

Factorized Inference in Deep Markov Models for Incomplete Multimodal Time Series

Tan Zhi-Xuan,^{1,2} Harold Soh,³ Desmond C. Ong^{1,4}

¹A*STAR Artificial Intelligence Initiative, A*STAR, Singapore

²Department of Electrical Engineering and Computer Science, MIT

³Department of Computer Science, National University of Singapore

⁴Department of Information Systems and Analytics, National University of Singapore

xuan@mit.edu, {harold, dco}@comp.nus.edu.sg

Abstract

Integrating deep learning with latent state space models has the potential to yield temporal models that are powerful, yet tractable and interpretable. Unfortunately, current models are not designed to handle missing data or multiple data modalities, which are both prevalent in real-world data. In this work, we introduce a factorized inference method for Multimodal Deep Markov Models (MDMMs), allowing us to filter and smooth in the presence of missing data, while also performing uncertainty-aware multimodal fusion. We derive this method by factorizing the posterior $p(z|x)$ for non-linear state space models, and develop a *variational backward-forward algorithm* for inference. Because our method handles incompleteness over both time and modalities, it is capable of interpolation, extrapolation, conditional generation, label prediction, and weakly supervised learning of multimodal time series. We demonstrate these capabilities on both synthetic and real-world multimodal data under high levels of data deletion. Our method performs well even with more than 50% missing data, and outperforms existing deep approaches to inference in latent time series.

Introduction

Virtually all sensory data that humans and autonomous systems receive can be thought of as multimodal time series—multiple sensors each provide streams of information about the surrounding environment, and intelligent systems have to integrate these information streams to form a coherent yet dynamic model of the world. These time series are often asynchronously or irregularly sampled, with many time-points having missing or incomplete data. Classical time series algorithms, such as the Kalman filter, are robust to such incompleteness: they are able to infer the state of the world, but only in linear regimes (Kalman 1960). On the other hand, human intelligence is robust *and* complex: we infer complex quantities with *nonlinear dynamics*, even from incomplete temporal observations of multiple modalities—for example, intended motion from both eye gaze and arm movements (Dragan, Lee, and Srinivasa 2013), or desires and emotions from actions, facial expressions, speech, and language (Baker et al. 2017; Ong, Zaki, and Goodman 2015).

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

There has been a proliferation of attempts to learn these nonlinear dynamics by integrating deep learning with traditional probabilistic approaches, such as Hidden Markov Models (HMMs) and latent dynamical systems. Most approaches do this by adding random latent variables to each time step in an RNN (Chung et al. 2015; Bayer and Osendorfer 2014; Fabius and van Amersfoort 2014; Fraccaro et al. 2016). Other authors begin with latent sequence models as a basis, then develop deep parameterizations and inference techniques for these models (Krishnan, Shalit, and Sontag 2017; Archer et al. 2015; Johnson et al. 2016; Karl et al. 2016; Doerr et al. 2018; Lin, Khan, and Hubacher 2018; Chen et al. 2018). Most relevant to our work are the Deep Markov Models (DMMs) proposed by Krishnan, Shalit, and Sontag (2017), a generalization of HMMs and Gaussian State Space Models where the transition and emission distributions are parameterized by deep networks.

Unfortunately, none of the approaches thus far are designed to handle inference with both missing data and multiple modalities. Instead, most approaches rely upon RNN inference networks (Krishnan, Shalit, and Sontag 2017; Karl et al. 2016; Che et al. 2018), which can only handle missing data using ad-hoc approaches such as zero-masking (Lipton, Kale, and Wetzell 2016), update-skipping (Krishnan, Shalit, and Sontag 2017), temporal gating mechanisms (Neil, Pfeiffer, and Liu 2016; Che et al. 2018), or providing time stamps as inputs (Chen et al. 2018), none of which have intuitive probabilistic interpretations. Of the approaches that do not rely on RNNs, Fraccaro *et al.* handle missing data by assuming linear latent dynamics (Fraccaro et al. 2017), while Johnson *et al.* (Johnson et al. 2016) and Lin *et al.* (Lin, Khan, and Hubacher 2018) use hybrid message passing inference that is theoretically capable of marginalizing out missing data. However, these methods are unable to learn nonlinear transition dynamics, nor do they handle multiple modalities.

We address these limitations—the inability to handle missing data, over multiple modalities, and with nonlinear dynamics—by introducing a multimodal generalization of Deep Markov Models, as well as a factorized inference method for such models that handles missing time points and modalities by design. Our method allows for both fil-

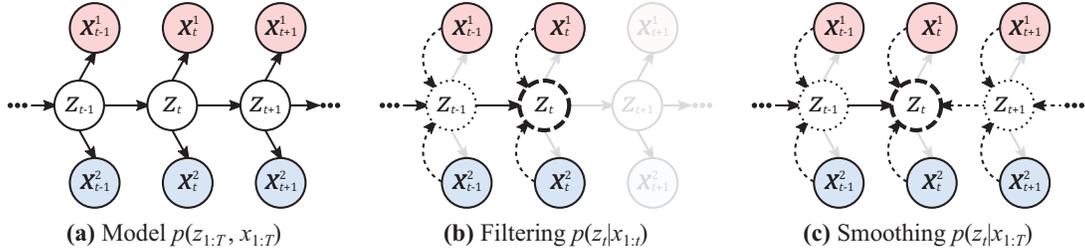


Figure 1: (a) A Multimodal Deep Markov Model (MDMM) with $M = 2$ modalities. Observations (filled) are generated from unobserved latent states (unfilled). (b) Filtering infers the current latent state z_t (bold dashed outline) given all observations up to t (solid outlines), and marginalizes (dotted outline) over past latent states. (c) Smoothing infers z_t given past, present, and future observations, and marginalizes over both past and future latent states.

tering and smoothing given incomplete time series, while also performing uncertainty-aware multimodal fusion *à la* the Multimodal Variational Autoencoder (MVAE) (Wu and Goodman 2018). Because our method handles incompleteness over both time and modalities, it is capable of (i) interpolation, (ii) forward / backward extrapolation, and (iii) conditional generation of one modality from another, including label prediction. It also enables (iv) weakly supervised learning from incomplete multimodal data. We demonstrate these capabilities on both a synthetic dataset of noisy bidirectional spirals, as well as a real world dataset of labelled human actions. Our experiments show that our method learns and performs excellently on each of these tasks, while outperforming state-of-the-art inference methods that rely upon RNNs.

Methods

We introduce Multimodal Deep Markov Models (MDMMs) as a generalization of Krishnan *et al.*'s Deep Markov Models (DMMs) (Krishnan, Shalit, and Sontag 2017). In a MDMM (Figure 1a), we model multiple sequences of observations, each of which is conditionally independent of the other sequences given the latent state variables. Each observation sequence corresponds to a particular data or sensor modality (e.g. video, audio, labels), and may be missing when other modalities are present. An MDMM can thus be seen as a sequential version of the MVAE (Wu and Goodman 2018).

Formally, let z_t and x_t^m respectively denote the latent state and observation for modality m at time t . An MDMM with M modalities is then defined by the transition and emission distributions:

$$z_t \sim \mathcal{N}(\mu_\theta(z_{t-1}), \Sigma_\theta(z_{t-1})) \quad (\text{Transition}) \quad (1)$$

$$x_t^m \sim \Pi(\kappa_\theta^m(z_t)) \quad (\text{Emission}) \quad (2)$$

Here, \mathcal{N} is the Gaussian distribution, and Π is some arbitrary emission distribution. The distribution parameters μ_θ , Σ_θ and κ_θ^m are functions of either z_{t-1} or z_t . We learn these functions as neural networks with weights θ . We also use $z_{t_1:t_2}$ to denote the time-series of z from t_1 to t_2 , and $x_{t_1:t_2}^{m_1:m_2}$ to denote the corresponding observations from modalities m_1 to m_2 . We omit the modality superscripts when all modalities are present (i.e., $x_t \equiv x_t^{1:M}$).

We want to jointly learn the parameters θ of the generative model $p_\theta(z_{1:T}, x_{1:T}) = p_\theta(x_{1:T}|z_{1:T})p_\theta(z_{1:T})$ and the parameters ϕ of a variational posterior $q_\phi(z_{1:T}|x_{1:T})$ which approximates the true (intractable) posterior $p_\theta(z_{1:T}|x_{1:T})$. To do so, we maximize a lower bound on the log marginal likelihood $L(x; \theta, \phi) \leq p_\theta(x_{1:T})$, also known as the evidence lower bound (ELBO):

$$L(x; \theta, \phi) = \mathbb{E}_{q_\phi(z_{1:T}|x_{1:T})}[\log p_\theta(x_{1:T}|z_{1:T})] - \mathbb{E}_{q_\phi(z_{1:T}|x_{1:T})}[\text{KL}(q_\phi(z_{1:T}|x_{1:T})||p_\theta(z_{1:T}))] \quad (3)$$

In practice, we can maximize the ELBO with respect to θ and ϕ via gradient ascent with stochastic backpropagation (Kingma and Welling 2014; Rezende, Mohamed, and Wierstra 2014). Doing so requires sampling from the variational posterior $q_\phi(z_{1:T}|x_{1:T})$. In the following sections, we derive a variational posterior that factorizes over time-steps and modalities, allowing us to tractably infer the latent states $z_{1:T}$ even when data is missing.

Factorized Posterior Distributions

In latent sequence models such as MDMMs, we often want to perform several kinds of inferences over the latent states. The most common of such latent state inferences are:

Filtering Inferring z_t given past observations $x_{1:t}$.

Smoothing Inferring some z_t given all observations $x_{1:T}$.

Sequencing Inferring the sequence $z_{1:T}$ from $x_{1:T}$

Most architectures that combine deep learning with state space models focus upon filtering (Fabius and van Amersfoort 2014; Chung et al. 2015; Hafner et al. 2018; Buesing et al. 2018), while Krishnan *et al.* optimize their DMM for sequencing (Krishnan, Shalit, and Sontag 2017). One of our contributions is to demonstrate that we can learn the filtering, smoothing, and sequencing distributions within the same framework, because they all share similar factorizations (see Figures 1b and 1c for the shared inference structure of filtering and smoothing). A further consequence of these factorizations is that we can naturally handle inferences given missing modalities or time-steps.

To demonstrate this similarity, we first factorize the sequencing distribution $p(z_{1:T}|x_{1:T})$ over time:

$$p(z_{1:T}|x_{1:T}) = p(z_1|x_{1:T}) \prod_{t=2}^T p(z_t|z_{t-1}, x_{t:T}) \quad (4)$$

This factorization means that each latent state z_t depends only on the previous latent state z_{t-1} , as well as all current and future observations $x_{t:T}$, and is implied by the graphical structure of the MDMM (Figure 1a). We term $p(z_t|z_{t-1}, x_{t:T})$ the *conditional smoothing posterior*, because it is the posterior that corresponds to the *conditional prior* $p(z_t|z_{t-1})$ on the latent space, and because it combines information from both past and future (hence ‘smoothing’).

Given one or more modalities, we can show that the conditional smoothing posterior $p(z_t|z_{t-1}, x_{t:T})$, the *backward filtering distribution* $p(z_t|x_{t:T})$, and the smoothing distribution $p(z_t|x_{1:T})$ all factorize almost identically:

Backward Filtering (5)

$$p(z_t|x_{t:T}) \propto p(z_t|x_{t+1:T}) \left[\prod_m \frac{p(z_t|x_t^m)}{p(z_t)} \right]$$

Forward Smoothing (6)

$$p(z_t|x_{1:T}) \propto p(z_t|x_{t+1:T}) \left[\prod_m \frac{p(z_t|x_t^m)}{p(z_t)} \right] \frac{p(z_t|x_{1:t-1})}{p(z_t)}$$

Conditional Smoothing Posterior (7)

$$p(z_t|z_{t-1}, x_{t:T}) \propto p(z_t|x_{t+1:T}) \left[\prod_m \frac{p(z_t|x_t^m)}{p(z_t)} \right] \frac{p(z_t|z_{t-1})}{p(z_t)}$$

Equations 5–7 show that each distribution can be decomposed into (i) its dependence on future observations $p(z_t|x_{t+1:T})$, (ii) its dependence on each modality m in the present, $p(z_t|x_t^m)$, and, excluding filtering, (iii) its dependence on the past $p(z_t|z_{t-1})$ or $p(z_t|x_{1:t-1})$. Their shared structure is due to the conditional independence of $x_{t:T}$ given z_t from all prior observations or latent states. Here we show only the derivation for Equation 7, because the others follow by either dropping z_{t-1} (Equation 5), or replacing z_{t-1} with $x_{1:t-1}$ (Equation 6):

$$\begin{aligned} & p(z_t|z_{t-1}, x_{t:T}^{1:M}) \\ &= p(x_{t+1:T}^{1:M}|z_t) p(x_t^{1:M}|z_t) \frac{p(z_t|z_{t-1})}{p(x_{t:T}^{1:M}|z_{t-1})} \\ &\propto p(x_{t+1:T}^{1:M}|z_t) \left[\prod_{m=1}^M p(x_t^m|z_t) \right] p(z_t|z_{t-1}) \\ &= \frac{p(z_t|x_{t+1:T}^{1:M}) p(x_{t+1:T}^{1:M})}{p(z_t)} \left[\prod_{m=1}^M \frac{p(z_t|x_t^m) p(x_t^m)}{p(z_t)} \right] p(z_t|z_{t-1}) \\ &\propto p(z_t|x_{t+1:T}^{1:M}) \left[\prod_{m=1}^M \frac{p(z_t|x_t^m)}{p(z_t)} \right] \frac{p(z_t|z_{t-1})}{p(z_t)} \end{aligned}$$

The factorizations in Equations 5–7 lead to several useful insights. First, they show that any missing modalities $\bar{m} \in [1, M]$ at time t can simply be left out of the product over modalities, leaving us with distributions that correctly condition on only the modalities $[1, M] \setminus \{\bar{m}\}$ that are present. Second, they suggest that we can compute all three distributions if we can approximate the dependence on the future, $q(z_t|x_{t+1:T}) \simeq p(z_t|x_{t+1:T})$, learn approximate posteriors $q(z_t|x_t^m) \simeq p(z_t|x_t^m)$ for each modality m , and know the model dynamics $p(z_t), p(z_t|z_{t-1})$.

Multimodal Fusion via Product of Gaussians

However, there are a few obstacles to performing tractable computation of Equations 5–7. One obstacle is that it is not tractable to compute the product of generic probability distributions. To address this, we adopt the approach used for

the MVAE (Wu and Goodman 2018), making the assumption that each term in Equations 5–7 is Gaussian. If each distribution is Gaussian, then their products or quotients are also Gaussian and can be computed in closed form. Since this result is well-known, we state it in the supplement (see Wu and Goodman (2018) for a proof).

This Product-of-Gaussians approach has the added benefit that the output distribution is dominated by the input Gaussian terms with lower variance (higher precision), thereby fusing information in a way that gives more weight to higher-certainty inputs (Cao and Fleet 2014; Ong, Zaki, and Goodman 2015). This automatically balances the information provided by each modality m , depending on whether $p(z_t|x_t^m)$ is high or low certainty, as well as the information provided from the past and future through $p(z_t|z_{t-1})$ and $p(z_t|x_{t+1:T})$, thereby performing multimodal temporal fusion in a manner that is uncertainty-aware.

Approximate Filtering with Missing Data

Another obstacle to computing Equations 5–7 is the dependence on future observations, $p(z_t|x_{t+1:T})$, which does not admit further factorization, and hence does not readily handle missing data among those future observations. Other approaches to approximating this dependence on the future rely on RNNs as recognition models (Krishnan, Shalit, and Sontag 2017; Che et al. 2018), but these are not designed to work with missing data.

To address this obstacle in a more principled manner, our insight was that $p(z_t|x_{t+1:T})$ is the expectation of $p(z_t|z_{t+1})$ under the backwards filtering distribution, $p(z_{t+1}|x_{t+1:T})$:

$$p(z_t|x_{t+1:T}) = \mathbb{E}_{p(z_{t+1}|x_{t+1:T})} [p(z_t|z_{t+1})] \quad (8)$$

For tractable approximation of this expectation, we use an approach similar to assumed density filtering (Huber, Beutler, and Hanebeck 2011). We assume both $p(z_t|x_{t+1:T})$ and $p(z_t|z_{t+1})$ to be multivariate Gaussian with diagonal covariance, and sample the parameters μ, Σ of $p(z_t|z_{t+1})$ under $p(z_{t+1}|x_{t+1:T})$. After drawing K samples, we approximate the parameters of $p(z_t|x_{t+1:T})$ via empirical moment-matching:

$$\hat{\mu}_{z_t|x_{t+1:T}} = \frac{1}{K} \sum_{k=1}^K \mu_k \quad (9)$$

$$\hat{\Sigma}_{z_t|x_{t+1:T}} = \frac{1}{K} \sum_{k=1}^K (\Sigma_k + \mu_k^2) - \hat{\mu}_{z_t|x_{t+1:T}}^2 \quad (10)$$

Approximating $p(z_t|x_{t+1:T})$ by $p(z_t|z_{t+1})$ led us to three important insights. First, by substituting the expectation from Equation 8 into Equation 5, the backward filtering distribution becomes:

$$p(z_t|x_{t:T}) \propto \mathbb{E}_{p(z_t|z_{t+1})} \left[\prod_{m=1}^M \frac{p(z_t|x_t^m)}{p(z_t)} \right] \quad (11)$$

In other words, by sampling under the filtering distribution for time $t+1$, $p(z_{t+1}|x_{t+1:T})$, we can compute the filtering distribution for time t , $p(z_t|x_{t:T})$. We can thus recursively compute $p(z_t|x_{t:T})$ backwards in time, starting from $t=T$.

Second, once we can perform filtering backwards in time, we can use this to approximate $p(z_t|x_{t+1:T})$ in the smoothing distribution (Equation 6) and the conditional smoothing

posterior (Equation 7). Backward filtering hence allows us to approximate both smoothing and sequencing.

Third, this approach removes the explicit dependence on all future observations $x_{t+1:T}$, allowing us to handle missing data. Suppose the data points $X_{\#} = \{x_{t_i}^{m_i}\}$ are missing, where t_i and m_i are the time-step and modality of the i th missing point respectively. Rather than directly compute the dependence on an incomplete set of future observations, $p(z_t|x_{t+1:T} \setminus X_{\#})$, we can instead sample z_{t+1} under the filtering distribution conditioned on incomplete observations, $p(z_{t+1}|x_{t+1:T} \setminus X_{\#})$, and then compute $p(z_t|z_{t+1})$ given the sampled z_{t+1} , thereby approximating $p(z_t|x_{t+1:T} \setminus X_{\#})$.

Backward-Forward Variational Inference

We now introduce factorized variational approximations of Equations 5–7. We replace the true posteriors $p(z_t|x_t^m)$ with variational approximations $q(z_t|x_t^m) := \tilde{q}(z_t|x_t^m)p(z_t)$, where $\tilde{q}(z_t|x_t^m)$ is parameterized by a (time-invariant) neural network for each modality m . As in the MVAE, we learn the Gaussian quotients $\tilde{q}(z_t|x_t^m) := q(z_t|x_t^m)/p(z_t)$ directly, so as to avoid the constraint required for ensuring a quotient of Gaussians is well-defined. We also parameterize the transition dynamics $p(z_t|z_{t-1})$ and $p(z_t|z_{t+1})$ using neural networks for the quotient distributions. This gives the following approximations:

$$\text{Backward Filtering} \quad (12)$$

$$q(z_t|x_{t:T}) \propto \mathbb{E}_{\leftarrow} [p(z_t|z_{t+1})] \prod_m \tilde{q}(z_t|x_t^m)$$

$$\text{Forward Smoothing} \quad (13)$$

$$q(z_t|x_{1:T}) \propto \mathbb{E}_{\leftarrow} [p(z_t|z_{t+1})] \prod_m \tilde{q}(z_t|x_t^m) \frac{\mathbb{E}_{\rightarrow} [p(z_t|z_{t-1})]}{p(z_t)}$$

$$\text{Conditional Smoothing Posterior} \quad (14)$$

$$q(z_t|z_{t-1}, x_{t:T}) \propto \mathbb{E}_{\leftarrow} [p(z_t|z_{t+1})] \prod_m \tilde{q}(z_t|x_t^m) \frac{p(z_t|z_{t-1})}{p(z_t)}$$

Here, \mathbb{E}_{\leftarrow} is shorthand for the expectation under the approximate backward filtering distribution $q(z_{t+1}|x_{t+1:T})$, while \mathbb{E}_{\rightarrow} is the expectation under the forward smoothing distribution $q(z_{t-1}|x_{1:T})$.

To calculate the backward filtering distribution $q(z_t|x_{t:T})$, we introduce a *variational backward algorithm* (Algorithm 1) to recursively compute Equation 12 for all time-steps t in a single pass. Note that simply by reversing time in Algorithm 1, this gives us a *variational forward algorithm* that computes the forward filtering distribution $q(z_t|x_{1:t})$.

Unlike filtering, smoothing and sequencing require information from both past ($p(z_t|z_{t-1})$) and future ($p(z_t|z_{t+1})$). This motivates a *variational backward-forward algorithm* (Algorithm 2) for smoothing and sequencing. Algorithm 2 first uses Algorithm 1 as a backward pass, then performs a forward pass to propagate information from past to future. Algorithm 2 also requires knowing $p(z_t)$ for each t . While this can be computed by sampling in the forward pass, we avoid the instability (of sampling T successive latents with no observations) by instead assuming $p(z_t)$ is constant with time, i.e., the MDMM is stationary when nothing is observed. During training, we add $\text{KL}(p(z_t)||\mathbb{E}_{z_{t-1}}p(z_t|z_{t-1}))$ and $\text{KL}(p(z_t)||\mathbb{E}_{z_{t+1}}p(z_t|z_{t+1}))$ to the loss to ensure that the transition dynamics obey this assumption.

Algorithm 1 A variational backward algorithm for approximate backward filtering.

```

function BACKWARDFILTER( $x_{1:T}, K$ )
  Initialize  $q(z_t|x_{t+1:T}) \leftarrow p(z_T)$ 
  for  $t = T$  to 1 do
    Let  $\mathcal{M} \subseteq [1, M]$  be the observed modalities at  $t$ 
     $q(z_t|x_{t:T}) \leftarrow q(z_t|x_{t+1:T}) \prod_{\mathcal{M}} \tilde{q}(z_t|x_t^m)$ 
    Sample  $K$  particles  $z_t^k \sim q(z_t|x_{t:T})$  for  $k \in [1, K]$ 
    Compute  $p(z_{t-1}|z_t^k)$  for each particle  $z_t^k$ 
     $q(z_{t-1}|x_{t:T}) \leftarrow \frac{1}{K} \sum_{k=1}^K p(z_{t-1}|z_t^k)$ 
  end for
  return  $\{q(z_t|x_{t:T}), q(z_t|x_{t+1:T})$  for  $t \in [1, T]\}$ 
end function

```

Algorithm 2 A variational backward-forward algorithm for approximate forward smoothing.

```

function FORWARDSMOOTH( $x_{1:T}, K_b, K_f$ )
  Initialize  $\tilde{p}(z_t|x_{1:0}) \leftarrow 1$ 
  Collect  $q(z_t|x_{t+1:T})$  from BACKWARDFILTER( $x_{1:T}, K_b$ )
  for  $t = 1$  to  $T$  do
    Let  $\mathcal{M} \subseteq [1, M]$  be the observed modalities at  $t$ 
     $q(z_t|x_{1:T}) \leftarrow q(z_t|x_{t+1:T}) \prod_{\mathcal{M}} [\tilde{q}(z_t|x_t^m)] \frac{q(z_t|x_{1:t-1})}{p(z_t)}$ 
    Sample  $K_f$  particles  $z_t \sim q(z_t|x_{1:T})$  for  $k \in [1, K_f]$ 
    Compute  $p(z_{t+1}|z_t^k)$  for each particle  $z_t^k$ 
     $q(z_{t+1}|x_{1:t}) \leftarrow \frac{1}{K_f} \sum_{k=1}^{K_f} p(z_{t+1}|z_t^k)$ 
  end for
  return  $\{q(z_t|x_{1:T}), q(z_t|x_{1:t-1})$  for  $t \in [1, T]\}$ 
end function

```

While Algorithm 1 approximates the filtering distribution $q(z_t|x_{t:T})$, by setting the number of particles $K = 1$, it effectively computes the (backward) conditional filtering posterior $q(z_t|z_{t+1}, x_t)$ and (backward) conditional prior $p(z_t|z_{t+1})$ for a randomly sampled latent sequence $z_{1:T}$. Similarly, while Algorithm 2 approximates smoothing by default, when $K_f = 1$, it effectively computes the (forward) conditional smoothing posterior $q(z_t|z_{t-1}, x_{t:T})$ and (forward) conditional prior $p(z_t|z_{t-1})$ for a random latent sequence $z_{1:T}$. These quantities are useful not only because they allow us to perform sequencing, but also because we can use them to compute the ELBO for both backward filtering and forward smoothing:

$$L_{\text{filter}} = \sum_{t=1}^T [\mathbb{E}_{q(z_t|x_{t:T})} \log p(x_t|z_t) - \mathbb{E}_{q(z_{t+1}|x_{t+1:T})} \text{KL}(q(z_t|z_{t+1}, x_t) || p(z_t|z_{t+1}))] \quad (15)$$

$$L_{\text{smooth}} = \sum_{t=1}^T [\mathbb{E}_{q(z_t|x_{1:T})} \log p(x_t|z_t) - \mathbb{E}_{q(z_{t-1}|x_{1:T})} \text{KL}(q(z_t|z_{t-1}, x_{t:T}) || p(z_t|z_{t-1}))] \quad (16)$$

L_{filter} is the filtering ELBO because it corresponds to a ‘backward filtering’ variational posterior $q(z_{1:T}|x_{1:T}) = \prod_t q(z_t|z_{t+1}, x_t)$, where each z_t is only inferred using the current observation x_t and the future latent state z_{t+1} . L_{smooth} is the smoothing ELBO because it corresponds to the correct factorization of the posterior in Equation 4, where each term combines information from both past and future. Since L_{smooth} corresponds to the correct factorization, it should theoretically be enough to minimize L_{smooth} to learn

Method	Recon.	Drop Half	Fwd. Extra.	Bwd. Extra.	Cond. Gen.	Label Pred.
Spirals Dataset: MSE (SD)						
BFVI (ours)	0.02 (0.01)	0.04 (0.01)	0.12 (0.10)	0.07 (0.03)	0.26 (0.26)	–
F-Mask	0.02 (0.01)	0.06 (0.02)	0.10 (0.08)	0.18 (0.07)	1.37 (1.39)	–
F-Skip	0.04 (0.01)	0.10 (0.05)	0.13 (0.11)	0.19 (0.06)	1.51 (1.54)	–
B-Mask	0.02 (0.01)	0.04 (0.01)	0.18 (0.14)	0.04 (0.01)	1.25 (1.23)	–
B-Skip	0.05 (0.01)	0.19 (0.05)	0.32 (0.22)	0.37 (0.15)	1.64 (1.51)	–
Weizmann Video Dataset: SSIM or Accuracy* (SD)						
BFVI (ours)	.85 (.03)	.84 (.04)	.84 (.04)	.83 (.05)	.85 (.03)	.69 (.33)*
F-Mask	.68 (.18)	.66 (.18)	.68 (.18)	.66 (.17)	.60 (.15)	.33 (.33)*
F-Skip	.70 (.12)	.68 (.14)	.70 (.12)	.67 (.16)	.63 (.12)	.21 (.26)*
B-Mask	.79 (.04)	.79 (.04)	.79 (.04)	.79 (.04)	.76 (.06)	.46 (.34)*
B-Skip	.80 (.04)	.79 (.04)	.80 (.04)	.79 (.04)	.74 (.08)	.29 (.37)*

Table 1: Evaluation metrics on both datasets across inference methods and tasks. Best performance per task (column) in bold. (Top) *Spirals Dataset*: MSE (lower is better) per time-step between reconstructions and ground truth spirals. For scale, the average squared spiral radius is about 5 sq. units. (Bottom) *Weizmann Video Dataset*: SSIM or label accuracy (higher is better) per time-step with respect to original videos. Means and Standard Deviations (SD) are across the test set.

good MDMM parameters θ, ϕ . However, in order to compute L_{smooth} , we must perform a backward pass which requires sampling under the backward filtering distribution. Hence, to accurately approximate L_{smooth} , the backward filtering distribution has to be reasonably accurate as well. This motivates learning the parameters θ, ϕ by jointly maximizing the filtering and smoothing ELBOs as a weighted sum. We call this paradigm **backward-forward variational inference (BFVI)**, due to its use of variational posteriors for both backward filtering and forward smoothing.

Experiments

We compare **BFVI** against state-of-the-art RNN-based inference methods on two multimodal time series datasets over a range of inference tasks. **F-Mask** and **F-Skip** use forward RNNs (one per modality), using zero-masking and update skipping respectively to handle missing data. They are thus multimodal variants of the ST-L network in (Krishnan, Shalit, and Sontag 2017), and similar to the variational RNN (Chung et al. 2015) and recurrent SSM (Hafner et al. 2018). **B-Mask** and **B-Skip** use backward RNNs, with masking and skipping respectively, and correspond to the Deep Kalman Smoother in (Krishnan, Shalit, and Sontag 2017). The underlying MDMM architecture is constant across inference methods. Architectural and training details can be found in the supplement. Code is available at <https://git.io/Jeeze>.

Datasets

Noisy Spirals. We synthesized a dataset of 1000 noisy 2D spirals (600 train / 400 test) similar to Chen *et al.* (Chen et al. 2018), treating the x and y coordinates as two separate modalities. Spiral trajectories vary in direction (clockwise or counter-clockwise), size, and aspect ratio, and Gaussian noise is added to the observations. We used 5 latent dimensions, and two-layer perceptrons for encoding $q(z_t|x_t^m)$ and decoding $p(x_t^m|z_t)$. For evaluation, we compute the mean

squared error (MSE) per time step between the predicted trajectories and ground truth spirals.

Weizmann Human Actions. This is a video dataset of 9 people each performing 10 actions (Gorelick et al. 2007). We converted it to a trimodal time series dataset by treating silhouette masks as an additional modality, and treating actions as per-frame labels, similar to He et al. (2018). Each RGB frame was cropped to the central 128×128 window and resized to 64×64 . We selected one person’s videos as the test set, and the other 80 videos as the training set, allowing us to test action label prediction on an unseen person. We used 256 latent dimensions, and convolutional / deconvolutional neural networks for encoding and decoding. For evaluation, we compute the Structural Similarity (SSIM) between the input video frames and the reconstructed outputs.

Inference Tasks

We evaluated all methods on the following suite of temporal inference tasks for both datasets: **reconstruction**: reconstruction given complete observations; **drop half**: reconstruction after half of the inputs are randomly deleted; **forward extrapolation**: predicting the last 25% of a sequence when the rest is given; and **backward extrapolation**: inferring the first 25% of a sequence when the rest is given.

When evaluating these tasks on the Weizmann dataset, we provided only video frames as input (i.e. with neither the silhouette masks nor the action labels), to test whether the methods were capable of unimodal inference after multimodal training.

We also tested cross-modal inference using the following *conditional generation / label prediction* tasks: **conditional generation (Spirals)**: given x - and initial 25% of y -coordinates, generate rest of spiral; **conditional generation (Weizmann)**: given the video frames, generate the silhouette masks; and **label prediction (Weizmann)**: infer action labels given only video frames.

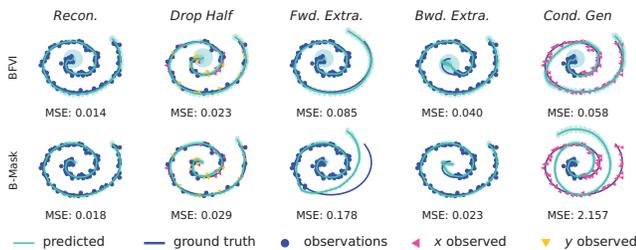


Figure 2: Reconstructions for all 5 spiral inference tasks for BFVI and the next best method, B-Mask. BFVI outperforms B-Mask significantly on both forward extrapolation and conditional generation.

Table 1 shows the results for the inference tasks, while Figure 2 and 3 show sample reconstructions from the Spirals and Weizmann datasets respectively. On the Spirals dataset, BFVI achieves high performance on all tasks, whereas the RNN-based methods only perform well on a few. In particular, all methods besides BFVI do poorly on the conditional generation task, which can be understood from the right-most column of Figure 2. BFVI generates a spiral that matches the provided x -coordinates, while the next-best method, B-Mask, completes the trajectory with a plausible spiral, but ignores the x observations entirely in the process.

On the more complex Weizmann video dataset, BFVI outperforms all other methods on every task, demonstrating both the power and flexibility of our approach. The RNN-based methods performed especially poorly on label prediction, and this was the case even on the training set (not shown in Table 1). We suspect that this is because the RNN-based methods lack a principled approach to multimodal fusion, and hence fail to learn a latent space which captures the mutual information between action labels and images. In contrast, BFVI learns to both predict one modality from another, and to propagate information across time, as can be seen from the reconstruction and predictions in Figure 3.

Weakly Supervised Learning

In addition to performing inference with missing data test time, we compared the various methods ability to learn with missing data at training time, amounting to a form of weakly supervised learning. We tested two forms of weakly supervised learning on the Spirals dataset, corresponding to different conditions of data incompleteness. The first was learning with data missing uniformly at random. This condition can arise when sensors are noisy or asynchronous. The second was learning with missing modalities, or semi-supervised learning, where a fraction of the sequences in the dataset only has a single modality present. This condition can arise when a sensor breaks down, or when the dataset is partially unlabelled by annotators. We also tested learning with missing modalities on the Weizmann dataset.

Results for these experiments are shown in Figure 4, which compare BFVI’s performance on increasing levels of missing data against the next best method, averaged across 10 trials. Our method (BFVI) performs well, maintaining

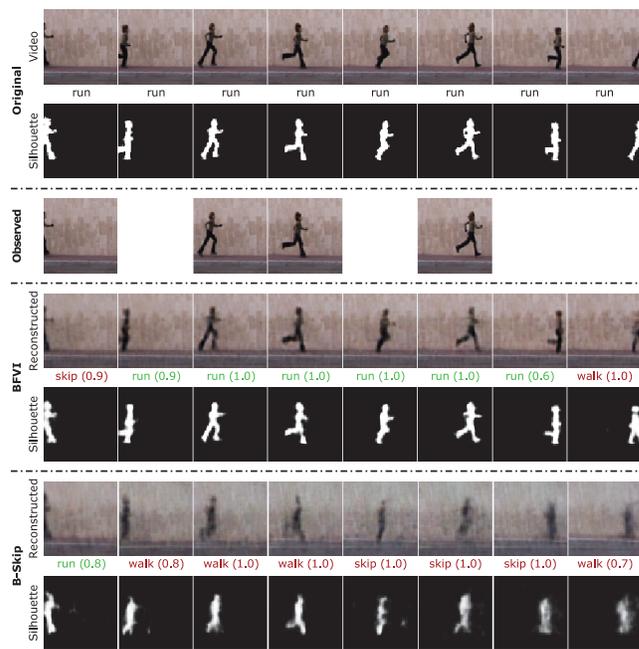


Figure 3: Snapshots of a ‘running’ video and silhouette mask from the Weizmann dataset (rows 1–2), with half of the frames deleted at random (14 out of 28 frames), and neither action labels nor silhouettes provided as observations (row 3). BFVI reconstructs the video and a running silhouette, and also correctly predicts the action (rows 4–5). By contrast, B-Skip (the next best method) creates blurred and wispy reconstructions, wrongly predicts the action label, and vacillates between different possible action silhouettes over time (rows 6–7).

good performance on the Spirals dataset even with 70% uniform random deletion (Figure 4a) and 60% uni-modal examples (Figure 4b), while degrading gracefully with increasing missingness. This is in contrast to B-Mask, which is barely able to learn when even 10% of the spiral examples are uni-modal, and performs worse than BFVI on the Weizmann dataset at all levels of missing data (Figure 4c).

Conclusion

In this paper, we introduced backward-forward variational inference (BFVI) as a novel inference method for Multimodal Deep Markov Models. This method handles incomplete data via a factorized variational posterior, allowing us to easily marginalize over missing observations. Our method is thus capable of a large range of multimodal temporal inference tasks, which we demonstrate on both a synthetic dataset and a video dataset of human motions. The ability to handle missing data also enables applications in weakly supervised learning of labelled time series. Given the abundance of multimodal time series data where missing data is the norm rather than the exception, our work holds great promise for many future applications.

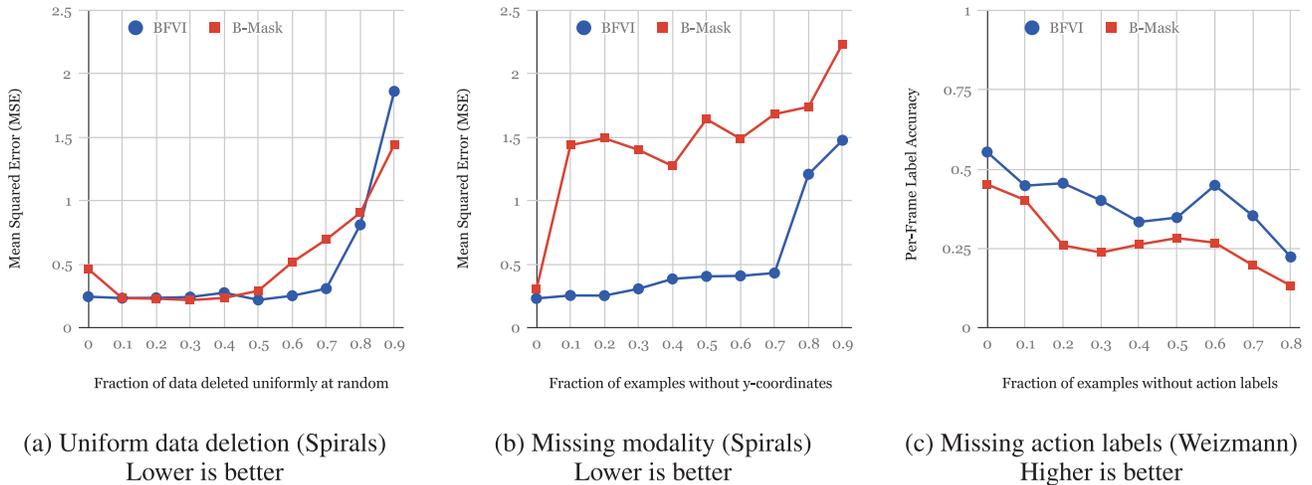


Figure 4: Learning curves under various forms of weak supervision. (a) Learning with randomly missing data on the Spirals dataset. (b) Semi-supervised learning on the Spirals dataset, where some sequences are entirely missing the y-coordinate. (c) Semi-supervised learning on the Weizmann dataset, where some sequences have no action labels.

Acknowledgements

This work was supported by the A*STAR Human-Centric Artificial Intelligence Programme (SERC SSF Project No. A1718g0048).

References

- Archer, E.; Park, I. M.; Buesing, L.; Cunningham, J.; and Paninski, L. 2015. Black box variational inference for state space models. *arXiv preprint arXiv:1511.07367*.
- Baker, C. L.; Jara-Ettinger, J.; Saxe, R.; and Tenenbaum, J. B. 2017. Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nature Human Behaviour* 1(4):0064.
- Bayer, J., and Osendorfer, C. 2014. Learning stochastic recurrent networks. In *NIPS 2014 Workshop on Advances in Variational Inference*.
- Buesing, L.; Weber, T.; Racaniere, S.; Eslami, S.; Rezende, D.; Reichert, D. P.; Viola, F.; Besse, F.; Gregor, K.; Hasabis, D.; et al. 2018. Learning and querying fast generative models for reinforcement learning. *arXiv preprint arXiv:1802.03006*.
- Cao, Y., and Fleet, D. J. 2014. Generalized product of experts for automatic and principled fusion of gaussian process predictions. In *Modern Nonparametrics 3: Automating the Learning Pipeline Workshop, NeurIPS 2014*.
- Che, Z.; Purushotham, S.; Li, G.; Jiang, B.; and Liu, Y. 2018. Hierarchical deep generative models for multi-rate multivariate time series. In Dy, J., and Krause, A., eds., *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 784–793. Stockholmsmässan, Stockholm Sweden: PMLR.
- Chen, T. Q.; Rubanova, Y.; Bettencourt, J.; and Duvenaud, D. K. 2018. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems*, 6571–6583.
- Chung, J.; Kastner, K.; Dinh, L.; Goel, K.; Courville, A. C.; and Bengio, Y. 2015. A recurrent latent variable model for sequential data. In *Advances in neural information processing systems*, 2980–2988.
- Doerr, A.; Daniel, C.; Schiegg, M.; Duy, N.-T.; Schaal, S.; Toussaint, M.; and Sebastian, T. 2018. Probabilistic recurrent state-space models. In Dy, J., and Krause, A., eds., *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 1280–1289. Stockholmsmässan, Stockholm Sweden: PMLR.
- Dragan, A. D.; Lee, K. C.; and Srinivasa, S. S. 2013. Legibility and predictability of robot motion. In *Proceedings of the 8th ACM/IEEE international conference on Human-robot interaction*, 301–308. IEEE Press.
- Fabius, O., and van Amersfoort, J. R. 2014. Variational recurrent auto-encoders. *arXiv preprint arXiv:1412.6581*.
- Fraccaro, M.; Sønderby, S. K.; Paquet, U.; and Winther, O. 2016. Sequential neural models with stochastic layers. In *Advances in Neural Information Processing Systems*, 2199–2207.
- Fraccaro, M.; Kamronn, S.; Paquet, U.; and Winther, O. 2017. A disentangled recognition and nonlinear dynamics model for unsupervised learning. In *Advances in Neural Information Processing Systems*, 3601–3610.
- Gorelick, L.; Blank, M.; Shechtman, E.; Irani, M.; and Basri, R. 2007. Actions as space-time shapes. *IEEE transactions on pattern analysis and machine intelligence* 29(12):2247–2253.
- Hafner, D.; Lillicrap, T.; Fischer, I.; Villegas, R.; Ha, D.; Lee, H.; and Davidson, J. 2018. Learning latent dynamics for planning from pixels. *arXiv preprint arXiv:1811.04551*.

- He, J.; Lehrmann, A.; Marino, J.; Mori, G.; and Sigal, L. 2018. Probabilistic video generation using holistic attribute control. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 452–467.
- Huber, M. F.; Beutler, F.; and Hanebeck, U. D. 2011. Semi-analytic gaussian assumed density filter. In *Proceedings of the 2011 American Control Conference*, 3006–3011. IEEE.
- Johnson, M.; Duvenaud, D. K.; Wiltchko, A.; Adams, R. P.; and Datta, S. R. 2016. Composing graphical models with neural networks for structured representations and fast inference. In *Advances in Neural Information Processing Systems*, 2946–2954.
- Kalman, R. E. 1960. A new approach to linear filtering and prediction problems. *Journal of basic Engineering* 82(1):35–45.
- Karl, M.; Soelch, M.; Bayer, J.; and van der Smagt, P. 2016. Deep variational bayes filters: Unsupervised learning of state space models from raw data. *arXiv preprint arXiv:1605.06432*.
- Kingma, D. P., and Welling, M. 2014. Auto-encoding variational bayes. In *International Conference on Learning Representations*.
- Krishnan, R. G.; Shalit, U.; and Sontag, D. 2017. Structured inference networks for nonlinear state space models. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Lin, W.; Khan, M. E.; and Hubacher, N. 2018. Variational message passing with structured inference networks. In *International Conference on Learning Representations*.
- Lipton, Z. C.; Kale, D.; and Wetzell, R. 2016. Directly modeling missing data in sequences with rnns: Improved classification of clinical time series. In *Machine Learning for Healthcare Conference*, 253–270.
- Neil, D.; Pfeiffer, M.; and Liu, S.-C. 2016. Phased lstm: Accelerating recurrent network training for long or event-based sequences. In *Advances in neural information processing systems*, 3882–3890.
- Ong, D. C.; Zaki, J.; and Goodman, N. D. 2015. Affective cognition: Exploring lay theories of emotion. *Cognition* 143:141–162.
- Rezende, D. J.; Mohamed, S.; and Wierstra, D. 2014. Stochastic backpropagation and approximate inference in deep generative models. In *International Conference on Machine Learning*, 1278–1286.
- Wu, M., and Goodman, N. 2018. Multimodal generative models for scalable weakly-supervised learning. In *Advances in Neural Information Processing Systems*, 5575–5585.