KI WISSEN
Automotive AI Powered by Knowledge

Stefan Rudolph

Towards Causality-Driven Reinforcement Learning for Autonomous Driving

20.1.2022 - Computational Science Lab Seminar

KI Wissen

- Part of *KI Familie*
- Next level autonomous driving
- Knowledge-based approaches to development of ML functions
  - Knowledge integration
  - Knowledge extraction
  - Knowledge conformity

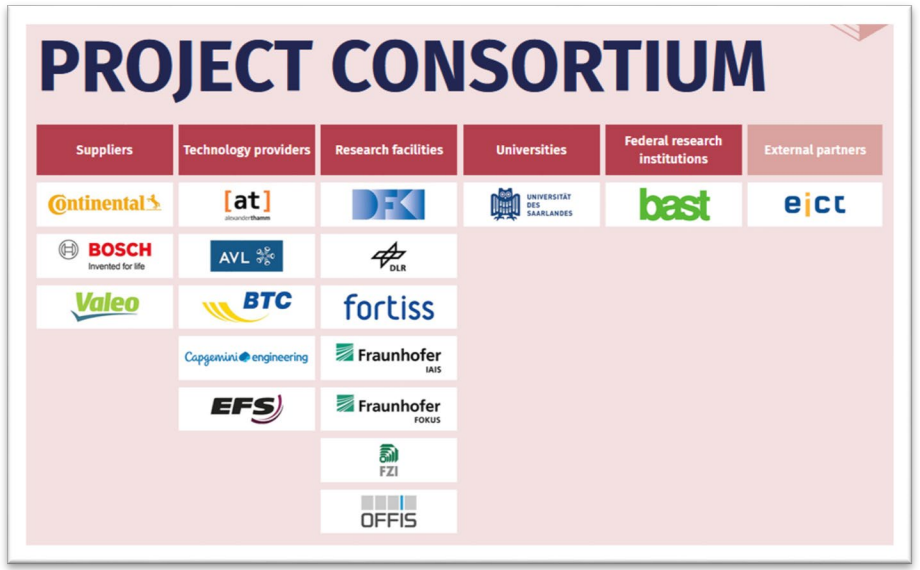

Motivation - Recent Achievements in Reinforcement Learning

- Several benchmarks “solved” over the last years

- Atari – “*Human-level control through deep reinforcement learning*” (Mnih et al., 2015)
  - Deep Q-learning (DQN)
  - Images as input

- Go – “*Mastering the game of Go with deep neural networks and tree search*” (Silver, 2016)
  - AlphaGo – Monte Carlo tree search (MCTS)
  - Very large search space

Motivation - Recent Achievements in Reinforcement Learning

- Dota 2 – "*Dota 2 with Large Scale Deep Reinforcement Learning*" (OpenAI, 2019)
  - Proximal policy optimization (PPO)
  - Long time horizons
  - Imperfect information
  - Complex, continuous state-action spaces

- Starcraft 2 – "*Grandmaster level in StarCraft II using multi-agent reinforcement learning*" (Vinyals et al., 2019)
  - AlphaStar – league training
  - Compete and coordinate with other agents in complex environments

Motivation – Current Research in Reinforcement Learning

- Model-based RL on continious control tasks
  - „*Dream to Control: Learning Behaviors by Latent Imagination*“ (Hafner et al., 2019)
  - Dreamer – Baysian model to train model-free RL

- MCTS with a learned model
  - „*Mastering Atari, Go, Chess and Shogi by Planning with a Learned Model*” (Schrittwieser et al., 2020)
  - MuZero – MCTS with a learned model

- Model-based RL in Autonomous Driving
  - “*Interpretable End-to-end Urban Autonomous Driving with Latent Deep Reinforcement Learning*” (Chen et al., 2020)

RL and Autonomous Driving

- Challenges
  - Data sparsity
  - Safety Guarantees
  - Collecting on-policy data is very expensive

- Idea: Use Causal Model
  - Generate/augment training data
  - Check conformity

Knowledge Integration & Conformity

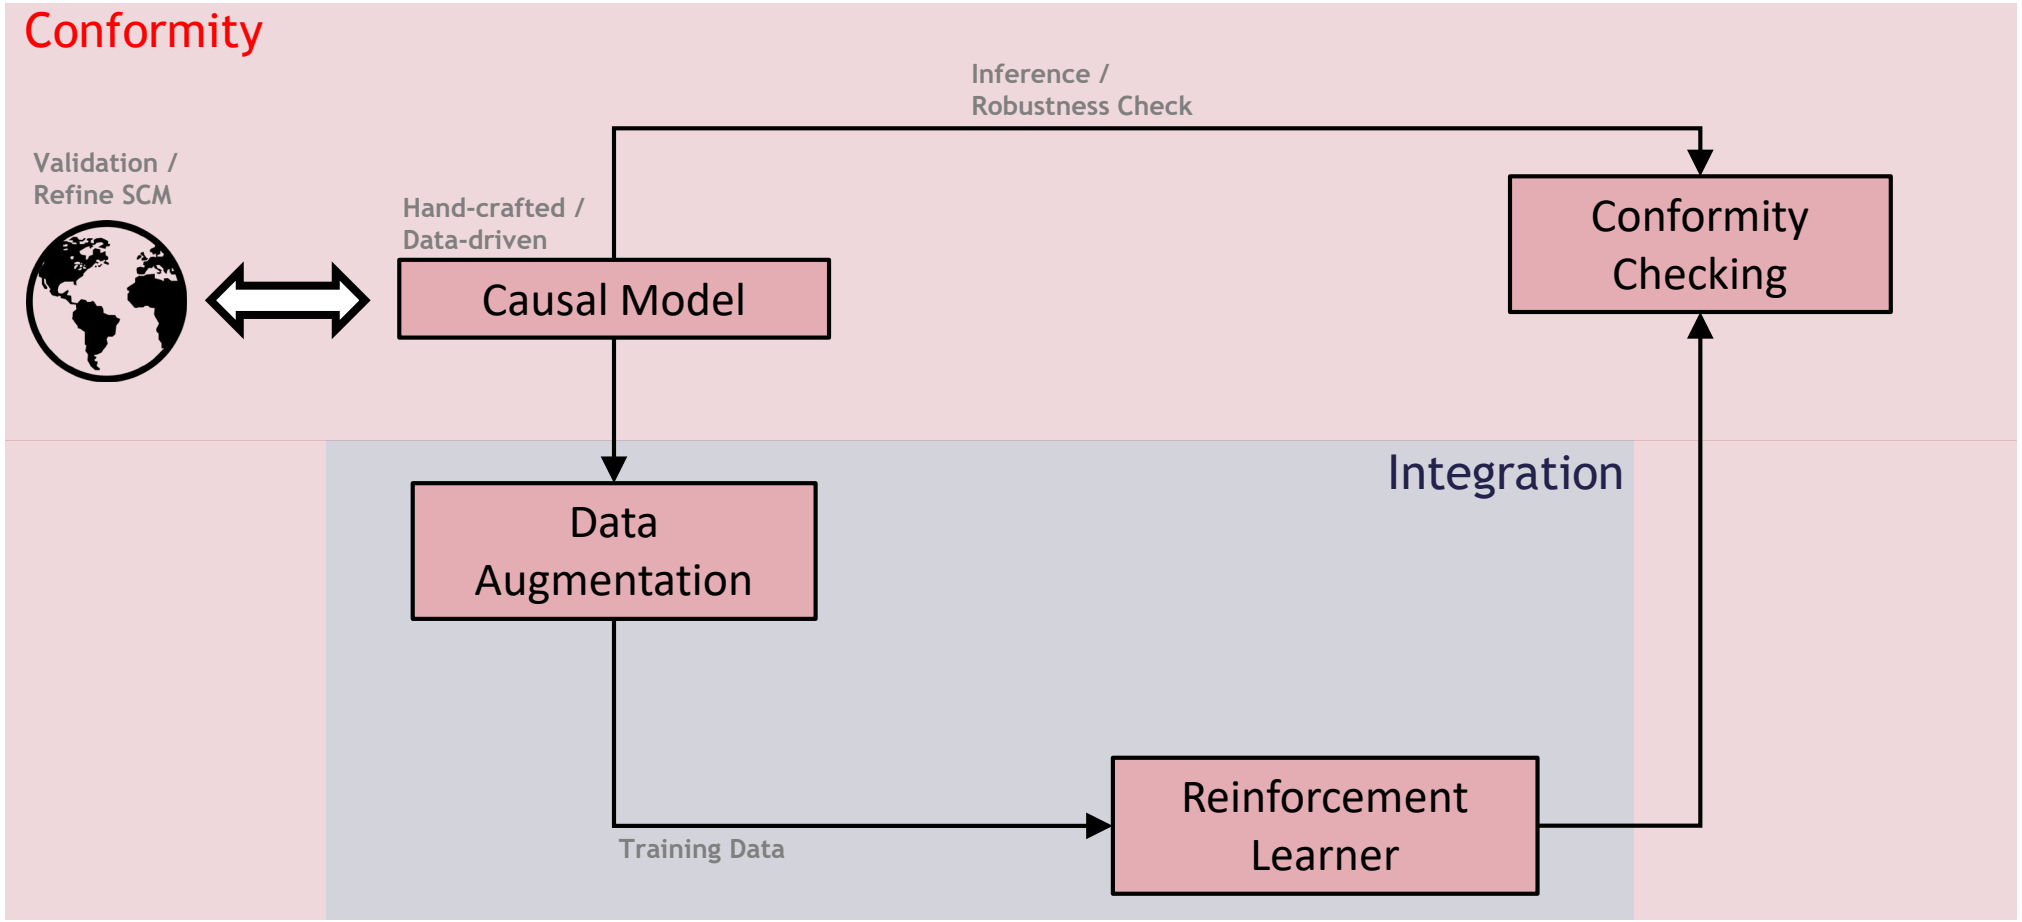

Exemplary Knowledge:
Vehicle Dynamics

The RL Setting

- **Agent** interacts with **environment**
  - Agent receives (partial) **observation** of **state**
  - Agent applies **action**
  - Agent receives **reward**
- Repeats over **multiple steps**
  - Sum of **rewards** is called the **return**
- Goal of the agent
  - Maximize expected **return** over time

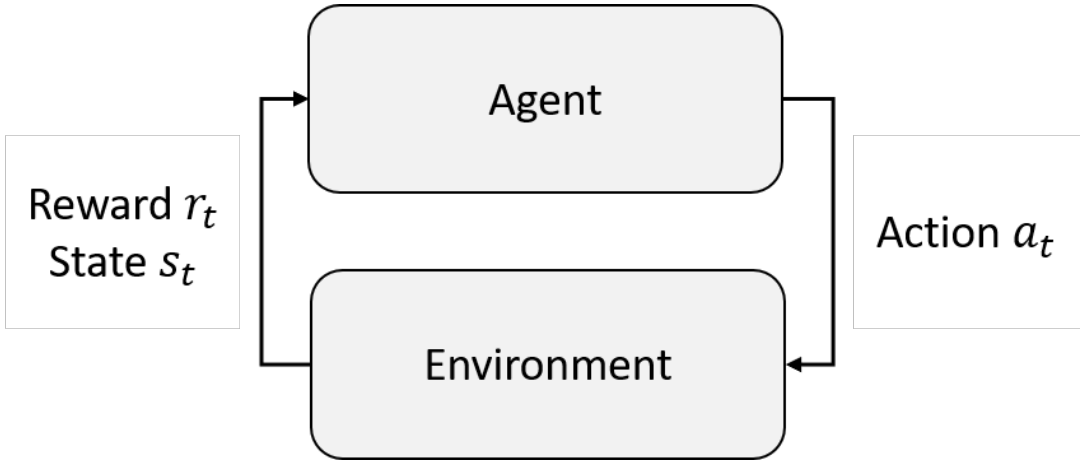

Dreamer v2 - Learning Flow

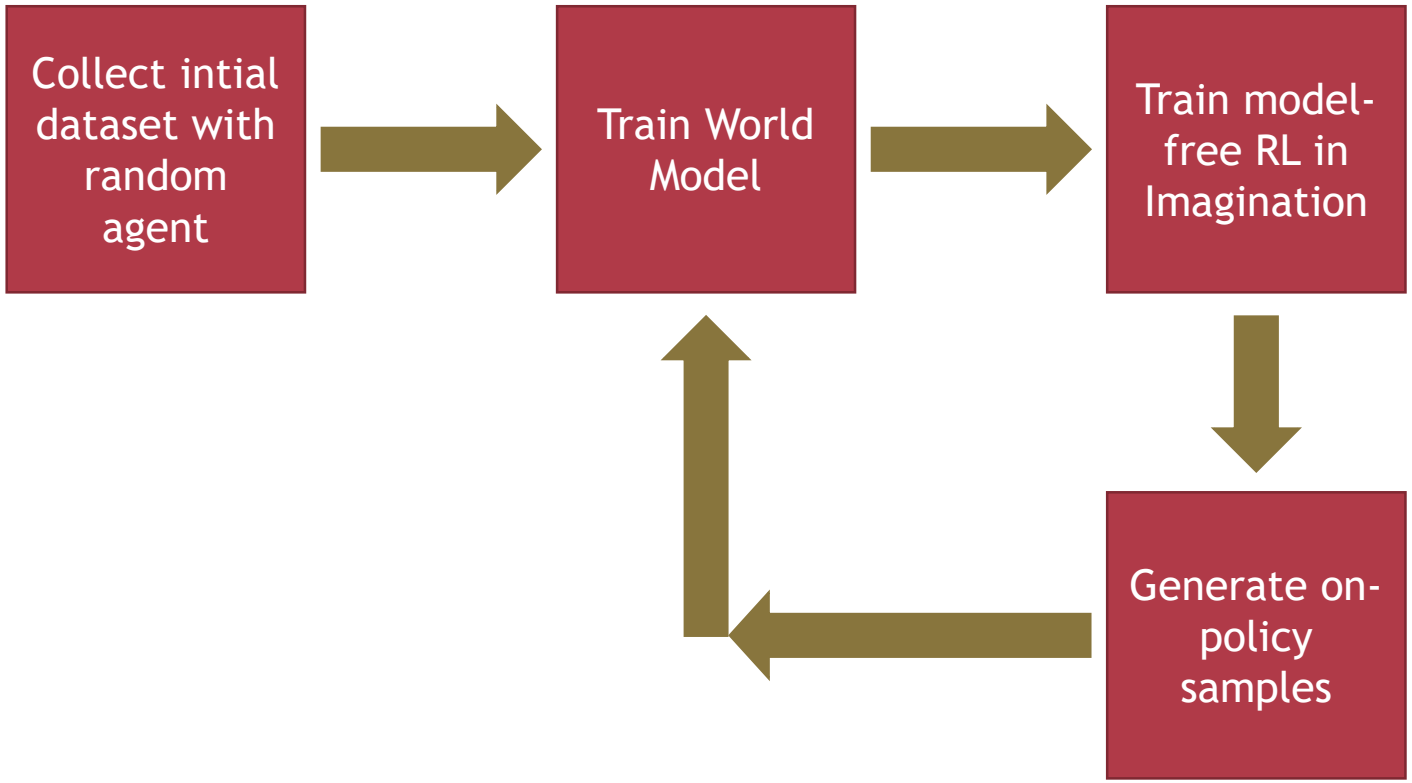

Dreamer v2 - Recurrent State Space Model

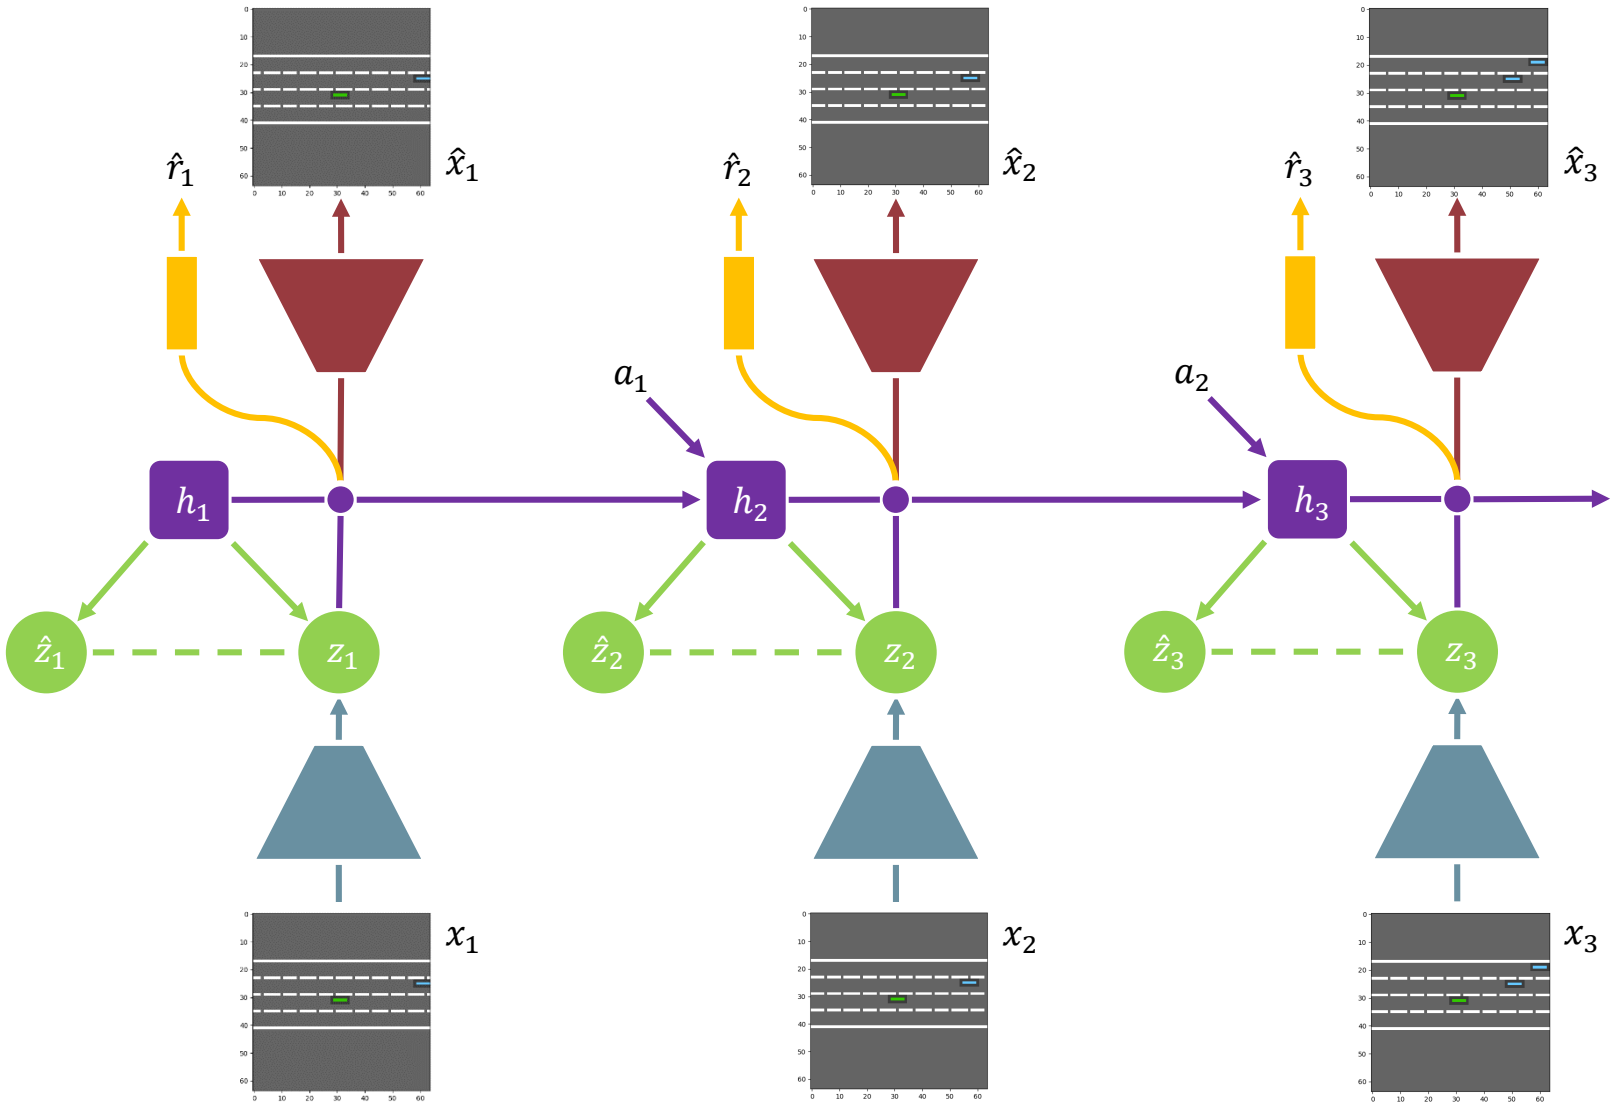

Dreamer v2 - Training of the model-free RL

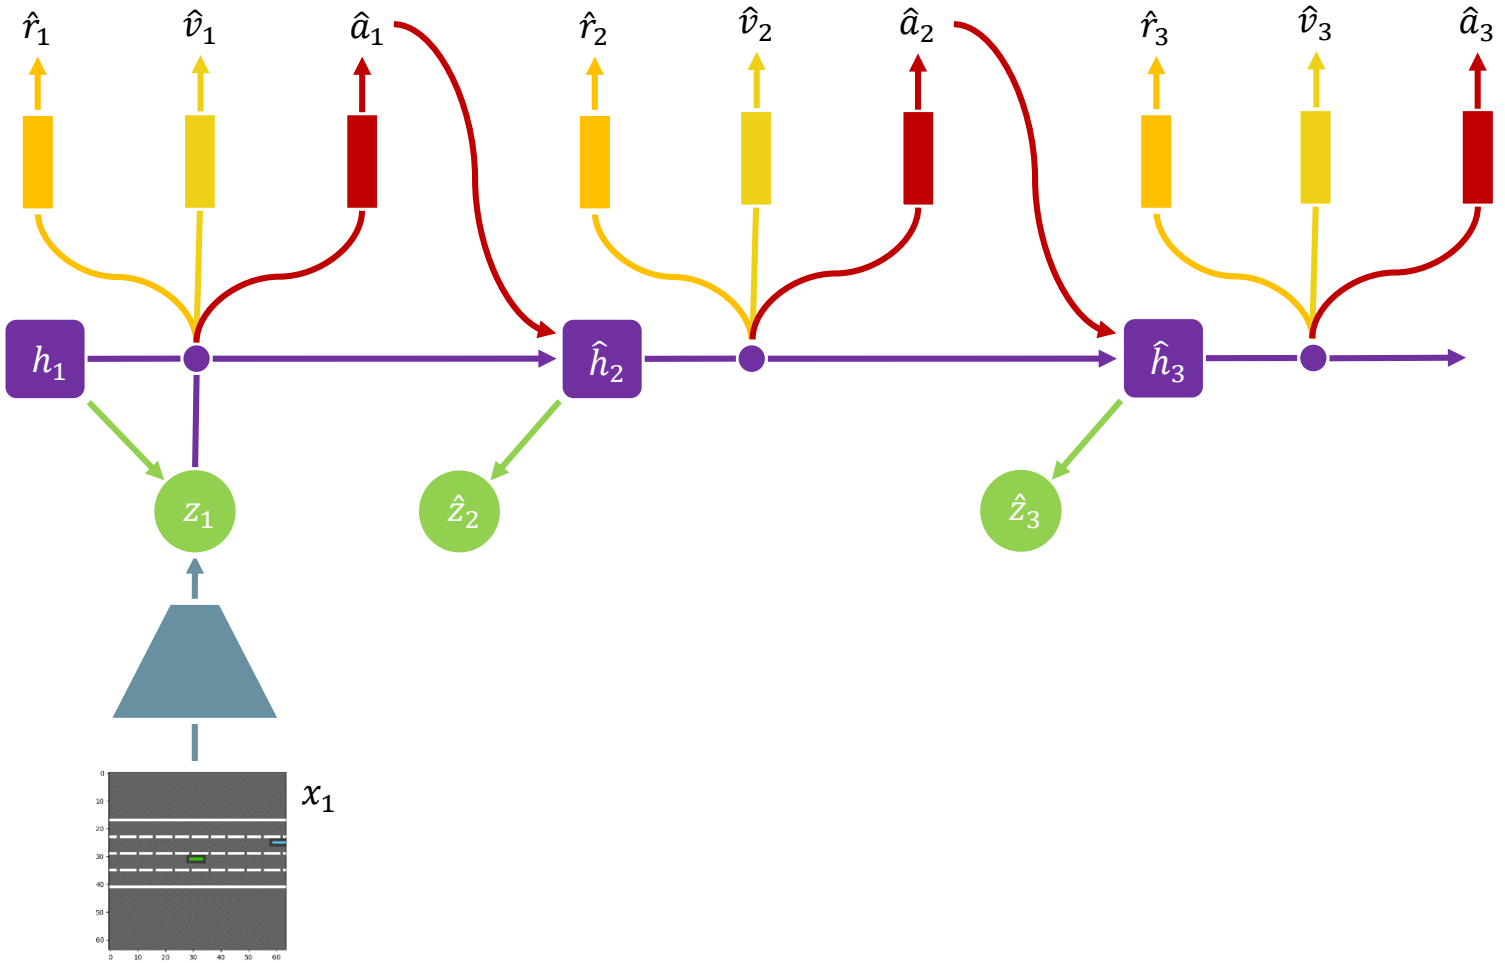

Structural Causal Models

- Structural causal model (SCM)
  - Graphical model
  - Causal relationships
  - Ladder of causation

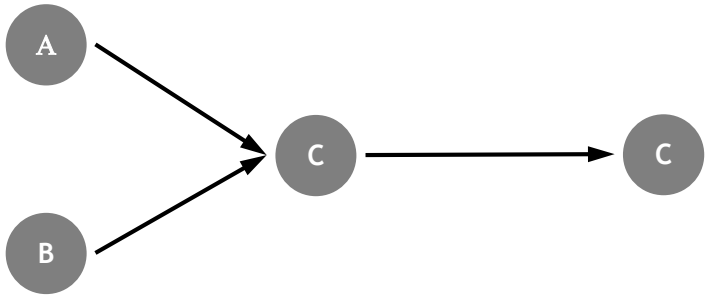

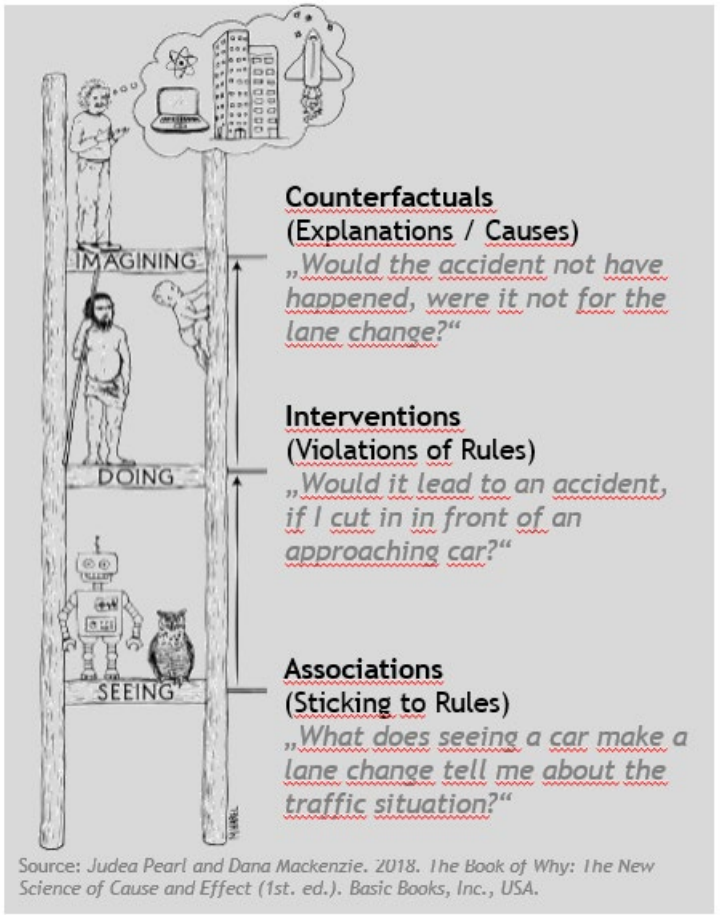

Source: Judea Pearl and Dana Mackenzie. 2018. The Book of Why: The New Science of Cause and Effect (1st. ed.). Basic Books, Inc., USA.

Vehicle Dynamics Model

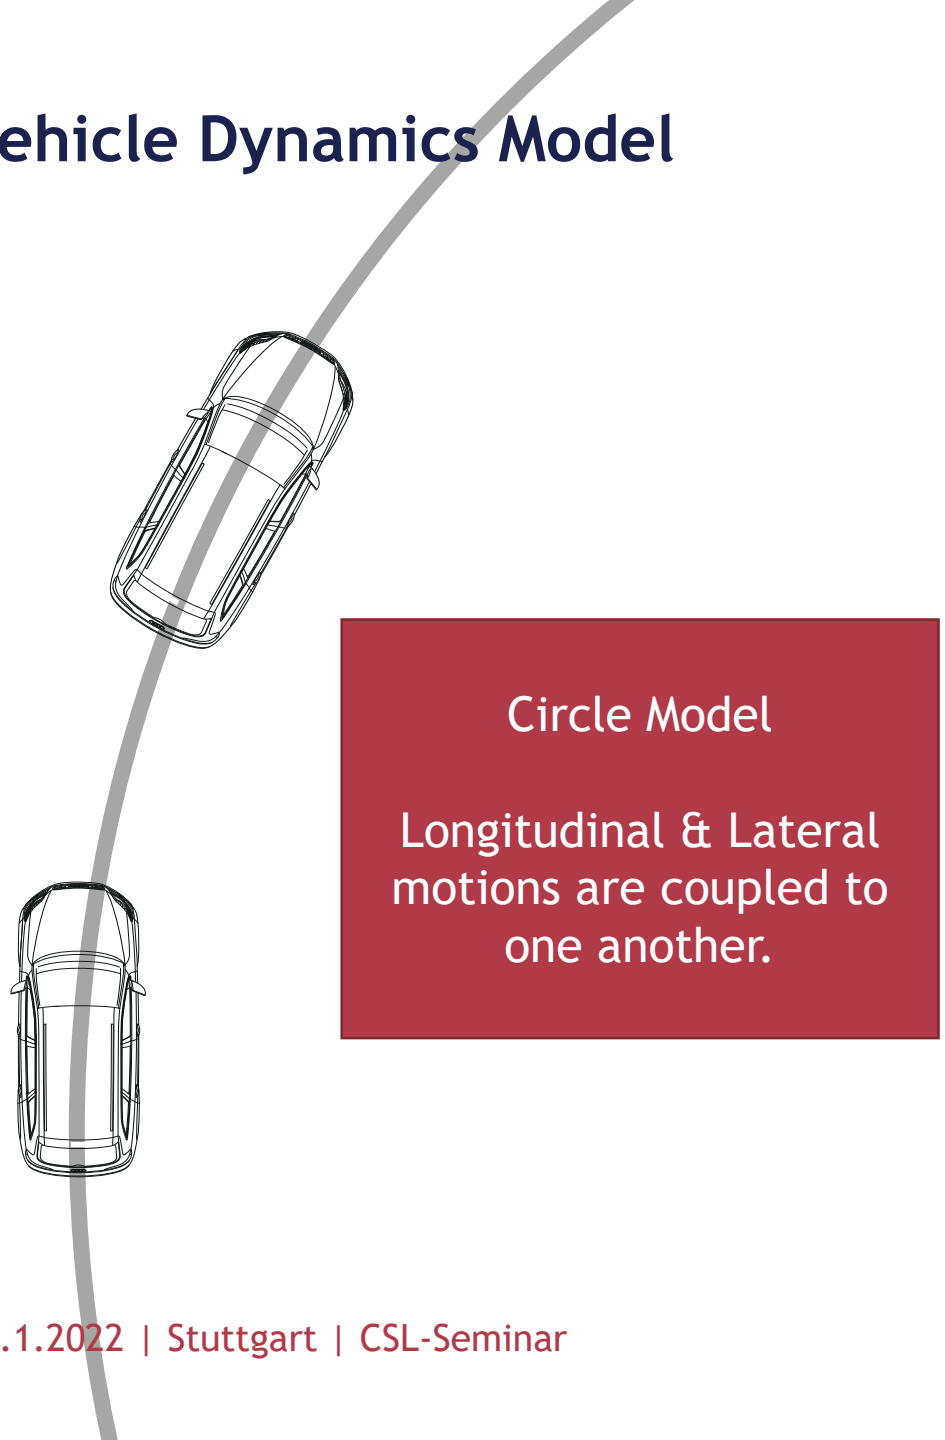

Circle Model

Longitudinal & Lateral motions are coupled to one another.

Representation:
Structural Causal Model

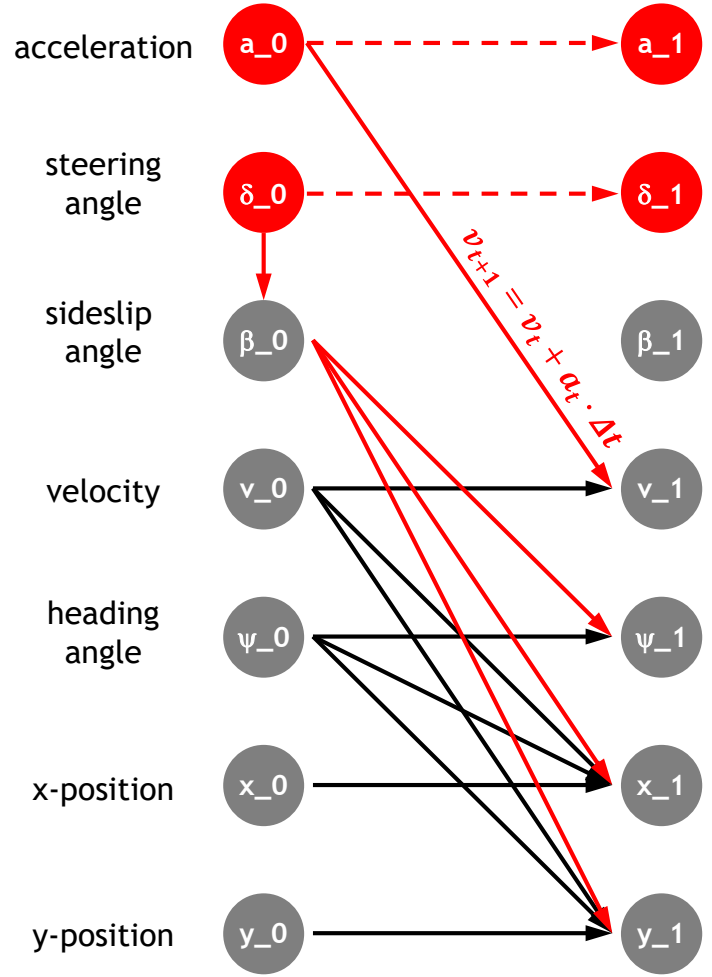

Dreamer v2 - Learning Flow with Causal Model

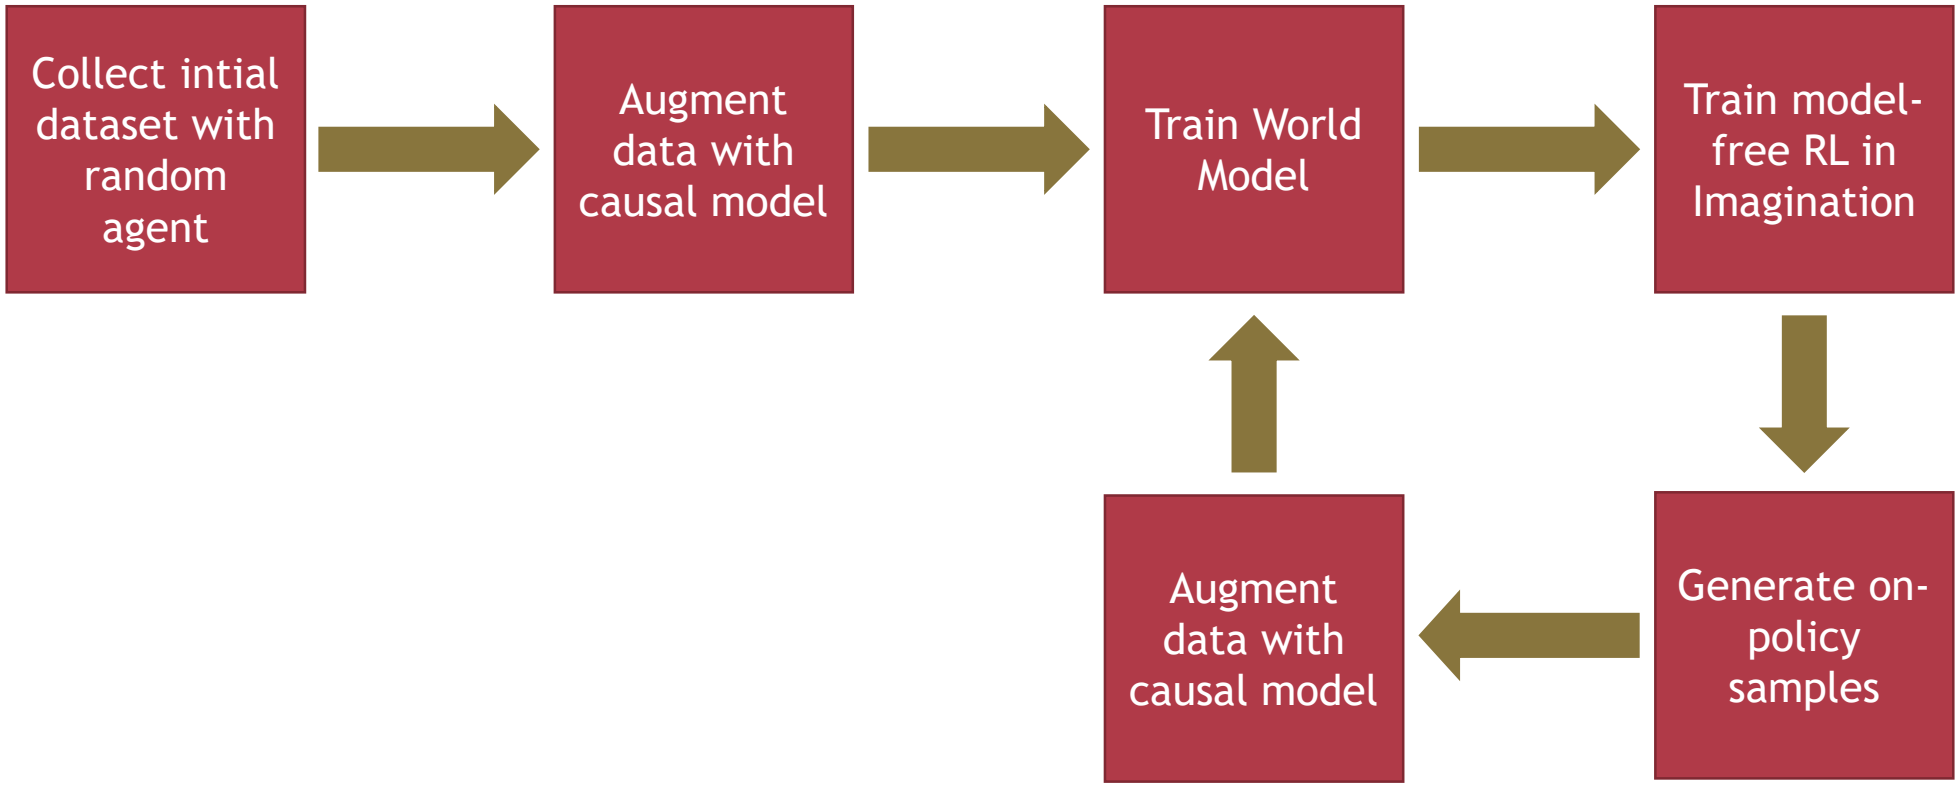

Knowledge Integration Concept

Goal:

- Improve lane change control

Basic idea:

- Use counterfactuals to generate/augment

Starting points:

- Counterfactually-Guided Policy Search [1]
- RL on latent space for autonomous driving [2]

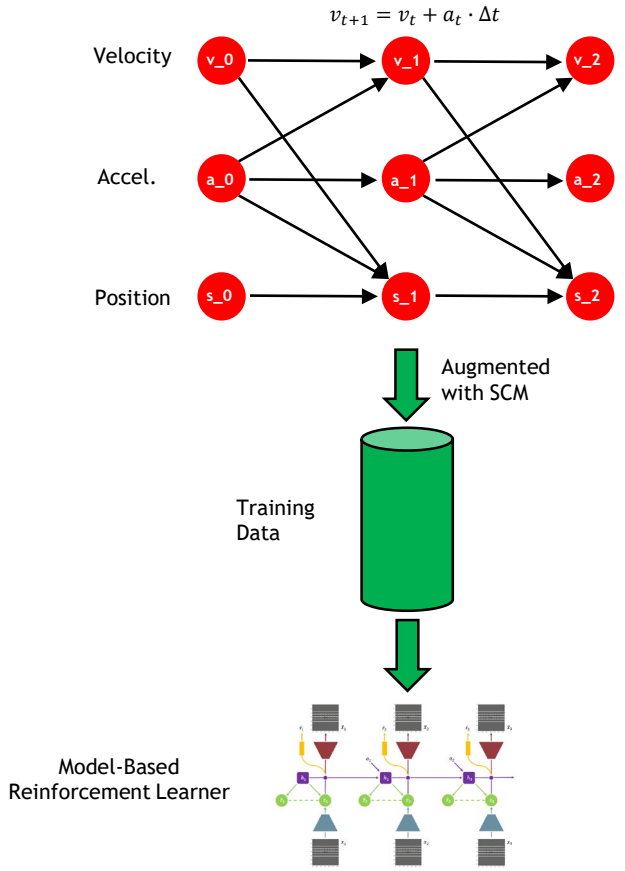

[1] Buesing *et al.*, „Woulda, Coulda, Shoulda: Counterfactually-Guided Policy Search“, *arXiv:1811.06272*, Nov. 2018, http://arxiv.org/abs/1811.06272
[2] Chen et al., „Interpretable End-to-end Urban Autonomous Driving with Latent Deep Reinforcement Learning”, *arXiv:2001.08726*, Jun. 2020, http://arxiv.org/abs/2001.08726

Dreamer v2

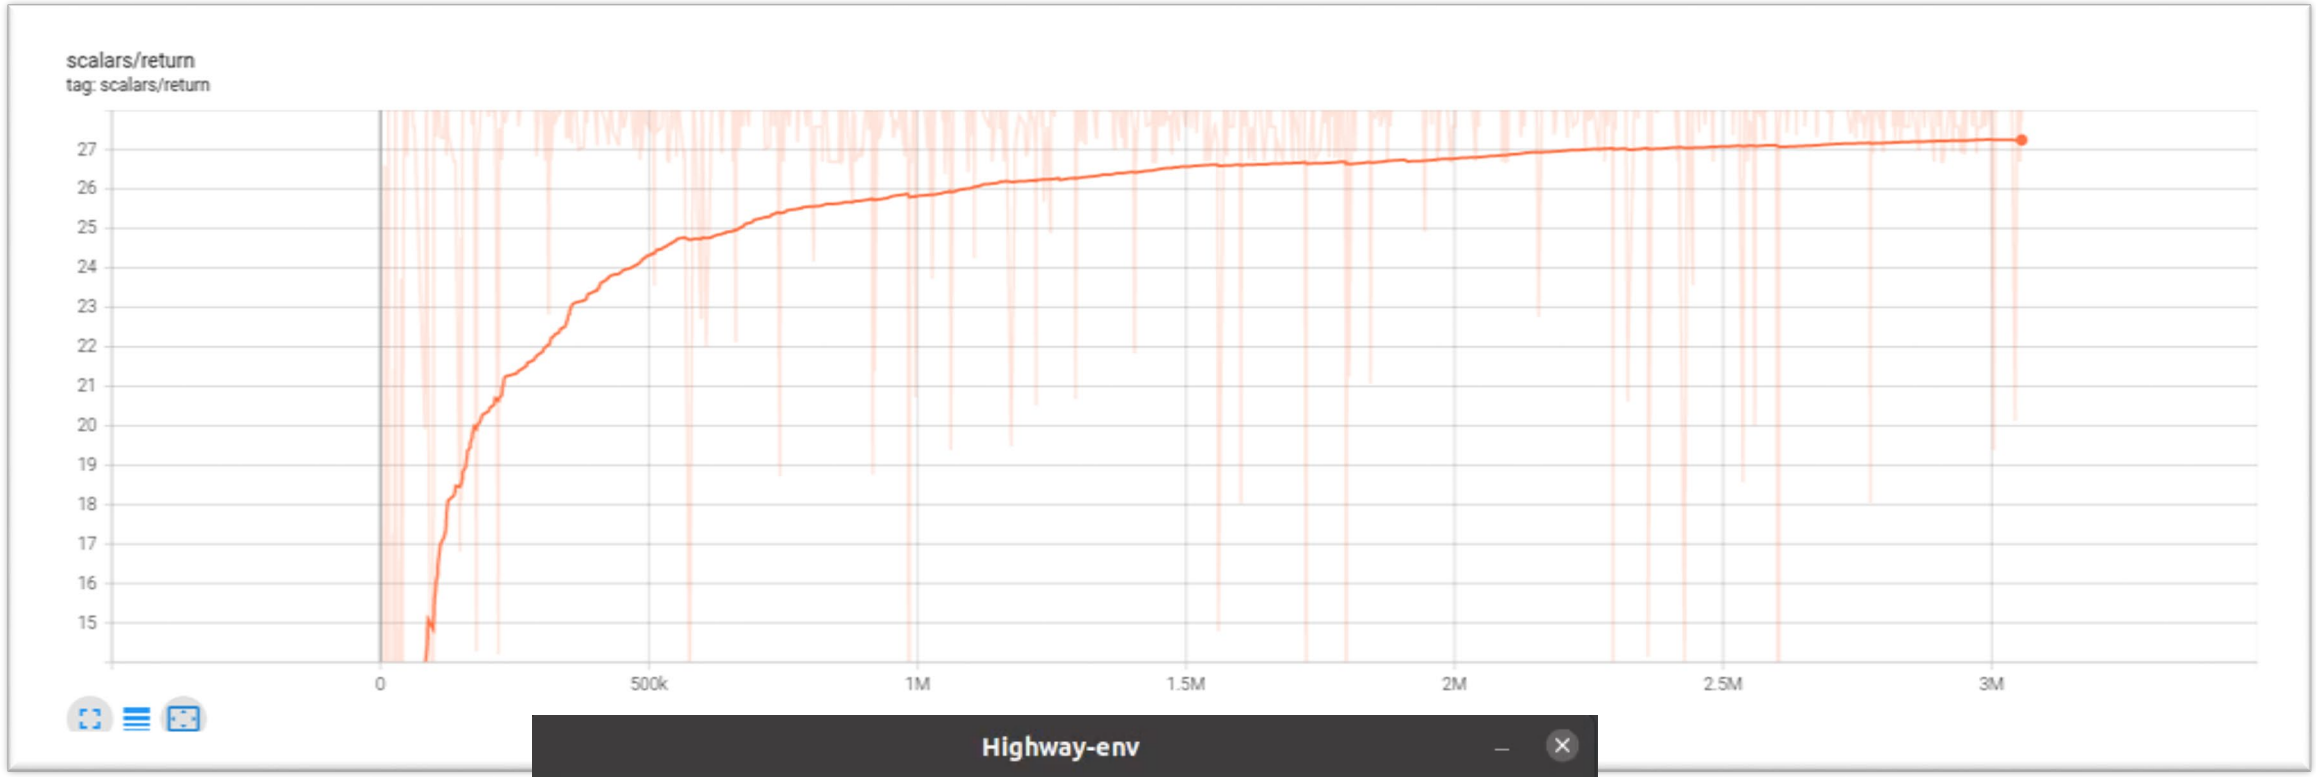
scalars/return
tag: scalars/return
27
26
25
24
23
22
21
20
19
18
17
16
15
0
500k
1M
1.5M
2M
2.5M
3M


Highway-env

Conclusion

- EFS in KI Wissen
- Recent Achievements in RL
- Dreamer v2
- Strucural Causal Models
- Concept on causality-driven RL

KI WISSEN

Automotive AI Powered by Knowledge

Dr. Stefan Rudolph (stefan.rudolph@efs-auto.com)

Dr. Tobias Latka (tobias.latka@efs-auto.com)

KI Wissen is a project of the KI Familie. It was initiated and developed by the VDA Leitinitiative autonomous and connected driving and is funded by the Federal Ministry for Economic Affairs and Energy.

KI FAMILIE

Gefördert durch:

Bundesministerium für Wirtschaft und Energie

aufgrund eines Beschlusses des Deutschen Bundestages

ki-wissen.vdali.de

@KI_Familie

KI Familie