

Incremental Learning for Scene Understanding

Karteek Alahari: karteek.alahari@inria.fr

Cordelia Schmid: cordelia.schmid@inria.fr

<http://thoth.inriaples.fr>

Location of the internship: The internship will be in the Thoth team at Inria Grenoble, and will be co-supervised by Karteek Alahari (Inria researcher) and Cordelia Schmid (Inria Research Director). The team is specialized in computer vision, in particular visual recognition.

Topic: In the context of learning a network-based representation of data, there is a growing need to perform incremental updates. This problem manifests itself in two scenarios: when additional data samples of existing classes or new data from unseen classes becomes available. Despite their success for several computer vision problems, CNNs are ill-equipped for such incremental learning [2, 4]. For example, adapting the original model trained on a set of classes to additionally represent samples from new classes, in the absence of the initial training data, leads to a phenomenon called “catastrophic forgetting” [1, 3]. This phenomenon is an abrupt degradation of performance on the original set of classes, when the training objective is adapted to the new classes. Our recent work [5] is a first step to address this issue. We proposed an approach for learning object detectors incrementally, with an optimization function to balance the interplay between predictions on the new classes and a new distillation loss which minimizes the discrepancy between responses for old classes from the original and the updated networks. This work now needs to be generalized to other end-to-end frameworks, beyond object detection, such as cross-modal learning. Furthermore, the incremental learning problem for the other challenging scenario with additional samples from known classes needs to be addressed.

Skills and profile: The student must have solid programming skills as well as solid mathematics knowledge (especially linear algebra and statistics). Knowledge of deep learning tools is a strong plus.

References

- [1] I. J. Goodfellow, M. Mirza, X. Da, A. C. Courville, and Y. Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks. In *ICLR*, volume abs/1312.6211, 2014.
- [2] Z. Li and D. Hoiem. Learning without forgetting. In *ECCV*, 2016.
- [3] M. McCloskey and N. J. Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. *Psychology of learning and motivation*, 24:109–165, 1989.
- [4] S.-A. Rebuffi, A. Kolesnikov, and C. H. Lampert. iCaRL: Incremental classifier and representation learning. In *CVPR*, 2017.
- [5] K. Shmelkov, C. Schmid, and K. Alahari. Incremental learning of object detectors without catastrophic forgetting. In *ICCV*, 2017.