Particle-Based Score Estimation for State Space Model Learning in Autonomous Driving

Anonymous Author(s) Affiliation Address email

Abstract: Multi-object state estimation is a fundamental problem for robotic ap-1 plications where a robot must interact with other moving objects. Typically, other 2 objects' relevant state features are not directly observable, and must instead be 3 inferred from observations. Particle filtering can perform such inference given 4 approximate transition and observation models. However, these models are often 5 unknown a priori, yielding a difficult parameter estimation problem since obser-6 vations jointly carry transition and observation noise. In this work, we consider 7 learning maximum-likelihood parameters using particle methods. Recent meth-8 ods addressing this problem typically differentiate through time in a particle filter, 9 which requires workarounds to the non-differentiable resampling step, that yield 10 biased or high variance gradient estimates. By contrast, we exploit Fisher's iden-11 tity to obtain a particle-based approximation of the score function (the gradient of 12 the log likelihood) that yields a low variance estimate while only requiring step-13 wise differentiation through the transition and observation models. We apply our 14 method to real data collected from autonomous vehicles (AVs) and show that it 15 learns better models than existing techniques and is more stable in training, yield-16 ing an effective smoother for tracking the trajectories of vehicles around an AV. 17

Keywords: Autonomous Driving, Particle Filtering, Self-supervised Learning

19 **1** Introduction

18

Multi-object state estimation is a fundamental problem in settings where a robot must interact with 20 other moving objects, since their state is directly relevant for decision making. Typically, other 21 objects' relevant state features are not directly observable. Instead, the robot must infer them from 22 a stream of observations it receives via a perception system. For example, an autonomous vehicle 23 (AV) selects actions based on the state of nearby road users. However, such road users are only 24 partially observed, owing to limited field of view, occlusions, and imperfections in the AV's sensors 25 and perception systems. Such partial observability negatively affects many downstream tasks in a 26 robot's behavioural stack that depend on observations, e.g., action planning. 27

Addressing partial observability requires sequential state estimation, to which Bayesian filtering of-28 fers a generic probabilistic approach. In particular, sequential Monte Carlo methods, also known as 29 particle filtering, have been successfully applied to state estimation in many robotics applications [1]. 30 However, Bayesian filters require models that reasonably approximate the transition and observation 31 models of a state-space model (SSM). In some special cases, these models can be derived analyti-32 cally from first principles, e.g., when the physical dynamics are well understood, or by modeling a 33 sensor's physical characteristics. In many real-world applications, however, these models cannot be 34 specified analytically. For example, the transition model may encode complicated motion dynamics 35 and environmental physics. In multi-agent settings, other agents' behaviour must also be modelled. 36 Modelling observations is also difficult. Modern perception systems often involve multiple stages 37 and combine information from multiple sensors, making observation models practically impossible 38 to specify by hand. By contrast, collecting observations from a robotic system is relatively easy and 39

^{*}These authors contributed equally to this work.

40 cheap. We are interested, therefore, in algorithms that can leverage such observations to learn tran-

41 sition and observation models in a self-supervised fashion, and yield an effective particle smoother.

42 Learned transition and observation models can also be independently useful for other applications,

⁴³ such as the evaluation of AVs by simulating realistic observations.

In this work, we propose Particle Filtering-Based Score Estimation using Fisher's Identity (PF-44 SEFI), a method for jointly learning maximum-likelihood parameters of both the transition and 45 observation models, that is applicable to a wide class of SSMs. Unlike many recently proposed 46 methods [2, 3, 4, 5, 6, 7], our approach avoids differentiable approximations of the resampling step. 47 We achieve this by revisiting a methodology originally proposed in statistics [8, 9] that relies on a 48 particle approximation of the score, i.e., the gradient of the log likelihood of observation sequences, 49 obtained through Fisher's identity. This only requires differentiating through the transition and 50 observation models. Unfortunately, a direct particle approximation of this identity provides a high 51 variance estimate of the score. While [8] propose an alternative low variance estimate, it admits a 52 $\mathcal{O}(N^2)$ cost, where N is the number of particles. Furthermore, these methods compute and store 53 the gradient of the marginal log-likelihood with respect to model parameters for each particle. This 54 requires computing Jacobian matrices, which are slow to compute using automatic differentiation 55 tools such as TensorFlow and PyTorch [10, 11] which rely on Jacobian-vector products. This makes 56 these methods impractical for large models. By contrast, PF-SEFI is a simple scalable $\mathcal{O}(N)$ variant 57 with only negligible bias. PF-SEFI marginalises over particles before computing gradients, allowing 58 automatic differentiation tools to make use of efficient Jacobian-vector product operations, making it 59 significantly faster and allowing us to scale to larger models. To the best of our knowledge, previous 60 particle methods estimating the score have been limited to SSMs with few parameters, whereas we 61 apply PF-SEFI to neural network models with thousands of parameters. 62

We apply PF-SEFI to jointly learn transition and observation models for tracking multiple objects 63 around an AV, using a large set of noisy trajectories, containing almost 10 hours of road-user tra-64 jectories observed by an AV. We show that PF-SEFI learns an SSM that yields an effective object 65 tracker as measured by average displacement and yaw errors. We compare the learned observation 66 model to one trained through supervised learning on a dataset of manually labelled trajectories, and 67 show that PF-SEFI yields a better model (as measured by log-likelihood on ground-truth labels) 68 even though it requires no labels for training. Finally, we compare PF-SEFI to a number of existing 69 particle methods for jointly learning transition and observation models and show that it learns better 70 models and is more stable to train. 71

72 2 Related Work

Particle filters are widely used for state estimation in non-linear non-Gaussian SSMs where no closed 73 74 form solution is available; see e.g., [12] for a survey. The original bootstrap particle filter [13] samples at each time step using the transition density particles that are then reweighted according to 75 their conditional likelihood, which measures their "fitness" w.r.t. to the available observation. Parti-76 cles with low weights are then eliminated while particles with high weights are replicated to focus 77 computational efforts into regions of high probability mass. Compared to many newer methods, 78 such as the auxiliary particle filter [14], the bootstrap particle filter only requires sampling from the 79 transition density, not its evaluation at arbitrary values, which is not possible for the compositional 80 81 transition density used in this work.

In most practical applications, the SSM has unknown parameters that must be estimated together with the latent state posterior (see [9] for a review). Simply extending the latent space to include the unknown parameters suffers from insufficient parameter space exploration [15]. While particle filters can estimate consistently the likelihood for fixed model parameters, a core challenge is that the such estimated likelihood function is discontinuous in the model parameters due to the resampling step, hence complicating its optimization; see e.g. [6, Figure 1] for an illustration.

Instead, the score vector can be computed using Fisher's identity [8]. However, as shown in [8], performance degrades quickly for longer sequences if a standard particle filter is used, due to the path degeneracy problem: repeated resampling of particles and their ancestors will leave few or even just one remaining ancestor path for earlier timesteps, resulting in unbiased, but very high variance estimates. Methods for overcoming this limitation exist [8, 16, 17], but with requirements making them unsuitable in this work. Poyiadjis et al. [8] store gradients separately for each particle, making

this approach infeasible for all but the smallest neural networks. Scibior and Wood [17] propose 94 an improved implementation with lower memory requirements by smartly using automatic differen-95 tiation. However, their approach still requires storing a computation graph whose size scales with 96 $\mathcal{O}(N^2)$ as the transition density for each particle *pair* must be evaluated during the forward pass. 97 Both previous methods' computational complexity also scales quadratically with the number of par-98 ticles, N, which is problematic for costly gradient backpropagation through large neural networks. 99 Lastly, Olsson and Westerborn [16] require evaluation of the transition density for arbitrary values, 100 which our compositional transition model does not allow. Instead, in this work, we show that fixed-101 lag smoothing [18, 19] is a viable alternative to compute the score function of large neural network 102 models in the context of extended object tracking. 103

There is extensive literature on combining particle filters with learning complex models such as 104 neural networks [2, 3, 4, 5, 6, 20, 21, 22, 23, 24]. In contrast to our work, they make use of a 105 learned, data-dependent proposal distribution. However, for parameter estimation, they rely on dif-106 ferentiation of an estimated lower bound (ELBO). Due to the non-differentiable resampling step, 107 this gradient estimation has either extremely high variance or is biased if the high variance terms 108 are simply dropped, as in [2, 3, 4]. As we show in Section 5, this degrades performance noticeably. 109 A second line of work proposes soft resampling [5, 20, 21], which interpolates between regular 110 and uniform sampling, thereby allowing to trade off variance reduction through resampling with 111 the bias introduced by ignoring the non-differentiable component of resampling. Lastly, Corenflos 112 et al. [6] make the resampling step differentiable by using entropy-regularized optimal transport, 113 also inducing bias and a $\mathcal{O}(N^2)$ cost. 114

Extended object tracking [25] considers how to track objects which, in contrast to "small" objects [26], generate multiple sensor measurements per timestep. Unlike in our work, transition and measurement models are assumed to be known or to depend on only a few learnable parameters. Similar to our work, the measurement model proposed in [27] assumes measurement sources lying on a rectangular shape. However, our model is more flexible, for example, allowing non-zero probability on all four sides simultaneously.

121 **3** State-Space Models and Particle Filtering

122 3.1 State-Space Models

A SSM is a partially observed discrete-time Markov process with initial density, $x_0 \sim \mu(\cdot)$, transition density $x_t | x_{t-1} \sim f_{\theta}(\cdot | x_{t-1})$, and observation density $y_t | x_t \sim g_{\theta}(\cdot | x_t)$, where x_t is the latent state at time t and y_t the corresponding observation. The joint density of $x_{0:T}$, $y_{0:T}$ satisfies:

$$p_{\theta}(x_{0:T}, y_{0:T}) = \mu(x_0)g_{\theta}(y_0|x_0) \prod_{t=1}^{I} f_{\theta}(x_t|x_{t-1})g_{\theta}(y_t|x_t).$$
(1)

Given this model, we are typically interested in inferring the states from the data by computing the filtering and one-step ahead prediction distributions, $\{p(x_t|y_{0:t})\}_{t\in 0,...,T}$ and $\{p(x_{t+1}|y_{0:t})\}_{t\in 0,...,T-1}$ respectively, and more generally the joint distributions $\{p(x_{0:t}|y_{0:t})\}_{t\in 0,...,T}$ satisfying

$$p_{\theta}(x_{0:t}|y_{0:t}) = \frac{p_{\theta}(x_{0:t}, y_{0:t})}{p_{\theta}(y_{0:t})}, \qquad p_{\theta}(y_{0:T}) = \int p_{\theta}(x_{0:T}, y_{0:T}) \mathrm{d}x_{0:T}.$$
(2)

Additionally, to estimate parameters, we would also like to compute the marginal log likelihood:

$$\ell_T(\theta) = \log p_\theta(y_{0:T}) = \log p_\theta(y_0) + \sum_{t=1}^T \log p_\theta(y_t|y_{0:t-1}),$$
(3)

where $p_{\theta}(y_0) = \int g_{\theta}(y_0|x_0)\mu(x_0) dx_0$ and $p_{\theta}(y_t|y_{0:t-1}) = \int g_{\theta}(y_t|x_t)p_{\theta}(x_t|y_{0:t-1}) dx_t$ for $t \ge 1$.

For non-linear non-Gaussian SSMs, these posterior distributions and the corresponding marginal likelihood cannot be computed in closed form.

134 3.2 Particle Filtering

Particle methods provide non-parametric and consistent approximations of these quantities. They rely on the combination of importance sampling and resampling steps of a set of N weighted particles (x_t^i, w_t^i) , where x_t^i denotes the values of the *i*th particle at time *t* and w_t^i is corresponding weight satisfying $\sum_{i=1}^{N} w_t^i = 1$. We focus on the bootstrap particle filter, shown in Algorithm 1, which samples particles according to the transition density.

Algorithm 1 Bootstrap Particle Filter

Sample $X_0^i \stackrel{\text{i.i.d.}}{\sim} \mu(\cdot)$ for $i \in [N]$ and set $\hat{\ell}_0(\theta) \leftarrow \log\left(\frac{1}{N}\sum_{i=1}^N g_\theta(y_0|x_0^i)\right)$. For t = 1, ..., T1. Compute weights $w_{t-1}^i \propto g_\theta(y_{t-1}|x_{t-1}^i)$ with $\sum_{i=1}^N w_{t-1}^i = 1$. 2. Sample $a_{t-1}^i \sim \text{Cat}(w_{t-1}^1, ..., w_{t-1}^N)$ then $x_t^i \sim f_\theta(\cdot|x_{t-1}^{a_{t-1}^i})$ for $i \in [N]$. 3. Set $x_{0:t}^i \leftarrow (x_{0:t-1}^{a_{t-1}^i}, x_t^i)$ for $i \in [N]$ and $\hat{\ell}_t(\theta) \leftarrow \hat{\ell}_{t-1}(\theta) + \log\left(\frac{1}{N}\sum_{i=1}^N g_\theta(y_t|x_t^i)\right)$.

Let $k \sim \text{Cat}(\alpha_1, ..., \alpha_N)$ denote the categorical distribution with N categories, where the probability of the k taking the *i*th category is α_i . At any time t, this algorithm produces particle approximations

$$\hat{p}_{\theta}(x_{0:t}|y_{0:t}) = \sum_{i=1}^{N} w_t^i \delta_{x_{0:t}^i}(x_{0:t}), \qquad \hat{\ell}_t(\theta) = \sum_{t=0}^{T} \log\left(\frac{1}{N} \sum_{i=1}^{N} g_{\theta}\left(y_t|x_t^i\right)\right), \tag{4}$$

of $p_{\theta}(x_{0:t}|y_{0:t})$ and $\ell_t(\theta) = \log p_{\theta}(y_{0:t})$, where δ_{α} is the Dirac delta distribution centred at α . Step 2 resamples, discarding particles with small weights while replicating those with large weights before evolving according to the transition density. This focuses computational effort on the "promising" regions of the state space. Unfortunately, resampling involves sampling N discrete random variables at each time step and as such produces estimates of the log likelihood that are not differentiable w.r.t. θ as illustrated in [6, Figure 1].

While the resulting estimates are consistent as $N \to \infty$ for any fixed time t [28], this does not guarantee good practical performance. Fortunately, under regularity conditions the approximation error for the estimate $\hat{p}_{\theta}(x_t|y_{0:t})$ and more generally $\hat{p}_{\theta}(x_{t-L+1:t}|y_{0:t})$ for a fixed lag $L \ge 1$ as well as $\log p_{\theta}(y_{0:t})/t$ does not increase with t for fixed N. However, this is not the case for the joint smoothing approximation because successive resampling means that $\hat{p}_{\theta}(x_{0:L}|y_{0:t})$ is eventually approximated by a single unique particle for large enough t, a phenomenon known as path degeneracy; see e.g. [12, Section 4.3].

156 4 Score Estimation using Particle Methods

To estimate the parameters θ of a given SSM (1) along with a dataset of observations $y_{0:T}$, we want to maximise via gradient ascent the marginal log likelihood in (3). However, the gradient of the marginal log likelihood, i.e., the *score function*, is intractable. As explained in Section 2, automatic differentiation through the filter is difficult due to the non-differentiable resampling step.

161 4.1 Score Function Using Fisher's Identity

We leverage here instead Fisher's identity [8] for the score to completely side-step the nondifferentiability problem. This identity shows that

$$\nabla_{\theta} \ell_T(\theta) = \int \nabla_{\theta} \log p_{\theta}(x_{0:T}, y_{0:T}) p_{\theta}(x_{0:T} | y_{0:T}) \mathrm{d}x_{0:T},$$
(5)

i.e., the score is the expectation of $\nabla_{\theta} \log p_{\theta}(x_{0:T}, y_{0:T})$ under the joint smoothing distribution $p_{\theta}(x_{0:T}|y_{0:T})$. Plugging in (1), the score function can be simplified to

$$\nabla_{\theta} \ell_{T}(\theta) = \sum_{t=0}^{T} \int \nabla_{\theta} \log g_{\theta}(y_{t}|x_{t}) p_{\theta}(x_{t}|y_{0:T}) \mathrm{d}x_{t}$$
$$+ \sum_{t=1}^{T} \int \nabla_{\theta} \log f_{\theta}(x_{t}|x_{t-1}) p_{\theta}(x_{t-1:t}|y_{0:T}) \mathrm{d}x_{t-1:t}.$$
(6)

166 4.2 Particle Score Approximation

The identity (6) shows that we can simply estimate the score by plugging particle approximations 167 of the marginal smoothing distributions $p(x_{t-1:t}|y_{0:T})$ into (6). This identity makes differentiating 168 through time superfluous and thereby renders the use of differentiable approximations of resampling 169 unnecessary. However, as discussed in Section 3.2, naive particle approximations of the smoothing 170 distribution's marginals, $p_{\theta}(x_t|y_{0:T})$ and $p_{\theta}(x_{t-1:t}|y_{0:T})$, suffer from path degeneracy. To bypass 171 this problem, [8, 17] propose an $\mathcal{O}(N^2)$ method inspired by dynamic programming. We propose 172 here a simpler and computationally cheaper method that relies on the following fixed-lag approxi-173 mation of the fixed-interval smoothing distribution, which states that for $L \ge 1$ large enough, 174

$$p_{\theta}(x_{t-1:t}|y_{0:T}) \approx p_{\theta}\left(x_{t-1:t}|y_{0:\min\{t+L,T\}}\right). \tag{7}$$

This approximation simply assumes that observations after time t + L do not bring further infor-175 mation about the states x_{t-1}, x_t . This is satisfied for most models and the resulting approximation 176 error decreases geometrically fast with L [19]. The benefit of this approximation is that the particle 177 approximation of $p_{\theta}(x_{t-1:t}|y_{0:\min\{t+L,T\}})$ does not suffer from path degeneracy and is a simple 178 byproduct of the bootstrap particle filtering of Algorithm 1; e.g., for t + L < T we consider the particle approximation $\hat{p}_{\theta}(x_{0:t+L}|y_{0:t+L}) = \sum_{i=1}^{N} w_{t+L}^i \delta_{x_{0:t+L}^i}(x_{0:t+L}^i)$ obtained at time t + L and 179 180 use its corresponding marginals in x_{t-1}, x_t and x_t to integrate respectively $\nabla_{\theta} \log f_{\theta}(x_t | x_{t-1})$ and 181 $\nabla_{\theta} \log g_{\theta}(y_t | x_t)$. For $t + L \ge T$, we just consider the marginals in x_{t-1}, x_t and x_t of $\hat{p}_{\theta}(x_{0:T} | y_{0:T})$. 182 So finally, we consider the estimate, 183

$$\widehat{\nabla_{\theta}\ell_{T}}(\theta) = \sum_{t=0}^{T} \int \nabla_{\theta} \log g_{\theta}(y_{t}|x_{t}) \, \hat{p}_{\theta}(x_{t}|y_{0:\min\{t+L,T\}}) \mathrm{d}x_{t} + \sum_{t=1}^{T} \int \nabla_{\theta} \log f_{\theta}(x_{t}|x_{t-1}) \, \hat{p}_{\theta}(x_{t-1:t}|y_{0:\min\{t+L,T\}}) \mathrm{d}x_{t-1:t}.$$
(8)

4.3 Score Estimation with Deterministic, Differentiable, Injective Motion Models

We have described a generic method to approximate the score using particle filtering techniques. For many applications, however, the transition density function, $f_{\theta}(x_t|x_{t-1})$, is the composition of a *policy*, $\pi_{\theta}(a_t|x_{t-1})$, which characterises the action distribution conditioned on the state, and a potentially complex but deterministic, differentiable, and injective *motion model*, $\tau : \mathbb{R}^{n_x} \times \mathbb{R}^{n_a} \to \mathbb{R}^{n_x}$ where $n_a < n_x$, which characterises kinematic constraints such that $x_t = \tau(x_{t-1}, a_t) = \bar{\tau}_{x_{t-1}}(a_t)$. Under such a composition, the transition density function on the induced manifold $\mathcal{M}_{x_{t-1}} =$ $\{\bar{\tau}_{x_{t-1}}(a_t) : a_t \in \mathbb{R}^{n_a}\}$ is thus obtained by marginalising out the latent action variable, i.e.,

$$f_{\theta}(x_t|x_{t-1}) = \mathbb{I}(x_t \in \mathcal{M}_{x_{t-1}}) \int \delta(x_t - \bar{\tau}_{x_{t-1}}(a_t)) \pi_{\theta}(a_t|x_{t-1}) \mathrm{d}a_t.$$
(9)

It is easy to sample from this density but it is intractable analytically if the motion model is only 192 available through a complex simulator or if it is not invertible. This precludes the use of sophisticated 193 proposal distributions within the particle filter. Additionally, even if it were known, one cannot use 194 the $\mathcal{O}(N^2)$ smoothing type algorithms developed in [8, 16] as the density is concentrated on a low-195 dimensional manifold [29]. This setting is common in mobile robotics, in which controllers factor 196 into policies that select actions and motion models that determine the next state. Indeed, this is 197 precisely the case in our application setting of estimating the state of observed road users around 198 an AV (see Section 5). Learning the corresponding SSM reduces to learning the parameters θ of 199 the policy, $\pi_{\theta}(a_t|x_{t-1})$, and the observation model, $q_{\theta}(y_t|x_t)$. Thankfully, even if the explicit form 200 of the motion model is unknown, we can still compute $\nabla \log f_{\theta}(x_t|x_{t-1})$ as required by the score 201 estimate (8). 202

Lemma 4.1. For any $x \in \mathbb{R}^{n_x}$, let $\tau_x : \mathbb{R}^{n_a} \to \mathbb{R}^{n_x}$ where $n_a < n_x$ be a smooth and injective mapping. Then, for any fixed x_{t-1} and $x_t \in \mathcal{M}_{x_{t-1}}$, the gradient of the transition log density, i.e., $\nabla_{\theta} \log f_{\theta}(x_t|x_{t-1})$, reduces to the gradient of the policy log density, i.e., $\nabla_{\theta} \log \pi_{\theta}(a_t|x_{t-1})$, where a_t is the unique action that takes x_{t-1} to x_t .

Proof. For x_{t-1} and $x_t \in \mathcal{M}_{x_{t-1}}$, we denote by $J[\bar{\tau}_{x_{t-1}}](\bar{\tau}_{x_{t-1}}^{-1}(x_t)) \in \mathbb{R}^{n_x \times n_a}$ the rectangular Jacobian matrix and write $a_t = \bar{\tau}_{x_{t-1}}^{-1}(x_t)$, i.e., this is the unique action such $\bar{\tau}_{x_{t-1}}(a_t) = x_t$. By a



(a) Observations from the real data.

(b) Observations from the synthetic data.

Figure 1: Observations from real and synthetic data. In each plot, the AV (orange) observes a set of 2D points (blue) forming a convex polygon around the observed road users. The actual bounding boxes corresponding to each observed road user (green) were manually labelled in the real data, and are determined while constructing the synthetic data.

standard result from differential geometry [30, 31], the transition density (9) satisfies

$$f_{\theta}(x_t|x_{t-1}) = \pi_{\theta}\left(a_t|x_{t-1}\right) \left| \det J[\bar{\tau}_{x_{t-1}}]^{\mathrm{T}}(a_t) J[\bar{\tau}_{x_{t-1}}](a_t) \right|^{-1/2} \mathbb{I}(x_t \in \mathcal{M}_{x_{t-1}}).$$
(10)

It follows directly that $\nabla_{\theta} \log f_{\theta}(x_t | x_{t-1}) = \nabla_{\theta} \log \pi_{\theta} (a_t | x_{t-1}).$

- Indeed for the marginals $\hat{p}_{\theta}(x_{t-1:t}|y_{0:\min\{t+L,T\}})$, we can store the actions corresponding to tran-
- sitions $x_{t-1} \to x_t$ during filtering, and it follows that for the class of SSMs described above, the
- 213 score estimate reduces to:

$$\widehat{\nabla_{\theta}\ell_{T}}(\theta) = \sum_{t=0}^{T} \int \nabla_{\theta} \log g_{\theta}(y_{t}|x_{t}) \, \hat{p}_{\theta}(x_{t}|y_{0:\min\{t+L,T\}}) \mathrm{d}x_{t} + \sum_{t=1}^{T} \int \nabla_{\theta} \log \pi_{\theta}(a_{t}|x_{t-1}) \, \hat{p}_{\theta}(x_{t-1:t}|y_{0:\min\{t+L,T\}}) \mathrm{d}x_{t-1:t}, \qquad (11)$$

where we use Lemma 4.1 to replace the gradient of the transition log density with the gradient of the policy log density in (8), and where a_t is the action sampled to go from x_{t-1} to x_t .

216 **5** Experiments

Problem Setting. Our experiments focus on the problem of state estimation of observed road users 217 (in particular other vehicles) from the viewpoint of an AV, which involves the estimation of 2D 218 poses from an observed sequence of 2D convex polygons in a "bird's eye view" (BEV) constructed 219 from LiDAR point clouds at each time step. For these experiments, we assume that the size of the 220 observed objects, the pose of the AV, and the association of observations with their corresponding 221 222 objects are known a priori. Some observations (and their corresponding states) are shown in Figure 1a. Here, the observation model must learn to describe the likelihood of 2D points around the pe-223 riphery of the observed road user (see [25] for a review on such models), while the transition model 224 225 must learn to describe driving behaviour. We use a feed-forward neural network to parameterise our observation model, where we provide it with range, bearing, and relative bearing from the viewpoint 226 of the corresponding AV as features (Appendix A), and factor our transition model into a determin-227 istic and differentiable motion model based on Ackermann dynamics [32] (Appendix B.1), and a 228 policy parameterised by another feed-forward neural network (Appendix B.2). 229

Baselines, Datasets, and Metrics. We compare the quality of the models learned using PF-SEFI (our method), DPF-SGR [17], PFNET [5], and differentiating through a vanilla PF (ignoring the bias introduced by resampling). We evaluate our method (and the baselines) on real data collected from an AV in an urban environment, equipped with LiDARs, cameras, and radar sensors. All sensors were used to associate LiDAR points to their corresponding objects, and the observations shown in Figure 1a were obtained via a convex hull computation of the associated LiDAR points.

In addition to using real data, we generate two synthetic datasets (with 25 and 50 step trajectories), using a hand-crafted policy, and an observation model trained using supervised learning on manually labelled trajectories (see Appendix A and B for more details). Example observations are shown in Figure 1b. Unlike with real data, where the true models are unknown, synthetic datasets allow us to compare the learned models against a known ground truth. We measure the quality of learned models using the following metrics:

- Marginal Log Likelihood (MLL): The marginal log likelihood $\ell_T(\theta)$ given by filtering observations $y_{0:T}$ using the learned models.
- Average Displacement Error (ADE) and Average Yaw Error (AYE): The average error in the positions and yaws respectively of the smoothed state estimates $\mathbb{E}_{\theta}(x_{0:T}|y_{0:T})$ against the true poses, $\bar{x}_{0:T}$. For the synthetic data, the true poses are sampled while generating the data; for the real data, the true poses are obtained by humans manually labelling object trajectories from videos. These measure the quality of the learned models for the purposes of state estimation.
- Average Observation True Log Likelihood (AOTLL): The average log likelihood of observations conditioned on the corresponding true states under the learned observation model, i.e., $\frac{1}{T+1}\sum_{t=0}^{T} \log g_{\theta}(y_t|\bar{x}_t)$. This measures the standalone quality of the learned observation model.
- Average Policy True Log Likelihood (APTLL): The average log likelihood of true actions, $\bar{a}_{1:T}$, (only available for experiments with synthetic data since it is not possible to manually label latent actions) conditioned on the corresponding true states under the learned policy, i.e., $\frac{1}{T}\sum_{t=1}^{T} \log \pi_{\theta}(\bar{a}_t | \bar{x}_{t-1})$. This measures the standalone quality of the learned policy.

Results. Figure 2 shows the progress of the learned models by tracking MLL of held out test data 257 for each of the three datasets (synthetic data with 25 steps, synthetic data with 50 steps, and real data 258 with 60 steps), and for each of the four methods (PF-SEFI, DPF-SGR, PFNET, and PF). Table 1 sum-259 marise the performance of the learned models at convergence. We pick the best hyper-parameters, 260 smoothing lag L for PF-SEFI, and trade-off parameter α for PFNET, in each of the experiments. 261 Appendix C includes analysis of the training sensitivity of each of the hyper-parameters (see Fig-262 ures 5, 6, and 7). We find that PF-SEFI improves with increasing L up to a point, past which it is 263 insensitive to the choice of L (see Figure 6, Appendix C). 264

In our experiments with synthetic data with 25 steps (Figure 2a and Experiment A in Table 1), we 265 observe a clear gap in performance of PF-SEFI and DPF-SGR relative to PFNET and PF. The im-266 provements over PF are likely due to the bias in PF's score estimates due to the non-differentiable 267 resampling step, while the improvements over PFNET are likely due to adverse effects of not re-268 sampling with the correct distribution at each time step. While PF-SEFI and DPF-SGR perform 269 similarly on this dataset, the difference is stark in the case of synthetic data with 50 steps (Figure 2b 270 271 and Experiment B in Table 1). PF-SEFI is invariant to the length of the trajectories used, converging stably; however, all other methods, struggle to learn useful models. We postulate that since each of 272 the baselines, in one way or another, differentiate through all time steps of the filter, the variance in 273 their score estimates is too high for good learning through gradient ascent.¹ 274

The results of our experiments with *real data* with 60 steps (Figure 2c and Experiment C in Table 275 276 1) are consistent with Experiment B (i.e., with experiments on synthetic data with 50 steps) and 277 show that PF-SEFI is able to learn useful models. The learned observation model using PF-SEFI 278 performs even better than the model that was trained offline through supervision with manually labelled data (see AOTLL in Table 1 for Experiment C). We also find that sampling from the learned 279 model produces observations that are qualitatively similar to the real data (Appendix D). While the 280 supervised model is trained only on the subset of the observations that are labeled (labelling only a 281 subset is common in practical applications due to the cost of labelling), PF-SEFI, by contrast, can 282 leverage *all* observations in a self-supervised fashion. Moreover, we speculate that the labels contain 283 noise and that the labelling distribution is biased towards observations that are easy to label. Both 284 limitations hinder supervised learning. 285

¹The authors of DPF-SGR [17] recommend the use of stop gradients not only for particle weights after resampling, i.e., $\tilde{v}_t^i = \bar{v}_t^{a^i} / \pm \bar{v}_t^{a^i}$ [17, Algorithm 1], but also, in the case of bootstrap particle filters, while computing the likelihood ratio $v_t^i = \tilde{v}_{t-1}^i p_\theta(x_t^i, y_t | x_{t-1}^{a^i}) / \pm q_\theta(x_t^i | x_{t-1}^{a^i})$ before resampling, and while sampling from $x_t^i \sim q_\theta(\cdot | x_{t-1}^{a^i})$ [17, Section 4.1]. While these additional stop gradients significantly reduce variance, our experiments with them yielded extremely poor overall performance (even with synthetic data with 25 steps). The results we report here thus make use of stop-gradients only for particle weights after resampling.



(a) MLL of test synthetic data with (25 step trajectories. 55

(b) MLL of test synthetic data with (c) MLL of test real data with 60 50 step trajectories. step trajectories.

Figure 2: Marginal Log Likelihood (MLL) on synthetic and real test data for models trained using PF-SEFI (us), DPF-SGR, PFNET, and PF, plotted against the corresponding training steps. For synthetic data we also show the MLL of the true models.

Table 1: Metrics computed on held out synthetic test data comparing PF-SEFI (us) against baselines DPF-SGR, PFNET, and vanilla PF, on three experiments - (A) learning from synthetic data with 25 steps, (B) learning from synthetic data with 50 steps, and (C) learning from real data with 60 steps. For experiments (A) and (B), we also compare against the performance of the true models. For experiment (C), we compare against the supervised observation model trained using manually labelled trajectories. For MLL, AOTLL, and APTLL, higher values imply better models, while for ADE and AYE, lower values imply better models.

Exp.	Method	MLL	AOTLL	APTLL	ADE (m)	AYE (rad)
A	TRUE PF-SEFI (us) DPF-SGR PFNET PF	$\begin{array}{c} -3.161 \pm 0.003 \\ \textbf{-3.147} \pm \textbf{0.004} \\ -3.159 \pm 0.004 \\ \textbf{-3.225} \pm 0.004 \\ \textbf{-3.229} \pm 0.006 \end{array}$	$\begin{array}{c} -2.128 \\ -2.285 \pm 0.028 \\ \textbf{-2.265} \pm \textbf{0.010} \\ -2.487 \pm 0.026 \\ -2.484 \pm 0.026 \end{array}$	$\begin{array}{c} 2.674 \\ \textbf{2.661} \pm \textbf{0.014} \\ 2.594 \pm 0.027 \\ 2.621 \pm 0.013 \\ 2.576 \pm 0.021 \end{array}$		$\begin{array}{c} 0.014 \pm 0.000 \\ 0.016 \pm 0.000 \\ \textbf{0.014} \pm \textbf{0.000} \\ 0.017 \pm 0.000 \\ 0.017 \pm 0.000 \end{array}$
В	TRUE PF-SEFI (us) DPF-SGR PFNET PF	$\begin{array}{c} -3.145 \pm 0.002 \\ \textbf{-3.141} \pm \textbf{0.005} \\ -3.966 \pm 0.050 \\ \textbf{-4.169} \pm 0.046 \\ \textbf{-4.118} \pm 0.038 \end{array}$	$\begin{array}{c} -2.165 \\ \textbf{-2.283} \pm \textbf{0.015} \\ -2.636 \pm 0.031 \\ -2.901 \pm 0.039 \\ -2.841 \pm 0.025 \end{array}$	$\begin{array}{c} 2.693 \\ \textbf{2.505} \pm \textbf{0.042} \\ 0.811 \pm 0.130 \\ 0.539 \pm 0.077 \\ 0.681 \pm 0.122 \end{array}$		$\begin{array}{c} 0.012 \pm 0.000 \\ \textbf{0.014} \pm \textbf{0.000} \\ 0.142 \pm 0.016 \\ 0.148 \pm 0.008 \\ 0.137 \pm 0.007 \end{array}$
С	SUPERVISED PF-SEFI (us) DPF-SGR PFNET PF		$\begin{vmatrix} -2.224 \pm 0.006 \\ \textbf{-1.973} \pm \textbf{0.029} \\ -2.236 \pm 0.218 \\ -2.752 \pm 0.079 \\ -2.639 \pm 0.140 \end{vmatrix}$	N/A N/A N/A N/A N/A		

286 6 Discussion, Limitations, and Future Work

In this work, we proposed an efficient particle-based approach for estimating the score function to 287 learn a wide class of SSMs in a completely self-supervised way. Compared to previous particle-288 based methods that estimate the score, our method is more computationally efficient, allowing us to 289 scale to learning models with many parameters. Unlike alternative methods, PF-SEFI is applicable 290 to SSMs where the transition distribution is concentrated on a low-dimensional manifold, allowing 291 us to apply it to a real-world AV object tracking problem. We also showed empirically that our 292 method learns better models and is more stable in training than recent methods that use automatic 293 differentiation to estimate the score, and that we can learn an observation model that outperforms 294 one trained through supervised learning, without using any labels. 295

While this solution is ideal for our problem, it does have a number of limitations. Most notably, it is restricted to maximising the marginal log-likelihood of the data, while differentiating through the filter allows for arbitrary differentiable loss functions. Furthermore, our method is not suitable for estimating the parameters of a proposal distribution. Beyond these algorithmic limitations, in our application, the models that we used were not very expressive. For the observation model, we did not model important phenomena that affect partial observability such as occlusions and we restricted our states and observations to 2D. For the policy, we used a simplified policy with only basic features that are insufficient for controlling an agent in simulation. Furthermore, the policy and the motion model are both specific to vehicles, and currently exclude other road users such as pedestrians.

In future work, we aim to scale up our problem setting, by making both models more expressive, and to estimate more state dimensions, such as full 3D poses and sizes of objects. We also believe that learning policies as components of an SSM to explicitly account for observation noise is, in practice, critical for learning good driving behaviour from demonstrations. Such policies could be used as models for predicting the behaviour of other road-users, or to control agents in simulation, and the method we proposed in this work offers an ideal starting point to explore this.

311 References

- [1] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. MIT Press, 2005.
- [2] T. A. Le, M. Igl, T. Rainforth, T. Jin, and F. Wood. Auto-encoding sequential Monte Carlo. In
 International Conference on Learning Representations, 2018.
- [3] C. J. Maddison, J. Lawson, G. Tucker, N. Heess, M. Norouzi, A. Mnih, A. Doucet, and Y. Teh.
 Filtering variational objectives. *Advances in Neural Information Processing Systems*, 30, 2017.
- [4] C. Naesseth, S. Linderman, R. Ranganath, and D. Blei. Variational sequential Monte Carlo. In *International Conference on Artificial Intelligence and Statistics*, pages 968–977, 2018.
- [5] P. Karkus, D. Hsu, and W. S. Lee. Particle filter networks with application to visual localization.
 In *Conference on Robot Learning*, pages 169–178, 2018.
- [6] A. Corenflos, J. Thornton, G. Deligiannidis, and A. Doucet. Differentiable particle filtering
 via entropy-regularized optimal transport. In *International Conference on Machine Learning*,
 pages 2100–2111, 2021.
- [7] J. Lai, J. Domke, and D. Sheldon. Variational marginal particle filters. In *International Conference on Artificial Intelligence and Statistics*, pages 875–895, 2022.
- [8] G. Poyiadjis, A. Doucet, and S. S. Singh. Particle approximations of the score and observed information matrix in state space models with application to parameter estimation. *Biometrika*, 98(1):65–80, 2011.
- [9] N. Kantas, A. Doucet, S. S. Singh, J. Maciejowski, and N. Chopin. On particle methods for parameter estimation in state-space models. *Statistical science*, 30(3):328–351, 2015.
- [10] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis,
 J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray,
 C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu,
 and X. Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015.
 URL http://tensorflow.org/. Software available from tensorflow.org.
- [11] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin,
 N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.
 cc/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf.
- [12] A. Doucet and A. M. Johansen. A tutorial on particle filtering and smoothing: Fifteen years
 later. *Handbook of Nonlinear Filtering*, 12(656-704):3, 2009.
- [13] N. J. Gordon, D. J. Salmond, and A. F. Smith. Novel approach to nonlinear/non-Gaussian
 Bayesian state estimation. *IEE Proceedings F (Radar and Signal Processing)*, 140(2):107–113, 1993.
- [14] M. K. Pitt and N. Shephard. Filtering via simulation: Auxiliary particle filters. *Journal of the American Statistical Association*, 94(446):590–599, 1999.
- [15] E. L. Ionides, C. Bretó, and A. A. King. Inference for nonlinear dynamical systems. *Proceedings of the National Academy of Sciences*, 103(49):18438–18443, 2006.
- Ifo] J. Olsson and J. Westerborn. Efficient particle-based online smoothing in general hidden
 markov models: the PaRIS algorithm. *Bernoulli*, 23(3):1951–1996, 2017.
- [17] A. Ścibior and F. Wood. Differentiable particle filtering without modifying the forward pass.
 arXiv preprint arXiv:2106.10314, 2021.

- [18] G. Kitagawa and S. Sato. Monte Carlo smoothing and self-organising state-space model. In
 A. Doucet, N. De Freitas, and N. Gordon, editors, *Sequential Monte Carlo Methods in Practice*,
 pages 177–195. Springer, 2001.
- I9] J. Olsson, O. Cappé, R. Douc, and E. Moulines. Sequential Monte Carlo smoothing with
 application to parameter estimation in nonlinear state space models. *Bernoulli*, 14(1):155–179, 2008.
- [20] X. Ma, P. Karkus, D. Hsu, and W. S. Lee. Particle filter recurrent neural networks. In AAAI
 Conference on Artificial Intelligence, 2020.
- X. Ma, P. Karkus, D. Hsu, W. S. Lee, and N. Ye. Discriminative particle filter reinforcement
 learning for complex partial observations. In *International Conference on Learning Representations*, 2020.
- [22] R. Jonschkowski, D. Rastogi, and O. Brock. Differentiable particle filters: End-to-end learning
 with algorithmic priors. In *Proceedings of Robotics: Science and Systems*, 2018.
- [23] M. Zhu, K. Murphy, and R. Jonschkowski. Towards differentiable resampling. *arXiv preprint arXiv:2004.11938*, 2020.
- A. Kloss, G. Martius, and J. Bohg. How to train your differentiable filter. *Autonomous Robots*, 45(4):561–578, 2021.
- ³⁷⁴ [25] K. Granstrom, M. Baum, and S. Reuter. Extended object tracking: Introduction, overview and ³⁷⁵ applications. *arXiv preprint arXiv:1604.00970*, 2016.
- [26] R. S. S. Blackman, Popoli. *Design and Analysis of Modern Tracking Systems*. Artech House,
 Boston, 1999.
- [27] K. Granström, S. Reuter, D. Meissner, and A. Scheel. A multiple model PHD approach to
 tracking of cars under an assumed rectangular shape. In *17th International Conference on Information Fusion (FUSION)*. IEEE, 2014.
- ³⁸¹ [28] P. Del Moral. *Feynman-Kac Formulae*. Springer, 2004.
- [29] F. Lindsten, P. Bunch, S. S. Singh, and T. B. Schön. Particle ancestor sampling for neardegenerate or intractable state transition models. *arXiv preprint arXiv:1505.06356*, 2015.
- [30] S. G. Krantz and H. R. Parks. *Geometric Integration Theory*. Springer Science & Business Media, 2008.
- [31] A. L. Caterini, G. Loaiza-Ganem, G. Pleiss, and J. P. Cunningham. Rectangular flows for
 manifold learning. *Advances in Neural Information Processing Systems*, 34, 2021.
- [32] K. M. Lynch and F. C. Park. *Modern Robotics*. Cambridge University Press, 2017.



(a) Input features provided to the observation model's (b) Output representation of an observed point y_t^i using neural network. (b) Output representation of an observed point y_t^i using the triplet $(e_t^i, \alpha_t^i(e_t^i), \beta_t^i(e_t^i))$.

Figure 3: Input and output representations for the observation model.

389 A Observation Model

For the problem setting described in Section 5, our observation model is concerned with the distribution of 2D points around the peripheries of observed road users. Such models are reviewed in [25]. Let the *i*th point from an observation y_t be y_t^i . We assumed independence across different points, i.e.

$$g_{\theta}(y_t|x_t) = \prod_i g_{\theta}\left(y_t^i|x_t\right).$$
(12)

In order to conveniently sample y_t^i or to measure its likelihood, we represented each point via a triplet, $(e_t^i, \alpha_t^i(e_t^i), \beta_t^i(e_t^i))$, where $e_t^i \in \{0, 1, 2, 3\}$ is a categorical variable encoding the edge of the road user's bounding box, $\alpha_t^i(e_t^i)$ is the corresponding parallel offset, and $\beta_t^i(e_t^i)$ is the corresponding perpendicular offset (see Figure 3b). Such a triplet uniquely determines y_t^i , and we used following generative model to sample such points: $e_t^i \sim \text{Categorical}(\phi_{\theta}^0(x_t), \phi_{\theta}^1(x_t), \phi_{\theta}^2(x_t), \phi_{\theta}^3(x_t))$,

399

400

401

402

$$\begin{split} \text{if } e_t^i &= 0, \quad \alpha_t^i \quad \sim \text{Uniform}(0, w) \text{ and} \\ \beta_t^i \quad \sim \text{Laplace } \left(\mu = \phi_\theta^4\left(x_t\right), b = \phi_\theta^5\left(x_t\right)\right), \\ \text{if } e_t^i &= 1, \quad \alpha_t^i \quad \sim \text{Uniform}(0, l) \text{ and} \\ \beta_t^i \quad \sim \text{Laplace } \left(\mu = \phi_\theta^6\left(x_t\right), b = \phi_\theta^7\left(x_t\right)\right), \\ \text{if } e_t^i &= 2, \quad \alpha_t^i \quad \sim \text{Uniform}(0, w) \text{ and} \\ \beta_t^i \quad \sim \text{Laplace } \left(\mu = \phi_\theta^8\left(x_t\right), b = \phi_\theta^9\left(x_t\right)\right), \\ \text{if } e_t^i &= 3, \quad \alpha_t^i \quad \sim \text{Uniform}(0, l) \text{ and} \\ \beta_t^i \quad \sim \text{Laplace } \left(\mu = \phi_\theta^{10}\left(x_t\right), b = \phi_\theta^{11}\left(x_t\right)\right), \end{split}$$

 $\begin{array}{l} \beta_t &\sim \text{Laplace} \left(\mu = \phi_{\theta}\left(x_t\right), b = \phi_{\theta}\left(x_t\right)\right),\\ \text{if } e_t^i = 3, \quad \alpha_t^i &\sim \text{Uniform}(0,l) \text{ and}\\ \beta_t^i &\sim \text{Laplace} \left(\mu = \phi_{\theta}^{10}\left(x_t\right), b = \phi_{\theta}^{11}\left(x_t\right)\right),\\ \text{where } l \text{ and } w \text{ are the length and the width of the road user's bounding box respectively,}\\ \text{and } \phi_{\theta}^{0:11}\left(x_t\right) \text{ are the outputs from a neural network with weights } \theta. \quad \text{While the mapping}\\ \left(e_t^i, \alpha_t^i\left(e_t^i\right), \beta_t^i\left(e_t^i\right)\right) \rightarrow y_t^i \text{ is unique, the reverse mapping has four representations depending on } \end{array}$

the choice of e_t^i , i.e., e_t^i is a latent variable. The likelihood of a point, $g_\theta(y_t^i|x_t)$, was therefore obtained by:-

$$g_{\theta}\left(y_{t}^{i}|x_{t}\right) = \sum_{e_{t}^{i}=0}^{3} \phi_{\theta}^{e_{t}^{i}}\left(x_{t}\right) p_{\theta}\left(\alpha_{t}^{i}\left(e_{t}^{i}\right), \beta_{t}^{i}\left(e_{t}^{i}\right)|e_{t}^{i}, x_{t}\right),$$

$$(13)$$

where we project $y_t^i \rightarrow (e_t^i, \alpha_t^i(e_t^i), \beta_t^i(e_t^i))$ for each $e_t^i \in \{0, 1, 2, 3\}$, and marginalise over it. We use a feed-forward neural network with 4 hidden layers, each with 16 Tanh units, which outputs 12 parameters, $\phi_{\theta}^{0:11}(x_t)$. We provide the network with 5 input features - range, bearing, relative bearing, length, and width - of the road-user's bounding box as measured from the observing AV's viewpoint (see Figure 3a).

For each of the experiments in Section 5 (with synthetic and real data), we used the same observation model design. Moreover, we trained an observation model using supervision from a dataset of manually labelled trajectories, $\bar{x}_{0:T}$, by maximising the AOTLL. This model was then used to generate the synthetic data and also as a baseline for the experiments with real data.

414 **B** Transition Model

As described in Section 4, the class of SSMs that we are concerned with involves a transition function, $f_{\theta}(x_t|x_{t-1})$, that factorises into a policy which produces actions, $\pi_{\theta}(a_t|x_{t-1})$, and a deterministic, differentiable, and injective motion model, such that $x_t = \tau(x_{t-1}, a_t)$. In this section, we provide more details on the motion model and on the policies used for the experiments described in Section 5.

420 B.1 Motion Model

All experiments used a **state space** that consists of the 2D Pose (x, y, θ) of an observed road user, in addition to its instantaneous linear speed v and curvature κ . Moreover, the **action space** used consists of linear acceleration a and pinch p, i.e., the instantaneous rate of change of curvature. This choice of actions, i.e. acceleration and pinch, naturally maps to the controls exercised by road users (vehicles users in particular), i.e., to gas and rate of change of steering respectively. The calculations below compute the next state, $(x(t), y(t), \theta(t), v(t), \kappa(t))$, from the previous state $(x_0, y_0, \theta_0, v_0, \kappa_0)$ under the influence of constant actions (a, p) for $t \in [0, \Delta t]$.

Clearly
$$\dot{\kappa}(t) = p \Rightarrow \kappa(t) = \kappa_0 + pt,$$
 (14)

and
$$\dot{v}(t) = a \Rightarrow v(t) = v_0 + at.$$
 (15)

428

Since
$$\dot{\theta}(t) = v(t)\kappa(t)$$
, we have (16)

$$\dot{\theta}(t) = v_0 \kappa_0 + (v_0 p + a\kappa_0)t + apt^2, \tag{17}$$

$$\Rightarrow \quad \theta(t) = \theta_0 + v_0 \kappa_0 t + (v_0 p + a \kappa_0) \frac{t^2}{2} + a p \frac{t^3}{3}.$$
(18)

429

Finally, using
$$\dot{x}(t) = v(t)\cos\theta(t)$$
, and $\dot{y}(t) = v(t)\sin\theta(t)$, we have (19)

$$x(t) = x_0 + \int_0^t v(s) \cos \theta(s) \, \mathrm{d}s,$$
 (20)

and
$$y(t) = y_0 + \int_0^t v(s) \sin \theta(s) \, ds$$
 (21)

$$= y_0 + \int_0^t v(s) \cos\left(\frac{\pi}{2} - \theta(s)\right) \mathrm{d}s.$$
(22)

To make these integrals analytically tractable, we drop the cubic term in θ_t , i.e. $ap \frac{t^3}{3}$, and use the integral² $\int (a + bs) \cos(c + ds + es^2) ds$ with appropriate coefficients to compute x(t) and y(t). This approximation is justified as for the experiments described in Section 5, we use a small Δt of ≈ 0.33 seconds.

434 B.2 Policy

For *synthetic data*, we designed a simple state-dependent stochastic policy that modulated its mean acceleration and pinch as a function of speed. The mapping from state to actions is described in Figure 4. The policy then had 2 learnable parameters - the standard deviations of its acceleration and pinch. The advantage of having only 2 learnable parameters is that it allowed us to easily verify if a method was converging to the correct values or not.

On *real data*, we used a more expressive policy which produces a multivariate Gaussian distribution over acceleration and pinch conditioned on its inputs. The policy's inputs are the instantaneous speed and curvature of an object, which are then fed to a 3-layer feedforward neural network, with hidden layers with 32 ReLU units, outputting 5 parameters - the means of acceleration and pinch, and the 3 elements of the lower triangular matrix representation of the covariance of the two. This gives a total of 1,317 learnable parameters.

²The closed form integral was obtained using Wolfram Alpha.



Figure 4: State to action mapping of the simple policy. The solid line represents the mean action taken at a given speed, while the shaded regions represent one standard deviation of Gaussian noise around that action. At a high-level, the policy accelerates at low speeds and decelerates at higher speeds. The policy also applies less pinch at higher speeds.



with 25 step trajectories using PF-SEFI with different values for the smoothing lag (L).

(b) MLL of test synthetic data with 50 step trajectories using PF-SEFI with different values for the smoothing lag (L).

(c) MLL of test real data with 60 step trajectories using PF-SEFI with different values for the smoothing lag (L) using 2048 training particles.

Figure 5: Marginal Log Likelihood (MLL) of synthetic test data for models trained using PF-SEFI with different values for the smoothing lag (L) plotted against the corresponding training steps.

446 C Training Setup

In this section, we present our training setup for the experiments described in Section 5. We show 447 sample trajectories from each of the datasets (synthetic and real), and provide the set of hyper-448 parameters that were used for each of the experiments. All experiments (including the ones used 449 for picking the best set of hyper-parameters) were repeated 10 times to obtain the median and in-450 terquartile ranges that are shown in Figures 2, 5, and 7, and in Table 1. All experiments used the 451 default settings of the Adam optimiser in TensorFlow with a learning rate of 0.01. We found that 452 smaller learning rates yield similar results, but require proportionately longer training time, while 453 larger learning rates cause instability. 454

455 C.1 Training on Synthetic Data

456 Sample trajectories from the generated synthetic data are shown in Figure 8. These samples were 457 generated using a hand-crafted policy (see Section B.2), the motion model derived in Section B.1, 458 and an observation model trained with supervised learning (see Section A). We generated two 459 datasets (with 25 and 50 steps respectively), each containing 10 scenes with 100 objects each for 460 training, and 2 scenes also with 100 objects each for evaluation. For every train step, we used all 461 100 objects from a single randomly sampled training scene, while for every evaluation step, we used 462 all 100 objects from both evaluation scenes.

Table 2 tabulates the set of hyper-parameters that were used for the experiments discussed in Section 5. Figures 5a and 5b show the effect of different values for the smoothing lag hyper-parameter (L)



Figure 6: Maximum MLL of test real data with 60 step trajectories using PF-SEFI sweeping over different values for the smoothing lag (L) using different numbers of particles for training.



25 step trajectories using PFNET 50 step trajectories using PFNET step trajectories using PFNET with with different values for the trade- with different values for the trade- different values for the trade-off paoff parameter (α).

off parameter (α).

(a) MLL of test synthetic data with (b) MLL of test synthetic data with (c) MLL of test real data with 60 rameter (α).

Figure 7: Marginal Log Likelihood (MLL) of synthetic and real test data for models trained using PFNET with different values for the trade-off parameter (α) plotted against the corresponding training steps.

for PF-SEFI on synthetic data, while 5c shows the effect of L on real data. We observed that for 465 synthetic data, the performance started to plateau at L = 8, and hence picked L = 8 for the final 466 experiments. On real data, performance continues to improve at higher values of L up to around 467 L = 14 and we also note that variance in training is high when L is too small. Moreover, Figures 7a 468 and 7b show the effect of different values for the trade-off parameter (α) for PFNET. While learning 469 failed on synthetic data with 50 steps, we observed that $\alpha = 0.8$ marginally outperformed $\alpha = 1.0$ 470 and significantly outperformed $\alpha = 0.6$. 471

Table 2: Hyper-parameters used for experiments A (synthetic data with 25 step trajectories) and B (synthetic data with 50 step trajectories). Smoothing lag (L) is only relevant for PF-SEFI, and the trade-off parameter (α) is only relevant for PFNET.

Hyper-Parameter	Value
Learning Rate	0.01
Number of Epochs	100
Smoothing Lag (L) for PF-SEFI	8
Trade-off Parameter (α) for PFNET	0.8
Number of Particles for Training	1024
Number of Particles for Evaluation	4096



Figure 8: Examples of synthetic trajectories. The figure shows 3 snapshots of the state of multiple objects, each occurring 10 steps apart from each other. The green boxes are the true states of the objects at different timesteps, while the blue polygons are the observations that the AV (orange) makes. The green dots are the true (x, y) coordinates of the objects at all timesteps.

472 C.2 Training on Real Data

For training on real data, we used a dataset of real-world road-user trajectories observed by an AV. Many of the frames in this set were also labelled manually by human-labelers, allowing us to compute relevant tracking metrics such as ADE, AYE, and AOTLL under the different models that we learned.

All trajectories were approximately 20 seconds long, where each step corresponded to 0.33 seconds
in real time, giving 60 steps in discrete time. The dataset consisted of a training set with 1502
trajectories, and a test set containing 404 trajectories. Figure 9 shows some example trajectories.

We ran hyper-parameter sweeps over learning rates, the smoothing lag (L) for PF-SEFI, the tradeoff parameter (α) for PFNET, and the number of training particles. The final results are using the best settings of these parameters which we list in Table 3, though we note that results were quite insensitive to most of these settings. Figure 6 shows the effect of L on the maximum achieved test MLL on real data for different numbers of training particles. As can be seen, PF-SEFI does significantly better as L increases from 0 to 10, highlighting the added benefit of smoothing. On the other hand, it is quite insensitive to the smoothing lag L beyond this point, except that the variance



(a) The AV in a crowded multilane road surrounded by mostly stationary objects.

(b) The AV decelerating, yielding for other objects at an intersection.

Figure 9: Examples of real-data trajectories. The figure shows 3 snapshots of the state of multiple objects, each occurring 5 seconds apart from each other. The green boxes are the labelled bounding boxes of objects at different timesteps, while the blue polygons are the observations that the AV (orange) makes. The dots represent the labelled (x, y) coordinates of the objects over time.



Figure 10: Sampled observations on the same labelled states from the real data that were shown in Figure 9. Notice the qualitative similarity in the observations in this figure relative to Figure 9.

Table 3: Hyper-parameters used for experiment C (real data with 60 step trajectories). Smoothing lag (L) is only relevant for PF-SEFI, and the trade-off parameter (α) is only relevant for PFNET.

Hyper-Parameter	Value
Learning Rate	0.01
Global Grad Norm Clipping	0.5
Number of Epochs	15
Smoothing Lag (L) for PF-SEFI	15
Trade-off Parameter (α) for PFNET	0.8
Number of Particles for Training	4096
Number of Particles for Evaluation	4096



(a) MLL of test real data after training models using (b) MLL of test real data after training models using 30 step sequences. 15 step sequences.

Figure 11: Marginal Log Likelihood (MLL) of real test data for models trained on shorter sequences plotted against the corresponding training steps.

⁴⁸⁷ appears to increase when L gets too large. We also find that PF-SEFI consistently improves with ⁴⁸⁸ more particles.

When training on real data, all methods were subject to very large gradients at times, especially earlier on when failing to track objects is much more likely, leading to very high losses and corresponding gradients. In order to stabilise training, we clipped the maximum global norm of all gradients to 0.5. This is particularly necessary for the methods that require differentiating through the filter.

494 D Sampled Observations from Learned Observation Model

Figure 10 shows sampled observations from the observation model learned on real data using PF-SEFI. These observations were sampled using the checkpoint that produced the highest MLL, and for the same (labelled) states that were shown in Figure 9. The qualitative similarity between the real and sampled observations indicates the efficacy of our method for learning generative models that can be used to sample observations in closed-loop simulation.

500 E Training on Shorter Sequences

In the case of synthetic data, we note that some methods such as DPF-SGR performed poorly when trained on 50 step sequences (Figure 2b), however performed much better when trained on shorter 25 step sequences (Figure 2a). We conducted similar experiments on real data to see if a similar improvement can be attained. Figure 11 shows the results when training DPF-SGR and other baselines on 30 and 15 step sequences of real data (instead of 60). At 30 steps (Figure 11a), we find that all 3 baselines still fail to learn good models, while PF-SEFI performs almost as well as on length Table 4: Metrics computed on held out synthetic test data comparing PF-SEFI (us) against baselines DPF-SGR, PFNET, and vanilla PF, on two experiments - (A) learning from synthetic data with 25 steps, and (B) learning from synthetic data with 50 steps. The Smoothing ADE and AYE reported here are the same as in Table 1. They are contrasted with the Filtering ADE and AYE, which measure the performance of the learned models in the online setting of state estimation.

Exp.	Method	Smoothing ADE (m)	Filtering ADE (m)	Smoothing AYE (rad)	Filtering AYE (rad)
A	TRUE PF-SEFI (us) DPF-SGR PFNET PF	$\begin{array}{c} 0.090 \pm 0.001 \\ 0.186 \pm 0.021 \\ \textbf{0.165} \pm \textbf{0.008} \\ 0.264 \pm 0.019 \\ 0.245 \pm 0.021 \end{array}$	$\begin{array}{c} 0.144 \pm 0.001 \\ 0.233 \pm 0.009 \\ \textbf{0.222} \pm \textbf{0.008} \\ 0.333 \pm 0.030 \\ 0.345 \pm 0.024 \end{array}$	$ \begin{vmatrix} 0.014 \pm 0.000 \\ 0.016 \pm 0.000 \\ 0.014 \pm 0.000 \\ 0.017 \pm 0.000 \\ 0.017 \pm 0.000 \end{vmatrix} $	$\begin{array}{c} 0.034 \pm 0.000 \\ 0.039 \pm 0.000 \\ \textbf{0.035} \pm \textbf{0.001} \\ 0.047 \pm 0.001 \\ 0.047 \pm 0.001 \end{array}$
в	TRUE PF-SEFI (us) DPF-SGR PFNET PF		$\begin{array}{c} 0.154 \pm 0.001 \\ \textbf{0.239} \pm \textbf{0.007} \\ 2.888 \pm 0.232 \\ 3.133 \pm 0.108 \\ 3.207 \pm 0.106 \end{array}$	$ \begin{vmatrix} 0.012 \pm 0.000 \\ 0.014 \pm 0.000 \\ 0.142 \pm 0.016 \\ 0.148 \pm 0.008 \\ 0.137 \pm 0.007 \end{vmatrix} $	$\begin{array}{c} 0.038 \pm 0.000 \\ \textbf{0.043} \pm \textbf{0.000} \\ 0.189 \pm 0.008 \\ 0.213 \pm 0.009 \\ 0.212 \pm 0.008 \end{array}$

⁵⁰⁷ 60 sequences. At 15 steps (Figure 11b), however, the baselines do actually perform much better, ⁵⁰⁸ though the final performance for all methods is worse than PF-SEFI on 60 steps.

This highlights an advantage of PF-SEFI, which is that it is relatively invariant to the length of 509 sequences that it is trained on. Depending on the problem setting, there is usually a minimum 510 sequence length required to obtain enough information to learn the correct models. If that sequence 511 length is longer than the maximum sequence length for which a method such as DPF-SGR is stable 512 to train, then one must sacrifice model quality for stable learning by cutting the sequences to shorter 513 subsequences or by subsampling the sequences, throwing away some of the observations. In this 514 case, shortening the sequences to 15 steps allowed for reasonable (though suboptimal) models to 515 be learned using the baseline methods. In other problems it may well be the case that even more 516 trimming and/or subsampling would be needed. 517

518 F Performance on the Filtering Task

In Section 5, we considered metrics such as ADE and AYE that pertain to the task of state estimation in the offline setting (known as smoothing), i.e., where we assume access to the entire sequence of observations, i.e. $y_{1:T}$. In this section, we additionally consider the online setting (known as filtering) where the task is state estimation at each time step t, with observations from time step 0 up to time step t, i.e. $p(x_t|y_{0:t})$. This setting is relevant for the use of our learned models on-board the AV.

Table 4 reports ADE and AYE using both the smoothing and filtering distribution for the offline and online task of state estimation respectively. The patterns are unchanged relative to the ones observed in Section 5 and in Table 1. However, these results highlight the applicability of the learned in both settings.

529 G Training with Higher Dimensional Observations

In Section 5, we reported experiments with 32 dimensional observations (16 2D points). In this section, we report additional experiments with even higher dimensional observations (32 and 64 2D points, i.e., 64 and 128 dimensional observations respectively) on the synthetic dataset with 25 steps, trained using PF-SEFI. In Table 5 we summarise the empirical findings of these experiments. In each case, the learned models match the performance of the true models as measured by MLL. These results suggest that PF-SEFI scales well with higher dimensional observations.

536 H Effect of a Noisy AV State on Learning

In Section 5, we assume that the AV state is known precisely. In practice, we expect there to be some minimal errors in state estimation. To test our sensitivity to the same, we ran additional experiments



Table 5: MLL computed on held out synthetic test data comparing PF-SEFI (us) against the true models on synthetic dataset generated with 32 (A), 64 (D), and 128 (E) dimensional observations.

Figure 12: Marginal Log Likelihood (MLL) of synthetic test data using models trained on synthetic data with noisy AV states.

- ⁵³⁹ with the 25 steps synthetic dataset by injecting Gaussian noise (with a standard deviation of 0.5m
- in x and y, and 0.05rad in θ) in the AV's 2D pose. We retrained our models using PF-SEFI in the
- presence of such noise, and observe no change in MLL at convergence over the held out test data (see Figure 12), nor in training stability.