Supplemental Material for An Adversarial Hierarchical Hidden Markov Model for Human Pose Modeling and Generation

Rui Zhao and **Qiang Ji**
Rensselaer Polytechnic Institute
Troy NY, USA
{zhaor,jiq}@rpi.edu

Derivation of gradient of HMM

We derive the gradient of marginal loglikelihood of HMM with respect to the model parameters. A similar derivation in the context of computing MLE of HMM parameters is also provided by (Cappé, Buchoux, and Moulines 1998). First of all, the HMM is parameterized by a set of parameters $\theta = \{\pi, \mathbf{A}, \gamma\}$, where π is the initial state parameters, $\mathbf{A}$ is the state transition matrix and $\gamma = \{\mu, \Sigma\}$ is the Gaussian emission distribution. The joint distribution of HMM is as follows.

$$P(\mathbf{X}, \mathbf{Z}|\theta) = P(Z_1|\pi) \prod_{t=2}^{T} P(Z_t|Z_{t-1}, \mathbf{A}) \prod_{t=1}^{T} P(X_t|Z_t, \gamma) \tag{1}$$

The marginal loglikelihood of the model can be written as

$$\log P(\mathbf{X}|\pi, A, \gamma) = \log \sum_{\mathbf{Z}} P(\mathbf{X}, \mathbf{Z}|\pi, A, \gamma) \tag{2}$$

However, the number of summations needed to sum over $\mathbf{Z}$ is exponential to the sequence length. It is infeasible even for moderately long sequence. We exploit the Markov property of the state chain to develop a recursive algorithm to compute gradient of $\log P(\mathbf{X})$. To begin with, we write $\log P(\mathbf{X})$ as follows. We omit the dependency on parameters for clarity.

$$\begin{aligned}\log P(\mathbf{X}) &= \sum_{t=1}^{T} \log P(X_t|X_{1\ldots t-1}) \qquad (3)\\ &= \sum_{t=1}^{T} \log \sum_{i=1}^{Q} P(X_t, Z_t = i|X_{1\ldots t-1})\\ &= \sum_{t=1}^{T} \log \sum_{i=1}^{Q} P(X_t|Z_t = i) P(Z_t = i|X_{1\ldots t-1})\end{aligned}$$

We define

$$f_i(X_t) = P(X_t|Z_t = i) \tag{4}$$

$$\phi_t(i) = P(Z_t = i|X_{1\ldots t-1}) \in [0, 1] \tag{5}$$

Then

$$\log P(\mathbf{X}) = \sum_{t=1}^{T} \log \sum_{i=1}^{Q} f_i(X_t)\phi_t(i) \tag{6}$$

Given Eq. (6), it is straightforward to show the following identities.

$$\nabla_\pi \log P(\mathbf{X}) = \frac{1}{c_1} \sum_{i=1}^{Q} f_i(X_1) \nabla_\pi \phi_1(i) \tag{7}$$

$$\nabla_{\mathbf{A}} \log P(\mathbf{X}) = \frac{1}{c_t} \sum_{t=1}^{T} \sum_{i=1}^{Q} f_i(X_t) \nabla_{\mathbf{A}} \phi_t(i) \tag{8}$$

$$\begin{aligned}\nabla_\gamma \log P(\mathbf{X}) = \frac{1}{c_t} \sum_{t=1}^{T} \sum_{i=1}^{Q} [\phi_t(i) \nabla_\gamma f_i(X_t) \qquad (9)\\ + f_i(X_t) \nabla_\gamma \phi_t(i)]\end{aligned}$$

Now we show how to compute $\nabla \phi_t(i)$ and $\nabla f_i(X_t)$. First, exploiting the Markov assumption of the state chain, we have the following recursion on $\phi_t(i)$.

$$\phi_{t+1}(j) = \frac{1}{c_t} \sum_{i=1}^{Q} f_i(X_t) \phi_t(i) A_{ij} \tag{10}$$

where $c_t = \sum_{i=1}^{Q} f_i(X_t)\phi_t(i) = P(X_t|X_{1\ldots t-1})$ is the normalization term at time t. Using this recursion, we can compute $\nabla \phi_t(i)$ recursively as follows.

$$\nabla_\pi \phi_1(j) = \nabla_\pi P(Z_1 = j) = e_j \tag{11}$$

$$\begin{aligned}\nabla_{\mathbf{A}} \phi_{t+1}(j) = \frac{1}{c_t} \sum_{i=1}^{Q} f_i(X_t) [\phi_t(i) \nabla_{\mathbf{A}} A_{ij} \qquad (12)\\ + (A_{ij} - \phi_{t+1}(j)) \nabla_{\mathbf{A}} \phi_t(i)]\end{aligned}$$

$$\begin{aligned}\nabla_\gamma \phi_{t+1}(j) = \frac{1}{c_t} \sum_{i=1}^{Q} (A_{ij} - \phi_{t+1}(j)) \qquad (13)\\ [\phi_t(i) \nabla_\gamma f_i(X_t) + f_i(X_t) \nabla_\gamma \phi_t(i)]\end{aligned}$$

For initialization, we use $\nabla_{\mathbf{A}} \phi_1(j) = \mathbf{0}, \nabla_\gamma \phi_1(j) = \{\nabla_\mu \phi_1(j), \nabla_\Sigma \phi_1(j)\} = \mathbf{0}, j = 1, ..., Q$.

Second, we can compute gradient of emission parameters $\nabla_\gamma f_i(X_t) = \{\nabla_\mu f_i(X_t), \nabla_\Sigma f_i(X_t)\}, i = 1, ..., Q$ as follows

$$\nabla_{\mu_i} f_i(X_t) = f_i(X_t)\Sigma_i^{-1}(X_t - \mu_i) \tag{14}$$

$$\nabla_{\sigma_i} f_i(X_t) = f_i(X_t)[-\frac{1}{2}\Sigma_i^{-1} + \frac{1}{2}\Sigma_i^{-1}(X_t - \mu_i)(X_t - \mu_i)^T\Sigma_i^{-1}] \tag{15}$$

Assuming Σ_i is diagonal with diagonal entries be σ_{ij}^2. We can reparameterize $\kappa_{ij} = \log \sigma_{ij}$. Then

$$\begin{aligned}\nabla_{\kappa_i} f_i(X_t) &= \frac{\partial f_i(X_t)}{\partial \Sigma_i}\frac{\partial \Sigma_i}{\partial \kappa_i} \\ &= f_i(X_t)[\Sigma_i^{-1}(X_t - \mu_i)(X_t - \mu_i)^T - I]\end{aligned} \tag{16}$$

where I is identity matrix with proper dimension.

Maximum Likelihood Learning of HHMM

Given a set of training data $\{\mathbf{X}\}_n$ and HHMM parameterized by $\alpha = \{\hat{\alpha}, \hat{\theta}\}$ as described in the main paper, the problem of MLE is defined as follows.

$$\begin{aligned}\alpha^* &= \arg\max_{\alpha} \log P(\{\mathbf{X}\}_n|\alpha) \\ &= \arg\max_{\alpha} \log \int_\theta \prod_n \sum_{\mathbf{Z}_n} P(\mathbf{X}_n, \mathbf{Z}_n|\theta, \hat{\theta})P(\theta|\hat{\alpha})d\theta\end{aligned} \tag{17}$$

Exact computation of Eq. (17) is intractable due to the integration over θ introduces additional dependencies among $\mathbf{Z}$. We instead approximate MLE using the following estimate.

$$\alpha^* = \arg\max_{\alpha} \sum_n \log \sum_{\mathbf{Z}_n} P(\mathbf{X}_n, \mathbf{Z}_n|\theta^*, \hat{\theta}) + \log P(\theta^*|\hat{\alpha}) \tag{18}$$

where θ^* is one particular choice of θ. This leads to an alternating estimation process between θ and α. First, we compute the following MAP estimate of θ using MAP-EM algorithm (Gauvain and Lee 1994) given current estimate of $\hat{\alpha}, \hat{\theta}$.

$$\theta^* = \arg\max_{\theta} \sum_n \log \sum_{\mathbf{Z}_n} P(\mathbf{X}_n, \mathbf{Z}_n|\theta, \hat{\theta}) + \log P(\theta|\hat{\alpha}) \tag{19}$$

The E-step has complexity $O(Q^2T)$ and M-step has closed-form solution. We initialize the value θ using prior mean and the MAP-EM algorithm typically converges in 10 iterations.

Second, we update $\hat{\alpha}, \hat{\theta}$ by solving Eq. (18) given current estimate of $\theta = \theta^*$. This yields two separate optimization problems for $\hat{\alpha}$ and $\hat{\theta}$ as follows.

$$\hat{\alpha}^* = \arg\max_{\hat{\alpha}} \log P(\theta^*|\hat{\alpha}) \tag{20}$$

$$\hat{\theta}^* = \arg\max_{\hat{\theta}} \sum_n \log \sum_{\mathbf{Z}_n} P(\mathbf{X}_n, \mathbf{Z}_n|\theta^*, \hat{\theta}) \tag{21}$$

The overall process is summarized in Algorithm 1.

Algorithm 1 Maximum likelihood learning of HHMM

Require: $\{\mathbf{X}\}_n$: observation sequences
Ensure: Augmented hyperparameters $\alpha = \{\hat{\alpha}, \hat{\theta}\}$.
1: Initialization of α
2: **repeat**
3: Update θ by solving Eq. (19) using MAP-EM
4: Update α by solving Eq. (20)-(21)
5: **until** convergence or reach maximum iteration number
6: **return** α

Detailed Results

Data reconstruction

The reconstruction results on each action class versus hidden state change is shown in Figure 1. Exact values of reconstruction results of individual action class is provided in Table 1 and 2. The methods compared include GPDM (Wang, Fleet, and Hertzmann 2008), ERD (Fragkiadaki et al. 2015) and TSBN (Gan et al. 2015).

Data synthesis

SSIM results of individual action class is provided in Table 3 and 4. The methods compared include CRBM (Taylor, Hinton, and Roweis 2006), ERD (Fragkiadaki et al. 2015) and TSBN (Gan et al. 2015).

References

Cappé, O.; Buchoux, V.; and Moulines, E. 1998. Quasi-newton method for maximum likelihood estimation of hidden markov models. In *ICASSP*.

Fragkiadaki, K.; Levine, S.; Felsen, P.; and Malik, J. 2015. Recurrent network models for human dynamics. In *ICCV*.

Gan, Z.; Li, C.; Henao, R.; Carlson, D. E.; and Carin, L. 2015. Deep temporal sigmoid belief networks for sequence modeling. In *NIPS*.

Gauvain, J.-L., and Lee, C.-H. 1994. Maximum a posteriori estimation for multivariate gaussian mixture observations of markov chains. *TSAP*.

Taylor, G. W.; Hinton, G. E.; and Roweis, S. T. 2006. Modeling human motion using binary latent variables. In *NIPS*.

Wang, J. M.; Fleet, D. J.; and Hertzmann, A. 2008. Gaussian process dynamical models for human motion. *PAMI*.

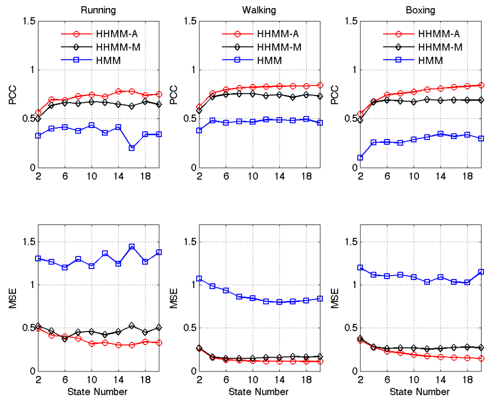

(a) Reconstruction on CMU

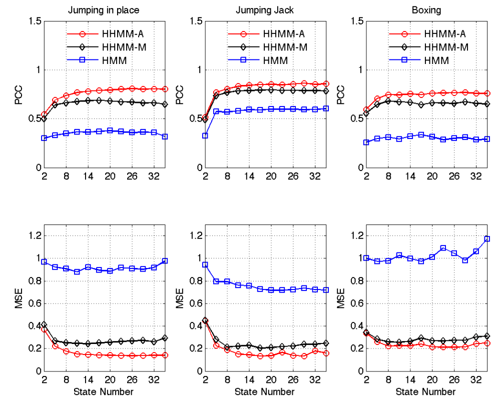

(b) Reconstruction on Berkeley

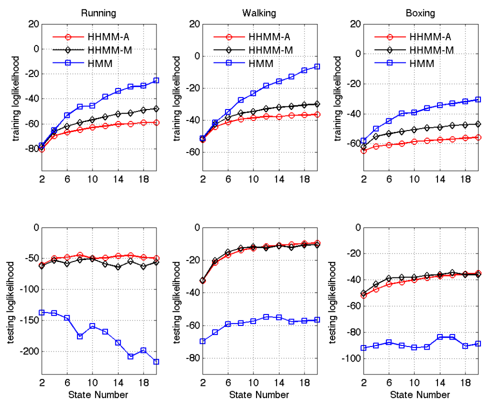

(c) Loglikelihood on CMU

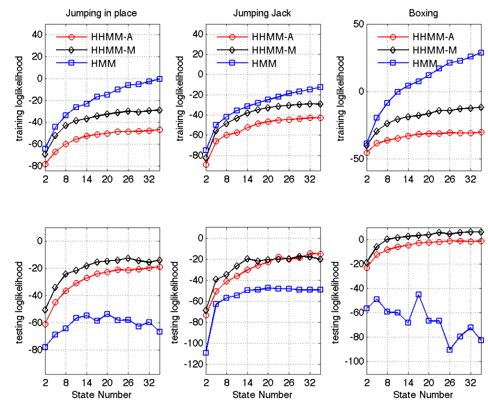

(d) Loglikelihood on Berkeley

Figure 1: Reconstruction experiment results versus change of number of hidden states. HHMM-A refers to adversarial learning variant and HHMM-M refers to maximum likelihood learning variant. (Best view in color)

Table 1: Motion reconstruction results on CMU.

Metric	PCC					
Method	HMM	GPDM	ERD	TSBN	HHMM-M	HHMM-A
Walking	0.5108	0.7018	0.5988	0.7991	0.7327	**0.8520**
Running	0.4291	0.6811	0.5716	0.6735	0.6472	**0.7576**
Boxing	0.2948	0.7089	0.8343	**0.8675**	0.6912	0.8413
Average	0.3637	0.6972	0.6574	0.7908	0.6904	**0.8170**
Metric	MSE					
Method	HMM	GPDM	ERD	TSBN	HHMM-M	HHMM-A
Walking	0.9233	0.1470	0.6050	0.1846	0.1690	**0.1037**
Running	1.2206	0.3723	0.8448	0.5181	0.5061	**0.3153**
Boxing	1.0714	0.2071	0.2474	**0.1259**	0.2719	0.1476
Average	1.1210	0.2421	0.6051	0.2724	0.3156	**0.1889**

Table 2: Motion reconstruction results on Berkeley.

Metric	PCC					
Method	HMM	GPDM	ERD	TSBN	HHMM-M	HHMM-A
Jumping	0.3762	0.4052	0.7323	0.7857	0.6819	**0.7910**
Jack	0.5961	0.4066	**0.8475**	0.8410	0.7964	0.8464
Boxing	0.3150	0.5916	0.6687	**0.7957**	0.6632	0.7876
Average	0.4291	0.4678	0.7495	0.8075	0.7138	**0.8084**
Metric	MSE					
Method	HMM	GPDM	ERD	TSBN	HHMM-M	HHMM-A
Jumping	0.8863	0.5355	0.3335	0.1497	0.2551	**0.1373**
Jack	0.7167	0.6546	0.2163	0.1662	0.2085	**0.1499**
Boxing	1.0115	0.3379	0.3910	0.2163	0.2702	**0.0645**
Average	0.8715	0.5093	0.3136	0.1774	0.2446	**0.1172**

Table 3: Average largest SSIM results on CMU.

Method	CRBM	ERD	TSBN	HHMM-M	HHMM-A	Training
Running	0.8295	0.3959	0.3390	0.4367	0.2326	0.0931
Walking	0.8276	0.1921	0.1342	0.3516	0.2114	0.0778
Boxing	0.4735	0.1301	0.1804	0.2562	0.1087	0.0657
Average	0.7102	0.2394	0.2179	0.3482	0.1842	0.0789

Table 4: Average largest SSIM results on Berkeley.

Method	CRBM	ERD	TSBN	HHMM-M	HHMM-A	Training
Jumping	0.5851	0.1886	0.2383	0.3245	0.1774	0.1031
Jack	0.6828	0.2764	0.2528	0.3962	0.2676	0.1260
Boxing	0.6558	0.2998	0.8497	0.3026	0.1514	0.1132
Average	0.6412	0.2550	0.4469	0.3411	0.1988	0.1141