Learning Bounds for Domain Adaptation [M27]

John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, Jenn Wortman
University of Pennsylvania

Setup: Plentiful but biased source data and relevant but scarce target data

Example task: identifying genes in DNA sequences

Domains: different organisms
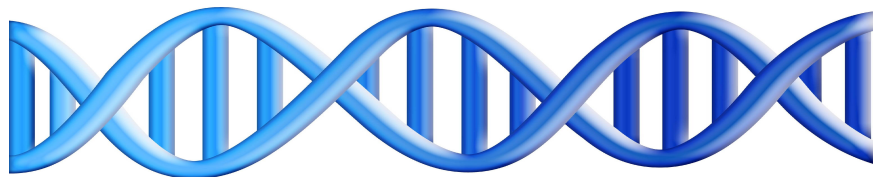
source

target

We develop a uniform convergence bound that is computable from finite labeled & unlabeled samples

Minimize convex combinations of source & target risk

Generalizes standard single-domain bounds as well as previous domain adaptation bounds