International Journal of Computer Vision (2023) 131:626–643
https://doi.org/10.1007/s11263-022-01723-4

Revisiting Consistency Regularization for Semi-Supervised Learning

Yue Fan[1] · Anna Kukleva[1] · Dengxin Dai[1] · Bernt Schiele[1]

Received: 14 February 2022 / Accepted: 18 November 2022 / Published online: 7 December 2022

Abstract

Consistency regularization is one of the most widely-used techniques for semi-supervised learning (SSL). Generally, the aim is to train a model that is invariant to various data augmentations. In this paper, we revisit this idea and find that enforcing invariance by decreasing distances between features from differently augmented images leads to improved performance. However, encouraging equivariance instead, by increasing the feature distance, further improves performance. To this end, we propose an improved consistency regularization framework by a simple yet effective technique, FeatDistLoss, that imposes consistency and equivariance on the classifier and the feature level, respectively. Experimental results show that our model defines a new state of the art across a variety of standard semi-supervised learning benchmarks as well as imbalanced semi-supervised learning benchmarks. Particularly, we outperform previous work by a significant margin in low data regimes and at large imbalance ratios. Extensive experiments are conducted to analyze the method, and the code will be published.

Keywords Semi-supervised learning · Consistency regularization · Representation learning · Classification

1 Introduction

Deep learning requires large-scale and annotated datasets to reach state-of-the-art performance (Russakovsky et al. 2015; Lin et al. 2014). As labels are not always available or expensive to acquire a wide range of semi-supervised learning (SSL) methods have been proposed to leverage unlabeled data (Tarvainen and Valpola 2017; Laine and Aila 2017; Miyato et al. 2018; Verma et al. 2019; Berthelot et al. 2019; Sohn et al. 2020; Xie et al. 2020; Berthelot et al. 2020; Arazo et al. 2020; Lee 2013; Pham et al. 2020; French et al. 2020; Bachman et al. 2019; Chen et al. 2020b).

Consistency regularization (Bachman et al. 2014; Laine and Aila 2017; Sajjadi et al. 2016) is one of the most widely-used SSL methods. Recent work (Sohn et al. 2020; Xie et al. 2020; Kuo et al. 2020) achieves strong performance by utilizing unlabeled data in a way that model predictions should be invariant to input perturbations. However, when using advanced and strong data augmentation schemes, we question if the model should be invariant to such strong perturbations. In Fig. 1 we illustrate that strong data augmentation leads to perceptually highly diverse images. Thus, we argue that improving equivariance on such strongly augmented images can provide even better performance rather than making the model invariant to all kinds of augmentations. Moreover, existing works apply consistency regularization either at the feature level or at the classifier level. We find empirically that it is more beneficial to introduce consistency on both levels. To this end, we propose a simple yet effective technique, Feature Distance Loss (FeatDistLoss), to improve data-augmentation-based consistency regularization.

We formulate our FeatDistLoss as to explicitly encourage invariance or equivariance between features from different augmentations while enforcing the same semantic class label. Figure 2 shows the intuition behind the idea. Specifically, encouragement of equivariance for the same image but different augmentations (increase distance between stars and circles of the same color) pushes representations apart from each other, thus, covering more space for the class. Impos-

Communicated by Christian Bauckhage.

✉ Yue Fan
yfan@mpi-inf.mpg.de

Anna Kukleva
akukleva@mpi-inf.mpg.de

Dengxin Dai
ddai@mpi-inf.mpg.de

Bernt Schiele
schiele@mpi-inf.mpg.de

[1] Max Planck Institute for Informatics, Saarland Informatics Campus, Saarbrücken, Germany