Meta Co-Training: Two Views are Better than One

Jay C. Rothenberger*, Dimitrios I. Diochnos, University of Oklahoma

In many practical computer vision scenarios unlabeled data is plentiful, but labels are scarce and difficult to obtain. As a result, semi-supervised learning which leverages unlabeled data to boost the performance of supervised classifiers have received significant attention in recent literature. One major class of semi-supervised algorithms is *Co-Training* [1]. In co-training two different models leverage different independent and sufficient "views" of the data to jointly make better predictions. During co-training each model creates pseudo labels on unlabeled points which are used to improve the other model. Our first contribution is that we show that in the common case when independent views are not available we can construct such views inexpensively using pre-trained models. Co-training on the constructed views yields a performance improvement over any of the individual views we construct and performance comparable with recent approaches in semi-supervised learning, but has some undesirable properties. To alleviate the issues present with co-training we present *Meta Co-Training* which is an extension of the successful *Meta Pseudo Labels* [3] approach to two views. Our method achieves 85.8% top-1 accuracy on ImageNet-10%, which outperforms the 85.1% of the current state-of-the-art (SotA) which is due to REACT [2]; see also https://paperswithcode.com/sota/semi-supervised-image-classification-on-2 for papers that have defined the SotA in recent years. Finally, we test our method in some additional computer vision datasets with equally encouraging results when comparisons were possible.

* Email address: jay.c.rothenberger@ou.edu

References

[1] Avrim Blum and Tom M. Mitchell. Combining Labeled and Unlabeled Data with Co-Training. In Peter L. Bartlett and Yishay Mansour, editors, *Proceedings of the Eleventh Annual Conference on Computational Learning Theory, COLT 1998, Madison, Wisconsin, USA, July 24-26, 1998*, pages 92–100. ACM, 1998.

[2] Haotian Liu, Kilho Son, Jianwei Yang, Ce Liu, Jianfeng Gao, Yong Jae Lee, and Chunyuan Li. Learning Customized Visual Models with Retrieval-Augmented Knowledge. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pages 15148–15158. IEEE, 2023.

[3] Hieu Pham, Zihang Dai, Qizhe Xie, and Quoc V. Le. Meta Pseudo Labels. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pages 11557–11568. Computer Vision Foundation / IEEE, 2021.