Language Technologies Institute

Thesis Defense

GHC 4405|Tuesday, August 3, 2021| 1:00 pm

Learning Neural Models for Natural Language Processing in the Face of Distributional Shift

Paul Michel

COMMITTEE:

Graham Neubig
(Chair)

Zachary Lipton

Zico Kolter

Tatsunori Hashimoto,
(Stanford)

Abstract

The dominating NLP paradigm of training a strong neural predictor to perform one task on a specific dataset has led to state-of-the-art performance in a variety of applications (eg. sentiment classification, span-prediction based question answering or machine translation). However, it builds upon the assumption that the data distribution is stationary, ie. that the data is sampled from a fixed distribution both at training and test time. This way of training is inconsistent with how we as humans are able to learn from and operate within a constantly changing stream of information. Moreover, it is ill-adapted to real-world use cases where the data distribution is expected to shift over the course of a model's lifetime.

The first goal of this thesis is to characterize the different forms this shift can take in the context of natural language processing, and propose benchmarks and evaluation metrics to measure its effect on current deep learning architectures. We then proceed to take steps to mitigate the effect of distributional shift on NLP models. To this end, we develop methods based on parametric reformulations of the distributionally robust optimization framework. Empirically, we demonstrate that these approaches yield more robust models as demonstrated on a selection of realistic problems. In the third and final part of this thesis, we explore ways of efficiently adapting existing models to new domains or tasks. Our contribution to this topic takes inspiration from information geometry to derive a new gradient update rule which alleviate catastrophic forgetting issues during adaptation.

https://www.cs.cmu.edu/~pmichel1/assets/Paul_Michel___PhD_Thesis_Draft.pdf