

Master Thesis Topics in Machine Learning

Master Thesis at RISE SICS in Kista, working on fast inference, uncertainty and online learning.

We are looking for students with a strong background in Machine Learning (ML) to work on state of the art research issues. The topics on offer deal with using ML for large scale data. More specifically we will investigate the following three topics:

Fast inference in deep networks: As neural networks grow deeper and more complex fast inference is becoming increasingly harder. In this thesis we plan to investigate architectural and system improvements to deep networks with the aim of offering faster inference. See [1] for relevant work.

Uncertainty in decision trees: Decision trees have been one of the most successful algorithms for classification and regression. However, research has lagged behind in determining the uncertainty over predictions in trees. This thesis will investigate various approaches in calculating uncertainty for decision trees. See [2] for related work.

Distributed Online Learning: As streaming data analytics become more popular, we are faced with the lack of a modern software platform aimed specifically at learning from streaming data in a distributed framework. For this thesis we aim to design and implement a proof of concept framework for distributed online learning. See [3] for related work.

Students are encouraged to propose their own topics as well.

Please contact

Theodore Vasiloudis
PhD Candidate, KTH
Researcher, RISE SICS
e-mail: tvas@sics.se

- [1] Bolukbasi, T., Wang, J., Dekel, O. & Saligrama, V. (2017). *Adaptive Neural Networks for Efficient Inference*. Proceedings of the 34th International Conference on Machine Learning, in PMLR 70:527-536
- [2] Balaji L., Daniel M. R., and Yee Whye Teh. 2014. *Mondrian forests: efficient online random forests*. In Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2 (NIPS'14),
- [3] De Francisci Morales, G. (2013, May). *SAMOA: A platform for mining big data streams*. In Proceedings of the 22nd International Conference on World Wide Web (pp. 777-778). ACM.