



## Developing Neural Memory

There is a distinct appetite for rational, deductive systems in natural language processing, yet we seem to miss even the first ingredient necessary – reliable, learnable memory. Neural networks are perfectly capable of unlearning things previously learned if they have not seen them for long enough in training. Furthermore, our experimentation with autoencoders indicates heavy dependency of latent representations on the distribution of input samples – samples that are more likely get more prominent positioning in the latent vectors regardless of their objective importance to the sample in terms of the ground-truth factors.



We aim to eradicate this issue with a meta-learning technique weighing weights during training. Think of it as dynamic multi-task learning if you wish, but with continuous information compression and hierarchical generalisation.

**About the thesis.** You are going to be working very closely with DISCO group members, advancing the research of the group. In the first few weeks you will get familiar with the related work, and you will have the option to decide whether you want to pursue the topic further. Going forward, we will formulate a project – based on your skills and interests – that we will work on for the rest of your time with us.

**Candidate Profile.** Generally speaking, a good candidate is a competent programmer in the language of his/her choice, has good knowledge of or solid experience with TensorFlow or (Py)Torch, and is interested in one or more of the following fields: meta-learning, transfer learning, natural language inference, program induction, program synthesis.

**Interested? Please contact us to learn more!**

### Contact

- Peter Belcak: [belcak@ethz.ch](mailto:belcak@ethz.ch), ETZ G61.3