
Recurrent Kalman Networks: Factorized Inference in High-Dimensional Deep Feature Spaces

Philipp Becker^{1,2,3} Harit Pandya⁴ Gregor Gebhardt¹ Cheng Zhao⁵ James Taylor⁶ Gerhard Neumann^{4,2,3}

Abstract

In order to integrate uncertainty estimates into deep time-series modelling, Kalman Filters (KFs) (Kalman et al., 1960) have been integrated with deep learning models, however, such approaches typically rely on approximate inference techniques such as variational inference which makes learning more complex and often less scalable due to approximation errors. We propose a new deep approach to Kalman filtering which can be learned directly in an end-to-end manner using backpropagation without additional approximations. Our approach uses a high-dimensional factorized latent state representation for which the Kalman updates simplify to scalar operations and thus avoids hard to backpropagate, computationally heavy and potentially unstable matrix inversions. Moreover, we use locally linear dynamic models to efficiently propagate the latent state to the next time step. The resulting network architecture, which we call *Recurrent Kalman Network* (RKN), can be used for any time-series data, similar to a LSTM (Hochreiter & Schmidhuber, 1997) but uses an explicit representation of uncertainty. As shown by our experiments, the RKN obtains much more accurate uncertainty estimates than an LSTM or Gated Recurrent Units (GRUs) (Cho et al., 2014) while also showing a slightly improved prediction performance and outperforms various recent generative models on an image imputation task.

1. Introduction

State-estimation in unstructured environments is a very challenging task as observations or measurements of the environment are often high-dimensional and only provide partial information about the state. Images are a good example: Even for low resolution, the number of pixels can quickly exceed tens or hundreds of thousands and it is impossible to obtain any information about the dynamics, such as velocities, from a single image. Additionally, the observations may be noisy or may not contain useful information for the task at hand. Such noise can, for example, be introduced by poor illumination or motion blur and occlusions can prevent us from observing some or all relevant aspects of the scene. In addition to state estimation, it is also often desirable to predict future states or observations, for example, in order to assess the consequences of future actions. To this end, an initial estimate of the current state is necessary which again has to be inferred from observations. In such environments, we typically also have to deal with high uncertainties in the state estimates. Being able to model this uncertainty is crucial in many decision making scenarios, e.g., if we need to decide to perform an action now or wait until more information about the scene is available.

Deep learning models have been very successful for time-series modelling in unstructured environments. Classical models such as LSTMs (Hochreiter & Schmidhuber, 1997) or GRUs (Cho et al., 2014) perform well but fail to capture the uncertainty of the state estimate. Recent probabilistic deep learning approaches have used the Kalman Filter (KF) as a tool to integrate uncertainty estimates into deep time-series modelling (Haarnoja et al., 2016; Watter et al., 2015; Archer et al., 2015; Fraccaro et al., 2017; Krishnan et al., 2017). These approaches use the KF to perform inference in a low-dimensional (latent) state space that is typically defined by a deep encoder. However, using KF in such a state space comes with two main limitations. In order to be usable for non-linear dynamics, we have to introduce approximations such as the extended KF (Haarnoja et al., 2016) and variational inference methods (Krishnan et al., 2017; Fraccaro et al., 2017). Moreover, the KF equations require computationally expensive matrix inversions that are hard to scale to high dimensional latent spaces for more

¹Computational Learning for Autonomous Systems, TU Darmstadt, Darmstadt, Germany. ²Bosch Center for Artificial Intelligence, Renningen, Germany. ³University of Tübingen, Tübingen, Germany. ⁴Lincoln Center for Autonomous Systems, University of Lincoln, Lincoln, UK. ⁵Extreme Robotics Lab, University of Birmingham, Birmingham, UK. ⁶Engineering Department, Lancaster University, Lancaster, UK.. Correspondence to: Philipp Becker <philippbecker93@googlemail.com>.

complex systems and computationally demanding to fully backpropagate in an end-to-end manner. Most of these methods are implemented as (variational) auto-encoders and are therefore also limited to predicting future observations or imputing missing observations and can not be directly be applied to state estimation.

We introduce the *Recurrent Kalman Network*, an end-to-end learning approach for Kalman filtering and prediction. While Kalman filtering in the original state space requires approximations due to the non-linear models, the RKN uses a learned high-dimensional latent state representation that allows for efficient inference using locally linear transition models and a factorized belief state representation. Exploiting this representation allows us to avoid the expensive and numerically problematic matrix inversions involved in the KF equations.

Conceptually, this idea is related to kernel methods which use high-dimensional feature spaces to approximate non-linear functions with linear models (Gebhardt et al., 2017). However, in difference to kernel feature spaces, our feature space is given by a deep encoder and learned in an end-to-end manner.

The RKN can be used for any time-series data set for which LSTMs and GRUs are currently the state of the art. In contrast to those, the RKN uses an explicit representation of uncertainty which governs the gating between keeping the current information in memory or updating it with the current observation. While the RKN shows a slightly improved performance in terms of state estimation errors, both LSTMs and GRUs struggle with estimating the uncertainty of the prediction while the RKN can provide accurate uncertainty estimates. In relation to existing KF-based approaches, our approach can be used for state estimation as well as for generative tasks such as image imputation. We also show that we outperform state of the art methods on a complex image imputation task.

1.1. Related Work

Using encoders for time-series modelling of high-dimensional data such as images is a common approach. Such encoders can also be easily integrated with well known deep time-series models such as LSTMs (Hochreiter & Schmidhuber, 1997) or GRUs (Cho et al., 2014). These models are very effective but do not provide good uncertainty estimates as shown in our experiments.

Pixel to Torques (P2T) (Wahlström et al., 2015) employs an autoencoder to obtain low dimensional latent representations from images together with a transition model. They subsequently use the models to perform control in the latent space. Embed to Control (E2C) (Watter et al., 2015) can be seen as an extension of the previous approach with the dif-

ference that a variational autoencoder (Kingma & Welling, 2013) is used. However, both of these approaches are not recurrent and rely on observations which allow inferring the whole state from a single observation. They can therefore not deal with noisy or missing data.

The BackpropKF (Haarnoja et al., 2016) applies a CNN to estimate the observable parts of the true state given the observation. Similar to our approach, this CNN additionally outputs a covariance matrix indicating the model’s certainty about the estimate and allows the subsequent use of an (extended) Kalman filter with known transition model. In contrast, we let our model chose the feature space that is used for the inference such that locally linear models can be learned and the KF computations can be simplified due to our factorization assumptions.

Another family of approaches interprets encoder-decoder models as latent variable models that can be optimized efficiently by variational inference. They derive a corresponding lower bound and optimize it using the stochastic gradient variational Bayes approach (Kingma & Welling, 2013). Black Box Variational Inference (BB-VI) (Archer et al., 2015) proposes a structured Gaussian variational approximation of the posterior, which simplifies the inference step at the cost of maintaining a tri-diagonal covariance matrix of the full state. To circumvent this issue, Structured Inference Networks (SIN) (Krishnan et al., 2017) employ a flexible recurrent neural network to approximate the dynamic state update. Deep Variational Bayes Filters (DVBF) (Karl et al., 2016) integrate general Bayes filters into deep feature spaces while the Kalman Variational Autoencoder (KVAE) (Fracaro et al., 2017) employs the classical Kalman Filter and allows not only filtering but also smoothing. Variational Sequential Monte Carlo (VSMC) (Naesseth et al., 2017) uses particle filters instead, however, they are only learning the proposal function and are not working in a learned latent space. Disentangled Sequential Autoencoder (DSA) (Yingzhen & Mandt, 2018) explicitly partitions latent variables for separating static and dynamic content from a sequence of observations however, they do not model state space noise. Yet, these models cannot be directly used for state estimation as they are formulated as generative models of the observations without the notion of the real state of the system. Moreover, the use of variational inference introduces an additional approximation that might affect the performance of the algorithms.

Recently, Rangapuram et al. (2018) introduced Deep State Space Models for Time Series Forecasting (DSSM). They employ a recurrent network to learn the parameters of a time varying linear state space model. The emissions of that state space model are the model’s final output, i.e. there is no decoder. While this makes the likelihood of the targets given the predicted state space model parameters tractable it limits

Table 1. Qualitative comparison of our approach to recent related work.

	scale-able	state est.	uncertainty	noise	direct opt.
LSTM	✓	✓	×/✓	✓	✓
GRU	✓	✓	×/✓	✓	✓
P2T	✓	✓	×/✓	×	✓
E2C	✓	×	✓	×	×
BB-VI	×	×	✓	✓	×
SIN	✓	×	✓	✓	×
KVAE	×	×	✓	✓	×
DVBF	✓	×	✓	✓	×
VSMC	✓	×	✓	✓	×
DSA	✓	×	✓	×	×
DSSM	×	✓ (1D)	✓	×	✓
PRSSM	×	✓	✓	✓	×
RKN	✓	✓	✓	✓	✓

them to simple latent spaces. Additionally, their approach is only derived for 1 dimensional output data. Probabilistic Recurrent State-Space Models (PRSSM) (Doerr et al., 2018) use Gaussian processes to capture state uncertainties, however, their approach is not demonstrated for high dimensional observations such as images.

A summary of all the approaches and their basic properties can be seen in Table 1. We compare whether the approaches are scalable to high dimensional latent states, whether they can be used for state estimation, whether they can provide uncertainty estimates, whether they can handle noise or missing data, and whether the objective can be optimized directly or via a lower bound. All probabilistic generative models rely on variational inference which optimizes a lower bound, which potentially affects the performance of the algorithms. P2T and E2C rely on the Markov assumption and therefore can not deal with noise or need very large window sizes. The approaches introduced in (Archer et al., 2015; Fraccaro et al., 2017; Haarnoja et al., 2016) directly use the Kalman update equations in the latent state, which limits these approaches to rather low dimensional latent states due to the expensive matrix inversions.

Traditional recurrent models such as LSTMs or GRUs can be trained directly by backpropagation through time and therefore typically yield very good performance but are lacking uncertainty estimates. However, as in our experiments, they can be artificially added. Our RKN approach combines the advantages of all methods above as it can be learned by direct optimization without the use of a lower bound and it provides a principled way of representing uncertainty inside the neural network.

2. Factorized Inference in Deep Feature Spaces

Lifting the original input space to a high-dimensional feature space where linear operations are feasible is a common approach in machine learning, e.g., in kernel regression and SVMs. The *Recurrent Kalman Network* (RKN) transfers this idea to state estimation and filtering, i.e., we learn a high dimensional deep feature space that allows for efficient inference using the Kalman update equations even for complex systems with high dimensional observations. To achieve this, we assume that the belief state representation can be factorized into independent Gaussian distributions as described in the following sections.

2.1. Latent Observation and State Spaces

The RKN encoder learns a mapping to a high-dimensional latent observation space $\mathcal{W} = \mathbb{R}^m$. The encoder also outputs a vector of uncertainty estimates σ_t^{obs} , one for each entry of the latent observation \mathbf{w}_t . The latent state space $\mathcal{Z} = \mathbb{R}^n$ of the RKN is related to the observation space \mathcal{W} by the linear latent observation model $\mathbf{H} = \begin{bmatrix} \mathbf{I}_m & \mathbf{0}_{m \times (n-m)} \end{bmatrix}$, i.e., $\mathbf{w} = \mathbf{H}\mathbf{z}$ with $\mathbf{w} \in \mathcal{W}$ and $\mathbf{z} \in \mathcal{Z}$. \mathbf{I}_m denotes the $m \times m$ identity matrix and $\mathbf{0}_{m \times (n-m)}$ denotes a $m \times (n-m)$ matrix filled with zeros.

The idea behind this choice is to split the latent state vector \mathbf{z}_t into two parts, a vector \mathbf{p}_t for holding information that can directly be extracted from the observations and a vector \mathbf{m}_t to store information inferred over time, e.g., velocities. We refer to the former as the observation or upper part and the latter as the memory or lower part of the latent state. For an ordinary dynamical system and images as observations the former may correspond to positions while the latter corresponds to velocities. Clearly, this choice only makes sense for $m \leq n$ and in this work we assume $n = 2m$, i.e., for each observation unit p_i , we also represent a memory unit m_i that stores its velocity information.

We initialize \mathbf{z}_0^+ with an all zeros vector and Σ_0^+ with $10 \cdot \mathbf{I}$. In practice, it is beneficial to normalize \mathbf{w}_t since the statistics of noisy and noise free images differ considerably.

2.2. The Transition Model

To obtain a locally linear transition model we learn K constant transition matrices $\mathbf{A}^{(k)}$ and combine them using state dependent coefficients $\alpha^{(k)}(\mathbf{z}_t)$, i.e.,

$$\mathbf{A}_t = \sum_{k=1}^K \alpha^{(k)}(\mathbf{z}_t) \mathbf{A}^{(k)}.$$

A small neural network with softmax output is used to learn $\alpha^{(k)}$. Similar approaches are used in (Fraccaro et al., 2017; Karl et al., 2016).

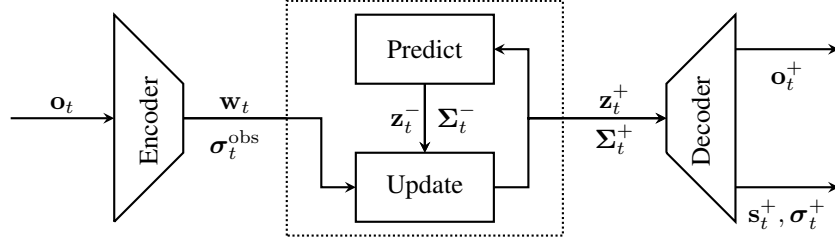


Figure 1. The *Recurrent Kalman Network*. An encoder network extracts latent features \mathbf{w}_t from the current observation \mathbf{o}_t . Additionally, it emits an estimate of the uncertainty in the features via the variance σ_t^{obs} . The transition model \mathbf{A}_t is used to predict the current latent prior $(\mathbf{z}_t^-, \Sigma_t^-)$ using the last posterior $(\mathbf{z}_{t-1}^+, \Sigma_{t-1}^+)$ and subsequently update the prior using the latent observation $(\mathbf{w}_t, \sigma_t^{\text{obs}})$. As we use a factorized representation of Σ_t , the Kalman update simplifies to scalar operations. The current latent state \mathbf{z}_t consists of the observable units \mathbf{p}_t as well as the corresponding memory units \mathbf{m}_t . Finally, a decoder produces either $(\mathbf{s}_t^+, \sigma_t^+)$, a low dimensional observation and an element-wise uncertainty estimate, or \mathbf{o}_t^+ , a noise free image.

Using a dense transition matrix in high-dimensional latent spaces is not feasible as it contains too many parameters and causes numerical instabilities and overfitting, as preliminary experiments showed. Therefore, we design each $\mathbf{A}^{(k)}$ to consist of four band matrices

$$\mathbf{A}^{(k)} = \begin{bmatrix} \mathbf{B}_{11}^{(k)} & \mathbf{B}_{12}^{(k)} \\ \mathbf{B}_{21}^{(k)} & \mathbf{B}_{22}^{(k)} \end{bmatrix}$$

with bandwidth b . Here, $\mathbf{B}_{11}^{(k)}, \mathbf{B}_{12}^{(k)}, \mathbf{B}_{21}^{(k)}, \mathbf{B}_{22}^{(k)} \in \mathbb{R}^{m \times m}$ since we assume $n = 2m$. This choice reduces the number of parameters while not affecting performance since the network is free to choose the state representation.

We assume the covariance of the transition noise to be diagonal and denote the vector containing the diagonal values by σ^{trans} . The noise is learned and independent of the state. Moreover, it is crucial to correctly initialize the transition matrix. Initially, the transition model should focus on copying the encoder output so that the encoder can learn how to extract good features if observations are available and useful. Additionally, it is crucial that \mathbf{A} does not yield an unstable system. We choose $\mathbf{B}_{11}^{(k)} = \mathbf{B}_{22}^{(k)} = \mathbf{I}$ and $\mathbf{B}_{12}^{(k)} = -\mathbf{B}_{21}^{(k)} = 0.2 \cdot \mathbf{I}$.

2.3. Factorized Covariance Representation

Since the RKN learns high-dimensional representations, we can not work with the full state covariance matrices Σ_t^+ and Σ_t^- . We can also not fully factorize the state covariances by diagonal matrices as this neglects the correlation between the memory and the observation parts. As the memory part is excluded from the observation model \mathbf{H} , the Kalman update step would not change the memory units nor their uncertainty if we would only use a diagonal covariance matrix Σ_t^+ . Hence, for each observation unit p_i , we compute the covariance with its corresponding memory unit m_i . All the other covariances are neglected. This might be a crude approximation for many systems, however, as our network

is free to choose its own state representation it can find a representation where such a factorization works well in practice. Thus, we use matrices of the form

$$\Sigma_t = \begin{bmatrix} \Sigma_t^u & \Sigma_t^s \\ \Sigma_t^s & \Sigma_t^l \end{bmatrix}$$

where each of $\Sigma_t^u, \Sigma_t^s, \Sigma_t^l \in \mathbb{R}^{m \times m}$ is a diagonal matrix. Again, we denote the vectors containing the diagonal values by σ_t^u, σ_t^s and σ_t^l .

2.4. Factorized Inference in the Latent Space

Inference in the latent state space can now be implemented, similar to a standard KF, using a prediction and an observation update.

Prediction Update. Equivalently to the classical Kalman Filter, the next prior $(\mathbf{z}_{t+1}^-, \Sigma_{t+1}^-)$ is obtained from the current posterior $(\mathbf{z}_t^+, \Sigma_t^+)$ by

$$\mathbf{z}_{t+1}^- = \mathbf{A}_t \mathbf{z}_t^+ \quad \text{and} \quad \Sigma_{t+1}^- = \mathbf{A}_t \Sigma_t^+ \mathbf{A}_t^T + \mathbf{I} \cdot \sigma^{\text{trans}}. \quad (1)$$

However, the special structure of the covariances enables us to significantly simplify the equation for the covariance. While being straight forward, the full derivations are rather lengthy, thus, we refer to the supplementary material where the equations are given in Eqs. 10,11 and 12.

Observation Update. Next, the prior is updated using the latent observation $(\mathbf{w}_t, \sigma_t^{\text{obs}})$. Similar to the state, we split the Kalman gain matrix \mathbf{Q}_t into an upper \mathbf{Q}_t^u and a lower part \mathbf{Q}_t^l . Both \mathbf{Q}_t^u and \mathbf{Q}_t^l are squared matrices. Due to the simple latent observation model $\mathbf{H} = [\mathbf{I}_m \quad \mathbf{0}_{m \times (n-m)}]$ and the factorized covariances all off-diagonal entries of \mathbf{Q}_t^u and \mathbf{Q}_t^l are zero and we can work with vectors representing the diagonals, i.e., \mathbf{q}_t^u and \mathbf{q}_t^l . Those are obtained by

$$\mathbf{q}_t^u = \sigma_t^{u,-} \oslash (\sigma_t^{u,-} + \sigma_t^{\text{obs}}) \quad (2)$$

$$\text{and} \quad \mathbf{q}_t^l = \sigma_t^{s,-} \oslash (\sigma_t^{u,-} + \sigma_t^{\text{obs}}), \quad (3)$$

where \odot denotes an elementwise vector division. With this the update equation for the mean simplifies to

$$\mathbf{z}_t^+ = \mathbf{z}_t^- + \begin{bmatrix} \mathbf{q}_t^u \\ \mathbf{q}_t^l \end{bmatrix} \odot \begin{bmatrix} \mathbf{w}_t - \mathbf{z}_t^{u,-} \\ \mathbf{w}_t - \mathbf{z}_t^{l,-} \end{bmatrix} \quad (4)$$

where \odot denotes the elementwise vector product. The update equations for the individual covariance parts are given by

$$\sigma_t^{u,+} = (\mathbf{1}_m - \mathbf{q}_t^u) \odot \sigma_t^{u,-}, \quad (5)$$

$$\sigma_t^{s,+} = (\mathbf{1}_m - \mathbf{q}_t^u) \odot \sigma_t^{s,-} \quad (6)$$

$$\text{and } \sigma_t^{l,+} = \sigma_t^{l,-} - \mathbf{q}_t^l \odot \sigma_t^{s,-}, \quad (7)$$

where $\mathbf{1}_m$ denotes the m dimensional vector consisting of ones. Again, we refer to the supplementary material for a more detailed derivations.

Besides avoiding the matrix inversion in the Kalman gain computation, the factorization of the covariance matrices reduces the total amount of numbers to store per matrix from n^2 to $3m$. Additionally, working with σ^s makes it trivial to ensure that the symmetry and positive definiteness of the state covariance are not affected by numerical issues.

2.5. Loss and Training

We consider two different potential output distributions, Gaussian distributions for estimating low dimensional observations and Bernoulli distributions for predicting high dimensional observations such as images. Both distributions are computed by decoders that use the current latent state estimate \mathbf{z}_t^+ as well its uncertainty estimates $\sigma_t^{u,+}$, $\sigma_t^{s,+}$, $\sigma_t^{l,+}$.

Inferring states. Let $\mathbf{s}_{1:T}$ be the ground truth sequence with dimension D_s , the Gaussian log-likelihood for a single sequence is then computed as

$$\mathcal{L}(\mathbf{s}_{(1:T)}) = \frac{1}{T} \sum_{t=1}^T \log \mathcal{N} \left(\mathbf{s}_t \mid \text{dec}_\mu(\mathbf{z}_t^+), \text{dec}_\Sigma(\sigma_t^{u,+}, \sigma_t^{s,+}, \sigma_t^{l,+}) \right), \quad (8)$$

where $\text{dec}_\mu(\cdot)$ and $\text{dec}_\Sigma(\cdot)$ denote the parts of the decoder that are responsible for decoding the latent mean and latent variance respectively.

Inferring images. Let $\mathbf{o}_{1:T}$ be images with D_o pixels. The Bernoulli log-likelihood for a single sequence is then given by

$$\mathcal{L}(\mathbf{o}_{(1:T)}) = \frac{1}{T} \sum_{t=1}^T \sum_{i=0}^{D_o} o_t^{(i)} \log(\text{dec}_{o,i}(\mathbf{z}_t^+)) \dots + (1 - o_t^{(i)}) \log(1 - \text{dec}_{o,i}(\mathbf{z}_t^+)), \quad (9)$$

where $o_t^{(i)}$ is the i th pixel of the t th image. The pixels are in this case represented by grayscale values in the range of $[0; 1]$. The i th dimension of the decoder is denoted by $\text{dec}_{o,i}(\mathbf{z}_t^+)$, where we use a sigmoid transfer function as output units for the decoder.

Gradients are computed using (truncated) backpropagation through time (BPTT) (Werbos, 1990) and clipped. We optimize the objective using the Adam (Kingma & Ba, 2014) stochastic gradient descent optimizer with default parameters.

2.6. The Recurrent Kalman Network

The prediction and observation updates result in a new type of recurrent neural network, that we call *Recurrent Kalman Network*, which allows working in high dimensional state spaces while keeping numerical stability, computational efficiency and (relatively) low memory consumption.

Similar to the input gate in LSTMs (Hochreiter & Schmidhuber, 1997) and GRUs (Cho et al., 2014) the Kalman gain can be seen as a gate controlling how much the current observation influences the state. However, this gating explicitly depends on the uncertainty estimates of the latent state and observation and is computed in a principled manner. While the sparse transition models and factorization assumptions may seem restrictive, they allow stable and efficient computations in high dimensional spaces. Since the high dimensional representation is learned jointly with the dynamics, this can yield very powerful models, as shown in our experiments.

In comparison to LSTMs and GRUs the number of parameters is considerably smaller. For a fixed bandwidth b and number of basis matrices k it scales linearly in the state size for the RKN while it scales quadratically for LSTMs and GRUs.

Moreover, the RKN provides a principled method to deal with absent inputs by just omitting the update step and setting the posterior to the prior.

3. Evaluation and Experiments

A full listing of hyperparameters and data set specifications can be found in the supplementary material. Code is available online¹. We compare to LSTM and GRU baselines for which we replaced the RKN transition layer with generic LSTM and GRU layers. Those were given the encoder output as inputs and have an internal state size of $2n$. The internal state was split into two equally large parts, the first part was used to compute the mean and the second to compute the variance. We additionally executed most of the following experiments using the root mean squared error to

¹<https://github.com/LCAS/RKN>

Table 2. Our approach outperforms the generic LSTM and GRU baselines. The GRU with $m = 8$ and the LSTM with $m = 6$ were designed to have roughly the same amount of parameters as the RKN with $b = 3$. In the case where $m = b$ the RKN uses a full transition matrix. fc stands for full covariance matrix, i.e., we do not use factorization of the belief state.

Model	Log Likelihood
RKN ($m = 15, b = 3, K = 15$)	6.182 ± 0.155
RKN $m = b = 15, K = 3$)	6.248 ± 0.1715
RKN ($m = 15, b = 3, K = 15, fc$)	6.161 ± 0.23
RKN ($m = b = 15, K = 15, fc$)	6.197 ± 0.249
LSTM $m = 50$	5.773 ± 0.231
LSTM $m = 6$	6.019 ± 0.122
GRU $m = 50$	5.649 ± 0.197
GRU $m = 8$	6.051 ± 0.145

illustrate that our approach is also competitive in prediction accuracy. The RMSE results can be found in the appendix.

3.1. Pendulum

We evaluated the RKN on a simple simulated pendulum with images of size 24×24 pixels as observations. Gaussian transition noise with standard deviation of $\sigma = 0.1$ was added to the angular velocity after each step. In the first experiment, we evaluated filtering in the presence of high observation noise. We compare against LSTMs and GRUs as these are the only methods that can also perform state estimation. The amount of noise varies between no noise at all and the whole image consisting of pure noise. Furthermore, the noise is correlated over time, i.e., the model may observe pure noise for several consecutive time steps. Details about the noise sampling process can be found in the appendix. We represent the joint angle θ_t as a two dimensional vector $\mathbf{s}_t = \boldsymbol{\theta}_t = (\sin(\theta_t), \cos(\theta_t))^T$ to avoid discontinuities. We compared different instances of our model to evaluate the effects of assuming sparse transition models and factorized state covariances. The results are given in Table 2. The results show that our assumptions do not affect the performance, while we need to learn fewer parameters and can use much more efficient computations using the factorized representation. Note that for the more complex experiments with a higher dimensional latent state space, we were not able to experiment with full covariance matrices due to lack of memory and massive computation times. Moreover, the RKN outperforms LSTM and GRU in all settings.

In a second experiment, we evaluated the image prediction performance and compare against existing variational inference approaches. We randomly removed half of the images from the sequences and tasked the models with imputing

Table 3. Comparison on the image imputation task. The informed models were given a mask of booleans indicating which images are valid and which not. The uninformed models were given a black image whenever the image was not valid. E2C and SIN only work in the informed case. Since the KVAE is capable of smoothing in addition to filtering, we evaluated both. Our approach outperforms all models. Only the informed KVAE yields comparable, but still slightly worse results while E2C and SIN fail to capture the dynamics. The uninformed KVAE fails at identifying the invalid images.

Model	Log Likelihood
RKN (informed)	-12.782 ± 0.0160
RKN (uninformed)	-12.788 ± 0.0142
KVAE (informed, filter)	-14.383 ± 0.229
KVAE (informed, smooth)	-13.337 ± 0.236
KVAE (uninformed, filter)	-46.320 ± 6.488
KVAE (uninformed, smooth)	-38.170 ± 5.399
E2C (informed)	-95.539 ± 1.754
SIN (informed)	-101.268 ± 0.567

those missing frames, i.e., we train the models to predict images instead of the position. We compared our approach to the Kalman Variational Autoencoder (KVAE) (Fraccaro et al., 2017), Embed to Control (E2C) (Watter et al., 2015) and Structured Inference Networks (SIN) (Krishnan et al., 2017). The results can be found in Table 3. Again, our RKN outperforms all other models. This is surprising as the variational inference models use much more complex inference methods and in some cases even more information such as in the KVAE smoothing case. Sample sequences can be found in the supplementary material. All hyperparameters are the same as for the previous experiment.

3.2. Multiple Pendulums

Dealing with uncertainty in a principled manner becomes even more important if the observation noise affects different parts of the observation in different ways. In such a scenario the model has to infer which parts are currently observable and which parts are not. To obtain a simple experiment with that property, we repeated the pendulum experiments with three colored pendulums in one image. The image was split into four equally sized square parts and the noise was generated individually for each part such that some pendulums may be occluded while others are visible. Exemplary images are shown in the supplementary material. A comparison to LSTM and GRU baselines can be found in Figure 2 and Table 4 as well as an exemplary trajectory in Figure 3. The RKN again clearly outperforms the competing methods. We also computed the quality of the uncertainty prediction by showing the histograms of

Table 4. Results of the multiple pendulum experiments.

Model	Log Likelihood
RKN $m = 45$ $b = 3, k = 15$	11.51 ± 1.703
LSTM $m = 50$	7.5224 ± 1.564
LSTM $m = 12$	7.429 ± 1.307
GRU $m = 50$	7.541 ± 1.547
GRU $m = 12$	5.602 ± 1.468

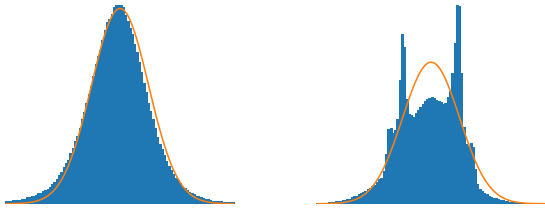


Figure 2. To evaluate the quality of our uncertainty prediction we compute the normalized error $\frac{s^{(j)} - s^{(j),+}}{\sigma^{(j),+}}$ for each entry j of \mathbf{s} for all time steps in all test sequences. This normalized error should follow a Gaussian distribution with unit variance if the prediction is correct. We compare the resulting error histograms with a unit variance Gaussian. The left histogram shows the RKN, the right one the LSTM. The RKN perfectly fits the normal distribution while the LSTM’s normalized error distribution has several modes. Again we designed the smaller LSTM and GRU to have roughly the same amount of parameters as the RKN.

the normalized prediction errors. While this histogram has clearly a Gaussian shape for the RKN, it looks like a less regular distribution for the LSTMs.

3.3. Quad Link

We repeated the filtering experiment on a system with much more complicated dynamics, a quad link pendulum on images of size 48×48 pixels. Since the individual links of the quad link may occlude each other different amounts of noise are induced for each link. Two versions of this experiment were evaluated. One without additional noise and one where we added noise generated by the same process used in the pendulum experiments. You can find the results in Table 5.

Furthermore, we repeated the imputation experiment with the quad link. We compared only to the informed KVAE, since it was the only model producing competitive results for the pendulum. Our approach achieved -44.470 ± 0.208 (informed) and -44.584 ± 0.236 (uninformed). The KVAE achieved -52.608 ± 0.602 for smoothing and -59.0218 ± 0.580 for filtering. Sample images can be found in the appendix.

Table 5. Comparison of our approach with the LSTM and GRU Baselines on the Quad Link Pendulum. Again the RKN performs significantly better than LSTM and GRU who fail to perform well.

Model	without noise Log Likelihood	with noise Log Likelihood
RKN ($m = 100,$ $b = 3, K = 15$)	14.534 ± 0.176	6.259 ± 0.412
LSTM ($m = 50$)	11.960 ± 1.24	5.21 ± 0.305
LSTM ($m = 100$)	7.858 ± 4.680	3.87 ± 0.938
GRU ($m = 50$)	10.346 ± 2.70	4.696 ± 0.699
GRU ($m = 100$)	5.82 ± 2.80	1.2 ± 1.105

3.4. KITTI Dataset for Visual Odometry

Next, we evaluate the RKN on the task of learning visual odometry from monocular images on the KITTI visual odometry dataset (Geiger et al., 2012). It consists of 11 labeled image sequences collected by driving a vehicle in an urban environment.

We use the unsupervised approach proposed by (Zhou et al., 2017) to extract features from the images. Using these features, we first evaluate a simple forward network, which we refer to as sfmlearner-s. We then augment this architecture with LSTMs, GRUs, and our proposed RKN. Additionally, we compare to DeepVO (Wang et al., 2017), an approach tailored to the problem of visual odometry. In Table 6 we show results using the common KITTI evaluation scheme. The results show that the RKN shows better performance in comparison to LSTMs and GRUs, and even performs comparably to the tailored DeepVO approach.

3.5. Pneumatic Brook Robot Arm

We consider another real world task, learning the dynamics of a pneumatically actuated joint of a Hydro-Lek HLK-7W robot arm. The data consists of measured joint positions and the input current to the controller of the joint sampled at 100Hz. Pneumatic robots are very hard to model as there is a large amount of hysteresis, requiring recurrent prediction models.

During training, we work with sequences of length 500. For the first 300 time steps those sequences consist of the full observation, i.e. joint position and current. The remaining 200 time steps only the current is given. The models have to impute the missing joint positions in an uninformed fashion, i.e., the absence of a position is only indicated by unrealistically high values.

In order to evaluate the learned dynamics, we used them to predict future joint positions. After an initial burn-in phase,

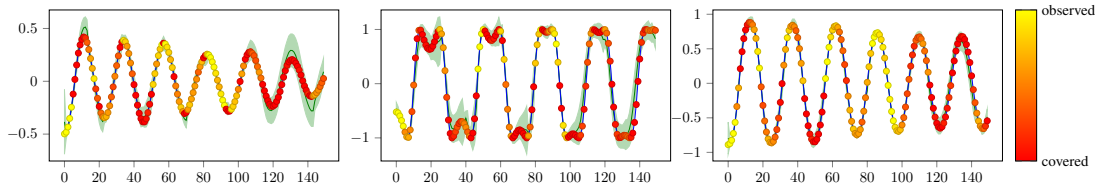


Figure 3. Predicted sine value of the tree links with 2 times standard deviation (green). Ground truth displayed in blue. The crosses visualize the current visibility of the link with yellow corresponding to fully visible and red to fully occluded. If there is no observation for a considerable time the predictions become less precise due to transition noise, however, the variance also increases.

Table 6. Comparison on the KITTI dataset using the common evaluation scheme. The RKN performs better than the non-recurrent sfmlerner-s baseline as well as the recurrent LSTM and GRU baselines. It performs comparably to DeepVO (Wang et al., 2017), an approach tailored to the task of visual odometry.

Sequence	Deep VO		sfmlerner-s		LSTM		GRU		RKN	
	$t_{rel}(\%)$	$r_{rel}(\circ)$	$t_{rel}(\%)$	$r_{rel}(\circ)$	$t_{rel}(\%)$	$r_{rel}(\circ)$	$t_{rel}(\%)$	$r_{rel}(\circ)$	$t_{rel}(\%)$	$r_{rel}(\circ)$
3	8.49	6.89	13.21	6.70	8.99	4.55	9.34	3.81	7.83	3.57
4	7.19	6.97	13.56	4.56	11.88	3.44	12.36	2.89	11.61	2.61
5	2.62	3.61	13.06	5.70	8.96	3.43	10.02	3.43	7.29	2.77
6	5.42	5.82	10.87	4.47	9.66	2.80	10.99	3.22	8.08	2.32
7	3.91	4.60	13.47	8.41	9.83	5.48	13.70	6.52	9.38	4.83
10	8.11	8.83	16.69	6.26	13.85	3.49	13.37	3.25	12.71	3.09
mean	5.96	6.12	13.48	6.02	10.53	3.87	11.63	3.85	9.48	3.20

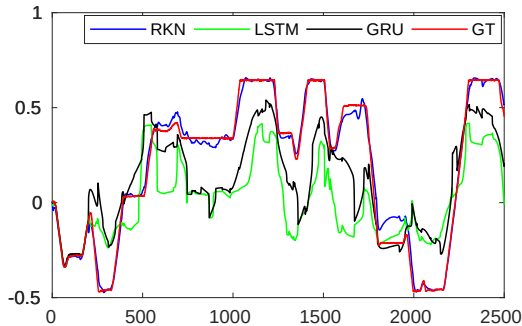


Figure 4. Predicted joint values for the pneumatic joint prediction task. The initial burn-in phase is 200 time steps long, afterwards only the predictions are displayed. GT denotes the ground truth. As can be seen, only our approach manages to give feasible predictions.

the model was tasked with predicting the joint position 2 seconds into the future, given only the future control inputs. Afterwards, the next observation was given to the model and the prediction process repeated.

Using a 60 dimensional latent space, 32 basis matrices and a bandwidth of 3, the RKN achieved a log likelihood of 4.930. This is considerably better than the LSTM and GRU baselines, which achieved 0.952 and 1.186 respectively. Figure 4 shows exemplary predictions of all approaches.

4. Conclusion

In this paper, we introduced the *Recurrent Kalman Network* that jointly learns high-dimensional representations of the system in a latent space with locally linear transition models and factorized covariances. The update equations in the high-dimensional space are based on the update equations of the classical Kalman filter, however, due to the factorization assumptions, they simplify to scalar operations that can be performed much faster and with greater numerical stability. Our model outperforms generic LSTMs and GRUs on various state estimation tasks while providing reasonable uncertainty estimates. Additionally, it outperformed several generative models on an image imputation task and we demonstrated its applicability to real world scenarios. Training is straight forward and can be done in an end-to-end fashion.

In future work, we want to exploit the principled notion of a variance provided by our approach in scenarios where such a notion is beneficial, e.g. reinforcement learning. Similar to (Fraccaro et al., 2017) we could expand our approach to not just filter but smooth over trajectories in offline post-processing scenarios which could potentially increase the estimation performance significantly. Additionally, different, more complex, factorization assumptions can be investigated in the future.

Acknowledgements

The research leading to these results has received funding from EPSRC under grant agreement EP/R02572X/1 (National Center for Nuclear Robotics). Calculations for this research were conducted on the Lichtenberg high performance computer of the TU Darmstadt.

References

- Archer, E., Park, I. M., Buesing, L., Cunningham, J., and Paninski, L. Black box variational inference for state space models. *arXiv preprint arXiv:1511.07367*, 2015.
- Ba, J. L., Kiros, J. R., and Hinton, G. E. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- Cho, K., Van Merriënboer, B., Bahdanau, D., and Bengio, Y. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*, 2014.
- Clevert, D.-A., Unterthiner, T., and Hochreiter, S. Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289*, 2015.
- Doerr, A., Daniel, C., Schiegg, M., Duy, N.-T., Schaal, S., Toussaint, M., and Sebastian, T. Probabilistic recurrent state-space models. In *International Conference on Machine Learning*, pp. 1279–1288, 2018.
- Fraccaro, M., Kamronn, S., Paquet, U., and Winther, O. A disentangled recognition and nonlinear dynamics model for unsupervised learning. In *Advances in Neural Information Processing Systems*, pp. 3601–3610, 2017.
- Gebhardt, G., Kupcsik, A., and Neumann, G. The kernel kalman rule - efficient nonparametric inference with recursive least squares. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- Geiger, A., Lenz, P., and Urtasun, R. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- Haarnoja, T., Ajay, A., Levine, S., and Abbeel, P. Backprop kf: Learning discriminative deterministic state estimators. In *Advances in Neural Information Processing Systems*, pp. 4376–4384, 2016.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- Kalman, R. E. et al. A new approach to linear filtering and prediction problems. *Journal of basic Engineering*, 82(1):35–45, 1960.
- Karl, M., Soelch, M., Bayer, J., and van der Smagt, P. Deep variational bayes filters: Unsupervised learning of state space models from raw data. *arXiv preprint arXiv:1605.06432*, 2016.
- Kingma, D. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Krishnan, R. G., Shalit, U., and Sontag, D. Structured inference networks for nonlinear state space models. In *AAAI*, pp. 2101–2109, 2017.
- Naesseth, C. A., Linderman, S. W., Ranganath, R., and Blei, D. M. Variational sequential monte carlo. *arXiv preprint arXiv:1705.11140*, 2017.
- Rangapuram, S. S., Seeger, M. W., Gasthaus, J., Stella, L., Wang, Y., and Januschowski, T. Deep state space models for time series forecasting. In *Advances in Neural Information Processing Systems*, pp. 7796–7805, 2018.
- Wahlström, N., Schön, T. B., and Deisenroth, M. P. From pixels to torques: Policy learning with deep dynamical models. *arXiv preprint arXiv:1502.02251*, 2015.
- Wang, S., Clark, R., Wen, H., and Trigoni, N. Deepvo: Towards end-to-end visual odometry with deep recurrent convolutional neural networks. In *Robotics and Automation (ICRA), 2017 IEEE International Conference on*, pp. 2043–2050. IEEE, 2017.
- Watter, M., Springenberg, J., Boedecker, J., and Riedmiller, M. Embed to control: A locally linear latent dynamics model for control from raw images. In *Advances in neural information processing systems*, pp. 2746–2754, 2015.
- Werbos, P. J. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, 78(10):1550–1560, 1990.
- Yingzhen, L. and Mandt, S. Disentangled sequential autoencoder. In *International Conference on Machine Learning*, pp. 5656–5665, 2018.
- Zhou, T., Brown, M., Snavely, N., and Lowe, D. G. Unsupervised learning of depth and ego-motion from video. In *CVPR*, volume 2, pp. 7, 2017.