

Dr. Christian Gumbsch

Learning Event-Based Temporal Abstractions for Hierarchical Prediction and Planning

Over the last decade, deep reinforcement learning systems have made remarkable progress in various domains, partially reaching superhuman performance when trained extensively. However, no artificial system has yet reached the flexibility and efficiency with which intelligent animals learn to solve novel problems. The goal of this thesis is to close this gap to some extent by enhancing the goal-directed behavior of artificial agents through the ability to hierarchically decompose sensorimotor experience into events.

The central hypothesis of this work is that model-based temporal abstractions of events not only play a crucial role for human decision making but by learning such structures artificial agents can also acquire more adaptive, far-reaching, goal-directed behavior. To investigate this, a framework is introduced for the event-based learning of hierarchical models with nested time scales. Throughout this talk, I will present three concrete deep learning systems that not only illustrate how to implement the framework in practice, but also demonstrate several strengths of the overall approach.

Learning hierarchical predictions requires a mechanism that suitably decomposes activity into events. As a first step, a recurrent neural network is introduced, which learns in a self-supervised way to compress dynamics into latent states that are sparsely updated over time. Including this mechanism in different prediction and planning systems improves their generalization abilities, their sample efficiency, and the explainability of the learned representations.

Next, the cognitive plausibility of the approach is investigated by modeling human anticipatory behavior. In these modeling experiments, an agent learns a hierarchy of models based on the previously introduced segmentation mechanism. When the agent selects its gaze to minimize uncertainty across its hierarchical predictions, goal-anticipatory gaze behavior emerges similarly to the eye fixations that develop in infants during their first year of life.

Finally, the segmentation and abstraction mechanisms are extended to learn a hierarchy of world models from scratch. On the high-level, meaningful temporal abstractions develop that guide low-level control. In challenging problems with long task horizons and pixel-based inputs, this hierarchical learning scheme can seamlessly enhance existing approaches of model-based reinforcement learning and planning.

Taken together, this thesis not only provides practical methods for learning event-based temporal abstractions, but also demonstrates how such structures can explain human behavior and enhance sequential decision making in artificial agents.