Improving Generalizability in Exemplar-Free Class-Incremental Learning

Horikawa Strakovsky, Arimitsu (School: Paul Laurence Dunbar High School)

While pre-training with supervised or self-supervised learning has recently shown great promise in continual learning (CL) by learning generalizable features, it requires a large, external dataset, similar to the target dataset. These types of datasets are available for popular benchmarking datasets but are often not available in real-world applications. As a result, many works still train randomly initialized models from scratch. However, no such works have investigated alternative measures to introduce generalizable features. Thus, self-supervised label augmentation (SSLA) is revisited in this work in a subset of CL, exemplar-free class-incremental learning (EFCIL). While SSLA has been previously proposed, it was utilized in an extremely naive manner, which is found to be harmful to performance with modern EFCIL algorithms. Due to this, two modifications based in theoretical observations are proposed to better fit SSLA to the EFCIL context: performing SSLA on only the initial task and performing knowledge distillation on only non-augmented class logits. Furthermore, SSLA is limited due to its reliance on either rotation or color-variant features, which may not be present in all datasets. Therefore, a novel augmentation for SSLA using binary low-pass and high-pass filters on the frequency domain of images is proposed for such settings. The proposed methods are evaluated on CIFAR-100, Tiny-ImageNet, and Skin23 and achieved significant gains in last task accuracy over the state-of-the-art with +2.5%. +2.7%, and over the baseline with +0.4% respectively.

Awards Won: