RECURRENT KALMAN NETWORKS

Factorized Inference in High-Dimensional Deep Feature Spaces

Philipp Becker[1 2 3] Harit Pandya[4] Gregor Gebhardt[1] Chen Zhao[5] James Taylor[6] Gerhard Neumann[4 2 3]

1: Computational Learning for Autonomous Systems, TU Darmstadt, Darmstadt, Germany
2: Bosch Center for Artificial Intelligence, Renningen, Germany
3: University of Tübingen, Tübingen, Germany
4: Lincoln Center for Autonomous Systems, University of Lincoln, Lincoln, UK
5: Extreme Robotics Lab, University of Birmingham, Birmingham, UK
6: Engineering Department, Lancaster University, Lancaster, UK

Motivation

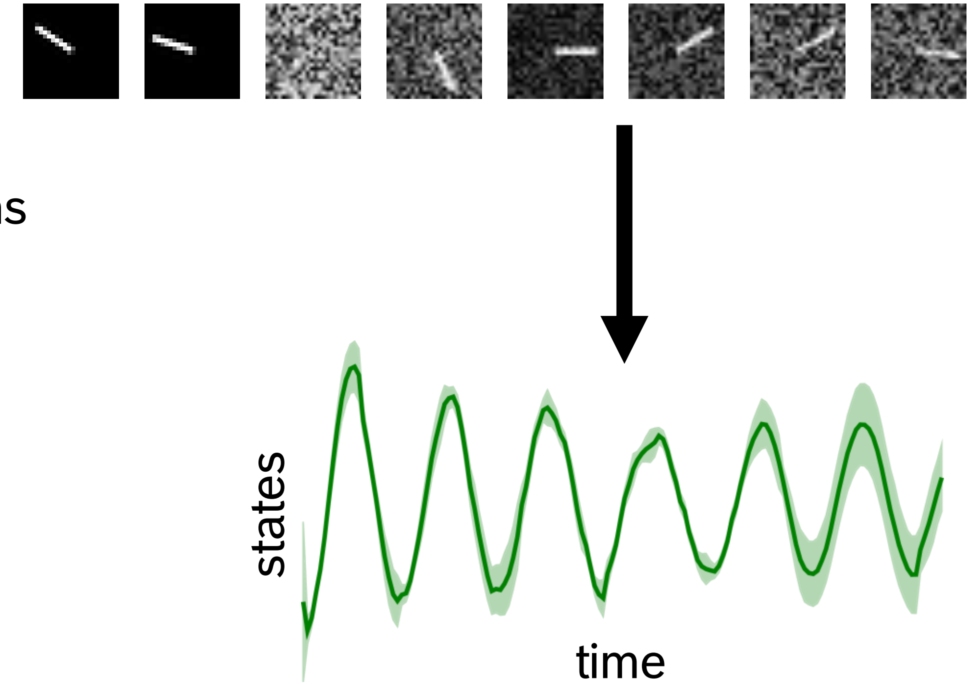

Goal: State estimation from high dimensional observations

- Filtering
- Prediction

Motivation

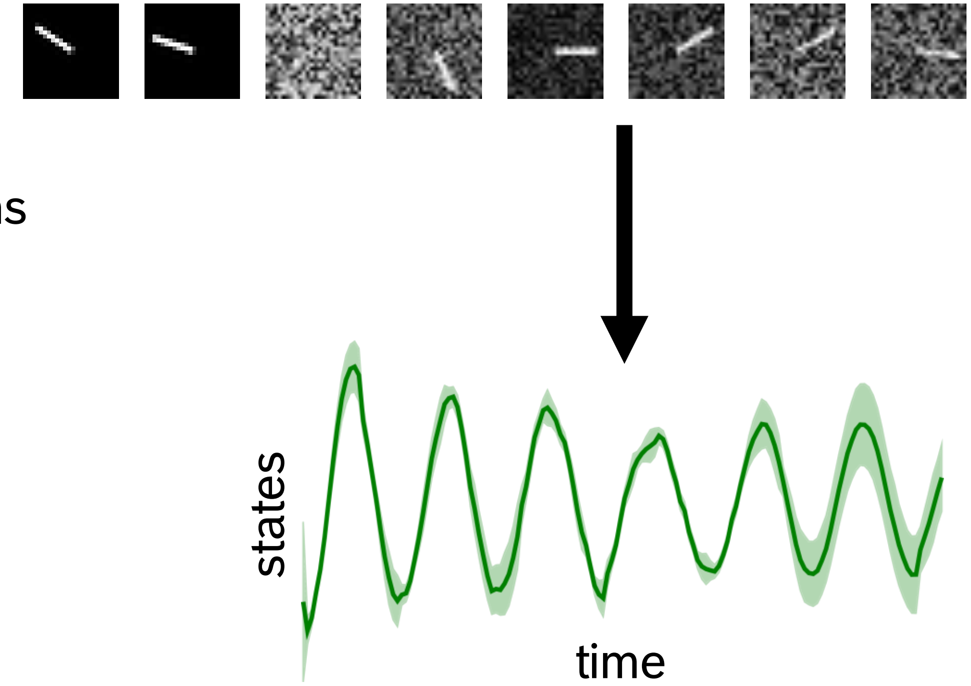

Goal: State estimation from high dimensional observations
- Filtering
- Prediction

Challenges:
- High dimensional observations
- Partially observable
- Nonlinear dynamics
- Uncertainty

Motivation

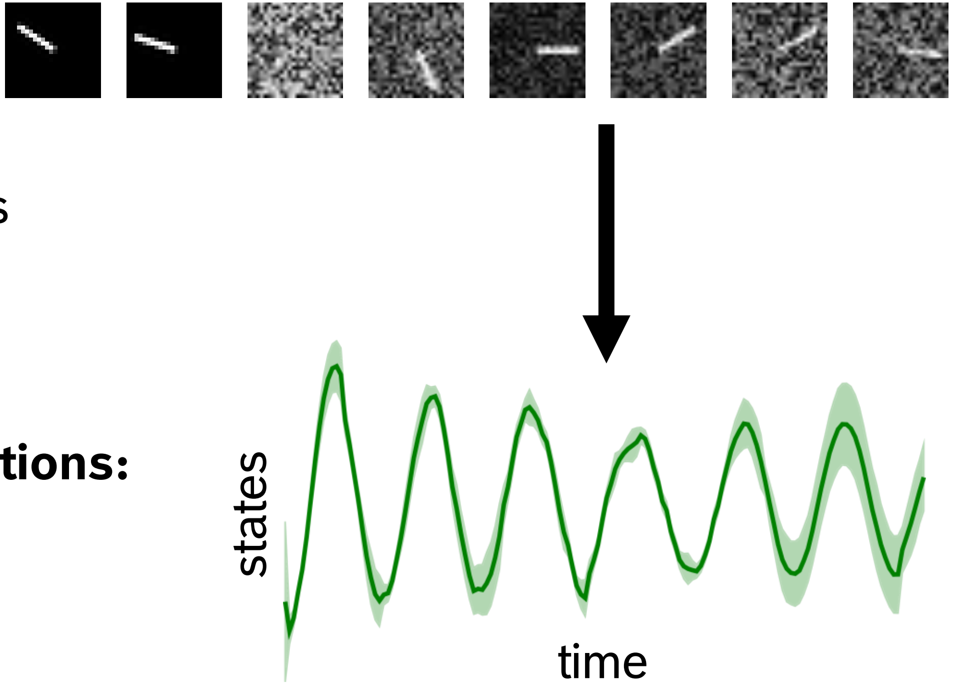

Goal: State estimation from high dimensional observations

- ➢ Filtering
- ➢ Prediction

Challenges:

- ➢ High dimensional observations
- ➢ Partially observable
- ➢ Nonlinear dynamics
- ➢ Uncertainty

(Deep Learning) Solutions:

- ✓ CNNs
- ✓ RNNs
- x Variational Inference (approximation errors)

BOSCH

Motivation

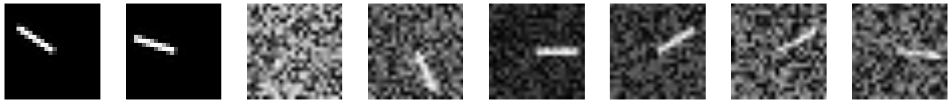

Goal: State estimation from high dimensional observations

- Filtering
- Prediction

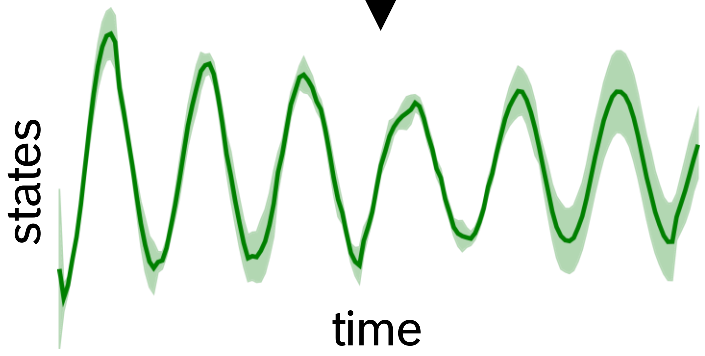

Challenges:

- High dimensional observations
- Partially observable
- Nonlinear dynamics
- Uncertainty

(Deep Learning) Solutions:

✓ CNNs

✓ RNNs

x Variational Inference (approximation errors)

How can we propagate uncertainty through RNNs without approximations?

→ Recurrent Kalman Networks (RKN): Recurrent cell based on Kalman filter

Overview

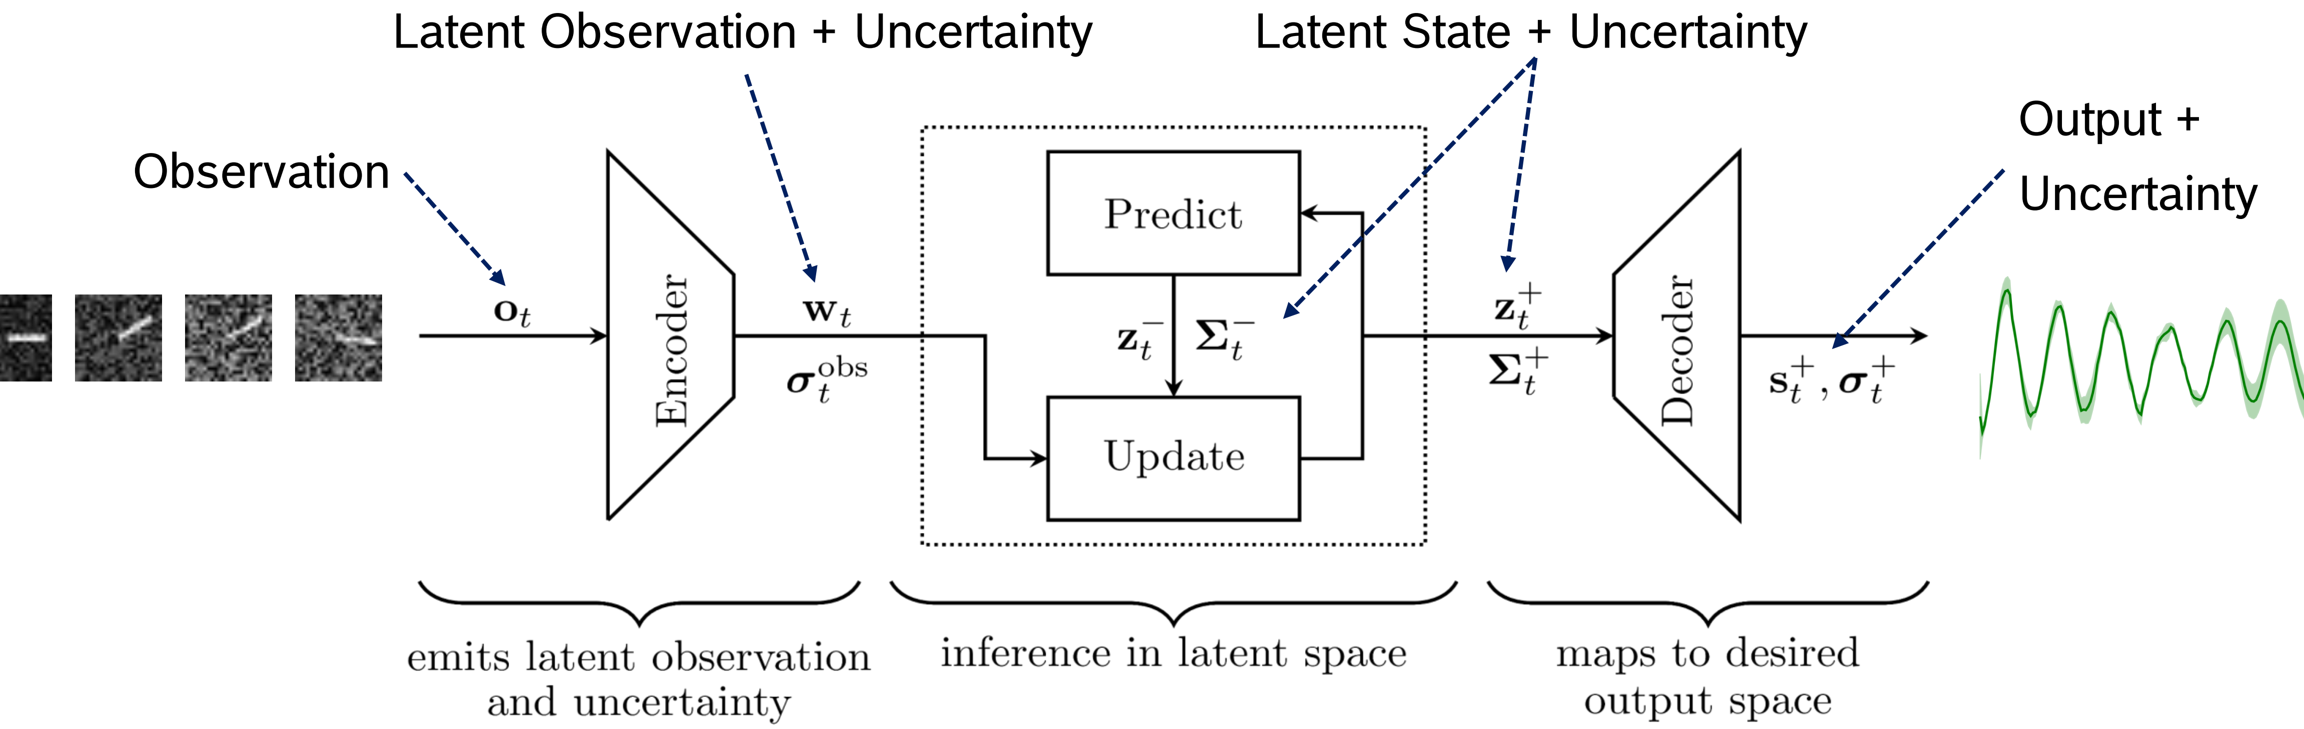

BOSCH

Overview

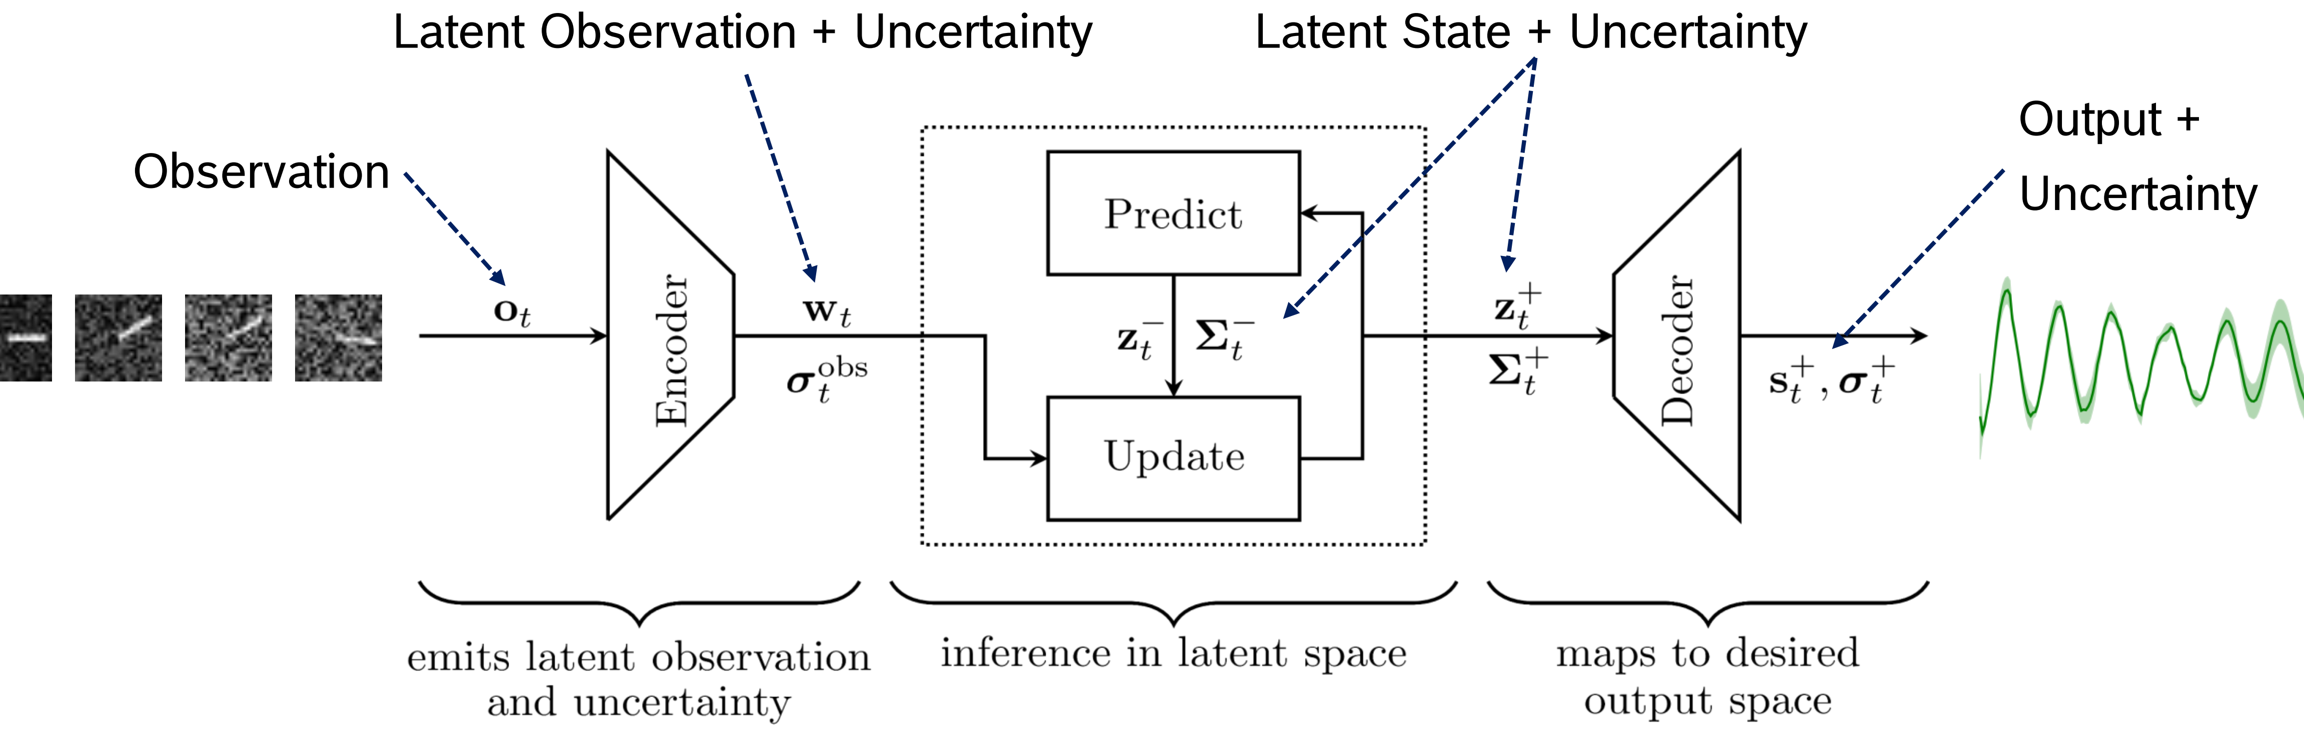

Make backpropagation through Kalman filter feasible?

- Locally linear transition models, even for highly nonlinear systems
- High dimensional latent spaces
- Factorized state representation to avoid expensive and unstable matrix inversions

BOSCH

Factorized State Representation

Observation Model

- $\mathbf{H} = \begin{pmatrix} \mathbf{I}_m & \mathbf{0}_{m \times m} \end{pmatrix}$
- Splits latent state
    1. Observable part
    2. Memory part

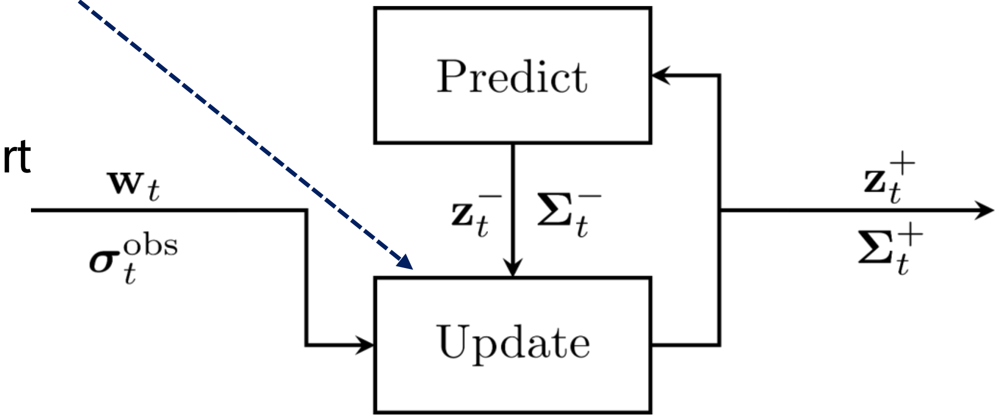

Factorized State Representation

Observation Model

- $\mathbf{H} = \begin{pmatrix} \mathbf{I}_m & \mathbf{0}_{m \times m} \end{pmatrix}$
- Splits latent state
  1. Observable part
  2. Memory part

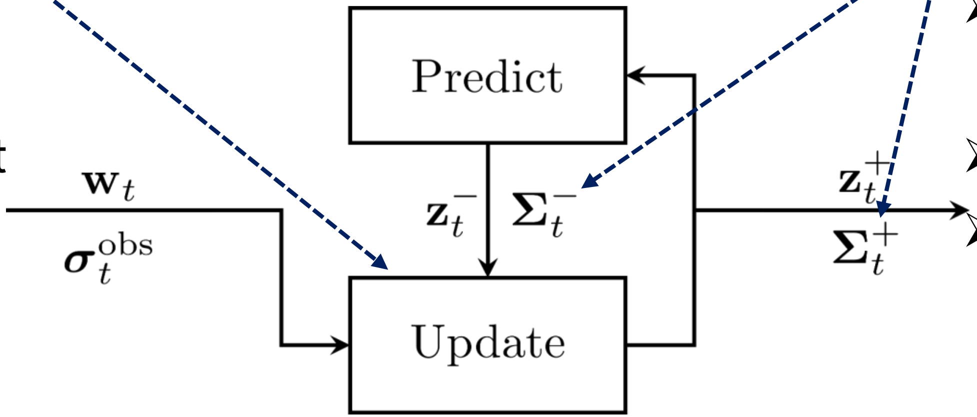

Factorized Representation

- $\boldsymbol{\Sigma} = \begin{pmatrix} \boldsymbol{\sigma}_u & \boldsymbol{\sigma}_s \\ \boldsymbol{\sigma}_s & \boldsymbol{\sigma}_l \end{pmatrix}$
- $\boldsymbol{\sigma}_u, \boldsymbol{\sigma}_l, \boldsymbol{\sigma}_s$ diagonal matrices
- $\boldsymbol{\sigma}_s$ correlates parts

Factorized State Representation

Observation Model

- $\mathbf{H} = \begin{pmatrix} \mathbf{I}_m & \mathbf{0}_{m \times m} \end{pmatrix}$
- Splits latent state
  1. Observable part
  2. Memory part

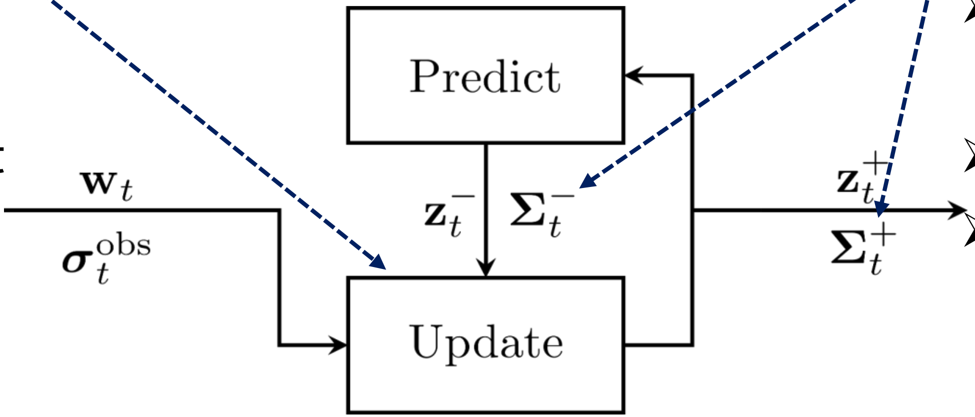

Factorized Representation

- $\boldsymbol{\Sigma} = \begin{pmatrix} \boldsymbol{\sigma}_u & \boldsymbol{\sigma}_s \\ \boldsymbol{\sigma}_s & \boldsymbol{\sigma}_l \end{pmatrix}$
- $\boldsymbol{\sigma}_u, \boldsymbol{\sigma}_l, \boldsymbol{\sigma}_s$ diagonal matrices
- $\boldsymbol{\sigma}_s$ correlates parts

Results in simplified Kalman Update

- No matrix inversion
- Instead only pointwise operations

$\}$ Makes inference and back-propagation feasible

- Assumptions not restrictive since latent space is learned

Quad Link Pendulum

- State (4 joint angles + velocity)
- Highly nonlinear dynamics
- Links occlude each other
- Estimate joint angles of all 4 links
- Observations: 48x48 pixel images

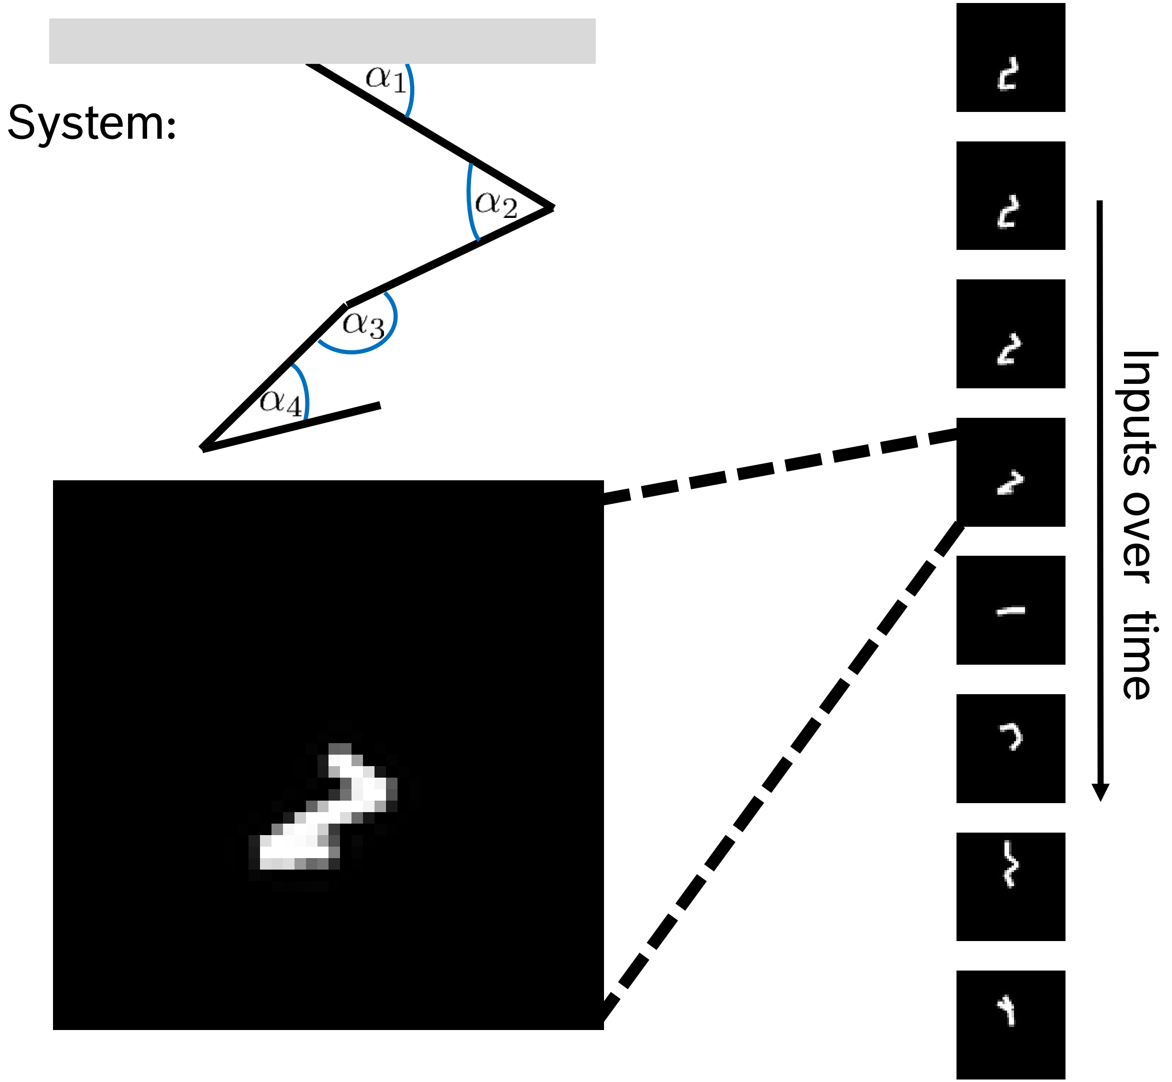

Quad Link Pendulum

- State (4 joint angles + velocity)
- Highly nonlinear dynamics
- Links occlude each other
- Estimate joint angles of all 4 links
- Observations: 48x48 pixel images

	RKN	LSTM	GRU
Log Likelihood	**14.534**	11.960	10.346
RMSE	**0.103**	0.118	0.121

- Significantly better uncertainty estimate (higher log-likelihood)
- Better prediction (smaller RMSE)

System:

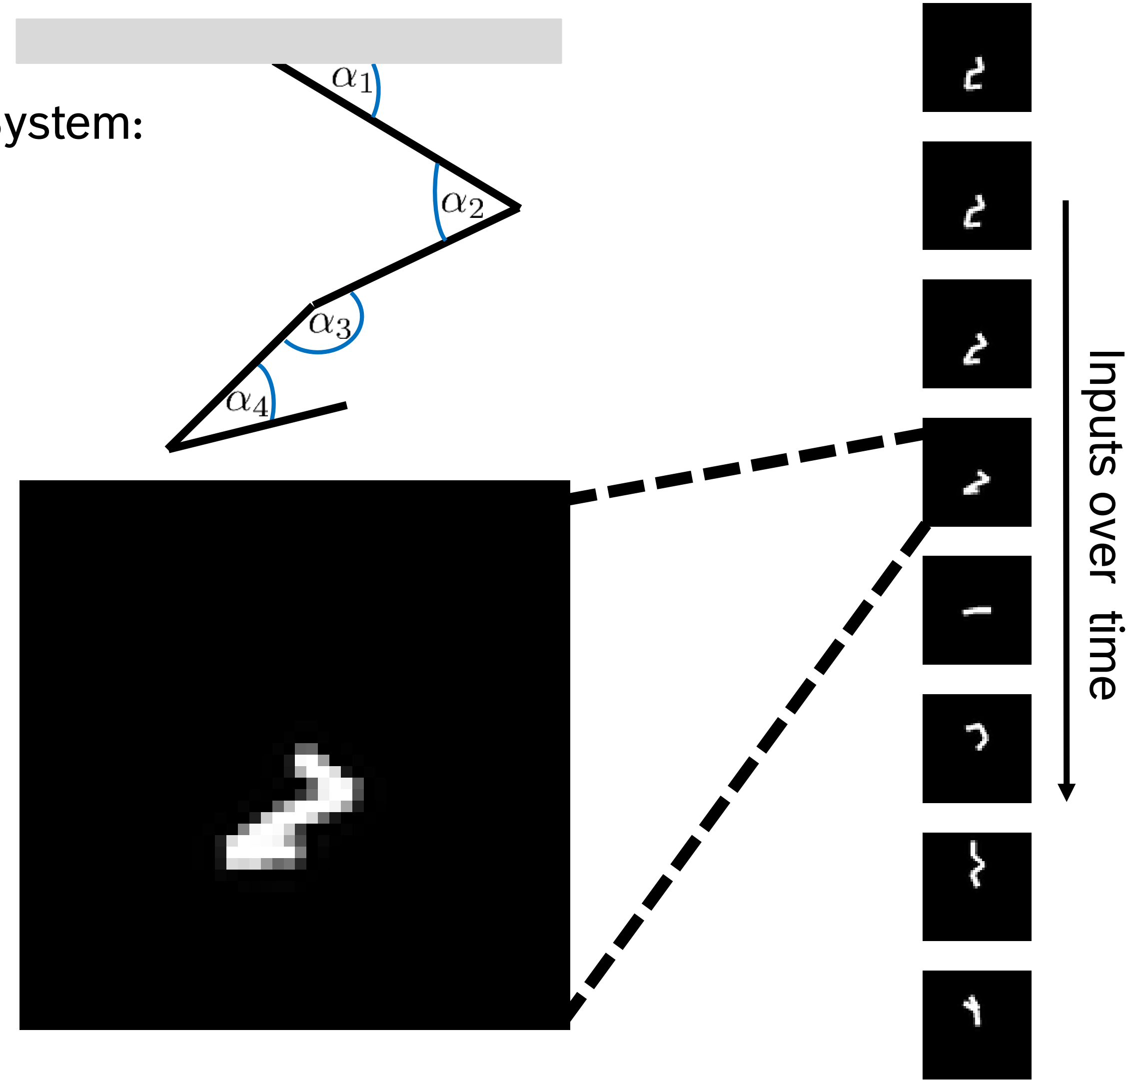

BOSCH

Summary & Conclusion

Recurrent Kalman Networks…

- … scale to real world systems
- … allow direct state estimation from images
- … use uncertainty in a principled manner to handle noise
- … can be trained end-to-end without approximations

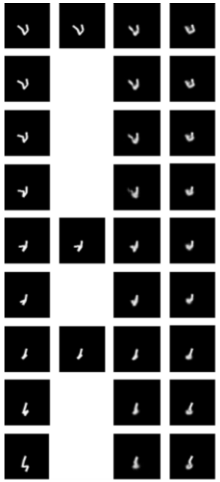

Additional Experiments

- Pendulum
- Image Imputation
- KITTI-Dataset for visual odometry

- Prediction for real pneumatic joint

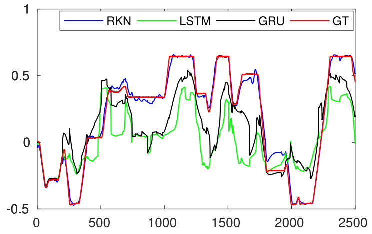

- Comparison to recent approaches
  - KVAE [1], E2C [2], Structured Inference Networks [3]
- Code available

[1]: Fraccaro et al. A disentangled recognition and nonlinear dynamics model for unsupervised learning. NIPS 2017
[2]: Watter et al. Embed to control: A locally linear latent dynamics model for control from raw images. NIPS, 2015
[3]: Krishnan et al. Structured inference networks for nonlinear state space models. AAAI, 2017