
Stochastic smoothing of the top-K calibrated hinge loss for deep imbalanced classification

Camille Garcin^{*1}, Maximilien Servajean , Alexis Joly , and Joseph Salmon

¹IMAG – CNRS, Université de Montpellier – France

Abstract

In modern classification tasks, the number of labels is getting larger and larger, as is the size of the datasets encountered in practice. As the number of classes increases, class ambiguity and class imbalance become more and more problematic to achieve high top-1 accuracy. Meanwhile, Top-K metrics (metrics allowing K guesses) have become popular, especially for performance reporting. Yet, proposing top-K losses tailored for deep learning remains a challenge, both theoretically and practically. In this paper we introduce a stochastic top-K hinge loss inspired by recent developments on top-K calibrated losses. Our proposal is based on the smoothing of the top-K operator building on the flexible "perturbed optimizer" framework. We show that our loss function performs very well in the case of balanced datasets, while benefiting from a significantly lower computational time than the state-of-the-art top-K loss function. In addition, we propose a simple variant of our loss for the imbalanced case. Experiments on a heavy-tailed dataset show that our loss function significantly outperforms other baseline loss functions

Keywords: deep learning, top, k, loss, plants

^{*}Speaker