

Model Learning With Recurrent Kalman Networks



Project Manager
Dr. Oleg Arenz

Researchers
Vaishak Shaj

Principal Investigator
Prof. Jan Peters (PhD)

Project Term
2019 - 2020

Clusters
Lichtenberg Cluster Darmstadt

Software
TensorFlow

Institute
Intelligent Autonomous Systems

University
Technische Universität Darmstadt

Introduction

Recurrent State-space models (RSSMs) are highly expressive models for learning patterns in time series data and system identification. However, these models assume that the dynamics are fixed and unchanging, which is rarely the case in real-world scenarios. Many control applications often exhibit tasks with similar but not identical dynamics which can be modeled as a latent variable. Our goal is to learn a state space model that can model the dynamics of partially observable robotic systems under scenarios where the dynamics changes over time. Often, the dynamics of real systems may differ in significant ways from the system our models were trained on. However, it is unreasonable to train a model across all possible conditions an agent may encounter. Instead, we propose a state space model that learns to account for the causal factors of variation observed across tasks at training time, and then infer at test time the model that best describe the system. In order to compare this approach with previous methods, we run extensive experiments using robot simulations.

Methods

We introduce a new formulation namely Hidden Parameter State Space Models (HiP-SSMs), a framework for modeling families of SSMs with different but related dynamics using lowdimensional latent task embeddings. We perform inference and learning in the HiP-SSM borrowing concepts from both deep learning and graphical model communities following recent works on recurrent neural network models, where the architecture of the network is informed by the structure of the probabilistic state

estimator. Aside from the simplicity of the training procedure, one of the key advantages of this approach is the ability to incorporate arbitrary nonlinear components into the observation and transition functions. We evaluate our method in simulation for robotic table tennis, locomotion and manipulation.

Results

We observe that the HiP-RSSM can outperform the previous state of the art predictions obtained by recurrent models. We hypothesise that the latent task variable captures multiple unobserved causal factors of variation that affect the dynamics in the latent space, which are not modelled in contemporary recurrent models.

Discussion

We proposed HiP-RSSM, a probabilistically principled recurrent neural network architecture for modelling changing dynamics scenarios. We start by formalizing a new framework, HiP-SSM, to address the multi-task state-space modelling setting. HiP-SSM assumes a shared latent state and action space across tasks but additionally assumes latent structure in the dynamics. We exploit the structure of the resulting Bayesian network to learn a universal dynamics model via forward inference algorithm and backpropagation through time. The resulting recurrent neural network, namely HiP-RSSM, learns to cluster SSM instances with similar dynamics together in an unsupervised fashion. Our experimental results on various robotic benchmarks show that HiP-RSSMs significantly outperform state of the art recurrent neural network architectures on dynamics modelling tasks. We believe that modelling the dynamics in the latent space which disentangles the state, action and task representations can benefit multiple future applications including planning/control in the latent space and causal factor identification.

Publications

Shaj, V., Büchler, D., Sonker, R., Becker, P., & Neumann, G.: "Hidden Parameter Recurrent State Space Models For Changing Dynamics Scenarios", In International Conference on Learning Representations, 2022 <https://openreview.net/forum?id=ds8yZOUsea>

Last Update: 2022-04-17 16:45