Learning to Teach: Scaffolded Meta-Learner for Efficient Teaching of Human Learners

Presented at NeurIPS 2020 (WIML workshop)

Alexandra Moringen

An appropriate choice of practice task w.r.t. their difficulty, and monitoring of progress are central for successful learning of most skills [3]. Motivated by the wish to accelerate and assist a human who is learning a new skill, we propose an extension of MAML for reinforcement learning [1] with a scaffold [3]. We exemplify our approach on the most basic but newly shaped teacher-learner interaction scenario during piano practice. The main idea is to have a model that on one hand guides the human learning process, but on the other hand, adapts to the user and their individual learning needs.

By monitoring human performance and its improvement, scaffolded MAML (sMAML) manages task complexity levels, represented by task sets provided for practice. Similar to the MAML concept, for a given complexity level, the meta-learner part of sMAML samples tasks out of the current task set, while the composition of the tasks and their number in the set is monitored by the scaffolding part of sMAML. For each sampled task within the inner loop, the reinforcement learning (RL) model generates concrete score pieces with practice instructions for the learner. During the RL model parameter optimization, the loss is evaluated based of the performance of the human learner. Larger rewards correspond to a more accurate performance of the practice example. Concurrent to the human learning process, the policy update seeks to generate practice examples that maximize reward for a given task level. The scaffolding part of sMAML decides to either stay on the same complexity level, or to change the task set if the performance of the human learner on the test task rapidly improves. The complexity is not raised as long as the human learner is still improving and is not capable to learn a previously unpracticed and slightly more challenging practice piece quickly enough. Analysis of midi produced by the learner and/or capture of movement data may allow us a very detailed error analysis.

sMAML is constantly coupled with the human learner through a reciprocal connection, defined by the task set selection and evolution, the quality of human performance and its dynamics. It provides basic implementation of scaffolding functions to the human learner, including performance analysis and monitoring, fading support, gradual task set modification, etc. The main difference to the standard application of MAML is that in our case, the sequences of practice examples corresponding to the action instructions generated by the MAML policy rollouts for a given task $\mathcal{T}_i \in \mathcal{T}$ are executed by the human learner, the reward is generated based on the learner's performance online. If, for example, on the current level, the learner has to play two black keys with two fingers at a given speed and temporal precision, then the sampled tasks may be defined by different speeds, left of right hand, different rhythmic patterns, or different volume instructions. The test task may be to play three black keys with three fingers.

To sum up, the meta-learner is responsible for training the learner on one skill level. Once the test phase of the meta-learner shows good adaptation of the human learner to a new task, the scaffold is responsible for changing the task set employed by the meta-learner. This should enable the human learner to move to the next level of complexity. The scaffold has to have an access to a large set of tasks (e.g. [2]), which are bit by bit provided to the meta-learner as the human learner progresses.

Acknowledgement

This work is supported by a Grant from the GIF, the German-Israeli Foundation for Scientific Research and Development.

References

[1] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. volume 70 of *Proceedings of Machine Learning Research*, pages 1126–1135, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR.

[2] Willard A. Palmer Palmer, Morton Manus, and Amanda Vick Lethco. *Alfred's basic piano library, Prep course for the young beginner*. Alfred Publiahing Co.; Inc., 2002.

[3] K. Ann Renninger and Alexandra List. *Scaffolding for Learning*, pages 2922–2926. Springer US, Boston, MA, 2012.