
T-testing: A Linear-time Adaptive Nonparametric Technique*

Zoltán Szabó[†] (CMAP, École Polytechnique)

Abstract

In my talk, I am going to focus on the two-sample testing problem: we are given two sets of observations, and our goal is to check whether the two sets are statistically indistinguishable (this may be thought of as a generalization of a t-test). We propose to choose features underlying the test in an adaptive way so as to maximize the distinguishability of the distributions by optimizing a lower bound on the test power. The constructed nonparametric t-test is based on a novel representation of probability distributions by analytic functions, and runs in linear time with respect to the number of samples. The efficiency of the method is illustrated in natural language processing (distinguishing articles from two categories) and computer vision (differentiating positive and negative emotions). The presented linear-time test achieves comparable performance to state-of-the-art quadratic-time alternatives, while returning human-interpretable features that explain the test results.

Paper: <http://www.cmap.polytechnique.fr/~zoltan.szabo/publications/jitkrittum16interpretable.pdf>, NIPS-2016 (full oral presentation = top 1.84%),

Code: <https://github.com/wittawatj/interpretable-test>,

Presenter: <http://www.cmap.polytechnique.fr/~zoltan.szabo/>.

*Machine Learning Seminar, Télécom ParisTech, Paris, France, 2 February 2017; abstract.

[†]Joint work with Wittawat Jitkrittum, Kacper Chwiałkowski and Arthur Gretton.