

Evaluating and Improving Recurrent Kalman Networks



Project Manager
Dr. Oleg Arenz

Researchers
Philipp Becker

Principal Investigator
Prof. Jan Peters (PhD)

Project Term
2019 - 2020

Clusters
Lichtenberg Cluster Darmstadt

Software
PyTorch

Additional Software
MuJoCo

Institute
Intelligent Autonomous Systems

University
Technische Universität Darmstadt

Introduction

Precise models of the system dynamics are crucial for model-based control and reinforcement learning (RL) in autonomous systems applications. In partially observable settings, practitioners commonly use state space models (SSMs) to formalize such systems. SSMs consist of a dynamics model, describing how one state relates to the next, and an observation model, which describes how system states generate observations. Given these models, probabilistic filters, such as the Kalman filter, allow efficient inference of state beliefs. Yet, these models are unknown for most relevant problems, and exact inference is usually not tractable. Recurrent state space models (RSSMs) have found considerable interest in the model-based RL community. Yet, while RSSMs are clearly inspired by classical state space models, they use an inference scheme that builds on simplified assumptions. They assume the belief is independent of future observations instead of using the common smoothing inference scheme. Moreover, RSSMs use nonlinear parametrizations of the state space model, which render closed-form inference intractable. Thus, they must rely on sample-based inference, which induces further noise in the training process. In this project, we introduce a novel approach called Variational Recurrent Kalman Network (VRKN), which allows exact smoothing inference in latent spaces by using a linear Gaussian parametrization of the state space model, and evaluate it for prediction of robotic movements in simulation.

Methods

We introduce a new parametrization of the latent dynamics

based on a linear Gaussian state space model (LGSSM) embedded in a latent space. The linear Gaussian assumptions allow for efficient inference and rigorous treatment of uncertainties while working in a learned latent space allows modeling high-dimensional and non-linear systems. We compare our method with a Bayesian variant (Bay-VRKN) and with the originally proposed RSSM on a variety of tasks and data sets recorded using the MuJoCo physics simulator.

Results

The VRKN cannot compete with the original RSSM-approaches. Yet, when combined with a simple form of modeling epistemic uncertainty, i.e., Monte-Carlo Dropout, the Bayesian version of the VRKN performs very similar to the original RSSM. We hypothesize that the noise introduced by the sample-based inference used by RSSM is beneficial for reinforcement learning as it improves exploration. Further, we find the Bay-VRKN converges faster to the final performance in several environments. Arguably, this again is related to the inference assumptions as the RSSM has to learn a more complex transition model, accounting for the aleatoric uncertainty.

Discussion

We propose an alternative approach, building on exact smoothing inference in linear Gaussian state space models embedded in a latent space. When combined with explicit modeling of epistemic uncertainty, this approach matches the RSSMs performance on the standard benchmarks and improves it in noisy settings while building on well-understood, theoretically founded components. A key limitation of this work is the usage of Monte-Carlo Dropout to represent the epistemic uncertainty. We chose it due to its simplicity, and revising this choice in future work could improve the approach. Additionally, while we provide a detailed analysis of RSSMs and a better understood alternative, there still remain open questions about the workings of RSSMs in practice.

Last Update: 2022-03-28 15:40