

Information Theoretic Co-Training

David McAllester
TTI-Chicago

January 5, 2018

Abstract

This paper introduces an information theoretic co-training objective for unsupervised learning. We consider the problem of predicting the future. Rather than predict future sensations (image pixels or sound waves) we predict “hypotheses” to be confirmed by future sensations. More formally, we assume a population distribution on pairs (x, y) where we can think of x as a past sensation and y as a future sensation. We train both a predictor model $P_\Phi(z|x)$ and a confirmation model $P_\Psi(z|y)$ where we view z as hypotheses (when predicted) or facts (when confirmed). For a population distribution on pairs (x, y) we focus on the problem of measuring the mutual information between x and y . By the data processing inequality this mutual information is at least as large as the mutual information between x and z under the distribution on triples (x, z, y) defined by the confirmation model $P_\Psi(z|y)$. The information theoretic training objective for $P_\Phi(z|x)$ and $P_\Psi(z|y)$ can be viewed as a form of co-training where we want the prediction from x to match the confirmation from y . Initial experiments applying information theoretic co-training to unsupervised learning of phonetics are presented.

Withdrawn