

Approximate Bayesian inference in spatial environments

Atanas Mirchev[†], Baris Kayalibay[†], Maximilian Soelch[†], Patrick van der Smagt[‡] and Justin Bayer[†]

Machine Learning Research Lab, Volkswagen Group
 80805, Munich, Germany

[†]{firstname.lastname}@argmax.ai

[‡]smagt@argmax.ai

Abstract—Model-based approaches bear great promise for decision making of agents interacting with the physical world. In the context of spatial environments, different types of problems such as localisation, mapping, navigation or autonomous exploration are typically addressed with specialised methods, often relying on detailed knowledge of the system at hand. We express these tasks as probabilistic inference and planning under the umbrella of deep sequential generative models. Using the frameworks of variational inference and neural networks, our method inherits favourable properties such as flexibility, scalability and the ability to learn from data. The method performs comparably to specialised state-of-the-art methodology in two distinct simulated environments.

I. INTRODUCTION

Sequential decision making is a framework to represent the interaction of an agent with its environment: an observation of the world is presented to the agent, upon which the informed agent picks an action, which in turn alters the world’s state. One instance of interest are spatial environments such as mobile robots on a factory floor, autonomous cars, robot arms or unmanned aerial vehicles. Various tasks are of interest in these scenarios. Localisation or pose estimation considers the relation of the agent itself to the environment. This is often combined with establishing a map of its surroundings and has been referred to as simultaneous localisation and mapping (SLAM) by the robotics community. If the aim is to obtain that map efficiently, autonomous exploration is about devising trajectories that uncover the map with as little effort as possible. Another goal is navigation, referring to the generation of plans that allow the agent to reach a pre-specified location.

We set out to address these problems in a unified framework. We augment a non-linear state space model, the deep variational Bayes filter [12] with a global latent variable representing a map (DVBF-LM). We propose the necessary learning algorithms that enable end-to-end learning and make it possible to express the aforementioned tasks as either inference or planning in that model. SLAM is performed by approximate, variational inference of the joint posterior over maps and pose trajectories. We rely on a standard formulation as probabilistic inference in a graphical model. For autonomous exploration, we optimise the expected information gain with respect to the control signals [25, 21]. This is possible due to the probabilistic treatment




Fig. 1: Example for the agent pose posterior in a LiDAR environment (pybox2d): (a) Scatter plot of observations predicted from the learned DVBF-LM map. (b) The current agent position and LiDAR readings. (c) Plot of the pose posterior over locations given the current observation. It is highly non-Gaussian and has several maxima.

of the map, where unexploredness is related to the remaining uncertainty in the respective region. Navigation is implemented as planning in a discretisation of the learned model.

Our contributions are:

- a deep non-linear state-space model with an explicit map component that can be estimated from data;
- methods of performing SLAM, autonomous exploration and navigation in said model;
- a solution to train stochastic recurrent models on a single, long, consecutive time series;
- a variational posterior formulation that copes with the complex joint posterior prevalent in SLAM (cf. Figure 1).

We validate the claims in a series of extensive experiments where we perform comparably to baselines tailored specifically to the respective settings.

II. RELATED WORK

The problem of concurrent estimation of an agent’s pose and its surrounding has seen considerable attention in the last decades. We refer the interested reader to the survey of Cadena et al. [3]. A contribution of Murphy [17] is most similar to our approach: the map is a matrix-valued global latent variable inferred through Bayesian methods.




Fig. 2: Sequential graphical model with global map \mathcal{M} and local charts \mathbf{m}_t .

Mapping and localisation has been adopted in the machine-learning community mostly to solve reinforcement-learning or visual-navigation problems [18, 19]. Fraccaro et al. [7] proposed a generative model for spatial environments. While their approach is similar to ours, their focus was primarily on simulator performance over long time spans. Further, an external memory is used which does not directly represent a random variable as part of a graphical model.

Early discussions of the importance of exploration of model parameters can be found in [20, 23], and information-theoretic methods for spatial exploration can be traced back to [6, 26]. Our work follows the framework of *curiosity-driven exploration* [20]. In the spatial environment context, it represents an instance of active SLAM [25, 3].

A large body of recent spatial exploration methods [21, 1, 24, 27] is driven by information theory, but assumes particular discretisations of the state space (occupancy grids) or skeletonisations of the possible action paths. Variational information maximizing exploration (VIME) [11] is closely related to our method, but imposes an intentional Gaussian constraint for tractability and mutual information is estimated for one step t at a time. Our method has no such assumptions.

III. DEEP VARIATIONAL BAYES FILTER WITH A LATENT MAP

Our aim is to provide a unified model that covers all of the aforementioned tasks. We defer the discussion of their exact implementations to Section IV, Section V and Section VI. Here we will focus on the underlying probabilistic generative model. We step on the solid foundation of state space models to represent sequential agent interactions:

$$p(\mathbf{x}_{1:T}, \mathbf{z}_{1:T} \mid \mathbf{u}_{1:T-1}) = p(\mathbf{z}_1) \prod_{t=2}^T p(\mathbf{z}_t \mid \mathbf{z}_{t-1}, \mathbf{u}_{t-1}) \prod_{t=1}^T p(\mathbf{x}_t \mid \mathbf{z}_t),$$

where

$$\begin{aligned} \mathbf{x}_{1:T} \in \mathbb{R}^{T \times D_x} & \text{ is a sequence of observations,} \\ \mathbf{z}_{1:T} \in \mathbb{R}^{T \times D_z} & \text{ is a sequence of poses, and} \\ \mathbf{u}_{1:T-1} \in \mathbb{R}^{(T-1) \times D_u} & \text{ is a sequence of control inputs.} \end{aligned}$$




Fig. 3: *Illustration of pose inference.* (a) An agent (teal star) traverses a maze, (b) collecting sensor readings (top) and control signals (bottom). (c) A belief of the map is formed from observations (top), a belief of the trajectory is formed from observations and controls (bottom). (d) The two beliefs are fused. (e) The map is indexed with a pose-based attention mechanism. (f) The attended region of the map is used to reconstruct the observation.

In terms of spatial environments, the *transition model* $p(\mathbf{z}_t \mid \mathbf{z}_{t-1}, \mathbf{u}_{t-1})$ represents the dynamics and the *emission model* $p(\mathbf{x}_t \mid \mathbf{z}_t)$ simulates the formation of observations.

The traditional state space model above is a powerful tool for the analysis of stochastic dynamical systems [12]. In order to unlock further inferences specific to spatial environments, more transparency is required. We thus introduce additional structure to provide the necessary entry points. First, we identify a part of the latent space with the environment itself. To that end we extend the traditional graphical model, incorporating a global map latent variable $\mathcal{M} \sim p(\mathcal{M})$. We introduce a latent middle layer of local charts $\mathbf{m}_t \sim p(\mathbf{m}_t \mid \mathbf{z}_t, \mathcal{M})$. Intuitively, the chart \mathbf{m}_t represents the currently relevant attended region of the map. It shapes the transition of poses over time, $p(\mathbf{z}_{t+1} \mid \mathbf{z}_t, \mathbf{m}_t, \mathbf{u}_t)$. The observation emission model operates solely on these charts, $\mathbf{x}_t \sim p(\mathbf{x}_t \mid \mathbf{m}_t)$. In total, this yields the graphical model (cf. Figure 2)

$$p(\mathbf{x}_{1:T}, \mathbf{z}_{1:T}, \mathbf{m}_{1:T}, \mathcal{M} \mid \mathbf{u}_{1:T-1}) = p(\mathcal{M}) \rho(\mathbf{z}_1) \prod_{t=1}^T p(\mathbf{x}_t \mid \mathbf{m}_t) p(\mathbf{m}_t \mid \mathbf{z}_t, \mathcal{M}) \prod_{t=1}^{T-1} p(\mathbf{z}_{t+1} \mid \mathbf{z}_t, \mathbf{m}_t, \mathbf{u}_t).$$

A common assumption is that the factors are governed by a set of parameters θ , i.e. $\rho_{\theta_t}(\mathbf{z}_1)$, $p_{\theta_T}(\mathbf{z}_{t+1} \mid \mathbf{z}_t, \mathbf{u}_t, \mathbf{m}_t)$, $p_{\theta_E}(\mathbf{x}_t \mid \mathbf{m}_t)$. We will leave out the dependency for notational brevity in the remainder of this work. Further, in the conducted experiments we assume the transition parameters θ_T and the initial state distribution are learned *a priori* or engineered.

A. Approximation via Variational Inference

Exact inference in such models is typically intractable. We obtain variational approximations $q(\mathcal{M})$ and $q(\mathbf{z}_{1:T} \mid \mathbf{x}_{1:T}, \mathbf{u}_{1:T-1}, \mathcal{M})$ of the corresponding posteriors, collecting their learnable variational parameters in ϕ , where we rely on Bayes by backprop [2] for the former and on SGVB [14] for the latter. The negative *evidence lower bound* (ELBO)

is given as

$$\mathcal{L}_{\text{elbo}} = \underbrace{\mathbb{E}_q[-\log p(\mathbf{x}_{1:T} | \mathbf{z}_{1:T}, \mathcal{M})]}_{=: \ell^r} + \underbrace{\text{KL}(q(\mathcal{M}) || p(\mathcal{M}))}_{=: \ell^{\mathcal{M}}} + \underbrace{\mathbb{E}_q[\text{KL}(q(\mathbf{z}_{1:T} | \mathbf{x}_{1:T}, \mathcal{M}) || p(\mathbf{z}_{1:T} | \mathcal{M}))]}_{=: \ell^z}. \quad (1)$$

The conditioning on $\mathbf{u}_{1:T-1}$ is dropped for brevity. We call ℓ^r the *reconstruction loss*, ℓ^z the *pose KL penalty* and $\ell^{\mathcal{M}}$ the *map KL penalty*. Inference of poses and the map then comes down to the minimisation of Equation (1) with respect to ϕ .

B. Implementation of the Generative Model

In general, the geometric properties of the environment that need to be represented will determine the particular form of the map and the associated attention model. For the purposes of this work we follow [17], defining the map \mathcal{M} to be a finite grid of width w and height h . Each grid cell \mathcal{M}_{ij} is a real-valued vector of dimensionality D_m , i.e. $\mathcal{M} \in \mathbb{R}^{w \times h \times D_m}$. As prior for such a latent map cell we use a standard normal, $\mathcal{M}_{ij} \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$. Extracting *local charts* \mathbf{m}_t from the map is done through a convex combination of the memory cells:

$$\mathbf{m}_t = f_m(\mathbf{z}_t, \mathcal{M}) = \sum_{i,j} \alpha(\mathbf{z}_t)_{ij} \mathcal{M}_{ij}.$$

The result is then a point mass:

$$p(\mathbf{m}_t | \mathbf{z}_t, \mathcal{M}, \theta_M) \propto \mathbb{I}[\mathbf{m}_t = f_m(\mathbf{z}_t, \mathcal{M})].$$

In this implementation, we choose α to be a bilinear interpolation kernel, combining four cells at a time.

The *emission model* and *transition model* are conditional Gaussian distributions with fixed diagonal covariances. The respective means are given by neural networks parameterised by θ_E and θ_T :

$$p(\mathbf{x}_t | \mathbf{m}_t) = \mathcal{N}(\boldsymbol{\mu}_E(\mathbf{m}_t), \text{diag}(\boldsymbol{\sigma}_E^2)), \\ p(\mathbf{z}_{t+1} | \mathbf{z}_t, \mathbf{u}_t, \mathbf{m}_t) = \mathcal{N}(\boldsymbol{\mu}_T(\mathbf{z}_t, \mathbf{u}_t, \mathbf{m}_t), \sigma_T^2 \mathbf{1}).$$

C. Design of the Variational Posterior

Inference of poses is done through a variational approximation of the true posterior $q(\mathbf{z}_t | \mathbf{x}_{1:T}) \approx p(\mathbf{z}_t | \mathbf{x}_{1:T})$, where we left out the control signals $\mathbf{u}_{1:T-1}$ for brevity and will do so for the remainder of this section. The global variable \mathcal{M} poses an atypical challenge for stochastic recurrent models trained with amortised variational inference, for which an intuitive explanation is as follows. Consider the true posterior, which has to account for all possible maps:

$$p(\mathbf{z}_t | \mathbf{x}_{1:T}) = \int p(\mathbf{z}_t | \mathcal{M}, \mathbf{x}_{1:T}) p(\mathcal{M} | \mathbf{x}_{1:T}) d\mathcal{M}.$$

Any parameterised variational approximation $q(\mathbf{z}_t | \mathbf{x}_{1:T})$ will have to implement its own belief of the map implicitly. During training, this will prove difficult as it has to track the current belief of the generative model to conform to it, as it essentially

implements its inverse. The task of the inference model can be substantially eased by informing it of the current belief of the map $q(\mathcal{M})$ explicitly. We choose to do so by implementing q as a bootstrap particle filter [8] with the particle forwarding distribution from Section III-D2 as a proposal distribution:

$$q(\mathbf{z}_{1:T} | \mathbf{x}_{1:T}, \mathcal{M}) = \prod_{t=1}^T q(\mathbf{z}_t | \mathbf{x}_{1:t}, \mathcal{M}), \\ q(\mathbf{z}_t | \mathbf{x}_{1:t}, \mathcal{M}) \propto \mathbb{E}_{\mathbf{z}_t^{(k)} \sim \hat{q}(\mathbf{z}_t)} \left[\sum_{k=1}^K \hat{\omega}_k \mathbb{I}(\mathbf{z}_t = \mathbf{z}_t^{(k)}) \right], \\ \hat{\omega}_k = \frac{\omega_k}{\sum_j \omega_j}, \\ \omega_k = \frac{p(\mathbf{x}_t | \mathbf{z}_t^{(k)}, \mathcal{M}) p(\mathbf{z}_t^{(k)})}{\hat{q}(\mathbf{z}_t^{(k)})}.$$

This has two immediate consequences. First, the variational posterior used does not have any parameters and is hence not optimised directly. Second, the true posterior is recovered for $K \rightarrow \infty$. But most importantly, the importance weights explicitly reflect the map (sampled from an outer expectation over $q(\mathcal{M})$ in Equation (1)) and the proposals in conflict with it will be sorted out in a natural manner as they have lower weights.

The variational approximation of the posterior map $q(\mathcal{M})$ was chosen to follow a mean-field approach with a factorised Gaussian $q(\mathcal{M}) = \prod_i \prod_j \mathcal{N}(\boldsymbol{\mu}_{\mathcal{M}_{ij}}, \boldsymbol{\sigma}_{\mathcal{M}_{ij}}^2)$, with variational parameters $\boldsymbol{\mu}_{\mathcal{M}_{ij}}, \boldsymbol{\sigma}_{\mathcal{M}_{ij}}^2 \in \phi$.

D. Faster Training with Mini Batches

In practice, inference is typically performed on very long, continuous streams of data recorded from a moving agent. Evaluating the ELBO for the whole trajectory at once proves prohibitive for learning or is downright impossible due to memory limitations. We therefore seek to relax the optimisation while still respecting the underlying model.

1) Decomposing the Loss into a Sum over Time Steps:

Under the Markov assumptions, the evidence lower bound from Equation (1) can be written as a sum over time steps:

$$\mathcal{L}_{\text{elbo}} = \ell^r + \ell^z + \ell^{\mathcal{M}} = \mathbb{E}_q \left[\sum_{t=1}^T \ell_t^r + \ell_t^z + \ell_t^{\mathcal{M}} \right] \quad (2)$$

with, leaving out the control signals $\mathbf{u}_{1:T-1}$ for brevity:

$$\ell_t^r = -\log p(\mathbf{x}_t | \mathbf{z}_t, \mathcal{M}), \quad \ell^r = \mathbb{E}_q \left[\sum_{t=1}^T \ell_t^r \right], \\ \ell_t^z = \log \frac{q(\mathbf{z}_t | \mathbf{x}_{1:t}, \mathcal{M})}{p(\mathbf{z}_t | \mathbf{z}_{t-1}, \mathcal{M})}, \quad \ell^z = \mathbb{E}_q \left[\sum_{t=1}^T \ell_t^z \right], \\ \ell_t^{\mathcal{M}} = \frac{1}{T} \log \frac{q(\mathcal{M})}{p(\mathcal{M})}, \quad \ell^{\mathcal{M}} = \mathbb{E}_q \left[\sum_{t=1}^T \ell_t^{\mathcal{M}} \right].$$

Following [2], we distribute the contribution of the map KL penalty term over different time steps, reflected in $\ell_t^{\mathcal{M}}$. We denote the overall loss at time step t as \mathcal{L}_t .

If the loss function is a sum over independent terms, a gradient estimator using only a subset of those terms will be unbiased. Unfortunately the terms for each time step in Equation (2) are not independent, which requires ancestral sampling from the whole Markov chain.

2) *Approximate Asynchronous Particle Representation:*

To overcome this issue we maintain sets of N particles $\xi_t^{(n)}, n = 1, \dots, N; t = 1, \dots, T$ that cache samples for each step of the variational posterior over poses $q(\mathbf{z}_{1:T} | \mathbf{x}_{1:T}, \mathcal{M}) = \prod_{t=1}^T q(\mathbf{z}_t | \mathbf{x}_{1:t}, \mathcal{M})$ during training.

We define an estimator of the gradients from the complete ELBO, akin to stochastic gradient descent. The time steps we wish to use for gradient estimation are gathered in a minibatch \mathcal{B} . We then approximate the loss given in Equation (2) via

$$\tilde{\mathcal{L}} = \frac{T}{|\mathcal{B}|} \sum_{t \in \mathcal{B}} \mathbb{E}_{\mathcal{M} \sim q} [\mathbb{E}_{\mathbf{z}_t \sim \tilde{q}} [\ell_t^r + \ell_t^z + \ell_t^{\mathcal{M}}]], \quad (3)$$

where $\tilde{q}(\mathbf{z}_{1:T}) = \prod_{t=1}^T \tilde{q}(\mathbf{z}_t)$ is an approximation of $q(\mathbf{z}_{1:T} | \mathbf{x}_{1:T}, \mathcal{M})$ based on the cached particles that allows more efficient sampling of \mathbf{z}_t . In particular, every $\tilde{q}(\mathbf{z}_t)$ is importance-resampled from an underlying proposal distribution $\hat{q}(\mathbf{z}_t)$, which in turn is based on the particles $\xi_t^{(n)}, n = 1, \dots, N$. In this work, $\hat{q}(\mathbf{z}_t)$ is represented as a Normal random variable with moments matched from the set of N particles for time step t : $\hat{q}(\mathbf{z}_t) = \mathcal{N}(\boldsymbol{\mu}_{\xi_t}, \boldsymbol{\sigma}_{\xi_t}^2)$. The mean $\boldsymbol{\mu}_{\xi_t} = \frac{1}{N} \sum_{n=1}^N \xi_t^{(n)}$ and variance $\boldsymbol{\sigma}_{\xi_t}^2 = \frac{1}{N} \sum_{n=1}^N (\boldsymbol{\mu}_{\xi_t} - \xi_t^{(n)})^2$ are the empirical mean and variance of the particles respectively. The approximating particle sets are updated during gradient estimation: for any training iteration with $t \in \mathcal{B}$, we can update the particles at following time steps $t+k, k = 1, \dots$:

$$\xi_{t+k}^{(n)} \sim \tilde{q}(\mathbf{z}_t) \prod_{i=1}^k p(\mathbf{z}_{t+i} | \mathbf{z}_{t+i-1}, \mathcal{M}),$$

effectively performing importance resampling and moving particles forward through the transition model, refreshing the approximation $\hat{q}(\mathbf{z}_{t+k})$. This leads to an asynchronous procedure: expectations in Equation (2) w.r.t. the approximate posterior over agent poses are implemented through particles stemming from previous training iterations, potentially biasing the gradients. This bias can be controlled with small parameter updates (i.e. $\phi^{(i+1)} \approx \phi^{(i)}$), since we can then expect the expectations to be close as well. In practice, we will choose chunks of consecutive time steps to be the elements of mini batches, requiring to only update the particles at the beginning of each such chunk.

IV. DVBF-LM AS A METHOD FOR SLAM

We first investigate the capabilities of DVBF-LM as a solution to SLAM. Our model is evaluated in two simulated, precisely controlled environments—a 2D environment with laser range finder observations and VizDoom [13]. A detailed description of each, along with additional information regarding the experimental setup can be found in the supplementary

material ¹.

We randomised seven distinct 2D maze patterns and replicated them in both environments. Each maze was traversed multiple times by two human operators to collect data. For both environments, the transition model $p(\mathbf{z}_{t+1} | \mathbf{z}_t, \mathbf{u}_t, \mathbf{m}_t)$ is pretrained on a first maze that is not considered during evaluation. Performing SLAM then consists of approximating the posterior of the poses and the map $p(\mathbf{z}_{1:t}, \mathcal{M} | \mathbf{x}_{1:t}, \mathbf{u}_{1:t-1})$ for a single traversal through the optimisation of Equation (1) with respect to the variational posteriors $q(\mathbf{z}_{1:T})$ and $q(\mathcal{M})$.

We consider two cases, *offline* and *online SLAM*. In the first case we optimise for all time steps at once, whereas in the second case we sweep $t = 1, \dots, T$ to obtain time-step-wise estimates $q(\mathbf{z}_{1:t})$.

All distances in the following experiments are unit-less, the width and height of the considered mazes were set to 1.0.

A. Pybox2d Environment

For this environment we implemented our own 2D simulator using pybox2d, in which the agent’s sensors are laser range finders (LiDAR readings).

a) *Improving Path Integration:* The aim of this set of experiments is to test whether the proposed approximation of the graphical model improves upon direct path integration based on the pretrained transition $p(\mathbf{z}_{t+1} | \mathbf{z}_t, \mathbf{u}_t, \mathbf{m}_t)$ only. For both online and offline SLAM, using a map clearly outperforms the path integration baseline in terms of localisation error: 0.03 ± 0.02 and 0.04 ± 0.02 at time step 3000 for online and offline SLAM respectively. At this time step, the motion model has diverged for most of the sequences with an average error of 0.14 ± 0.1 . Most notably, the use of a map practically eliminates drift: after 3000 steps, a relative error of less than $20,000^{-1}$, effectively zero, is obtained. This shows that our method stabilises the motion model and keeps the location estimate from diverging. We illustrate the findings in Figure 4.

b) *Comparison to Cartographer:* Next we compare DVBF-LM’s online localisation performance to that of Google’s Cartographer [10], which we consider a representative baseline model for 2D LiDAR SLAM. Cartographer is a realtime SLAM system which operates on laser range finder data and is capable of detecting loop closures. In addition to the LiDAR observations collected from pybox2d, we provided Cartographer with the angular velocity of the agent at every time step in the form of IMU readings. In order to improve upon the default Cartographer configuration and tune it to our setup, we performed a hyperparameter search over more than 40 of Cartographer’s hyperparameters, with 6000 trials on a held out trajectory of 1000 steps. Both Cartographer and DVBF-LM manage to eliminate drift, with respective errors of 0.05 ± 0.04 and 0.03 ± 0.02 at time step 3000. The proposed graphical model approximation leads to localisation performance that is consistently on par in quality to that of Cartographer. The results from the comparison are depicted in Figure 5a.

¹Available at: <https://arxiv.org/abs/1805.07206>.




Fig. 4: We compare the online and offline DVBF-LM localisation error to path integration for all 24 test traversals across 6 mazes in the pybox2d environment. The plots show aggregate results, shaded regions contain 50% and 80% of the traversals. The second plot shows the localisation error relative to the distance travelled.

B. VizDoom Environment

The VizDoom experiments take place in the same set of mazes as pybox2d. Observations are now two-dimensional images taken from the perspective of the agent.

For the VizDoom environment we only investigate offline SLAM performance. The experiment procedure was identical to the pybox2d counterpart. All model components apart from the emission model are kept the same. The latter is modified to better fit visual observations.


We summarise quantitative localisation results in Figure 5b. The method performs on a similar level as in the laser scan-based environment: localisation error is 0.04 ± 0.03 after 5000 time steps. We can see that DVBF-LM is capable of correcting the drift resulting from path integration (an error of 0.11 ± 0.05). The final relative error is $0.06\% \pm 0.06\%$ of the trajectory length. The results indicate that our method is able to perform accurate localisation when applied to different observation modalities in an offline fashion by adapting only the architecture of the emission model.

V. NAVIGATION IN LEARNED ENVIRONMENTS


One incentive for learning a generative model of a spatial environment is that such a model can be used to plan interactions with that environment. As an example, we use DVBF-LM as a black-box environment simulator and provide its predictions to a classical path planning algorithm to solve navigation tasks.

A. Latent Hybrid- A^*

The latent map \mathcal{M} of DVBF-LM conserves the Euclidean geometry of the true environment through the inductive bias of the pretrained transition. Hence, we are able to use the *hybrid- A^** algorithm [5] for path planning. The goal of *hybrid- A^** is to find a path from a *continuous* starting pose to a *continuous* target pose. This is done by discretising the search space into a grid of N cells $\{c_n\}_{n \in [N]}$, which the conventional A^* -search [9] can operate on.



(a) Pybox2d



(b) VizDoom

Fig. 5: (a) Online DVBF-LM localisation error compared to that of Cartographer on test mazes in the pybox2d environment. We omit the offline DVBF-LM results from Figure 4 for the sake of legibility. (b) Offline DVBF-LM localisation error compared to path integration on test maze traversals in the VizDoom environment. The plots show aggregate results, shaded regions contain 50% and 80% of the traversals. The plots on the right show the localisation error relative to the distance travelled.

To obtain smooth navigation trajectories, every discrete state c_n is associated with a continuous state z_n —the agent’s state when that cell was explored for the first time. New cells c_{n+1} are explored by picking random sequences of controls $\mathbf{u}_{1:K}^n$ and predicting a following continuous state by applying the controls to the current z_n . The successor states are found by applying the transition model of DVBF-LM and picking the mean of the resulting Gaussian distribution. When the target state is found, we backtrack to obtain a consistent sequence of controls $\mathbf{u}_{1:T}$, which can be executed by the agent to reach the target.

One issue with using an approximate transition model is that collisions with obstacles cannot always be modelled accurately. The problem is exacerbated by the fact that shortest paths in spatial environments tend to stay close to obstacles. In order to alleviate the negative effect of these two factors on navigation success, we introduce a safety term which penalises closeness to obstacles: $\mathcal{L}_{\text{safe}} = \sum_i s(l_i)$. Here, l_i is the reading of the i -th range sensor of the agent, as predicted by DVBF-LM based on the learned map $q(\mathcal{M})$, and $s(\cdot)$ is a sigmoidal function mirrored along the y -axis.

We add $\mathcal{L}_{\text{safe}}$ to the travel distance when calculating edge weights. The penalty term assumes that we have access to depth readings, which is true in the case of laser scan-based settings but not when the agent only has access to visual observations.

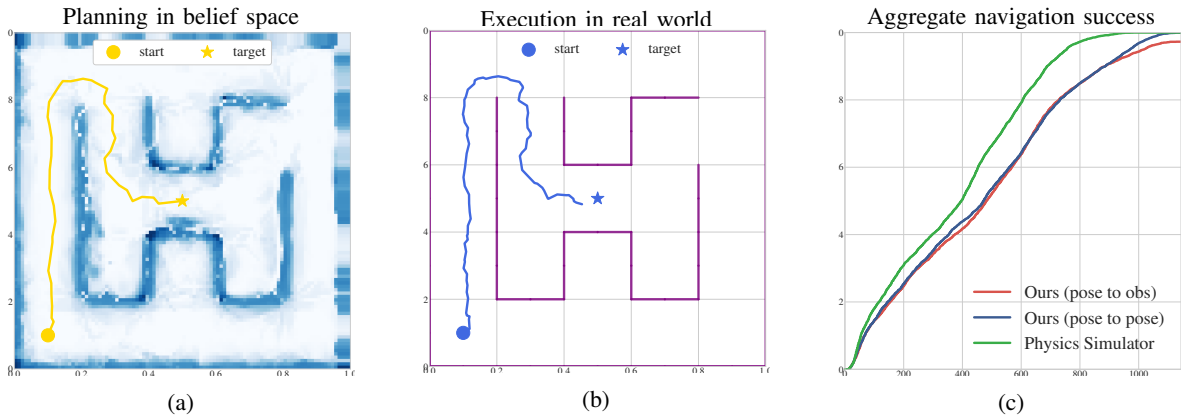


Fig. 6: (a) Navigation plan constructed by using the belief space of the learned environment. (b) Executed trajectory in the actual simulator. (c) Ratio of successful navigation attempts over the number of simulation steps in A^* . In green the upper bound, A^* with access to the ground truth map and transition.

Generalising this path planning framework to visual observations will be part of future work.

B. Results

In our experiments, we verify that the model described in Section III can be applied to navigation.

a) Generative Model Learning: Before DVBF-LM can be used for planning, the emission model and the map \mathcal{M} must first be learned. Here, we use the same models that were acquired as part of the SLAM experiments. For the transition we use an engineered model that allows the agent to move forward along its heading as long as the predicted LiDAR reading in that direction is greater than the desired step length. All navigation tasks are performed in the same set of mazes introduced earlier.

b) Pose-to-Pose Navigation with Hybrid- A^ :* For each maze, we exhaustively pick pairs of starting and target pose from a 5×5 grid over the map. For each pair of poses, we apply the hybrid- A^* as described in Section V-A to plan a trajectory. The obtained controls are then executed in the true physics simulator. The navigation task is considered successful if the agent lands in a proximity of 0.05 or less from the target pose. As a baseline, we consider the navigation performance when the planning algorithm is executed directly in the physics simulator, with access to the true transition and emission. This allows us to assess the drop in performance resulting from approximating the true environment with DVBF-LM. Figure 6c shows the results of the evaluation. Planning based on our generative model comes very close in terms of navigation efficiency to planning in the simulator, affirming the usability of the learned environment maps for navigation tasks. Furthermore, all planned trajectories are successful in reaching the target.

c) Pose-to-Observation Navigation with Hybrid- A^ :* In this scenario, the observation targets are sensor readings from the environment simulator. The corresponding starting poses are the same as in the previous case. Before we can apply the algorithm from V-A the observation targets must first be mapped


to pose targets. To that end, we train a separate variational auto-encoder (VAE, [14]) on the same data used for learning a map of the environment. We set the generative part of the VAE to the emission model of DVBF-LM, we condition on the learned map $q(\mathcal{M})$, and freeze the map and emission parameters. Thus, we obtain an approximation $q(\mathbf{z} | \mathbf{x})$ of the posterior over poses $p(\mathbf{z} | \mathbf{x})$ that conforms to the learned spatial map. This is done once for each of the six mazes. The obtained approximation can be reused for multiple navigation tasks in the given environment. The rest of the evaluation proceeds analogously to the *pose-to-pose* case, using the mode of the approximate posterior, $\mathbf{z}^* = \arg \max_{\mathbf{z}} q(\mathbf{z} | \mathbf{x})$, as a target. Figure 6c illustrates the results of the evaluation. Performance is very similar to the *pose-to-pose* case, with less than 2% of all trajectories failing to reach their actual target. The slight drop in performance can be attributed to *perceptual aliasing*—ambiguity in the pose given an observation—that is typical for spatial environments.

VI. EXPLORATION

Next, we tackle the problem of exploration, expressed in the efficient mapping of the environment. Fast inference of the map is a prerequisite for making informed decisions in spatial settings, as was demonstrated in the navigation task. Autonomous exploration amounts to the selection of control signals such that the data acquired makes inference progress fast. The control signals are then executed in the environment and the process repeats. We will now discuss how we use DVBF-LM to define an exploration policy.

A. Exploration via Active Learning

We choose to follow an information-theoretic approach—we define optimal exploration to be that which leads to the largest change in information in the map variable, a metric commonly referred to as infogain [15]. The change in information is quantified by the mutual information (MI) between future observations $\mathbf{x}_{1:T}$ predicted by DVBF-LM for a sequence of



(a) White means high MI (b) Black means high uncertainty.

Fig. 7: Generated candidates conform with the obstacles belief. Candidates leading into uncertain regions have high MI scores.

planned controls and the map \mathcal{M} :

$$\begin{aligned} & \mathbb{E}_{\mathbf{x}_{1:T} \sim p(\cdot)} [\text{KL}(p(\mathcal{M} | \mathbf{x}_{1:T}) || p(\mathcal{M}))] \\ &= \text{KL}(p(\mathbf{x}_{1:T}, \mathcal{M}) || p(\mathbf{x}_{1:T})p(\mathcal{M})) \\ &= I(\mathbf{x}_{1:T}; \mathcal{M}), \end{aligned} \quad (4)$$

omitting $\mathbf{u}_{1:T-1}$ for brevity. Intuitively, the goal is to select those control signals which maximise the information gained through them in expectation. Thus, we pose the following optimal-control problem:

$$\mathbf{u}_{1:T-1}^* = \arg \max_{\mathbf{u}_{1:T-1}} I(\mathbf{x}_{1:T}; \mathcal{M} | \mathbf{u}_{1:T-1}). \quad (5)$$

In light of this goal, the generative nature of DVBF-LM and the explicit modelling of a global latent map appear essential to the formulation of a principled exploration solution. To solve the posed problem, two issues need to be addressed: the intractability of computing mutual information from Equation (4) and conducting the optimisation in Equation (5) (w.r.t. $\mathbf{u}_{1:T-1}$).

B. Approximating Mutual Information

Equation (4) depends on the highly nonlinear intractable joint $p(\mathbf{x}_{1:T}, \mathcal{M} | \mathbf{u}_{1:T-1})$ and marginal $p(\mathbf{x}_{1:T} | \mathbf{u}_{1:T-1})$. For brevity, we will omit the conditioning on $\mathbf{u}_{1:T-1}$ until the end of this section. Approximation poses a challenge because of the double integration in the multi-dimensional spaces of $\mathbf{x}_{1:T}$ and \mathcal{M} . We resort to combining MC sampling with applying a black-box entropy estimator, as done in [4]. This results in the following estimation:

$$\begin{aligned} I(\mathcal{M}; \mathbf{x}_{1:T}) &= H[\mathbf{x}_{1:T}] - H[\mathbf{x}_{1:T} | \mathcal{M}] \\ &\approx \hat{H}(\bar{\mathbf{X}}) - \frac{1}{M} \sum_{m=1}^M \hat{H}(\bar{\mathbf{X}}^{(m)}). \end{aligned} \quad (6)$$

\hat{H} represents a black-box entropy estimator that works on sample sets. $\bar{\mathbf{X}}$ is a set of samples from the marginal $p(\mathbf{x}_{1:T})$ and each $\bar{\mathbf{X}}^{(m)}$, $m = 1, \dots, M$ is a set of samples from a conditional $p(\mathbf{x}_{1:T} | \mathcal{M}^{(m)})$ for $\mathcal{M}^{(m)} \sim q(\mathcal{M})$. All samples are obtained through ancestral sampling from DVBF-LM, which is only possible since DVBF-LM is a generative model.




Fig. 8: Qualitative exploration comparison: proposed method vs. LSTM baseline. The plot shows the parallel exploration progress of both agents over time.

In practice, exploration is performed in parallel with the inferences of $q(\mathbf{z}_{1:T})$ and $q(\mathcal{M})$, gradually adding new data points and increasing the overall data set size. Note that we do not maximise MI once w.r.t. the prior $p(\mathcal{M})$ for all future time steps, but we maximise it on-line multiple times for T steps ahead w.r.t. the current variational posterior $q(\mathcal{M}) \approx p(\mathcal{M} | \mathcal{D})$. This is well-grounded due to the validity of sequential Bayesian updates, i.e. $p(\mathcal{M} | \mathcal{D}, \mathbf{x}_{1:T}) \propto p(\mathbf{x}_{1:T} | \mathcal{M})p(\mathcal{M} | \mathcal{D})$. In this work, a k-NN black-box entropy estimator is used for \hat{H} [22].

C. Optimising Mutual Information

Optimisation of Equation (5) is performed in two stages. First, a set of proposal controls \mathcal{U} is generated by using the current belief of the map $q(\mathcal{M}) \approx p(\mathcal{M} | \mathcal{D})$. Second, the best candidate control sequence $\mathbf{u}_{1:T-1}^* = \arg \max_{\mathbf{u}_{1:T-1} \in \mathcal{U}} I(\mathbf{x}_{1:T}; \mathcal{M} | \mathbf{u}_{1:T-1}, \mathcal{D})$ is selected among the candidates and executed by the agent, following the scheme from the previous section.

Generating control sequences at random is very sample-inefficient, as the majority of sampled trajectories pass through obstacles in the environment. Instead we follow a heuristic approach, exploiting the laser range nature of our sensors. First we define an obstacle penalty $\mathcal{L}_{\text{obstacle}}(\mathbf{z})$ by building an occupancy map of the environment using $q(\mathcal{M})$ and the DVBF-LM emission model. We then draw F random control sequences and form a set of candidates, minimising $\mathcal{L}_{\text{obstacle}}$ in expectation over the model:

$$\bar{\mathbf{u}}_{1:T-1}^{(f)} = \arg \min_{\mathbf{u}_{1:T-1}} \mathbb{E}_{\mathbf{z}_{1:T} \sim p(\cdot | \mathbf{u}_{1:T-1}, \mathcal{D})} \left[\sum_{t=1}^T \mathcal{L}_{\text{obstacle}}(\mathbf{z}_t) \right], \quad (7)$$

with $f = 1, \dots, F$. We then approximate Equation (5) via

$$\mathbf{u}_{1:T-1}^* = \arg \max_{\mathbf{u}_{1:T-1} \in \mathcal{U}} I(\mathbf{x}_{1:T-1}; \mathcal{M} | \bar{\mathbf{u}}_{1:T-1}^{(f)}).$$

Figure 7a illustrates the described steps.

D. Results

The exploration experiments were carried out in our pybox2d simulator. All model components are learned during exploration, with the exception of the transition model, which is handled in

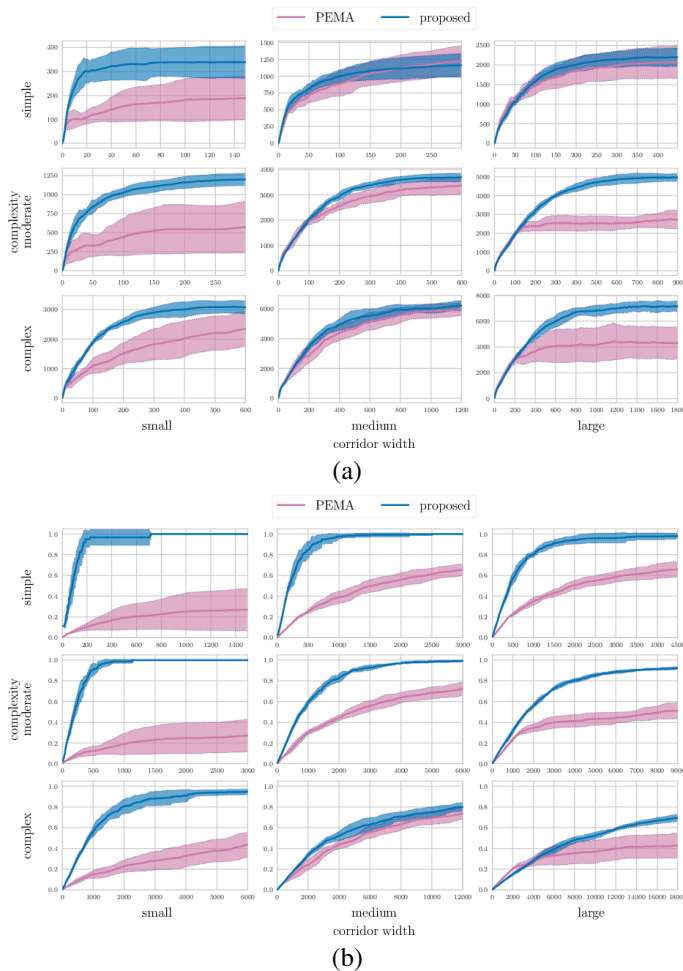


Fig. 9: Comparison of the proposed method (blue) against the LSTM baseline (violet): (a) infogain (b) exploration ratio. Higher values are better. The x axis shows the number of interaction steps divided by 10. Shaded area corresponds to one standard deviation over seven runs.

the same way as in the navigation experiments. We evaluate explorative performance based on two metrics—information gain over the course of exploration and the fraction of “tiles” visited by the agent, the exploration ratio. For the latter, we divide the mazes into a fixed number of equally sized tiles.

As a baseline we consider a method that directly optimises the second metric, the exploration ratio. The baseline agent is represented by a deep deterministic LSTM network. Given a sequence of observations as its input, it puts out a control signal at each time step. We refer to it as the pose-entropy maximising agent (PEMA). Since the exploration ratio objective is non-differentiable w.r.t. the LSTM parameters, we use Augmented Random Search [16] to optimise it.

Qualitatively, we see that the proposed method rapidly traverses the maze. DVBF-LM exploration exhibits nearly uniform coverage, being driven to places that are visited least and are thus most uncertain (see Figure 7). This holds even

for the most complex mazes we considered, as we show in Figure 8.

The quantitative evaluation shows that our method consistently and significantly outperforms PEMA, even though PEMA is directly trained on the exploration ratio evaluation criterion. Figure 9 summarises the comparison for the different metrics over time, aggregated over multiple runs in mazes with different complexity and corridor width.

VII. CONCLUSION AND FUTURE WORK

We have introduced a deep variational Bayes filter that integrates a global latent variable of spatial form. The novelty of our contribution lies in the flexibility that is inherited from neural networks and variational inference: contrary to most recent work in the area, our model still constitutes a generative model, which allows for a number of essential types of inference in spatial environments. We validated the proposed method by applying it to the problems of SLAM, autonomous exploration and navigation. Our model exhibits competitive localisation performance in comparison to an existing 2D LiDAR SLAM system, outperforms a strong baseline for exploration and can be used as a simulator for planning with virtually no loss in performance. The results bear promise for real world application, which we will address in upcoming studies, along with comparisons to state-of-the-art visual SLAM methods.

REFERENCES

- [1] Shi Bai, Jinkun Wang, Fanfei Chen, and Brendan Englot. Information-theoretic exploration with bayesian optimization. In *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on*, pages 1816–1822. IEEE, 2016.
- [2] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural networks. *CoRR*, abs/1505.05424, 2015. URL <http://arxiv.org/abs/1505.05424>.
- [3] Cesar Cadena, Luca Carlone, Henry Carrillo, Yasir Latif, Davide Scaramuzza, José Neira, Ian D. Reid, and John J. Leonard. Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age. *IEEE Trans. Robotics*, 32(6):1309–1332, 2016. doi: 10.1109/TRO.2016.2624754. URL <https://doi.org/10.1109/TRO.2016.2624754>.
- [4] Stefan Depeweg, Jose-Miguel Hernandez-Lobato, Finale Doshi-Velez, and Steffen Udluft. Decomposition of uncertainty in Bayesian deep learning for efficient and risk-sensitive learning. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 1192–1201. PMLR, 10–15 Jul 2018. URL <http://proceedings.mlr.press/v80/depeweg18a.html>.
- [5] Dmitri Dolgov, Sebastian Thrun, Michael Montemerlo, and James Diebel. Path planning for autonomous vehicles in unknown semi-structured environments. *The International Journal of Robotics Research*, 29(5):485–501, 2010. doi:

- 10.1177/0278364909359210. URL <https://doi.org/10.1177/0278364909359210>.
- [6] Alberto Elfes. Robot navigation: Integrating perception, environmental constraints and task execution within a probabilistic framework. In *Reasoning with Uncertainty in Robotics*, pages 91–130. Springer, 1996.
- [7] Marco Fraccaro, Danilo Jimenez Rezende, Yori Zwols, Alexander Pritzel, S. M. Ali Eslami, and Fabio Viola. Generative temporal models with spatial memory for partially observed environments. *CoRR*, abs/1804.09401, 2018. URL <http://arxiv.org/abs/1804.09401>.
- [8] N. J. Gordon, D. J. Salmond, and A. F. M. Smith. Novel approach to nonlinear/non-gaussian bayesian state estimation. *IEE Proceedings F - Radar and Signal Processing*, 140(2):107–113, April 1993. ISSN 0956-375X. doi: 10.1049/ip-f-2.1993.0015.
- [9] P. E. Hart, N. J. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2):100–107, July 1968. ISSN 0536-1567. doi: 10.1109/TSSC.1968.300136.
- [10] Wolfgang Hess, Damon Kohler, Holger Rapp, and Daniel Andor. Real-time loop closure in 2d lidar slam. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1271–1278, 2016.
- [11] Rein Houthoofd, Xi Chen, Yan Duan, John Schulman, Filip De Turck, and Pieter Abbeel. Vime: Variational information maximizing exploration. In *Advances in Neural Information Processing Systems*, pages 1109–1117, 2016.
- [12] Maximilian Karl, Maximilian Sölch, Justin Bayer, and Patrick van der Smagt. Deep variational bayes filters: Unsupervised learning of state space models from raw data. *CoRR*, abs/1605.06432, 2016. URL <http://arxiv.org/abs/1605.06432>.
- [13] Michał Kempka, Marek Wydmuch, Grzegorz Runc, Jakub Toczek, and Wojciech Jaśkowski. ViZDoom: A Doom-based AI research platform for visual reinforcement learning. In *IEEE Conference on Computational Intelligence and Games*, pages 341–348, Santorini, Greece, Sep 2016. IEEE. URL <http://arxiv.org/abs/1605.02097>. The best paper award.
- [14] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *Proceedings of the 2nd International Conference on Learning Representations (ICLR)*, 2014.
- [15] David JC MacKay. Information-based objective functions for active data selection. *Neural computation*, 4(4):590–604, 1992.
- [16] Horia Mania, Aurelia Guy, and Benjamin Recht. Simple random search provides a competitive approach to reinforcement learning. *CoRR*, abs/1803.07055, 2018. URL <http://arxiv.org/abs/1803.07055>.
- [17] Kevin P. Murphy. Bayesian map learning in dynamic environments. In *Advances in Neural Information Processing Systems 12, [NIPS Conference, Denver, Colorado, USA, November 29 - December 4, 1999]*, pages 1015–1021, 1999. URL <http://papers.nips.cc/paper/1716-bayesian-map-learning-in-dynamic-environments>.
- [18] Emilio Parisotto and Ruslan Salakhutdinov. Neural map: Structured memory for deep reinforcement learning. *CoRR*, abs/1702.08360, 2017. URL <http://arxiv.org/abs/1702.08360>.
- [19] Nikolay Savinov, Alexey Dosovitskiy, and Vladlen Koltun. Semi-parametric topological memory for navigation. *CoRR*, abs/1803.00653, 2018. URL <http://arxiv.org/abs/1803.00653>.
- [20] Jürgen Schmidhuber. Curious model-building control systems. In *Neural Networks. 1991 IEEE International Joint Conference on*, pages 1458–1463. IEEE, 1991.
- [21] Cyrill Stachniss, Giorgio Grisetti, and Wolfram Burgard. Information gain-based exploration using rao-blackwellized particle filters. In *Robotics: Science and Systems I, June 8-11, 2005, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA*, pages 65–72, 2005. URL <http://www.roboticsproceedings.org/rss01/p09.html>.
- [22] Zoltán Szabó. Information theoretical estimators toolbox. *Journal of Machine Learning Research*, 15:283–287, 2014.
- [23] S. Thrun. The role of exploration in learning control. In D.A. White and D.A. Sofge, editors, *Handbook for Intelligent Control: Neural, Fuzzy and Adaptive Approaches*. Van Nostrand Reinhold, Florence, Kentucky 41022, 1992.
- [24] Chaoqun Wang, Lili Meng, Teng Li, Clarence W De Silva, and Max Q-H Meng. Towards autonomous exploration with information potential field in 3d environments. In *Advanced Robotics (ICAR), 2017 18th International Conference on*, pages 340–345. IEEE, 2017.
- [25] Peter Whaite and Frank P Ferrie. Autonomous exploration: driven by uncertainty. In *Conference on Computer Vision and Pattern Recognition, CVPR 1994, 21-23 June, 1994, Seattle, WA, USA*, pages 339–346, 1994. doi: 10.1109/CVPR.1994.323849. URL <https://doi.org/10.1109/CVPR.1994.323849>.
- [26] Peter Whaite and Frank P Ferrie. Autonomous exploration: Driven by uncertainty. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(3):193–205, 1997.
- [27] Kai Xu, Lintao Zheng, Zihao Yan, Guohang Yan, Eugene Zhang, Matthias Niessner, Oliver Deussen, Daniel Cohen-Or, and Hui Huang. Autonomous reconstruction of unknown indoor scenes guided by time-varying tensor fields. *ACM Transactions on Graphics (TOG)*, 36(6):202, 2017.