
Divergence-Based Motivation for Online EM and Combining Hidden Variable Models

Ehsan Amid and Manfred K. Warmuth
UC Santa Cruz and Google Brain, Mountain View
{eamid, manfred}@google.com

Abstract

Expectation-Maximization (EM) is a prominent approach for parameter estimation of hidden (aka latent) variable models. Given the full batch of data, EM forms an upper-bound of the negative log-likelihood of the model at each iteration and updates to the minimizer of this upper-bound. We first provide a “model level” interpretation of the EM upper-bound as a sum of relative entropy divergences to a set of singleton models induced by the batch of observations. Our alternative motivation unifies the “observation level” and the “model level” view of the EM. As a result, we formulate an online version of the EM algorithm by adding an analogous inertia term which is a relative entropy divergence to the old model. Our motivation is more widely applicable than the previous approaches and leads to simple online updates for mixture of exponential distributions, hidden Markov models, and the first known online update for Kalman filters. Additionally, the finite sample form of the inertia term lets us derive online updates when there is no closed-form solution. Finally, we extend the analysis to the distributed setting where we motivate a systematic way of combining multiple hidden variable models. Experimentally, we validate the results on synthetic as well as real-world datasets.

1 INTRODUCTION

The goal of EM is to minimize¹ the negative log-likelihood (loss) of a hidden variable model given a set of iid observations from the data. Instead of directly mini-

¹We adopt the minimization view of the EM algorithm by considering the negative of the log-likelihood function. This will simplify our online EM motivation in the following.

mizing the negative log-likelihood, EM forms an upper-bound of the loss at each iteration and then updates to the minimizer² of the upper-bound. Due to the tightness of the upper-bound at the current estimate of the parameters, every EM update decreases the negative log-likelihood (or keeps it unchanged when the current estimate is at a local minimum).

We first provide an alternative view of the EM upper-bound as sum of joint relative entropy divergences between a set of “singleton models” and the new model. Each singleton model is induced by minimizing the EM upper-bound using only a single observation. This motivates the use of this divergence as an inertia term for an online variant of the EM algorithm. We add the relative entropy divergence between the joint distributions of the old and the new models to the EM upper-bound, and update to the minimizer of the sum. Curiously enough, our divergence based online EM updates coincide with the one given in (Sato, 2000) and its generalization in (Cappé & Moulines, 2009)³. However, we will show that the new formulation is more versatile and can be used to derive online updates for more complex models. In particular, our approach avoids having to identify the sufficient statistics of the joint distribution. Additionally, it can handle cases where there exists no closed-form solution for the EM update by using the Monte Carlo approximation of the inertia term. Finally, we extend the same methodology to minimizing a sum of relative entropy divergences between different models in a distributed setting. This results in an efficient way of combining hidden variable models. To summarize:

- We motivate the observation level view of the EM algorithm as combining models via minimizing a sum of relative entropy divergences.
- We formulate an online EM algorithm based on minimizing a sum of divergences to the singleton models

²Or the approximate solution of the minimization problem.

³Which is based on a stochastic approximation of the EM upper-bound.

(observations) plus a divergence to the old model (inertia term). The new formulation allows an approximate form of the online EM algorithm for cases where the updates do not have a closed-form.

- Using the new formulation, we obtain closed-form updates for mixtures of exponential distributions, hidden Markov models (HMMs), and Kalman filters, and approximate updates for the Compound Dirichlet distribution.
- Most importantly, we develop divergences between hidden variable models and provide a method for combining such models by minimizing convex combinations of divergences.

Here, we only consider models for which the EM upper-bound induces a closed-form. We omit extensions where other approximations are used for the upper bound (e.g. variational inference (Blei et al., 2017)).

Previous work EM is one of the most well-studied algorithms due to its simplicity and monotonic descent property (McLachlan & Krishnan, 2008; Do & Batzoglou, 2008; Gupta et al., 2011). It was also shown that EM converges to a stationary point of the negative log-likelihood under some mild conditions (Wu, 1983). EM is naturally a batch algorithm. Attempts for developing online versions of EM start with the work of Titterton (1984), who employs a second order method by approximating the complete data Fisher information matrix. This algorithm has been shown to almost surely converge to a local minimum of the negative log-likelihood (Wang & Zhao, 2006). However, deriving the updates for this method requires calculating sophisticated derivatives and matrix inversions, making it intractable for more complex models such as HMMs and Kalman filters. We show that our alternative divergence based motivation of the EM algorithm reduces to the work of Cappé & Moulines (2009), which substitutes the E-step by a stochastic approximation of the EM upper-bound while keeping the M-step unchanged. Cappé & Moulines (2009) showed that for models where the complete data likelihood belongs to an exponential family, the updates correspond to stochastic approximations of sufficient statistics. While this is intuitive, identifying the complete data sufficient statistic for more complex models becomes infeasible in practice. In contrast, our new formulation provides several advantages. First, we avoid characterizing the sufficient statistics by directly forming the inertia term between the current model and the updates. As a result, we can easily derive online EM updates for more complex models such as HMMs and Kalman filters. Additionally, we can apply the approximate form of the inertia term for problems where the minimization of the EM upper-bound does not have a closed-form solution. Finally, the methodology of minimizing sum of divergences between hidden vari-

able models leads to a natural way of combining multiple hidden variable models in a distributed settings.

Previous online EM algorithms for learning exponential family models have been mainly based on gradient ascent methods, or heuristic approaches for maximizing the likelihood or updating the sufficient statistics. Therefore, the resulting updates are commonly unstable and require careful tuning of the parameters. Generally, these updates also lack performance guarantees. Specifically, online methods have been developed for mixtures of exponential distributions (Neal & Hinton, 1998; Singer & Warmuth, 1999) and for online (aka block-wise) learning of HMMs (Baldi & Chauvin, 1994; Singer & Warmuth, 1997; Cappé et al., 1998; Mizuno et al., 2000). Also, inline (aka symbol-based) methods have been proposed for learning HMMs (Krishnamurthy & Moore, 1993; Collings et al., 1994; LeGland & Mevel, 1997; Garg & Warmuth, 2003; Florez-Larrahondo et al., 2005; Mongillo & Deneve, 2008; Cappé, 2011; Kontorovich et al., 2013). Our method falls into the category of block-wise updates for HMMs. To the best of our knowledge, no online updates were known for Kalman filters; all earlier training methods were based on either the batch EM algorithm via Kalman smoothing or inline updates via Kalman filtering (Ghahramani & Hinton, 1996).

2 BATCH EM MOTIVATION

Given a set of iid samples $\mathcal{V} = \{v_n\}_{n=1}^N$ from an underlying data distribution $p_d(v)$, the EM algorithm seeks to minimize the *negative log-likelihood* (as the loss)

$$\mathcal{L}(\tilde{\Theta} | \mathcal{V}) = -1/N \sum_n \log \underbrace{p(v_n | \tilde{\Theta})}_{\int_h p(h, v_n | \tilde{\Theta})},$$

wrt the parameters $\tilde{\Theta}$. Here, v_n is the n -th observation of some visible variable and h denotes the hidden variable. The above minimization problem, which involves logs of integrals (or sums in the discrete case) is typically non-convex and difficult to solve in practice. Adding a divergence to a loss can simplify the minimization. Batch EM employs the following *upper-bound* of the loss:

$$\begin{aligned} U_{\Theta}(\tilde{\Theta} | \mathcal{V}) &:= -1/N \sum_n \log \int_h p(h, v_n | \tilde{\Theta}) \\ &\quad + 1/N \sum_n \int_h p(h | v_n, \Theta) \log \frac{p(h | v_n, \Theta)}{p(h | v_n, \tilde{\Theta})} \\ &= -1/N \sum_n \mathbb{E}_{p(h | v_n, \Theta)} \left[\log p(h, v_n | \tilde{\Theta}) \right] \\ &\quad - 1/N \underbrace{\sum_n \mathbb{H}_{\Theta_n}(H | v_n)}_{\text{const.}}, \end{aligned} \tag{1}$$

where Θ denotes the current set of parameters and $\mathbb{H}_{\Theta_n}(H|v_n) := -\int_h p(h|v_n, \Theta) \log p(h|v_n, \Theta)$ is the conditional *differential entropy* of $H|v_n$. Batch EM algorithm proceeds by forming the upper-bound by calculating the posteriors $p(h|v_n, \Theta)$ based on the current estimate Θ (the E-step) and then minimizing (1) wrt $\tilde{\Theta}$ and updating Θ to the solution (the M-step). Minimizing the upper-bound is easier than minimizing the negative log-likelihood directly because logs of integrals are now replaced by logs of joints. Since the upper-bound is tight, i.e. $U_{\Theta}(\Theta|\mathcal{V}) = \mathcal{L}(\Theta|\mathcal{V})$, a decrement in the value of upper-bound amounts to a reduction in the negative log-likelihood, that is, $U_{\Theta}(\Theta^{\text{new}}|\mathcal{V}) < U_{\Theta}(\Theta|\mathcal{V}) \Rightarrow \mathcal{L}(\Theta^{\text{new}}|\mathcal{V}) < \mathcal{L}(\Theta|\mathcal{V})$.

We now rewrite the upper-bound as a sum of relative entropy divergences to a set of singleton models. Given the current model estimate Θ , let

$$p(h, v|\Theta_n) := \delta_{v_n}(v) p(h|v, \Theta), \quad (2)$$

where $\delta_{v_n}(v)$ is the *Dirac measure* centered at v_n . Note that Θ_n is an estimate of the model that minimizes the upper-bound (1) at Θ using a single observation v_n . Now the *relative entropy divergence* between the models Θ_n and $\tilde{\Theta}$ becomes

$$\begin{aligned} D_{\text{RE}}(\Theta_n, \tilde{\Theta}) &:= \int_{h,v} p(v, h|\Theta_n) \log \frac{p(v, h|\Theta_n)}{p(v, h|\tilde{\Theta})} \\ &= -\mathbb{H}_{\Theta_n}(H, V) - \int_{h,v} p(h, v|\Theta_n) \log p(h, v|\tilde{\Theta}) \\ &= \underbrace{-\mathbb{H}_{\Theta_n}(H, V)}_{\text{const.}} - \mathbb{E}_{p(h|v_n, \Theta)} \left[\log p(h, v_n|\tilde{\Theta}) \right]. \end{aligned} \quad (3)$$

where $\mathbb{H}_{\Theta_n}(H, V) := -\int_{h,v} p(h, v|\Theta_n) \log p(h, v|\Theta_n)$ is the joint differential entropy of H and V . Note that $\mathbb{H}_{\Theta_n}(H, V) = \mathbb{H}_{\Theta_n}(V) + \mathbb{H}_{\Theta_n}(H|V)$. Thus, the n -th term in the EM upper-bound (1) has the exact same form as (3) minus the $\mathbb{H}_{\Theta_n}(V)$ term. Although this term is unbounded for the singleton distribution in (2), it acts a constant wrt $\tilde{\Theta}$ and can be omitted from the upper-bound. Thus, one step of the EM algorithm can be seen as minimizing sum of the relative entropy divergences (3) to the singleton distributions (2), that is,

$$U_{\Theta}(\tilde{\Theta}|\mathcal{V}) = 1/N \sum_n D_{\text{RE}}(\Theta_n, \tilde{\Theta}) + 1/N \underbrace{\sum_n \mathbb{H}_{\Theta_n}(V)}_{\text{const.}}.$$

3 ONLINE EM MOTIVATION

EM is naturally a batch algorithm and requires the full set of observations to carry out each iteration. On the other hand, online algorithms only receive one example (or a

small batch \mathcal{V}^t) at every iteration t . The updates minimize the loss of the given batch (in our case an upper-bound of the loss) plus an *inertia* term (a second divergence) that ensures that the updates remain close to the current estimates Θ^t . From the model view of the EM upper-bound in (3), it is natural to choose the inertia term in the same form, i.e. a relative entropy divergence to the current model Θ^t . Thus, the online update minimizes

$$\Theta^{t+1} = \arg \min_{\tilde{\Theta}} \left\{ \underbrace{U_{\Theta^t}(\tilde{\Theta}|\mathcal{V}^t)}_{\text{loss}} + 1/\eta^t \underbrace{D_{\text{RE}}(\Theta^t, \tilde{\Theta})}_{\text{inertia}} \right\}, \quad (4)$$

where Θ^t and \mathcal{V}^t are the parameters and the given batch of observations at round t , respectively. Based on our discussion in the previous section, update (4) corresponds to combining $|\mathcal{V}^t| + 1$ models, therefore is guaranteed to have the same form as a batch EM update. The parameter η^t can be seen as a *learning rate* which controls the extent that the parameters are affected by the new observations. Note that $\eta^t \rightarrow \infty$ recovers the vanilla EM algorithm on the batch \mathcal{V}^t while $\eta^t \rightarrow 0$ keeps the parameters unchanged. Moreover, following the tightness of the upper-bound, the objective of (4) is equal to $\mathcal{L}(\Theta^t|\mathcal{V}^t)$ at $\tilde{\Theta} = \Theta^t$. Thus, every step of the online EM algorithm decreases the negative log-likelihood of the model over the batch \mathcal{V}^t .

A few remarks are in order. The objective (4) is equal to the objective function of the online EM algorithm of Cappé & Moulines (2009) up to additive constant terms wrt $\tilde{\Theta}$.

Theorem 1. *The objective function in (4) is equal to the objective function of the online EM algorithm of Cappé & Moulines (2009) up to additive constant terms wrt $\tilde{\Theta}$.*

The proof is given in Appendix A.

The formulation in (Cappé & Moulines, 2009) is specifically applied to models where the complete data likelihood belongs to an exponential family and the updates are shown to reduce to stochastic approximation of the sufficient statistics. While their approach is applicable to simpler models such as mixture of Poisson, identifying the sufficient statistics immediately becomes tedious for more complex models such as HMMs and Kalman filters. As a result, the corresponding updates for these models had not been discovered. Moreover, the decrement of the negative log-likelihood over the current observation \mathcal{V}^t is not evident in their formulation.

Corollary 2. *Under mild assumptions on the parameter space of the exponential family model and using a decaying learning rate $0 < \eta^t < 1$, $\sum_{t=1}^{\infty} \eta^t = \infty$ and $\sum_{t=1}^{\infty} \eta^t < \infty$, the update (4) almost surely converges to a stationary point that maximizes the expected log-likelihood wrt the data distribution.*

3.1 NATURAL GRADIENT APPROXIMATION

For small $d\tilde{\Theta} := \tilde{\Theta} - \Theta^t$, the inertia term can be approximated as

$$D_{\text{RE}}(\Theta^t, \tilde{\Theta}) \approx 1/2 d\tilde{\Theta}^\top I_{\text{F}}(\Theta^t) d\tilde{\Theta}, \quad (5)$$

where $I_{\text{F}}(\Theta^t) = -\mathbb{E}_{p(h,v|\Theta^t)}[\nabla^2 \log p(h,v|\Theta^t)]$ is the *Fisher information matrix*. Using (5), the update (4) can be approximated as

$$\Theta^{t+1} \approx \Theta^t - I_{\text{F}}^{-1}(\Theta^t) \nabla U_{\Theta^t}(\Theta^t | \mathcal{V}^t),$$

which is called the natural gradient update (Amari & Nagaoka, 2007) and resembles the gradient-based updates for EM proposed in (McLachlan & Krishnan, 2008) where the *observed Fisher* $I_{\text{O}}(\Theta^t) = -\mathbb{E}_{p(v|\Theta^t)}[\nabla^2 \log p(v|\Theta^t)]$ is used in place of $I_{\text{F}}(\Theta^t)$. Using the equality,

$$\begin{aligned} \nabla U_{\Theta^t}(\Theta^t | \mathcal{V}^t) &= -1/N \sum_n \mathbb{E}_{p(h|v_n, \Theta^t)} \left[\frac{\nabla p(h, v_n | \Theta^t)}{p(h, v_n | \Theta^t)} \right] \\ &= -1/N \sum_n \int_h \frac{\nabla p(h, v_n | \Theta^t)}{p(v_n | \Theta^t)} = \nabla \mathcal{L}(\mathcal{V}^t | \Theta^t), \end{aligned}$$

we have,

$$\Theta^{t+1} \approx \Theta^t - I_{\text{F}}^{-1}(\Theta^t) \nabla \mathcal{L}(\mathcal{V}^t | \Theta^t).$$

This connection was also observed in (Sato, 2000) and the extension of this work in (Cappé & Moulines, 2009).

3.2 FINITE-SAMPLE APPROXIMATION

In cases where $D_{\text{RE}}(\Theta^t, \tilde{\Theta})$ does not yield a closed-form, an approximate inertia term can be obtained via a finite number of samples as

$$\begin{aligned} D_{\text{RE}}(\Theta^t, \tilde{\Theta}) &= \mathbb{E}_{p(v|\Theta^t)} \left[\mathbb{E}_{p(h|v, \Theta^t)} \left[\log \frac{p(h, v | \Theta^t)}{p(h, v | \tilde{\Theta})} \right] \right] \quad (6) \\ &\approx \underbrace{-1/N' \sum_{n'} \mathbb{E}_{p(h|v_{n'}, \tilde{\Theta})} \left[\log p(v_{n'}, h | \tilde{\Theta}) \right]}_{U_{\Theta^t}(\tilde{\Theta} | \mathcal{V}') + \text{const.}} + \text{const.} \end{aligned}$$

where the samples $\mathcal{V}' = \{v_{n'}\}_{n'=1}^{N'}$ are drawn from the distribution $p(v|\Theta^t)$, not the data distribution $p_d(v)$. Thus the samples \mathcal{V}' may be seen as *pseudo-observations* that encourage the solution to remain close to the current model parameters Θ^t . We will use this sampled form of the inertia term to derive updates for the compound Dirichlet model. The sampled form is also similar to the update given in (Neal & Hinton, 1998). However, in their formulation, \mathcal{V}' is replaced with $\mathcal{V} - \{v_i\}$ where v_i is a random sample from \mathcal{V} and the upper-bound is formed at $\Theta^{(t-1)}$ instead of Θ^t .

4 UPDATES WITH CLOSED-FORM

The objective is easier to minimize when it reduces to a linear combination of negative log-likelihoods of exponential family distributions which includes mixtures of exponential families, HMMs and Kalman Filters. Note that the latter two cases are already hard to handle with the methodology of Cappé & Moulines (2009). We start with a brief background material. The *exponential family* (Wainwright et al., 2008) with vector of *sufficient statistics* $\phi(x)$ and *natural parameter* θ is defined as $p_G(x|\theta) = \exp(\theta \cdot \phi(x) - G(\theta))$, where $G(\theta) = \log \int_x \exp(\theta \cdot \phi(x))$ is called the *log partition function*, which ensures that $p_G(x|\theta)$ integrates to one. The expectation parameter $\mu = g(\theta) = \int_x \phi(x) p_G(x|\theta)$ is the *dual* (Hiriart-Urruty & Lemarchal, 2001) of the natural parameter θ where $g(\theta) := \nabla G(\theta)$. The duality implies $\theta = g^*(\mu) = g^{-1}(\mu)$ where $g^*(\mu) := \nabla G^*(\mu)$ and $G^*(\mu) = \sup_{\theta'} \{\mu \cdot \theta' - G(\theta')\}$ is the *convex conjugate* of G . It is easy to show that $G(\theta)$ is a convex function. In fact, the relative entropy divergence between two exponential distributions (of the same form) with parameters θ and $\tilde{\theta}$ yields

$$\int_x p_G(x|\theta) \log \frac{p_G(x|\theta)}{p_G(x|\tilde{\theta})} = D_G(\tilde{\theta}, \theta) = D_{G^*}(\mu, \tilde{\mu}),$$

where $D_G(\tilde{\theta}, \theta) = G(\tilde{\theta}) - G(\theta) - g(\theta) \cdot (\tilde{\theta} - \theta)$, is the *Bregman divergence* (Bregman, 1967) induced by the convex function $G(\cdot)$ and $\frac{\partial}{\partial \theta} D_G(\tilde{\theta}, \theta) = g(\tilde{\theta}) - g(\theta)$. The following lemma will be useful for deriving the updates. For more details, see Appendix B.

Lemma 3. For $\{\alpha_m\}_{m=1}^M$ s.t. $\alpha_m \in \mathbb{R}_+$ and $\sum_m \alpha_m > 0$,

$$\begin{aligned} \theta_{\text{opt}} &= \arg \min_{\tilde{\theta}} \sum_m \alpha_m (G(\tilde{\theta}) - \tilde{\theta} \cdot \mu_m) \\ &= g^{-1} \left(\frac{\sum_m \alpha_m \mu_m}{\sum_m \alpha_m} \right), \end{aligned}$$

$$\text{i.e. } \mu_{\text{opt}} = \frac{\sum_m \alpha_m \mu_m}{\sum_m \alpha_m}.$$

4.1 MIXTURE OF EXPONENTIAL FAMILY

In a k -mixture of exponential family, model $h \in [k]$ is chosen according to the probability $\omega_h := p(h|\Theta)$ and the observation is drawn from the corresponding distribution $P(v|h, \Theta) = p_G(v|\theta_h) = \exp(\theta_h \cdot \phi(v) - G(\theta_h))$, which belongs to an exponential family. Thus, the model parameters are $\Theta = \{\omega_h, \mu_h(\theta_h)\}_h$. The joint distribution becomes

$$p(v, h | \Theta) = \omega_h \exp(\theta_h \cdot \phi(v) - G(\theta_h)),$$

while the marginal is simply a sum over all models,

$$p(v|\Theta) = \sum_h \omega_h \exp(\theta_h \cdot \phi(v) - G(\theta_h)).$$

The EM upper-bound can be formed using the posterior distributions of each observation v_n , that is,

$$\begin{aligned} U_{\Theta}(\tilde{\Theta} | \mathcal{V}) &\cong -1/N \sum_n \sum_h \gamma_{n,h} \left(\log \tilde{\omega}_h + (G(\tilde{\theta}_h) - \tilde{\theta}_h \cdot \phi(v_n)) \right), \end{aligned}$$

where \cong means equal up to constants. The posteriors $\gamma_{n,h}$ are calculated as

$$\gamma_{n,h} = \frac{\omega_h \exp(\theta_h \cdot \phi(v_n) - G(\theta_h))}{\sum_{h'} \omega_{h'} \exp(\theta_{h'} \cdot \phi(v_n) - G(\theta_{h'}))}.$$

The inertia term for the online EM algorithm becomes

$$\begin{aligned} D_{\text{RE}}(\Theta, \tilde{\Theta}) &= \sum_h \int_v \omega_h p_G(v | \theta_h) \log \frac{\omega_h p_G(v | \theta_h)}{\tilde{\omega}_h p_G(v | \tilde{\theta}_h)} \\ &= \sum_h \omega_h \log \frac{\omega_h}{\tilde{\omega}_h} + \underbrace{\sum_h \omega_h \int_v p_G(v | \theta_h) \log \frac{p_G(v | \theta_h)}{p_G(v | \tilde{\theta}_h)}}_{D_G(\tilde{\theta}_h, \theta_h)}. \end{aligned}$$

Combining the inertia term with the upper-bound and applying Lemma 3, we have

$$\omega_h^{\text{new}} = \frac{1/\eta \omega_h + 1/N \sum_n \gamma_{n,h}}{1/\eta + 1}, \quad (7)$$

$$\mu_h^{\text{new}} = \frac{1/\eta \omega_h \mu_h + 1/N \sum_n \gamma_{n,h} \phi(v_n)}{1/\eta \omega_h + 1/N \sum_n \gamma_{n,h}}. \quad (8)$$

4.2 HIDDEN MARKOV MODELS

A Hidden Markov Model (HMM) (Rabiner, 1989) consists of an underlying finite-state (hidden) Markov chain with probability of transitioning from state h to h' equal to $a_{h,h'} := P(h | h', \Theta)$ and an initial state probability equal to $\pi_{h_1} := P(h_1 | \Theta)$. At every round, the model makes a transition to a new state according to the state transition probabilities, and given the new state h , generates an observation according to the state emission probability $P(v | h, \Theta)$. We make the assumption that the state emission probabilities are members of an exponential family, that is, $P(v | h, \Theta) = P_G(v | \theta_h)$. Thus, the model parameters are $\Theta = \{\pi_h, \{a_{h,h'}\}_{h,h'}, \mu_h(\theta_h)\}_h$. The joint distribution of the model can be written as

$$P(v, h | \Theta) = \prod_{t=1}^T a_{h_{t-1}, h_t} P_G(v_t | \theta_{h_t}),$$

in which, we define $a_{h_0, h_1} := \pi_{h_1}$. The marginal can be obtained by summing over all the possible hidden states:

$$P(v | \Theta) = \sum_{h_1, \dots, h_T} \prod_{t=1}^T a_{h_{t-1}, h_t} P_G(v_t | \theta_{h_t}).$$

Ignoring constant terms, the EM upper-bound can be written as

$$\begin{aligned} U_{\Theta}(\tilde{\Theta} | \mathcal{V}) &\cong 1/N \sum_n \left(\sum_h \gamma_h^{n,1} \log \frac{\gamma_{h_1}^n}{\tilde{\pi}_{h_1}} \right. \\ &\quad + \sum_{t=1}^{T-1} \sum_{h,h'} \gamma_{h,h'}^{n,t} \log \frac{\gamma_{h,h'}^{n,t}}{\tilde{a}_{hh'}} \\ &\quad \left. + \sum_{t=1}^T \sum_h \gamma_h^{n,t} (G(\tilde{\theta}_{h_t}) - \phi(v_{nt}) \cdot \tilde{\theta}_{h_t}) \right). \end{aligned}$$

The state posteriors are found using the Baum-Welch algorithm by performing a forward-backward pass (Rabiner, 1989). We define

$$\begin{aligned} \gamma_h^{n,t} &:= P(h_{n,t} = h | v_n, \Theta), \\ \gamma_{h,h'}^{n,t} &:= P(h_{n,t+1} = h', h_{n,t} = h | v_n, \Theta), \end{aligned}$$

where $h_{n,t}$ is the state at round t for the sample n . The inertia term can be written as

$$\begin{aligned} D_{\text{RE}}(\Theta, \tilde{\Theta}) &= \sum_h \pi_h \log \frac{\pi_h}{\tilde{\pi}_h} + \sum_h u_h \sum_{h'} a_{h,h'} \log \frac{a_{h,h'}}{\tilde{a}_{h,h'}} \\ &\quad + \sum_h u_h D_G(\tilde{\theta}_h, \theta_h), \end{aligned}$$

where $u_h = \sum_{t=1}^{\infty} \delta_h^t$, with $\delta_h^1 = \pi_h$, and $\delta_{h'}^{t+1} = \sum_h \delta_h^t a_{h,h'}$. That is, u_h is the *expected usage* of state h . Note that the usage u_h is not finite in general and should be instead approximated by a finite length sequence. However, for the class of *absorbing* HMMs, we can calculate the usages in the exact form. More specifically, the transition matrix $A = [a_{h,h'}]$ of an absorbing HMM with r absorbing states has the following form

$$A = \begin{bmatrix} Q & R \\ 0 & I_r \end{bmatrix},$$

where the Q entails the transition probabilities from a transient state to another while R describes the transition probabilities of from transient states to absorbing states. I_r is an identity matrix which describes the transitions from each absorbing state back to itself. The expected usages of the transient states can be calculated as

$$u^{\top} = \pi^{\top} + \pi^{\top} Q + \pi^{\top} Q^2 + \dots = \pi^{\top} (I - Q)^{-1}.$$

Additionally, note that for an absorbing state h , we always have $\sum_{h'} a_{h,h'} \log \frac{a_{h,h'}}{\tilde{a}_{h,h'}} = 0$. Therefore, the corresponding terms can be omitted from the inertia term.

Combining the EM upper-bound and the inertia term and

applying Lemma 3 gives the following updates

$$\begin{aligned}\pi_h^{\text{new}} &= \frac{1/\eta \pi_h + 1/N \sum_n \gamma_h^{n,1}}{1/\eta + 1}, \\ a_{h,h'}^{\text{new}} &= \frac{1/\eta u_h a_{h,h'} + 1/N \sum_n \sum_t \gamma_h^{n,t}}{1/\eta u_h + 1/N \sum_n \sum_t \gamma_h^{n,t}}, \\ \mu_h^{\text{new}} &= \frac{1/\eta u_h \mu_h + 1/N \sum_n \sum_t \gamma_h^{n,t} \phi(v_{n,t})}{1/\eta u_h + 1/N \sum_n \sum_t \gamma_h^{n,t}}.\end{aligned}$$

4.3 KALMAN FILTERS

Kalman filters (Welch & Bishop, 1995) can be described using the following two update equations

$$\begin{aligned}h_{t+1} &= A h_t + \rho_t, \\ v_t &= C h_t + \varepsilon_t,\end{aligned}$$

where h_t is the underlying (hidden) state at t and v_t is the corresponding output. Both state and observation noise, ρ_t and ε_t , are zero-mean Gaussian random variables with covariance matrices equal to Q and R , respectively. The initial state h_1 is generally assumed to be drawn from a Gaussian distribution with mean π_1 and covariance V . Thus the model parameters are $\Theta = \{\pi_1, V, A, C, Q, R\}$. In Kalman filters, only the output is observed, and thus the state as well as the noise variables are hidden.

The joint distribution of a Kalman filter can be written as

$$p(h, v | \Theta) = \prod_{t=1}^T p(h_t | h_{t-1}, \Theta) p(v_t | h_t, \Theta),$$

$$\begin{aligned}\text{where } p(h_1 | h_0, \Theta) &:= p(h_1 | \Theta) \\ p(v_t | h_t, \Theta) &\sim \mathcal{N}(C h_t, R), \\ p(h_{t+1} | h_t, \Theta) &\sim \mathcal{N}(A h_t, Q), \\ p(h_1 | \Theta) &\sim \mathcal{N}(\pi_1, V).\end{aligned}$$

Here, $\mathcal{N}(\xi, \Sigma)$ denotes a Gaussian probability density with mean and covariance equal to ξ and Σ , respectively. The marginal can be obtained by integrating over all the state variables, that is,

$$P(v | \Theta) = \int_{h_1, \dots, h_T} P(h, v | \Theta).$$

Forming the EM upper-bound requires calculating the posteriors. Note that because all the random variables are Gaussian, it suffices to keep track of the means and covariances. The posteriors depend on the following three expectations

$$\hat{h}_t := \mathbb{E}[h_t | v], \quad P_t := \mathbb{E}[h_t h_t^\top | v], \quad P_{t,t-1} := \mathbb{E}[h_t h_{t-1}^\top | v],$$

which can be calculated recursively using the *Kalman filtering* and *Kalman smoothing* equations (Ghahramani

& Hinton, 1996). Thus, the EM upper-bound can be written as

$$\begin{aligned}2 \times \text{U}_\Theta(\tilde{\Theta} | \mathcal{V}) &\cong 1/N \sum_n \left(\text{tr}(\tilde{V}^{-1} \hat{V}_n) + \log |\tilde{V}| \right. \\ &\quad \left. + \sum_{t=2}^T \text{tr}(\tilde{Q}^{-1} \hat{Q}_{n,t}) + (T-1) \log |\tilde{Q}| \right. \\ &\quad \left. + \sum_{t=1}^T \text{tr}(\tilde{R}^{-1} \hat{R}_{n,t}) + T \log |\tilde{R}| \right),\end{aligned}$$

where⁴

$$\begin{aligned}\hat{V}_n &= P_1^n - \tilde{\pi}_1 \hat{h}_{n,1}^\top - \hat{h}_{n,1} \tilde{\pi}_1^\top + \tilde{\pi}_1 \tilde{\pi}_1^\top, \\ \hat{Q}_{n,t} &= P_t^n - \tilde{A} P_{t-1,t}^n - P_{t,t-1}^n \tilde{A}^\top + \tilde{A} P_{t-1}^n \tilde{A}^\top, \\ \hat{R}_{n,t} &= v_{n,t} v_{n,t}^\top - \tilde{C} \hat{h}_{n,t} v_t^\top - v_{n,t} \hat{h}_{n,t}^\top \tilde{C}^\top + \tilde{C} P_t \tilde{C}^\top.\end{aligned}$$

Again, our inertia term for the online EM algorithm is relative entropy between the joints, assuming a fixed observation length equal to T , that is,

$$\begin{aligned}2 \times D_{\text{RE}}(\Theta, \tilde{\Theta}) &\cong \text{tr}(\tilde{V}^{-1}(\pi_1 - \tilde{\pi}_1)(\pi_1 - \tilde{\pi}_1)^\top) + D_{\text{ld}}(V, \tilde{V}) \\ &\quad + \text{tr}(\tilde{Q}^{-1}(A - \tilde{A}) \sum_{t=1}^{T-1} U_t (A - \tilde{A})^\top) + (T-1) D_{\text{ld}}(Q, \tilde{Q}) \\ &\quad + \text{tr}(\tilde{R}^{-1}(C - \tilde{C}) \sum_{t=1}^T U_t (C - \tilde{C})^\top) + T D_{\text{ld}}(R, \tilde{R}),\end{aligned}$$

where $U_1 = V + \pi_1 \pi_1^\top$ and $U_{t+1} = Q + A U_t A^\top$. Moreover $D_{\text{ld}}(X, Y) = \text{tr}(XY^{-1}) - \log |XY^{-1}| - d$, is the *log-determinant divergence* (Cichocki et al., 2009).

Combining the inertia term with the EM upper-bound and setting the derivatives wrt the parameters to zero yields

$$\begin{aligned}\pi_1^{\text{new}} &= \frac{1/\eta \pi_1 + 1/N \sum_n \hat{h}_{n,1}}{1/\eta + 1}, \\ V^{\text{new}} &= \frac{1/\eta (V + (\pi_1 - \tilde{\pi}_1)(\pi_1 - \tilde{\pi}_1)^\top) + \hat{V}}{1/\eta + 1}, \\ A^{\text{new}} &= \left(1/\eta \sum_{t=2}^{T-1} A U_t + \sum_{t=2}^T P_{t,t-1} \right) S_{T-1}^{-1}, \\ C^{\text{new}} &= \left(1/\eta \sum_{t=1}^T C U_t + 1/N \sum_n \sum_{t=1}^T v_{n,t} \hat{h}_{n,t}^\top \right) S_T^{-1}, \\ Q^{\text{new}} &= \frac{1/\eta (Q + \Delta A_U) + 1/(N(T-1)) \sum_n \sum_{t=2}^T \hat{Q}_{n,t}}{1/\eta + 1}, \\ R^{\text{new}} &= \frac{1/\eta (R + \Delta C_U) + 1/(NT) \sum_n \sum_{t=1}^T \hat{R}_{n,t}}{1/\eta + 1},\end{aligned}$$

⁴Note that $n \in [N]$ denotes the observation index.

where

$$S_T = \sum_{t=1}^T (1/\eta U_t + P_t),$$

$$\Delta A_U = 1/(T-1) (A - A^{\text{new}}) \sum_{t=2}^T U_t (A - A^{\text{new}})^\top,$$

$$\Delta C_U = 1/T (C - C^{\text{new}}) \sum_{t=1}^T U_t (C - C^{\text{new}})^\top.$$

Remark 4. For mixtures of exponential families, HMMs, and Kalman filters, the following holds: for a sufficiently small learning rate η , the negative log-likelihood wrt the underlying data distribution $p_d(v)$,

$$\mathbb{E}_{p_d} \left[\log p(v | \Theta) \right] = \int_v p_d(v) \log \int_h p(h, v | \Theta),$$

improves after each update.

This is a direct result of applying Proposition 3 in (Cappé & Moulines, 2009). See (Cappé & Moulines, 2009) for further details.

5 UPDATES WITHOUT CLOSED-FORM

The minimization problem in the M-step of the batch EM does not always have a closed-form solution. In those cases, it is likely that the divergence term between the models also does not have a closed-form either. An example of such a model is the compound Dirichlet distribution (Gupta et al., 2011). In this case, applying the online EM updates in form of (4) is infeasible. However, we can instead use the finite sample form of the inertia term. That is, in each iteration, we draw N' samples $\mathcal{V}' = \{v_{n'}\}_{n'=1}^{N'}$ from $P(v | \Theta^t)$ and form the corresponding EM upper-bound by treating the additional samples as pseudo-observations. The update can be achieved by (numerically) minimizing the combined upper-bounds,

$$\Theta^{t+1} \approx \arg \min_{\tilde{\Theta}} U_{\Theta^t}(\tilde{\Theta} | \mathcal{V}^t) + 1/\eta_t U_{\Theta^t}(\tilde{\Theta} | \mathcal{V}'),$$

where again $\eta_t > 0$ is a learning rate parameter. Note that this is fundamentally different than combining the samples $\mathcal{V}^t \cup \mathcal{V}'$ and forming a single upper-bound. In fact, combining the samples may require a larger number of model samples as we proceed with the online updates (which corresponds to a decaying learning rate) while our approach can be carried out by a fixed number of samples at every iteration. Note that the quality of the approximation of the inertia term depends on the size of the model samples N' . In fact, approximating the inertia

term leads to a higher variance: The variance decreases by increasing the number of model samples N' ; however, the computational cost also increases with larger N' . We will provide experimental results on the online parameter estimation with the compound Dirichlet distribution.

6 COMBINING MODELS

In many cases, multiple *local models* need to be combined to form a *global model*. For instance, due to the large amount of data, the model training is distributed over multiple machines where each machine only receives a subset of the dataset and performs updates on its local model. The local models are then combined into a single global model at the end of each iteration (synchronous) or the end of the training process (asynchronous). Our divergences between the hidden variable models provide a natural way of combining the local models in a distributed setting. More formally, let $\Theta^{(m)}$ denote the set of local parameters of model $m \in [M]$. We can define the combined model parameters $\Theta^{(\text{comb})}$ as

$$\Theta^{(\text{comb})} = \arg \min_{\tilde{\Theta}} \sum_{m \in [M]} \alpha_m D_{\text{RE}}(\Theta^{(m)}, \tilde{\Theta}), \quad (9)$$

where $\alpha_m \geq 0$ is the associated weight for combining model m (s.t. $\sum_m \alpha_m > 0$). The value of α_m can be tuned based on the amount of data seen by model m , likelihood of the model, etc. For the exponential family models, updates in (9) reduces to a convex combination of the complete data sufficient statistics of the models⁵. As an example, for hidden Markov models, Equation (9) yields

$$\begin{aligned} \pi_h^{(\text{comb})} &= \frac{\sum_m \alpha_m \pi_h^{(m)}}{\sum_m \alpha_m}, \\ a_{h,h'}^{(\text{comb})} &= \frac{\sum_m \alpha_m u_h^{(m)} a_{h,h'}^{(m)}}{\sum_m \alpha_m u_h^{(m)}}, \\ \mu_h^{(\text{comb})} &= \frac{\sum_m \alpha_m u_h^{(m)} \mu_h^{(m)}}{\sum_m \alpha_m u_h^{(m)}}. \end{aligned}$$

We experimentally show that combining the models via (9) provides improved results compared to the commonly used methods of combining the models via simple averaging (Sanderson & Curtin, 2017).

Again, for cases where the divergence between the models does not admit a closed-form, we can use the sampling form of the divergence to combine the models. That is, we draw N'_m samples \mathcal{V}'_m from $P(v | \Theta^{(m)})$ and form $U_{\Theta^{(m)}}(\tilde{\Theta} | \mathcal{V}'_m)$. The combined model can be obtained as

$$\Theta^{(\text{comb})} = \arg \min_{\tilde{\Theta}} \sum_m \alpha_m U_{\Theta^{(m)}}(\tilde{\Theta} | \mathcal{V}'_m).$$

⁵And not the sufficient statistics of the components.

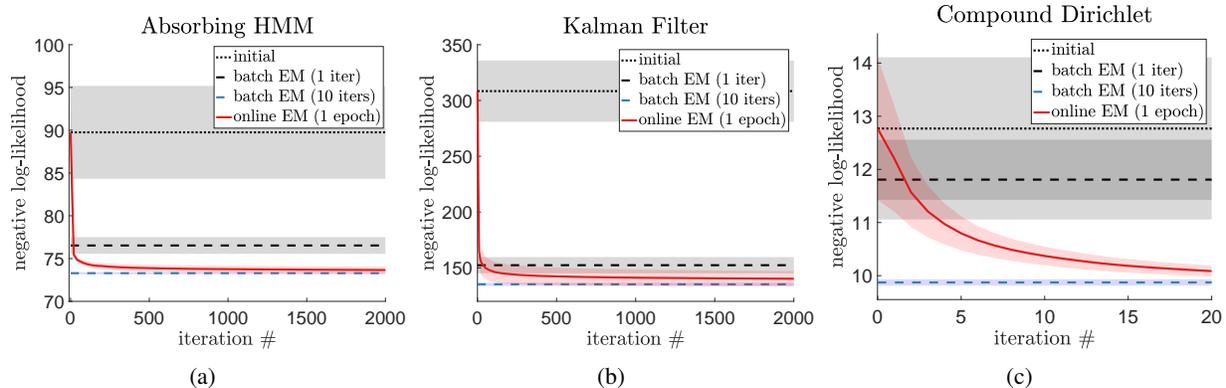


Figure 1: Online EM results: a) absorbing HMM, b) Kalman filter, and c) compound Dirichlet distribution.

7 EXPERIMENTS

In this section, we first conduct experiments on online learning of absorbing HMMs and Kalman filters. Next, we apply the approximate form of the inertia term for estimating a compound Dirichlet distribution in which the updates (as well as the inertia term) do not have a closed form solution. Finally, we consider learning of Gaussian mixture models in a distributed setting. In all experiments, we use a decaying learning rate of the form $\eta_t = \eta_0/t^\beta$ where t is the iteration number and $\eta_0 > 0$ and $0.5 < \beta < 1$ are specified for each case. We repeat each experiment over 20 random initializations.

7.1 ABSORBING HMM

We validate the derived updates for HMMs by conducting an experiment on estimating the parameters of an absorbing HMM with 3 transient and a single absorbing state (4 hidden states in total) and Gaussian emission probabilities of dimension 4. We consider 2000 samples from the model and apply batch EM updates as well as online updates with $(\eta_0, \beta) = (0.5, 0.9)$. The results are shown in Figure 1(a). The online algorithm processes one observation per iteration. Regarding processing time, one pass of the online update over the entire data set (called one epoch) is comparable to a single batch update. The online update rapidly outperforms the single batch EM update after around 30 iterations, and at the end of the first epoch converges to a value close to the loss of 10 batch EM iterations. Also, the online updates are stable to using lower or higher learning rates: The final loss values obtained for $\eta_0 = 0.1$ and $\eta_0 = 1.0$ are 77.17 and 72.27, respectively (not shown in the figure). For comparison, we also apply a gradient based update based on (Cappé, 2011). The gradient based updates are extremely unstable and best final result obtained is 83.84 (also not shown in the figure).

7.2 KALMAN FILTER

To validate the correctness of the updates for Kalman filters, we consider online estimation of the parameters of a Kalman filter with hidden state vector of dimension 5 and observation vector of dimension 10. We assume that the noise covariances Q and R are known and consider estimating the remaining parameters, i.e. $\{\pi_1, V, A, C\}$. We apply the batch EM updates as well as the online updates with parameters $(\eta_0, \beta) = (1.0, 0.9)$. The results are shown in Figure 1(b). Again, the online updates outperform the solution of one batch EM after around 40 iterations and converge to a solution with a loss close to 10 batch EM updates. Moreover, the updates are stable wrt the initial learning rate η_0 . The final value of the negative log-likelihood of the model obtained using $\eta_0 = 0.1$ and $\eta_0 = 10.0$ are 89.36 and 75.90, respectively.

7.3 COMPOUND DIRICHLET DISTRIBUTION

We consider online estimation of a compound Dirichlet distribution (Gupta et al., 2011). In this case, the EM updates for the model do not have a closed-form solution and therefore, numerical techniques such as Newton’s method (Nocedal & Wright, 2006) are used for performing the updates. The details are given in Appendix C. As a result, the relative entropy inertia term also does not admit a closed-form and thus, we use the sampling approximation of the inertia term. (See end of Section 6.) We consider 2000 samples from a 10 dimensional model and perform batch EM updates as well as online updates with mini-batch size equal to 100. In order to form the inertia term, we use 2000 samples from the model and use parameters $(\eta_0, \beta) = (1.0, 0.9)$ for the learning rate. We use Newton’s method for optimization. The result is shown in Figure 1(c). As can be seen, the online EM algorithm effectively learns the model parameters. The updates are stable for a lower or higher initial learning rate. The final negative log-likelihood values for $\eta_0 = 0.5$ and $\eta_0 = 2.0$ are 10.09 and 10.08, respectively (results not shown in the figure).

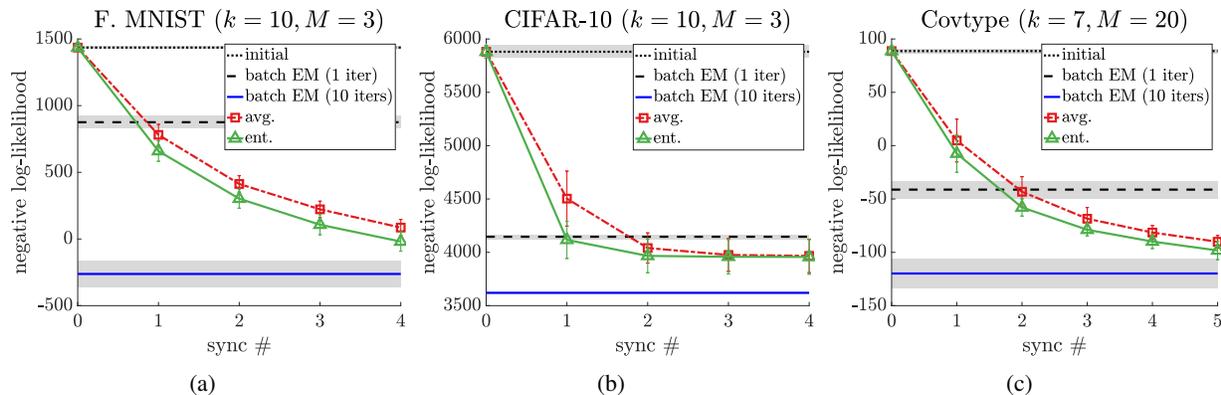


Figure 2: Combining models: results using simple averaging of the means and covariance matrices (avg.) compared to our proposed model averaging (ent.) on (a) Fashion MNIST, (b) CIFAR-10, and (c) Covtype datasets. The number of mixtures k as well as the number of machines M is shown on top of each plot. The sync step happens after every 5000 iterations of online EM on each machine.

7.4 DISTRIBUTED TRAINING OF GAUSSIAN MIXTURES

We conduct experiments on combining the parameters of Gaussian mixture models in a distributed setting. For this set of experiments, we consider the Fashion MNIST⁶ (dim=784), CIFAR-10⁷ (dim = 3072), and Covtype⁸ (dim = 54) datasets. The number of machines for each dataset is set to $M = 3, 3,$ and $20,$ respectively. To achieve an equal number of splits across machines, we consider a subset of 60K, 60K, and 500K points from each dataset, respectively. We set the number of mixtures k equal to the number of classes, which amounts to $k = 10$ for Fashion MNIST, $k = 10$ for CIFAR-10, and $k = 7$ for Covtype dataset. We use $(\eta_0, \beta) = (0.05, 0.5)$ for all datasets.

At each trial, all the machines are initialized with the same set of initial parameters. We consider synchronous updates where the parameters of all machines are combined into a single set of parameters at the end of each step and propagated back to each individual machine for the next step. Each machine receives a different set of 5000 observations at each step and the process is repeated until one pass over the whole dataset is achieved.

We compare two parameter combining strategies: 1) simple averaging where mixture weights as well as the expected values of the conditional sufficient statistics (i.e. means and covariances of each mixture component) are averaged over all machines (Sanderson & Curtin, 2017), and 2) our entropic combining of parameters as in (9) where we average the complete data sufficient statistics. The results are shown in Figure 2. As can be seen, the

⁶<https://github.com/zalandoresearch/fashion-mnist>

⁷<https://www.cs.toronto.edu/~kriz/cifar.html>

⁸<https://archive.ics.uci.edu/ml/datasets/covertype>

divergence based combining of the model provides a consistently better performance. Specifically, it shows faster convergence to a better solution. Additionally, on the Fashion MNIST and Covtype datasets, the final combined model has a lower negative log-likelihood using our divergence based combining.

8 CONCLUSION AND FUTURE WORK

We provided an alternative view of the online EM algorithm of Cappé & Moulines (2009) based on divergences between the models. Our new formulation casts new insights on the algorithm and facilitates finding the updates for more complex models without the need for identifying the sufficient statistics. The divergences between models that we use as inertia terms are interesting in their own right and are the most important outcome of this research. These divergences can be approximated in cases where the EM updates do not have a closed-form. Also, the divergences between the models lead to a new technique for combining models which is useful in distributed settings.

There are a number of intriguing open problems coming out of the current work. All our divergences are based on joint relative entropies where the new model is always in the second argument. In online learning, the new model is typically in the first argument (see e.g. (Kivinen & Warmuth, 1997)). Also in the context of reinforcement learning (Neu et al., 2017), the alternate joint entropies for HMMs (with the new parameters as the first argument) have been used effectively. The alternate relative entropies appear to be more stable. Therefore, the question is whether there is a use of the alternate for producing useful updates for minimizing the negative log-likelihood of hidden variable models.

Acknowledgment Ehsan Amid was partially supported by the NSF grant IIS 1546459 and a gift grant from the Intel Corporation.

References

- Amari, S.-i. and Nagaoka, H. *Methods of Information Geometry*, volume 191. American Mathematical Soc., 2007.
- Baldi, P. and Chauvin, Y. Smooth on-line learning algorithms for hidden Markov models. *Neural Computation*, 6(2):307–318, 1994.
- Blei, D. M., Kucukelbir, A., and McAuliffe, J. D. Variational inference: A review for statisticians. *Journal of the American statistical Association*, 112(518):859–877, 2017.
- Bregman, L. M. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR computational mathematics and mathematical physics*, 7(3):200–217, 1967.
- Cappé, O. Online EM algorithm for hidden Markov models. *Journal of Computational and Graphical Statistics*, 20(3): 728–749, 2011.
- Cappé, O. and Moulines, E. On-line expectation–maximization algorithm for latent data models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(3): 593–613, 2009.
- Cappé, O., Buchoux, V., and Moulines, E. Quasi-Newton method for maximum likelihood estimation of hidden Markov models. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 4, pp. IV–2265, 1998.
- Cichocki, A., Zdunek, R., Phan, A.-H., and Amari, S.-i. *Non-negative Matrix and Tensor Factorizations: Applications to Exploratory Multi-Way Data Analysis and Blind Source Separation*. Wiley, first edition, 2009.
- Collings, I., Krishnamurthy, V., and Moore, J. On-line identification of hidden Markov models via recursive prediction error techniques. *IEEE Transactions on Signal Processing*, 42:3535–3539, 1994.
- Do, C. B. and Batzoglou, S. What is the expectation maximization algorithm? *Nature biotechnology*, 26(8):897, 2008.
- Florez-Larrahondo, G., Bridges, S., and Hansen, E. A. Incremental estimation of discrete hidden Markov models based on a new backward procedure. In *Proceedings of the 20th national conference on Artificial intelligence-Volume 2*, pp. 758–763. AAAI Press, 2005.
- Garg, A. and Warmuth, M. K. Inline updates for HMMs. In *INTERSPEECH*, 2003.
- Ghahramani, Z. and Hinton, G. E. Parameter estimation for linear dynamical systems. Technical report, CRG-TR-96-2, 1996.
- Gupta, M. R., Chen, Y., et al. Theory and use of the EM algorithm. *Foundations and Trends® in Signal Processing*, 4(3):223–296, 2011.
- Hiriart-Urruty, J.-B. and Lemarchal, C. *Fundamentals of Convex Analysis*. Springer-Verlag Berlin Heidelberg, first edition, 2001.
- Kivinen, J. and Warmuth, M. K. Exponentiated gradient versus gradient descent for linear predictors. *Inf. Comput.*, 132(1): 1–63, 1997.
- Kontorovich, A., Nadler, B., and Weiss, R. On learning parametric-output HMMs. In *International Conference on Machine Learning*, pp. 702–710, 2013.
- Krishnamurthy, V. and Moore, J. B. On-line estimation of hidden Markov model parameters based on the Kullback-Leibler information measure. *IEEE Transactions on signal processing*, 41(8):2557–2573, 1993.
- LeGland, F. and Mevel, L. Recursive estimation in hidden Markov models. In *Decision and Control, 1997., Proceedings of the 36th IEEE Conference on*, volume 4, pp. 3468–3473. IEEE, 1997.
- McLachlan, G. J. and Krishnan, T. *The EM Algorithm and Extensions (Wiley Series in Probability and Statistics)*. Wiley-Interscience, 2 edition, 2008.
- Mizuno, J., Watanabe, T., Ueki, K., Amano, K., Takimoto, E., and Maruoka, A. On-line estimation of hidden Markov model parameters. In *International Conference on Discovery Science*, pp. 155–169. Springer, 2000.
- Mongillo, G. and Deneve, S. Online learning with hidden Markov models. *Neural computation*, 20(7):1706–1716, 2008.
- Neal, R. M. and Hinton, G. E. A view of the EM algorithm that justifies incremental, sparse, and other variants. In *Learning in graphical models*, pp. 355–368. Springer, 1998.
- Neu, G., Jonsson, A., and Gómez, V. A unified view of entropy-regularized markov decision processes. *arXiv preprint <http://arxiv.org/abs/1705.07798>*, 2017.
- Nocedal, J. and Wright, S. J. *Nonlinear Equations*. Springer, 2006.
- Rabiner, L. R. A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.
- Sanderson, C. and Curtin, R. An open source C++ implementation of multi-threaded Gaussian mixture models, k-means and expectation maximisation. In *IEEE Int. Conf. on Signal Processing and Communication Systems (ICSPCS)*, pp. 1–8, 2017.
- Sato, M. Convergence of on-line em algorithm. *7th International Conference on Neural Information Processing*, 1:476–481, 01 2000.
- Singer, Y. and Warmuth, M. K. Training algorithms for hidden Markov models using entropy based distance functions. In *Advances in Neural Information Processing Systems*, pp. 641–647, 1997.
- Singer, Y. and Warmuth, M. K. Batch and on-line parameter estimation of Gaussian mixtures based on the joint entropy. In *Proceedings of Advances in Neural Information Processing Systems*, pp. 578–584, 1999.
- Titterton, D. M. Recursive Parameter Estimation Using Incomplete Data. *Journal of the Royal Statistical Society, Series B*, 46(2):257–267, 1984.
- Wainwright, M. J., Jordan, M. I., et al. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1–2):1–305, 2008.
- Wang, S. and Zhao, Y. Almost sure convergence of Titterton’s recursive estimator for mixture models. *Statistics & probability letters*, 76(18):2001–2006, 2006.
- Welch, G. and Bishop, G. An introduction to the Kalman filter. Technical report, University of North Carolina at Chapel Hill, 1995.
- Wu, C. J. On the convergence properties of the EM algorithm. *The Annals of statistics*, pp. 95–103, 1983.

A PROOF OF THEOREM 1

Proof. The online E and M steps in (Cappé & Moulines, 2009) are defined respectively as

$$\begin{aligned}\hat{Q}_{t+1}(\tilde{\Theta}) &= \hat{Q}_t(\tilde{\Theta}) \\ &+ \tilde{\eta}_t \left(\mathbb{E}_{P(h|v_{t+1}, \Theta^t)} [\log p(h, v_{t+1} | \tilde{\Theta})] - \hat{Q}_t(\tilde{\Theta}) \right), \\ \Theta^{t+1} &= \arg \max_{\tilde{\Theta}} \hat{Q}_{t+1}(\tilde{\Theta}) = \arg \min_{\tilde{\Theta}} -\hat{Q}_{t+1}(\tilde{\Theta}),\end{aligned}\quad (10)$$

$$\Theta^{t+1} = \arg \max_{\tilde{\Theta}} \hat{Q}_{t+1}(\tilde{\Theta}) = \arg \min_{\tilde{\Theta}} -\hat{Q}_{t+1}(\tilde{\Theta}), \quad (11)$$

where

$$\hat{Q}_t(\tilde{\Theta}) = \mathbb{E}_{p(h, v | \Theta^t)} [\log p(h, v | \tilde{\Theta})].$$

The relative entropy divergence between the hidden variable models can be written as

$$\begin{aligned}D_{\text{RE}}(\Theta^t, \tilde{\Theta}) &= \int_{h, v} p(h, v | \Theta^t) \log \frac{p(h, v | \Theta^t)}{p(h, v | \tilde{\Theta})} \\ &= \int_{h, v} p(h, v | \Theta^t) \log p(h, v | \Theta^t) - \hat{Q}_t(\tilde{\Theta}).\end{aligned}\quad (12)$$

Note that the first term in (12) does not depend on $\tilde{\Theta}$. Additionally, for the EM upper-bound, we have

$$\begin{aligned}U_{\Theta}(\tilde{\Theta} | v_{t+1}) &= \mathbb{E}_{p(h|v_{t+1}, \Theta^t)} [\log p(h | v_{t+1}, \Theta)] \\ &- \mathbb{E}_{p(h|v_{t+1}, \Theta^t)} [\log p(h, v_{t+1} | \tilde{\Theta})],\end{aligned}\quad (13)$$

where again the first term does not depend on $\tilde{\Theta}$ and the second term corresponds to the negative of the expectation in (10). Comparing to Equation (4) with $\eta_t = \tilde{\eta}_t / (1 - \tilde{\eta}_t)$ and ignoring the constant yields the same updates given in (10) and (11). \square

B BREGMAN DIVERGENCE AND EXPONENTIAL FAMILY

In this section, we review Bregman divergence and exponential family as well as the required lemmas for deriving the updates.

For a real-valued continuously-differentiable and strictly convex function $F : \mathbb{R}^d \rightarrow \mathbb{R}$, the Bregman divergence (Bregman, 1967; Cichocki et al., 2009) $D_F(\tilde{\psi}, \psi)$ between $\tilde{\psi}$ and ψ in the domain of F is defined as

$$D_F(\tilde{\psi}, \psi) = F(\tilde{\psi}) - F(\psi) - f(\psi) \cdot (\tilde{\psi} - \psi),$$

where $f(\psi) := \nabla_{\psi} F(\psi)$. The gradient wrt the first arguments take the form

$$\nabla_{\tilde{\psi}} D_F(\tilde{\psi}, \psi) = f(\tilde{\psi}) - f(\psi),$$

while the gradient wrt the second argument becomes

$$\nabla_{\psi} D_F(\tilde{\psi}, \psi) = -\nabla^2 F(\psi) (\tilde{\psi} - \psi).$$

The *Fenchel dual* (Hiriart-Urruty & Lemarchal, 2001) of the function F is defined as

$$F^*(\psi^*) = \sup_{\psi'} \{ \psi' \cdot \psi^* - F(\psi') \}.$$

Assuming that the supremum is achieved at ψ , we have the following relation between variables ψ and ψ^*

$$\psi^* = f(\psi), \quad \psi = f^*(\psi^*), \quad \text{and} \quad f^* = f^{-1},$$

where $f^*(\psi^*) := \nabla_{\psi^*} F^*(\psi^*)$. Note that as a result of convexity of F^* , we can form the dual Bregman divergence using F^* as the generating convex function. The following equality holds for pairs of dual variables (ψ, ψ^*) and $(\tilde{\psi}, \tilde{\psi}^*)$

$$\begin{aligned}D_F(\tilde{\psi}, \psi) &= F(\tilde{\psi}) - F(\psi) - f(\psi) \cdot (\tilde{\psi} - \psi) \\ &\quad + f(\tilde{\psi}) \cdot \tilde{\psi} - f(\tilde{\psi}) \cdot \tilde{\psi} \\ &= \underbrace{-F(\psi) + f(\psi) \cdot \psi}_{F^*(\psi^*)} + \underbrace{F(\tilde{\psi}) - f(\tilde{\psi}) \cdot \tilde{\psi}}_{-F^*(\tilde{\psi}^*)} \\ &\quad - \underbrace{\tilde{\psi} \cdot (f(\psi) - f(\tilde{\psi}))}_{f^*(\tilde{\psi}^*) \cdot (\psi^* - \tilde{\psi}^*)} \\ &= D_{F^*}(\psi^*, \tilde{\psi}^*).\end{aligned}$$

Note that the order of variables is reversed when switching to the dual divergence. Additionally, using the definition of the dual function, we have

$$D_F(\tilde{\psi}, \psi) = D_{F^*}(\psi^*, \tilde{\psi}^*) = F(\tilde{\psi}) + F^*(\psi^*) - \tilde{\psi} \cdot \psi^*.$$

The following lemmas for combining Bregman divergences are useful for our discussion of the EM updates.

Lemma 5. Forward Combination Let $\{\alpha_i\}_{i=1}^N$ where $\alpha_i \in \mathbb{R}_+$ and $\sum_i \alpha_i > 0$. We have

$$\psi_{\text{opt}} = \arg \min_{\tilde{\psi}} \sum_i \alpha_i D_F(\tilde{\psi}, \psi_i) = f^* \left(\frac{\sum_i \alpha_i f(\psi_i)}{\sum_i \alpha_i} \right).$$

Proof. Taking the derivative of the objective function wrt $\tilde{\psi}$ and replacing for the gradient of the Bregman divergence with respect to the first argument, we have

$$\sum_i \alpha_i \left(f(\tilde{\psi}) - f(\psi_i) \right) = 0,$$

which yields

$$\left(\sum_i \alpha_i \right) f(\psi_{\text{opt}}) = \sum_i \alpha_i f(\psi_i),$$

$$\text{or} \quad \psi_{\text{opt}} = f^{-1} \left(\frac{\sum_i \alpha_i f(\psi_i)}{\sum_i \alpha_i} \right).$$

Using the fact that $f^{-1} = f^*$ completes the proof. \square

Corollary 6. Forward Triangular Equality

$$\begin{aligned} \sum_i \alpha_i D_F(\tilde{\psi}, \psi_i) - \sum_i \alpha_i D_F(\psi_{\text{opt}}, \psi_i) \\ = \left(\sum_i \alpha_i \right) D_F(\tilde{\psi}, \psi_{\text{opt}}). \end{aligned}$$

Lemma 7. Backward Combination Let $\{\alpha_i\}_{i=1}^N$ where $\alpha_i \in \mathbb{R}_+$ and $\sum_i \alpha_i > 0$. We have

$$\begin{aligned} \psi_{\text{opt}}^* &= \arg \min_{\tilde{\psi}^*} \left[\sum_i \alpha_i D_{F^*}(\psi_i^*, \tilde{\psi}^*) \right] \\ &= \sum_i \alpha_i \left(F^*(\psi_i^*) + F(\tilde{\psi}) - \tilde{\psi} \cdot \psi_i^* \right) = \frac{\sum_i \alpha_i \psi_i^*}{\sum_i \alpha_i}. \end{aligned}$$

Proof. Taking the derivative of the objective function wrt $\tilde{\psi}^*$ and replacing for the gradient of the Bregman divergence with respect to the second argument, we have

$$-\sum_i \alpha_i \left(\nabla^2 F^*(\tilde{\psi}^*) (\psi_i^* - \tilde{\psi}^*) \right) = 0.$$

Using the fact that $\nabla^2 F^*(\tilde{\psi}^*) \succeq 0$ and rearranging the terms concludes the proof. \square

Corollary 8. Backward Triangular Equality

$$\begin{aligned} \sum_i \alpha_i D_{F^*}(\psi_i^*, \tilde{\psi}^*) - \sum_i \alpha_i D_F(\psi_i^*, \psi_{\text{opt}}^*) \\ = \left(\sum_i \alpha_i \right) D_F(\psi_{\text{opt}}^*, \tilde{\psi}^*). \end{aligned}$$

In some cases, the value of $F^*(\psi_i^*)$ becomes unbounded (e.g. negative entropy of the Dirac measure). However, we can still apply Lemma 5 and 7 by dropping the $F^*(\psi_i^*)$ terms from the objective.

Lemma 9. Partial Combination Let $\{\alpha_i\}_{i=1}^N$ where $\alpha_i \in \mathbb{R}_+$ and $\sum_i \alpha_i > 0$. We have

$$\begin{aligned} \psi_{\text{opt}} &= \arg \min_{\tilde{\psi}} \sum_i \alpha_i (F(\tilde{\psi}) - \tilde{\psi} \cdot \psi_i^*) \\ &= f^* \left(\frac{\sum_i \alpha_i f(\psi_i)}{\sum_i \alpha_i} \right), \end{aligned}$$

i.e.

$$\psi_{\text{opt}}^* = \frac{\sum_i \alpha_i \psi_i^*}{\sum_i \alpha_i}.$$

Corollary 10.

$$\begin{aligned} \sum_i \alpha_i (F(\tilde{\psi}) - \tilde{\psi} \cdot \psi_i^*) + \left(\sum_i \alpha_i \right) F^*(\psi_{\text{opt}}^*) \\ = \left(\sum_i \alpha_i \right) D_F(\tilde{\psi}, \psi_{\text{opt}}^*). \end{aligned}$$

C COMPOUND DIRICHLET DISTRIBUTION

Compound Dirichlet distribution (also referred to as Pólya distribution) Gupta et al. (2011) is commonly used to model distribution over topics. A topic entails a distribution over words. More specifically, the compound Dirichlet distribution includes a non-negative parameter vector $\alpha > 0$ corresponding to a Dirichlet distribution over topics. The sampling process consists of sampling a topic h_n for the n -th document from the Dirichlet distribution. The component $h_{n,i}$ corresponds to the probability of sampling the i -th word. Next, a set of iid samples v_n are drawn from the topic. That is, $v_{n,i}$ denotes the frequency of the i -th word and $\sum_i v_{n,i}$ is the total number of words in the n -th document. Note that the sampled topics are hidden and only the set of documents are given. The set of model parameters equals to $\Theta = \{\alpha\}$.

The joint distribution over the hidden topics and visible documents can be written as

$$p(h, v | \Theta) = \prod_n \frac{\Gamma(\alpha_0)}{\prod_j \Gamma(\alpha_j)} \frac{(\sum_j v_{n,j})!}{\prod_j (v_{n,j})!} \prod_i h_{n,i}^{\alpha_i + v_{n,i} - 1},$$

where $\alpha_0 = \sum_j \alpha_j$ and $\Gamma(\cdot)$ is the gamma function.

The marginal probability of the documents can be calculated by integrating out the hidden topics, that is,

$$p(v | \Theta) = \prod_n \frac{(\sum_j v_{n,j})! \Gamma(\alpha_0) \prod_j \Gamma(\alpha_j + v_{n,j})}{(\prod_j v_{n,j})! (\prod_j \Gamma(\alpha_j)) \Gamma(\sum_j (\alpha_j + v_{n,j}))}.$$

The EM upper-bound can be written as

$$\begin{aligned} U_{\Theta}(\tilde{\Theta}) &= 1/N \sum_n p(h | v_n, \alpha) \log p(v_n, h | \tilde{\alpha}) \\ &= 1/N \sum_n p(h | v_n, \alpha) \log p(h | \tilde{\alpha}) \\ &= \log \left(\frac{\Gamma(\tilde{\alpha}_0)}{\prod_j \Gamma(\tilde{\alpha}_j)} \right) \\ &\quad + 1/N \sum_n \sum_j p(h | v_n, \alpha) (\tilde{\alpha}_j - 1) \log h_j \\ &= N \log \left(\frac{\Gamma(\tilde{\alpha}_0)}{\prod_j \Gamma(\tilde{\alpha}_j)} \right) \\ &\quad + \sum_n \sum_j (\tilde{\alpha}_j - 1) \left(\psi(v_{nj} + \alpha_j) - \psi \left(\sum_i v_{ni} + \alpha_0 \right) \right), \end{aligned}$$

where $\psi(\alpha) := \frac{\partial}{\partial \alpha} \log \Gamma(\alpha)$ is called the digamma func-

tion. The inertia term on the hand

$$D_{\text{RE}}(\Theta, \tilde{\Theta}) = \frac{\Gamma(\alpha_0)}{\Gamma(\tilde{\alpha}_0)} - \sum_j \frac{\Gamma(\alpha_j)}{\Gamma(\tilde{\alpha}_j)}$$

$$+ \left[\sum_j (\alpha_j - \tilde{\alpha}_j) \sum_v \frac{\Gamma(\alpha_0)}{\prod_i \Gamma(\alpha_i)} \frac{(\sum_i v_i)!}{\prod_i (v_i!)} \frac{\Gamma(\alpha_0 + \sum_i v_i)}{\prod_i \Gamma(\alpha_i + v_i)} \right. \\ \left. \times \left(\psi(v_j + \alpha_j) - \psi\left(\sum_i v_i + \alpha_0\right) \right) \right],$$

involves summing over all possible combinations of v and therefore, is intractable. Alternatively, we can use the approximate form of the upper-bound to perform the updates.

A standard approach to minimize the upper-bound is the Newton's method (Nocedal & Wright, 2006), which requires calculating the gradient and the Hessian matrix. The gradient of the upper-bound can be written as

$$\frac{\partial \mathbf{U}_{\Theta}(\tilde{\Theta} | \mathcal{V})}{\partial \alpha_i} = \psi(\tilde{\alpha}_0) - \psi(\tilde{\alpha}_i) \\ + 1/N \sum_n \left(\psi\left(\sum_j v_{nj} + \alpha_0\right) - \psi(v_{ni} + \alpha_i) \right).$$

The Hessian is

$$\nabla^2 \mathbf{U}_{\Theta}(\tilde{\Theta} | \mathcal{V}) = \psi_1(\tilde{\alpha}_0) \mathbf{1}\mathbf{1}^\top - \text{diag}[\psi_1(\tilde{\alpha}_1), \dots, \psi_1(\tilde{\alpha}_d)],$$

where $\psi_1(\alpha) := \frac{\partial}{\partial \alpha} \psi(\alpha)$ is called the trigamma function.