Karlsruhe Institute of Technology

Autonomous Learning Robots (ALR)
Prof. Gerhard Neumann

Project Type

- ☑ Master Thesis
- ☐ Bachelor Thesis
- ☑ Research Project

Supervisors

☞ Philipp Becker
@ philipp.becker@kit.edu

Difficulty

Algorithmic

Math

Application

State Space Models vs Transformers for Continuous Control

Description

Many recent approaches tackle reinforcement learning and control as sequence modeling problems using Transformers as their foundational architecture [1, 3, 4]. However, Transformers were originally developed for language processing and may not be ideally suited for continuous systems prevalent in robotics due to their inherent inductive biases. Recently state space model-based approaches have emerged as promising alternatives to Transformers for sequence modeling [2]. Besides beneficial scaling properties, their underlying theory builds on classical signal processing for continuous systems which makes them potentially much better suited as a backbone for continuous control tasks.

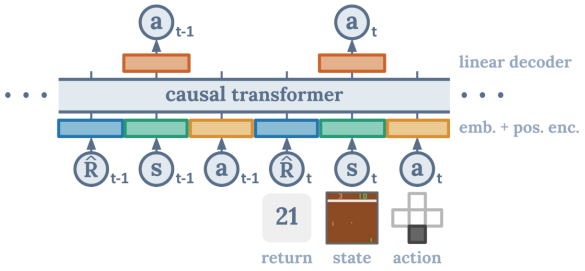

Figure 1: Left: The Decision Transformer [1] is one example of a recent sequence modeling approach to Reinforcement Learning. It is trained to predict the optimal actions given states, rewards, and previous actions. As the name implies, it builds on a Transformer - an architecture originally designed for discrete language modeling tasks. Right: Recent alternative approaches to sequence modeling build on continuous state space models and pose a more principled alternative to transformers for continuous control tasks.

In this thesis, we will investigate how to replace the backbone of transformer-based control approaches with modern state-space models and if their more principled inductive bias leads to practical improvements in continuous robot tasks.

Tasks

- Getting familiar with transformer-based sequence modeling approaches for RL as well as current state-space approaches.
- Replacing the transformer backbone in several of those approaches with the state space approach. Here, we can build on existing implementations.
- Investigate different design choices of how to include the state space approaches.
- Conducting a larger scale evaluation and comparison.

References

[1] Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Misha Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 2021.

[2] Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*, 2023.

[3] Yanchao Sun, Shuang Ma, Ratnesh Madaan, Rogerio Bonatti, Furong Huang, and Ashish Kapoor. Smart: Self-supervised multi-task pretraining with control transformers. In *International Conference on Learning Representations*, 2023.

[4] Philipp Wu, Arjun Majumdar, Kevin Stone, Yixin Lin, Igor Mordatch, Pieter Abbeel, and Aravind Rajeswaran. Masked trajectory models for prediction, representation, and control. In *International Conference on Machine Learning*, 2023.