested by its name, the solution to the DIB problem is a deterministic encoder, as opposed to the stochastic encoder that is optimal under the IB. We then compare the IB and DIB on synthetic data, showing that the IB and DIB perform similarly in terms of the IB cost function, but that the DIB vastly outperforms the IB in terms of the DIB cost function. Our derivation of the DIB also provides a family of models which interpolates between the DIB and IB by adding noise of a particular form. We discuss the role of this noise as a regularizer.

DJ Strouse
Princeton University

Date submitted: 06 Nov 2015

Electronic form version 1.4