

## A Derivation of the ELBO

This appendix contains three parts: we first explicate the joint distribution of the model and data  $p(\mathbf{X}, \mathbf{f}(\cdot), \mathbf{Y})$ ; then we describe the variational approximation to the model posterior  $q(\mathbf{X}, \mathbf{f}(\cdot))$ ; then we show how they combine to produce the ELBO. Table 2 provides some nomenclature.

Table 2: Nomenclature used in this derivation

$t \in \{0 \dots T\}$	time steps indexed $t$
$d \in \{1 \dots D\}$	dimension of hidden states $\mathbf{x}_t$ indexed $d$
$O$	dimension of the observed data
$m \in \{1 \dots M\}$	number of inducing variables indexed $m$
$\mathbf{x}_t$	hidden state at time $t$ , $\mathbf{x}_t \in \mathbb{R}^D$
$\mathbf{a}_t$	control input (action) at time $t$ , $\mathbf{a}_t \in \mathbb{R}^P$
$\tilde{\mathbf{x}}_t$	concatenation of control input and state at $t$
$\mathbf{y}_t$	observation at time $t$ , $\mathbf{y}_t \in \mathbb{R}^O$
$\tilde{\mathbf{y}}_t$	concatenation of control input and observation at $t$
$\mathbf{X}$	collection of hidden states, $\mathbf{X} = [\mathbf{x}_t]_{t=0}^T$ .
$\mathbf{Y}$	collection of observations, $\mathbf{Y} = [\mathbf{y}_t]_{t=0}^T$ .
$\sigma_f^2$	variance of state transition noise
$\sigma_g^2$	variance of observation noise
$f_d(\cdot)$	the $d^{\text{th}}$ Gaussian process (GP)
$\mathbf{f}(\cdot)$	collection of GPs, $= [f_d(\cdot)]_{d=1}^D$
$\eta_d(\cdot)$	prior mean function of the $d^{\text{th}}$ GP
$k_d(\cdot, \cdot)$	prior covariance function of the $d^{\text{th}}$ GP
$\mu_d(\cdot)$	posterior mean function of the $d^{\text{th}}$ GP
$v_d(\cdot, \cdot)$	posterior covariance function of the $d^{\text{th}}$ GP
$\mathbf{Z}$	Locations of variational pseudo-inputs
$\mathbf{u}_d$	evaluations of the $d^{\text{th}}$ GP at the pseudo-inputs: $\mathbf{u}_d = [f_d(\mathbf{z}_m)]_{m=1}^M$ .
$\mathbf{U}$	collection: $\mathbf{U} = [\mathbf{u}_d]_{d=1}^D$
$\boldsymbol{\mu}_d$	variational posterior mean of $\mathbf{u}_d$ : $\boldsymbol{\mu}_d = [\mu_d(\mathbf{z}_m)]_{m=1}^M$
$\boldsymbol{\Sigma}_d$	variational posterior covariance of $\mathbf{u}_d$
$\mathbf{A}_t$	variational transition matrix of $q(\mathbf{x}_t   \mathbf{x}_{t-1})$
$\mathbf{L}_t$	triangular-square-root of variational covariance of $q(\mathbf{x}_t   \mathbf{x}_{t-1})$
$\mathbf{m}_t$	variational mean of the marginal $q(\mathbf{x}_t)$
$\mathbf{S}_t$	variational covariance of the marginal $q(\mathbf{x}_t)$

### A.1 Model joint distribution

Here we define the joint distribution of the Gaussian processes  $f$ , the latent states  $\mathbf{x}$  and the data  $\mathbf{y}$ .

The Gaussian processes have prior mean  $\eta(\cdot)$  and prior covariances  $k(\cdot, \cdot)$ :

$$p(f_d(\cdot)) = \mathcal{GP}(\eta_d(\cdot), k_d(\cdot, \cdot)) \quad d = 1 \dots D. \quad (17)$$

We note that placing a measure  $p$  on the function  $f$  causes some measure-theoretic discrepancies. Nonetheless, the derivation holds following a more theoretical consideration of the problem (Mathews et al., 2016), and the intuition given by our derivation is correct.

The initial state is assumed to be drawn from a standard normal distribution

$$p(\mathbf{x}_0) = \mathcal{N}(\mathbf{0}, \mathbf{I}_D). \quad (18)$$

The state transition depends on the Gaussian processes:

$$p(\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{f}(\cdot)) = \mathcal{N}(\mathbf{x}_t | \mathbf{f}(\tilde{\mathbf{x}}_{t-1}), \sigma_f^2 \mathbf{I}_D), \quad (19)$$

We assume a linear-Gaussian observation model:

$$p(\mathbf{y}_t | \mathbf{x}_t) = \mathcal{N}(\mathbf{y}_t | \mathbf{W}_g \mathbf{x}_t + \mathbf{b}_g, \sigma_g^2 \mathbf{I}_O) \quad (20)$$

The joint density is then

$$p(\mathbf{f}, \mathbf{X}, \mathbf{Y}) = \prod_{d=1}^D p(f_d(\cdot)) p(\mathbf{x}_0) \prod_{t=1}^T p(\mathbf{y}_t | \mathbf{x}_t) \prod_{t=1}^T p(\mathbf{x}_t | \mathbf{f}, \mathbf{x}_{t-1}) \quad (21)$$

## A.2 Approximate posterior distribution

We will use variational Bayes to approximate the posterior distribution over  $\mathbf{f}$  and  $\mathbf{X}$ , whilst simultaneously obtaining a bound on the marginal likelihood (the ELBO) which will be used to train the parameters of the model, including covariance function parameters, noise variances and the parameters  $\mathbf{W}_g, \mathbf{b}_g$  of the linear output mapping.

The posterior over Gaussian processes takes the form of a sparse GP. We introduce a series of  $M$  variational inducing points  $\mathbf{Z} = [\mathbf{z}_m]_{m=1}^M$  which lie in the same domain as  $\tilde{\mathbf{x}}$ . Following convention, the values of the  $d^{\text{th}}$  function at those points are denoted  $\mathbf{u}_d = [f_d(\mathbf{z}_m)]_{m=1}^M$ , while evaluations from the prior as  $\boldsymbol{\eta}_d = [\eta_d(\mathbf{z}_m)]_{m=1}^M$ . Note that the variables  $\mathbf{u}$  are not *auxiliary* variables, but part of the original model specification, being part of the GP. We assume a variational posterior of the form

$$q(\mathbf{U}) = \prod_{d=1}^D \mathcal{N}(\mathbf{u}_d | \boldsymbol{\mu}_d, \boldsymbol{\Sigma}_d). \quad (22)$$

The remainder of the GPs conditioned on  $\mathbf{u}$  are assumed to take the same form as the GP prior conditional. That is

$$q(f_d(\cdot) | \mathbf{u}_d) = p(f_d(\cdot) | \mathbf{u}_d) = \mathcal{GP}(\eta_d(\cdot) + k(\cdot, \mathbf{Z}) \mathbf{K}_{zz}^{-1} (\mathbf{u}_d - \boldsymbol{\eta}_d), k(\cdot, \cdot) - k(\cdot, \mathbf{Z}) \mathbf{K}_{zz}^{-1} k(\mathbf{Z}, \cdot)). \quad (23)$$

Marginalising with respect to  $\mathbf{u}_d$  leads to our approximation to the GP:

$$q(f_d(\cdot)) = \mathcal{GP}(\mu_d(\cdot), v_d(\cdot, \cdot)), \quad (24)$$

with

$$\mu_d(\cdot) = \eta_d(\cdot) + k(\cdot, \mathbf{Z}) \mathbf{K}_{zz}^{-1} (\boldsymbol{\mu}_d - \boldsymbol{\eta}_d), \quad (25)$$

$$v_d(\cdot, \cdot) = k(\cdot, \cdot) - k(\cdot, \mathbf{Z}) \mathbf{K}_{zz}^{-1} [\mathbf{K}_{zz} - \boldsymbol{\Sigma}_d] \mathbf{K}_{zz}^{-1} k(\mathbf{Z}, \cdot). \quad (26)$$

The approximation to the posterior over state trajectories is given a Gauss-Markov structure of the form

$$q(\mathbf{X}) = q(\mathbf{x}_0) \prod_{t=1}^T q(\mathbf{x}_t | \mathbf{x}_{t-1}), \quad (27)$$

where

$$q(\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_0 | \mathbf{m}_0, \mathbf{L}_0 \mathbf{L}_0^\top) \quad (28)$$

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t | \mathbf{A}_t \mathbf{x}_{t-1}, \mathbf{L}_t \mathbf{L}_t^\top). \quad (29)$$

The complete set of variational parameters is then  $\mathbf{Z}, \{\boldsymbol{\mu}_d, \boldsymbol{\Sigma}_d\}_{d=1}^D, \mathbf{m}_0, \mathbf{L}_0, \{\mathbf{A}_t, \mathbf{L}_t\}_{t=1}^T$ . The parameters of  $q(\mathbf{X})$  are reconfigured to be the output of an RNN recognition model (see main text), whilst we optimise the parameters controlling  $\mathbf{f}(\cdot)$  directly.

The joint posterior then factors as

$$q(\mathbf{f}(\cdot), \mathbf{X}) = \prod_{d=1}^D q(f_d(\cdot)) q(\mathbf{X}). \quad (30)$$

### A.3 The ELBO

Having specified the forms of the model and the approximate posterior, we are ready to derive the ELBO. Following the standard variational Bayes methods, we write

$$\text{ELBO} = \mathbb{E}_{q(\mathbf{X})q(\mathbf{f}(\cdot))} \left[ \log \frac{p(\mathbf{Y} | \mathbf{X})p(\mathbf{X} | \mathbf{f}(\cdot))}{q(\mathbf{X})} \frac{p(\mathbf{f}(\cdot))}{q(\mathbf{f}(\cdot))} \right]. \quad (31)$$

We will split the ELBO into four parts, dealing with each in turn:

$$\begin{aligned} \text{ELBO} &= \underbrace{\mathbb{E}_{q(\mathbf{X})} [\log p(\mathbf{Y} | \mathbf{X})]}_{\text{part 1}} + \underbrace{\mathbb{E}_{q(\mathbf{X})q(\mathbf{f}(\cdot))} [\log p(\mathbf{X} | \mathbf{f}(\cdot))]}_{\text{part 2}} \\ &\quad - \underbrace{\mathbb{E}_{q(\mathbf{X})} [\log q(\mathbf{X})]}_{\text{part 3}} + \underbrace{\mathbb{E}_{q(\mathbf{f}(\cdot))} [\log \frac{p(\mathbf{f}(\cdot))}{q(\mathbf{f}(\cdot))}]}_{\text{part 4}}. \end{aligned} \quad (32)$$

**Part 1** This expression can be computed straight-forwardly in closed form due to our choice of a linear-Gaussian emission  $g(\mathbf{x})$ . Let  $\mathbf{m}_t, \Sigma_t$  be the marginals of  $q(\mathbf{x}_t)$  computed via the recursion, and recall the form of the linear emission function  $g(\mathbf{x}_t) = \mathbf{W}_g \mathbf{x}_t + \mathbf{b}_g$

$$\begin{aligned} \mathbb{E}_{q(\mathbf{X})} [\log p(\mathbf{Y} | \mathbf{X})] &= \mathbb{E}_{q(\mathbf{X})} \left[ \sum_{t=1}^T \log \mathcal{N}(\mathbf{y}_t | g(\mathbf{x}_t), \sigma_g^2) \right] \\ &= \sum_{t=1}^T \mathbb{E}_{q(\mathbf{x}_t)} \left[ \log \mathcal{N}(\mathbf{y}_t | \mathbf{W}_g \mathbf{x}_t + \mathbf{b}_g, \sigma_g^2) \right] \\ &= \sum_{t=1}^T \log \mathcal{N}(\mathbf{y}_t | \mathbf{W}_g \mathbf{m}_t + \mathbf{b}_g, \sigma_g^2) - \frac{1}{2\sigma_g^2} \text{tr}(\mathbf{W}_g^\top \mathbf{W}_g \Sigma_t). \end{aligned} \quad (33)$$

In practise we defer this simple computation to the `variational_expectations` functionality in GPflow (Matthews et al., 2017).

**Part 2** This expression cannot be computed in closed form without restriction to the RBF kernel as in (Frigola, 2015). We eliminate the integral with respect to  $\mathbf{f}$  here, and then use the reparameterisation trick to estimate the integral with respect to  $\mathbf{X}$  (see main text).

$$\begin{aligned} \text{part 2} &= \mathbb{E}_{q(\mathbf{X})q(\mathbf{f}(\cdot))} [\log p(\mathbf{X} | \mathbf{f}(\cdot))] \\ &= \mathbb{E}_{q(\mathbf{X})q(\mathbf{f}(\cdot))} \left[ \log p(\mathbf{x}_0) \prod_{t=1}^T \mathcal{N}(\mathbf{x}_t | \mathbf{f}(\tilde{\mathbf{x}}_{t-1}), \sigma_f^2 \mathbf{I}_D) \right] \\ &= \mathbb{E}_{q(\mathbf{x}_0)} [\log p(\mathbf{x}_0)] + \mathbb{E}_{q(\mathbf{X})q(\mathbf{f}(\cdot))} \left[ \sum_{t=1}^T \sum_{d=1}^D \log \mathcal{N}(\mathbf{x}_t^{(d)} | f_d(\tilde{\mathbf{x}}_{t-1}), \sigma_f^2) \right] \\ &= \mathbb{E}_{q(\mathbf{x}_0)} [\log p(\mathbf{x}_0)] + \mathbb{E}_{q(\mathbf{X})} \left[ \sum_{t=1}^T \sum_{d=1}^D \log \mathcal{N}(\mathbf{x}_t^{(d)} | \mu_d(\tilde{\mathbf{x}}_{t-1}), \sigma_f^2) - \frac{1}{2} \sigma_f^{-2} v_d(\tilde{\mathbf{x}}_{t-1}, \tilde{\mathbf{x}}_{t-1}) \right], \end{aligned} \quad (34)$$

which matches the term in the main text.

**Part 3** This corresponds to the entropy of  $q(\mathbf{X})$ . It is straightforward to derive:

$$-\mathbb{E}_{q(\mathbf{X})} [\log q(\mathbf{X})] = \mathbf{H}[q(\mathbf{X})] = \frac{(T+1)D}{2} \log(2\pi e) + \sum_{t=0}^T \log(\det(\mathbf{L}_t)). \quad (35)$$

**Part 4** This final part is the Kullback-Leibler divergence between the prior and (approximate) posterior GPs. We first note that it can be written as a sum across dimensions  $d$ , and then that each

GP  $f_d(\cdot)$  can be factored into two parts:  $p(f_d(\cdot) | \mathbf{u}_d)p(\mathbf{u}_d)$  and similarly for  $q$ . This results in

$$\begin{aligned} \mathbb{E}_{q(\mathbf{f}(\cdot))} \left[ \log \frac{p(\mathbf{f}(\cdot))}{q(\mathbf{f}(\cdot))} \right] &= \sum_{d=1}^D \mathbb{E}_{q(f_d(\cdot))} \left[ \log \frac{p(f_d(\cdot))}{q(f_d(\cdot))} \right] \\ &= \sum_{d=1}^D \mathbb{E}_{q(f_d(\cdot) | \mathbf{u}_d)q(\mathbf{u}_d)} \left[ \log \frac{p(f_d(\cdot) | \mathbf{u}_d)p(\mathbf{u}_d)}{q(f_d(\cdot) | \mathbf{u}_d)q(\mathbf{u}_d)} \right]. \end{aligned} \quad (36)$$

Since we have defined the posterior conditional  $q(f_d(\cdot) | \mathbf{u}_d)$  to match the prior conditional, the two terms cancel, resulting in

$$\begin{aligned} \mathbb{E}_{q(\mathbf{f}(\cdot))} \left[ \log \frac{p(\mathbf{f}(\cdot))}{q(\mathbf{f}(\cdot))} \right] &= \sum_{d=1}^D \mathbb{E}_{q(\mathbf{u}_d)} \left[ \log \frac{p(\mathbf{u}_d)}{q(\mathbf{u}_d)} \right] \\ &= - \sum_{d=1}^D \text{KL}[q(\mathbf{u}_d) || p(\mathbf{u}_d)]. \end{aligned} \quad (37)$$

Since the result is a Kullback-Leibler divergence between two finite-dimensional normal distributions, it is computed straightforwardly.

Although this notation is somewhat sloppy (since the sets of variables  $f_d(\cdot)$  and  $\mathbf{u}_d$  overlap), the result is correct. Matthews (2017) contains a more careful and significantly more technical derivation.

## B Full visualisation of synthetic 1D dataset

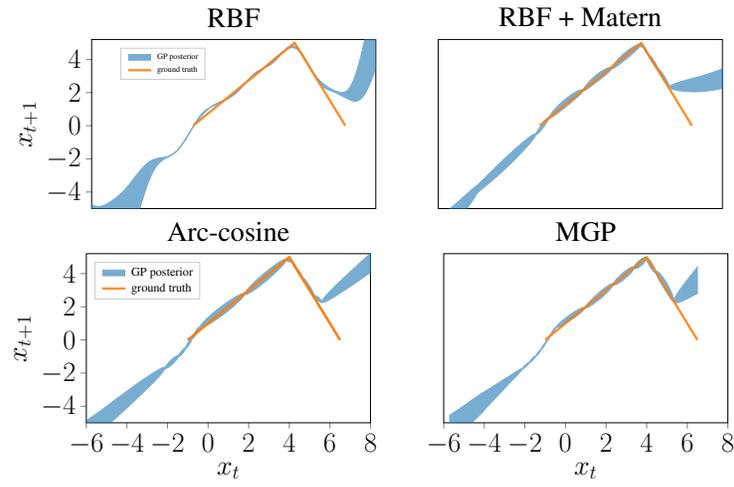


Figure 7: Visualisation of the learned GP transition functions across a greater domain of the function. It can be seen that all models revert to the mean function (defined as the identity function) away from the data. The short lengthscales of the Arc-cosine and MGP (compounded with a Matern kernel) that are used to fit the kink of the true transition function mean that they almost instantaneously revert to the mean function. The longer length scales of the RBF-containing kernels mean that we revert much more slowly to the mean.

## C Learnt latent states for cart-pole

Below we provide the learnt latent states for the cart-pole dataset with observed and hidden velocities. It is worth noting that the model has recovered similar structure for both cases.

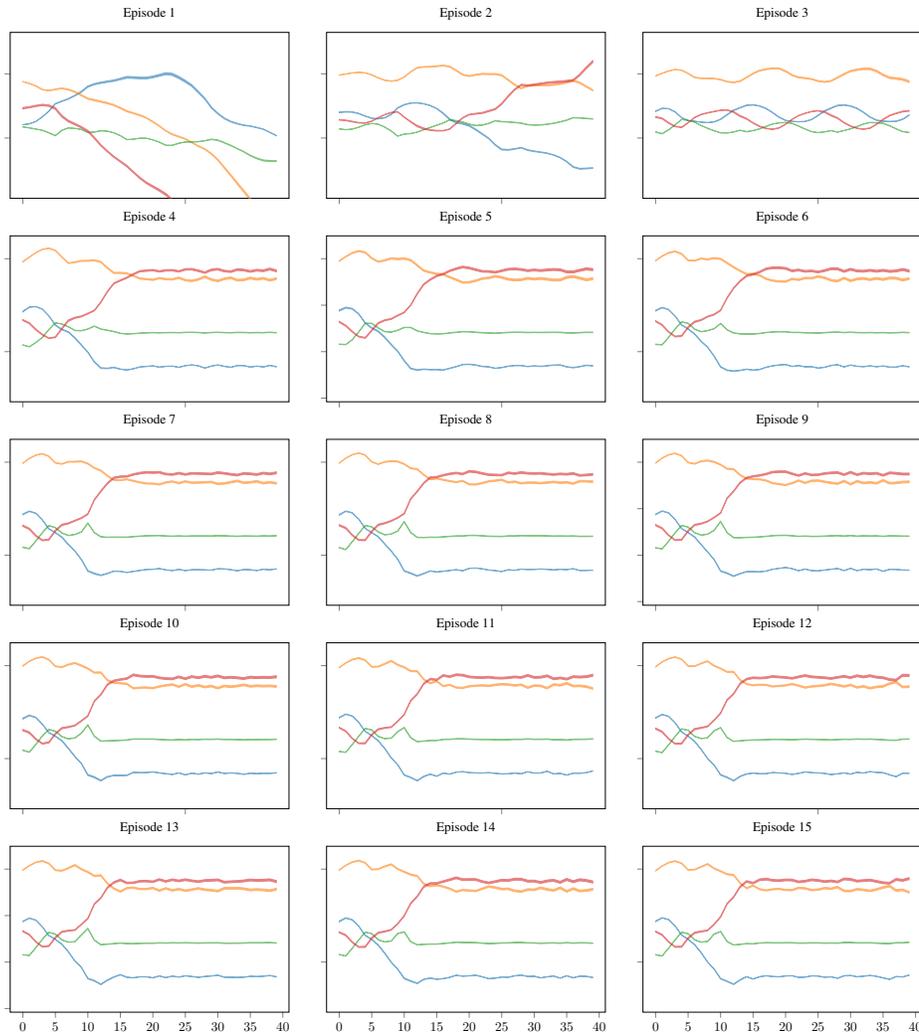


Figure 8: Learnt latent states for the cart-pole dataset with observed velocities.

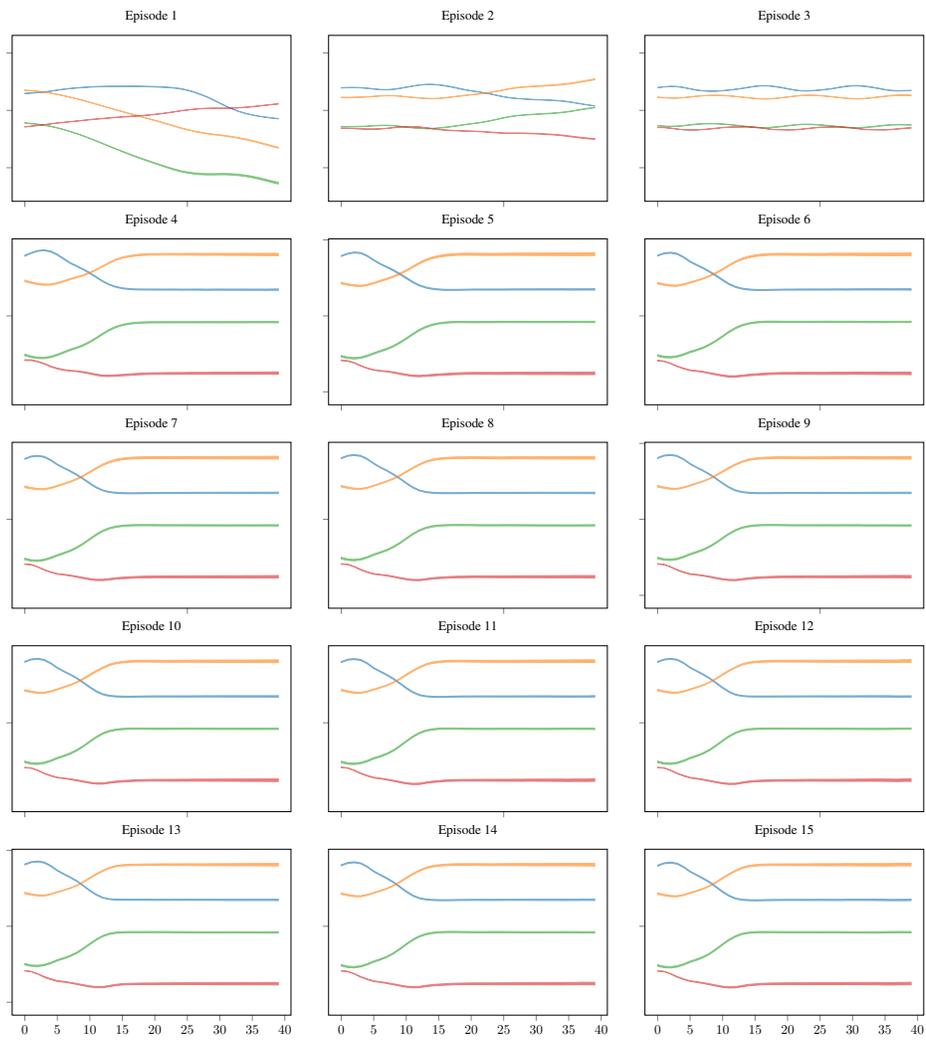


Figure 9: Learnt latent states for the cart-pole dataset with hidden velocities.

## D Cart-pole training data fitting

Below we provide detailed fittings on the training episodes for the cart-pole dataset.

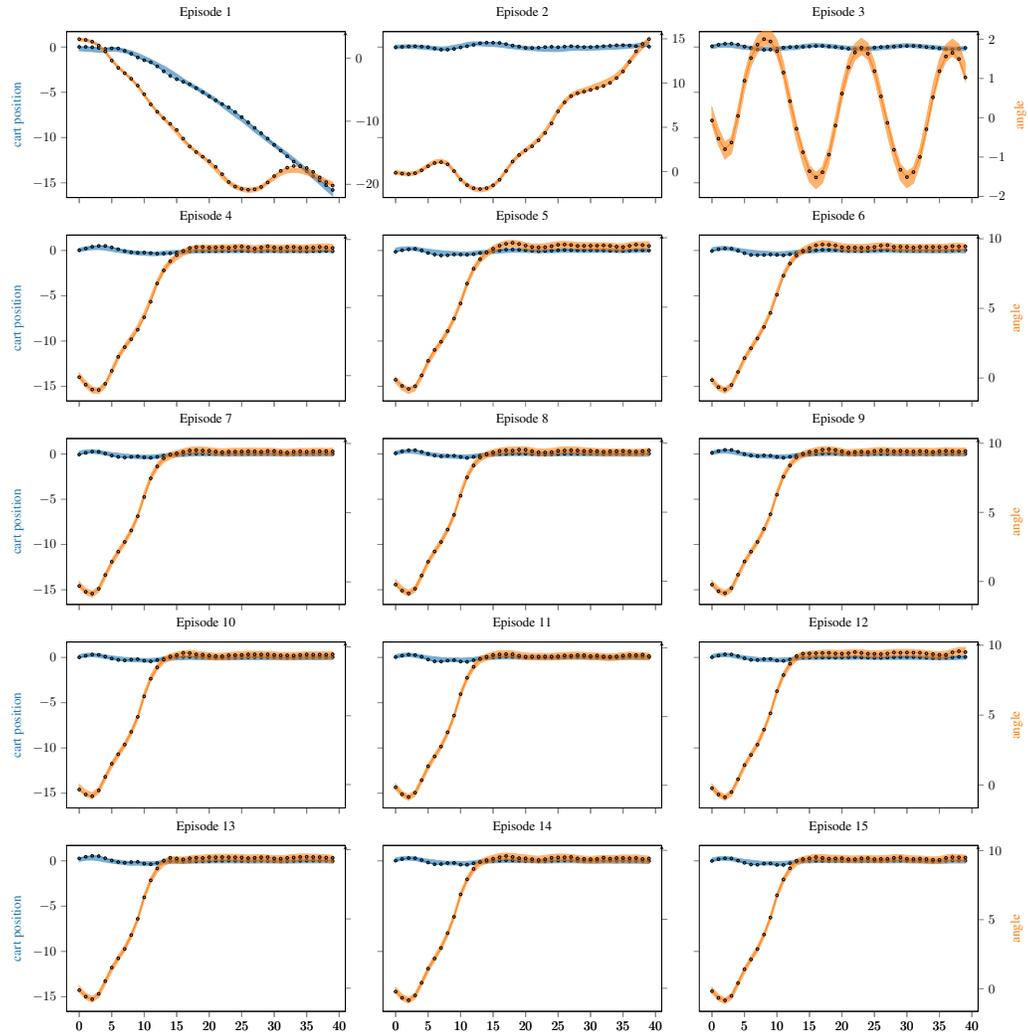


Figure 10: Detailed fittings per episode for the cart-pole dataset with observed velocities.

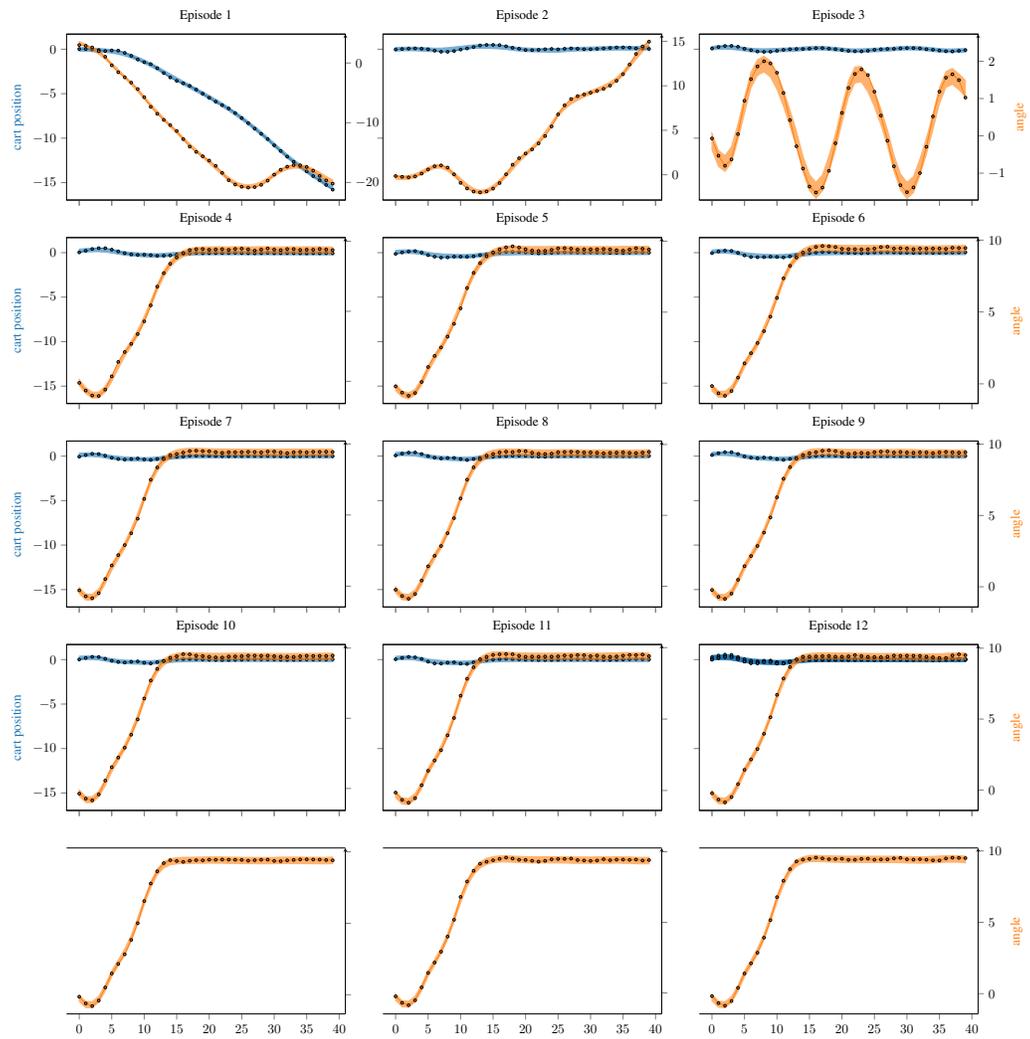


Figure 11: Detailed fittings per episode for the cart-pole dataset with hidden velocities.