



HAL
open science

Streaming Bayesian inference: theoretical limits and mini-batch approximate message-passing

Andre Manoel, Florent Krzakala, Eric W. Tramel, Lenka Zdeborová

► **To cite this version:**

Andre Manoel, Florent Krzakala, Eric W. Tramel, Lenka Zdeborová. Streaming Bayesian inference: theoretical limits and mini-batch approximate message-passing. 2017. cea-01553517

HAL Id: cea-01553517

<https://cea.hal.science/cea-01553517>

Preprint submitted on 3 Jul 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Streaming Bayesian inference: theoretical limits and mini-batch approximate message-passing

Andre Manoel¹, Florent Krzakala², Eric W. Tramel³, and Lenka Zdeborová⁴

¹Neurospin, CEA, Université Paris-Saclay

²LPS ENS, CNRS, PSL, UPMC & Sorbonne Univ.

³OWKIN

⁴IPhT, CNRS, CEA, Université Paris-Saclay

June 5, 2017

Abstract

In statistical learning for real-world large-scale data problems, one must often resort to “streaming” algorithms which operate sequentially on small batches of data. In this work, we present an analysis of the information-theoretic limits of mini-batch inference in the context of generalized linear models and low-rank matrix factorization. In a controlled Bayes-optimal setting, we characterize the optimal performance and phase transitions as a function of mini-batch size. We base part of our results on a detailed analysis of a mini-batch version of the approximate message-passing algorithm (Mini-AMP), which we introduce. Additionally, we show that this theoretical optimality carries over into real-data problems by illustrating that Mini-AMP is competitive with standard streaming algorithms for clustering.

1 Introduction

In current machine learning applications, one often faces the challenge of *scale*: massive data causes algorithms to explode in time and memory requirements. In such cases, when it is infeasible to process the full dataset simultaneously, one must resort to “online” or “streaming” methods which process only a small fraction of data points at a time — producing a step-by-step learning process. Such procedures are becoming more and more necessary to cope with massive datasets. For example, one can see the effectiveness of such approaches in deep learning via the stochastic gradient descent algorithm [1] or in statistical inference via the stochastic variational inference framework [2].

In this work, we treat streaming inference within a Bayesian framework where, as new data arrives, posterior beliefs are updated according to Bayes’ rule. One well known approach in this direction is assumed density filtering (ADF) [3, 4], which processes a single data point at a time, a procedure to which we refer to as *fully online*. A number of other works analyzed various related fully online algorithms [5, 6], especially in the statistical physics literature [7–11]. We are instead interested in the case where multiple samples – a *mini-batch* – arrive at once. Tuning the size of these mini-batches allows us to explore the trade-off between the precision and efficiency.

Our motivation and setting are very much along the lines of streaming variational Bayes (VB) inference [12]. With respect to existing works, we bring three main contributions. (i) We introduce a streaming algorithm based on approximate message passing (AMP) [13–15] that we call Mini-AMP. As AMP treats some of the correlations which VB neglects, it is expected that AMP either outperforms or matches VB. (ii) Unlike other general-purpose algorithms for Bayesian inference, such as Gibbs sampling or VB, AMP possesses the state evolution method which asymptotically describes the performance of the algorithm for a class of generative models. We extend this state evolution analysis to Mini-AMP. (iii) For these generative

models, we also analyze the optimal streaming procedure, within a class of procedures that retains only point-wise marginals from one step to another, and characterize regions of parameters where Mini-AMP reaches optimality.

2 Problem setting

Denoting the vector of N values to be estimated by \mathbf{x} , the data presented at step k by $\mathbf{y}^{(k)}$, and the collection of all previously presented data by $\mathcal{D}^{(k-1)} = \{\mathbf{y}^{(1)}, \dots, \mathbf{y}^{(k-1)}\}$, the posterior distribution at step k is given by

$$P(\mathbf{x}|\mathbf{y}^{(k)}, \mathcal{D}^{(k-1)}) = \frac{P(\mathbf{y}^{(k)}|\mathbf{x}) P(\mathbf{x}|\mathcal{D}^{(k-1)})}{\int d\mathbf{x} P(\mathbf{y}^{(k)}|\mathbf{x}) P(\mathbf{x}|\mathcal{D}^{(k-1)})}. \quad (1)$$

In other words, with each presentation of new data, the prior distribution is updated with the posterior distribution derived from the previously presented data. Directly implementing this strategy is seldom feasible as the normalizing integral in (1) is intractable in general. Additionally, to keep the memory requirements small, we would like consider only the case where $O(N)$ parameters are passed from one step to the next. With this restriction, one cannot carry over high-order correlations from previous steps. Instead, following the strategy of [12], we resort to a factorized approximation of the "prior" term of (1),

$$P(\mathbf{x}|\mathcal{D}^{(k-1)}) \approx Q^{(k-1)}(\mathbf{x}) = \prod_{i=1}^N \mathcal{P}_i[P(\mathbf{x}|\mathbf{y}^{(k-1)}, \mathcal{D}^{(k-2)})], \quad (2)$$

where $\mathcal{P}_i[\cdot]$ denotes the posterior marginals of parameter x_i at a given step. Computing the marginals exactly is still computationally intractable for most models of interest. In the present work, this program is carried out with a scheme that is asymptotically exact for a class of generative models *and* that has the advantage of being amenable to a rigorous analysis.

We leverage the analysis of these models already conducted in the offline setting using concepts and techniques from statistical physics [13–20] which have now been made almost entirely rigorous [21–26]. We show in particular that – just as for the offline setting – phase transitions exist for mini-batch learning problems, and that their description provides information about the learning error that is achievable information-theoretically or computationally efficiently.

3 Generative models and offline learning

In our theoretical analysis, we consider inference in popular models with synthetic data generated from a given distribution, such as the perceptron with random patterns [27, 28], sparse linear regression with a random matrix (compressed sensing) [29] and clustering random mixtures of Gaussians [19]. For clarity, we restrict our presentation to the generalized linear models (GLMs), focusing on sparse linear estimation. Our results, however, can be extended straightforwardly to any problem where AMP can be applied. For offline GLMs, the joint distribution of the observation $\mathbf{y} \in \mathbb{R}^M$ and the unknown $\mathbf{x} \in \mathbb{R}^N$ is given by

$$P(\mathbf{y}, \mathbf{x}|\Phi) = \prod_{\mu=1}^M P(y_\mu|z_\mu \equiv \Phi_\mu \cdot \mathbf{x}) \prod_{i=1}^N P_X(x_i). \quad (3)$$

where Φ_μ is the μ -th line of the $M \times N$ matrix Φ . We consider the situation where Φ is a random matrix where each element is taken i.i.d. from $\mathcal{N}(0, 1/N)$ and $\alpha = M/N$. Structured matrices have also been studied with AMP [30, 31]. Two situations of interest described by GLMs are (a) sparse linear regression (SLR) where the likelihood is Gaussian $P(y_\mu|z_\mu) = \mathcal{N}(y_\mu; z_\mu, \Delta)$ and the parameters are sparse, for instance drawn from a Gauss-Bernoulli distribution $P_X(x_i) = \rho \mathcal{N}(x_i; 0, 1) + (1 - \rho) \delta(x_i)$, and (b) the probit regression problem $P(y_\mu|z_\mu) = \frac{1}{2} \operatorname{erfc}(-\frac{y_\mu z_\mu}{\sqrt{2\Delta}})$ that reduces to the perceptron $P(y_\mu|z_\mu) = \theta(y_\mu z_\mu)$ when $\Delta \rightarrow 0$ [15, 28, 32]. We

first summarize the known relevant results for the fully *offline* learning problem, where one processes all data at once. Again, for clarity, we focus on the case of SLR.

The marginals estimated by AMP are given by [13, 14, 18]

$$P(x_i|\Phi, \mathbf{y}) \approx q(x_i|A, B_i) = P_X(x_i) e^{-\frac{1}{2}Ax_i^2 + B_ix_i} / Z(A, B_i), \quad (4)$$

where $Z(A, B_i)$ is a normalization factor. We shall denote the mean and variance of this distribution by $\eta(A, B) \equiv \frac{\partial}{\partial B} \log Z(A, B)$ and $\eta'(A, B) \equiv \frac{\partial}{\partial B} \eta(A, B)$. The mean, in particular, provides an approximation to the minimum mean-squared error (MMSE) estimate of \mathbf{x} . The AMP iteration reads

$$\mathbf{z}^{(t)} = \mathbf{y} - \Phi \hat{\mathbf{x}}^{(t)} + \alpha^{-1} \mathbf{z}^{(t-1)} A^{(t-1)} V^{(t)}, \quad (5)$$

$$\mathbf{B}^{(t)} = A^{(t)} \hat{\mathbf{x}}^{(t)} + A^{(t)} \alpha^{-1} \Phi^T \mathbf{z}^{(t)}, \quad A^{(t)} = \frac{\alpha}{\Delta + V^{(t)}}, \quad (6)$$

$$\hat{x}_i^{(t+1)} = \eta(A^{(t)}, B_i^{(t)}) \forall i, \quad V^{(t+1)} = \frac{1}{N} \sum_{i=1}^N \eta'(A^{(t)}, B_i^{(t)}). \quad (7)$$

One of the main strengths of AMP is that when the matrix Φ has i.i.d. elements, the ground truth parameters are generated i.i.d. from a distribution $P_0(x_i)$, and $P(y_\mu|z_\mu) = \mathcal{N}(y_\mu; z_\mu, \Delta_0)$, then the behavior and performance of AMP can be studied analytically in the large system limit ($N \rightarrow \infty$) using a technique called state evolution. This was proven by [21] who show that in the large N limit, $A^{(t)}$ and $B_i^{(t)}$ converge in distribution such that, defining $\mathcal{E}^{(t)} \equiv \mathbb{E}(\eta(\mathcal{A}^{(t-1)}, \mathcal{B}^{(t-1)}) - x)^2$, with $x \sim P_0(x)$, and $\mathcal{V}^{(t)} \equiv \mathbb{E} \eta'(\mathcal{A}^{(t-1)}, \mathcal{B}^{(t-1)})$, one has

$$A^{(t)} \rightsquigarrow \mathcal{A}^{(t)} = \frac{\alpha}{\Delta + \mathcal{V}^{(t)}}, \quad B_i^{(t)} \rightsquigarrow \mathcal{B}^{(t)} \sim \mathcal{N}\left(\mathcal{A}^{(t)} x, \alpha \frac{\Delta_0 + \mathcal{E}^{(t)}}{(\Delta + \mathcal{V}^{(t)})^2}\right). \quad (8)$$

The behavior of the algorithm is monitored by the computation of the scalar quantities $\mathcal{E}^{(t)}$ and $\mathcal{V}^{(t)}$.

The "Bayes-optimal" setting is defined as the case when the generative model is known and matches the terms in the posterior (1), i.e. when $P_X = P_0$, and $\Delta = \Delta_0$. One can show that in this case $\mathcal{E}^{(t)} = \mathcal{V}^{(t)}$ (the so-called Nishimori property [15]), so that the state evolution further reduces to

$$\mathcal{A}^{(t)} = \frac{\alpha}{\Delta + \mathcal{E}^{(t)}}, \quad \mathcal{E}^{(t)} = \mathbb{E} \eta'(\mathcal{A}^{(t-1)}, \mathcal{A}^{(t-1)} x + \sqrt{\mathcal{A}^{(t-1)}} z), \quad (9)$$

with $x \sim P_X(x)$, $z \sim \mathcal{N}(0, 1)$, and $\mathcal{E}^{(t)}$ is the mean-squared error (MSE) achieved at iteration t .

Another set of recent results [24, 25] allows for the exact computation of the Bayes-optimal MMSE and the mutual information between the observations and the unknown parameters. Given model (3) with Gaussian likelihood, the mutual information per variable is given by the minimum of the so-called replica mutual information: $\lim_{N \rightarrow \infty} I(\mathbf{X}, \mathbf{Y}) = \min i_{\text{RS}}(\mathcal{E})$ where, defining, $\Sigma^{-2}(\mathcal{E}) \equiv \frac{\alpha}{\Delta + \mathcal{E}}$,

$$i_{\text{RS}}(\mathcal{E}) = \frac{\alpha}{2} \left[\frac{\mathcal{E}}{\Delta + \mathcal{E}} + \log \left(1 + \frac{\mathcal{E}}{\Delta} \right) \right] - \mathbb{E}_{x, z} \left[\log \mathbb{E}_{\tilde{x}} e^{-\frac{(\tilde{x} - (x + z \Sigma(\mathcal{E})))^2}{2\Sigma^2(\mathcal{E})}} \right] - \frac{1}{2}, \quad (10)$$

with $x \sim P_X(x)$, $\tilde{x} \sim P_X(\tilde{x})$ and $z \sim \mathcal{N}(0, 1)$. The MMSE is then given by $\arg \min i_{\text{RS}}(\mathcal{E})$.

Comparisons between the MMSE and the MSE provided by AMP after convergence are very instructive, as shown in [15]. Typically, for large enough noise, $\mathcal{E}^{(t \rightarrow \infty)} = \mathcal{E}_{\text{AMP}} = \text{MMSE}$ and AMP achieves the Bayes-optimal result in polynomial time, thus justifying, *a posteriori*, the interest of such algorithms in this setting. In fact, since the fixed points of the state evolution are all extrema of the mutual information (10), it is useful to think of AMP as an algorithm that attempts to minimize (10). However, a *computational* phase transition can exist at low noise levels, where $i_{\text{RS}}(\mathcal{E})$ has more than a single minimum. In this case, it may happen that AMP does not reach the *global* minimum, and therefore $\mathcal{E}_{\text{AMP}} > \text{MMSE}$. It is a remarkable open problem to determine whether finding the MMSE in this region is computationally tractable. The results we have just described are not merely restricted to SLR, but appear *mutatis mutandis* in various cases of low-rank matrix and tensor factorization [17, 19, 20, 23, 26, 33] and also partly in GLMs [14, 15] (in GLMs the replica mutual information is so far only conjectured).

4 Main results

4.1 Mini-AMP

Our first contribution is the Mini-AMP algorithm, which adapts AMP to the streaming setting. Again, we shall restrict the presentation to the linear regression case. The adaptation to other AMP algorithms is straightforward. We consider a dataset of M samples with N features each, which we split into B mini-batches, each containing $M_b = M/B$ samples. We denote $\alpha = M/N$ and $\alpha_b = M_b/N$. Crucially, for each step, the posterior marginal given by AMP (4) is the prior multiplied by a quadratic form. Performing the program discussed in (2) is thus tractable as the \mathcal{P}_i are given by a Gaussian distribution multiplied by the original prior. The only modification w.r.t. the offline AMP at each step is thus to update the prior by multiplying the former one by the exponential in (4). In other words, we use the following "effective" prior when processing the k -th mini-batch:

$$P_{\Lambda_{k-1}, \Theta_{k-1}}^k(\mathbf{x}) = P_X(\mathbf{x}) \prod_{i=1}^N e^{-\frac{1}{2} \Lambda_{k-1} x_i^2 + \Theta_{k-1, i} x_i}, \quad \text{where } \Lambda_{k-1} = \sum_{\ell=1}^{k-1} A_\ell, \quad \Theta_{k-1, i} = \sum_{\ell=1}^{k-1} B_{\ell, i}. \quad (11)$$

In practice, the only change when moving from AMP to Mini-AMP is therefore the update of the arguments of the η function. After k mini-batches have been processed, one replaces (7) by

$$\begin{aligned} \hat{\mathbf{x}}_{k, i}^{(t+1)} &= \eta \left(\underbrace{\sum_{\ell=1}^{k-1} A_\ell}_{\Lambda_{k-1}} + A_k^{(t)}, \underbrace{\sum_{\ell=1}^{k-1} B_{\ell, i} + B_{k, i}^{(t)}}_{\Theta_{k-1, i}} \right), \\ V_k^{(t+1)} &= \frac{1}{N} \sum_{i=1}^N \eta' \left(\underbrace{\sum_{\ell=1}^{k-1} A_\ell}_{\Lambda_{k-1}} + A_k^{(t)}, \underbrace{\sum_{\ell=1}^{k-1} B_{\ell, i} + B_{k, i}^{(t)}}_{\Theta_{k-1, i}} \right). \end{aligned} \quad (12)$$

The corresponding pseudo-code is given as Algorithm 1. Each Mini-AMP iteration has a computational complexity proportional to $M_b \times N$. We note that, in the fully online scheme when $M_b = 1$, Mini-AMP with a single iteration performed per sample gives the same as ADF [3, 11].

Algorithm 1 Mini-AMP

- 1: initialize $\Lambda_0 = 0, \Theta_{0, i} = 0 \forall i$
 - 2: **for** $k \leftarrow 1$ **to** B **do**
 - 3: initialize $z_{k, \mu}^{(1)} = 0 \forall \mu$
 - 4: initialize $\hat{\mathbf{x}}_{k, i}^{(1)} = \eta(\Lambda_{k-1}, \Theta_{k-1, i}) \forall i, V_k^{(1)} = \frac{1}{N} \sum_{i=1}^N \eta'(\Lambda_{k-1}, \Theta_{k-1, i})$
 - 5: **for** $t \leftarrow 1$ **to** $t_{\max}^{(t)}$ **do**
 - 6: compute $\mathbf{z}_k^{(t)}$ using (5)
 - 7: compute $A_k^{(t)}, \mathbf{B}_k^{(t)}$ using (6)
 - 8: compute $V_k^{(t+1)}, \hat{\mathbf{x}}_k^{(t+1)}$ using (12)
 - 9: **end for**
 - 10: accumulate $\Lambda_k \leftarrow \Lambda_{k-1} + A_k$
 - 11: accumulate $\Theta_k \leftarrow \Theta_{k-1} + \mathbf{B}_k$
 - 12: **end for**
-

4.2 State evolution

Theorem 1 (State evolution of Mini-AMP). *For a random matrix Φ , where each element is taken i.i.d. from $\mathcal{N}(0, 1/N)$, the MSE of Mini-AMP can be monitored asymptotically ($N \rightarrow \infty$ while $\alpha_b = O(1)$) by iterating the*

following state evolution equations,

$$\begin{aligned}\lambda_k^{(t)} &= \lambda_{k-1} + \frac{\alpha_b}{\Delta + \mathcal{V}_k^{(t)}}, & \mathcal{V}_k^{(t+1)} &= \mathbb{E}_{x,z} \eta'(\lambda_k^{(t)}, \lambda_k^{(t)} x + \sqrt{\gamma_k^{(t)}} z), \\ \gamma_k^{(t)} &= \gamma_{k-1} + \alpha_b \frac{\Delta_0 + \mathcal{E}_k^{(t)}}{(\Delta + \mathcal{V}_k^{(t)})^2}, & \mathcal{E}_k^{(t+1)} &= \mathbb{E}_{x,z} \left(\eta(\lambda_k^{(t)}, \lambda_k^{(t)} x + \sqrt{\gamma_k^{(t)}} z) - x \right)^2,\end{aligned}\tag{13}$$

where $x \sim P_0(x)$, and $z \sim \mathcal{N}(0, 1)$. For each $k = 1, \dots, N_b$, these equations are iterated from $t = 1, \dots, t_{max}$, at which point we assign $\lambda_{k+1} = \lambda_{k+1}^{(t_{max})}$ and $\gamma_{k+1} = \gamma_{k+1}^{(t_{max})}$. The MSE given by AMP after the k -th mini-batch has been processed is given by \mathcal{E}_k . In the Bayes-optimal case, in particular, one can further show that the state evolution reduces to

$$\lambda_k^{(t)} = \lambda_{k-1} + \frac{\alpha_b}{\Delta + \mathcal{E}_k^{(t)}}, \quad \mathcal{E}_k^{(t+1)} = \mathbb{E}_{x,z} \left(\eta(\lambda_k^{(t)}, \lambda_k^{(t)} x + \sqrt{\lambda_k^{(t)}} z) - x \right)^2.\tag{14}$$

Proof sketch. We apply the proof of state evolution for AMP in [21] to each mini-batch step, each with its own denoiser function $\eta(\cdot)$. Each step is an instance of AMP with a new, independent matrix, and an effective denoiser given by (12). Using (8), the statistics of the denoisers are known, and the application of the standard AMP state evolution leads to (13). The Bayes-optimal case (14) then follows by induction, as in Sec. V.A.2 of [34]. \square

Note that the above Theorem holds for any value of t_{max} . Hence, even a stochastic version of the Mini-AMP algorithm, where for every mini-batch one only performs a few iterations without waiting for convergence in order to further speed up the algorithm, is analyzable using the above state evolution.

4.3 Optimal MMSE and mutual information under mini-batch setting

Theorem 2 (Mutual information for each mini-batch). *For a random matrix Φ , where each element is taken i.i.d. from $\mathcal{N}(0, 1/N)$, in the Bayes-optimal setting, assume one has been given, after $k - 1$ mini-batches, a noisy version \mathbf{r} of unknown signal \mathbf{x} with i.i.d noise $\mathcal{N}(0, \lambda^{-1})$. Given a new mini-batch with $\alpha_b = O(1)$, the mutual information per variable between the couple (\mathbf{r}, \mathbf{y}) and the unknown \mathbf{x} is asymptotically given by $i = \min i_{RS}^b(\mathcal{E}_k)$ where, defining $\Sigma_b^{-2}(\lambda, \mathcal{E}_k) \equiv \lambda + \frac{\alpha_b}{\Delta + \mathcal{E}_k}$,*

$$i_{RS}^b(\mathcal{E}_k) = \frac{\alpha_b}{2} \left[\frac{\mathcal{E}_k}{\Delta + \mathcal{E}_k} + \log \left(1 + \frac{\mathcal{E}_k}{\Delta} \right) \right] - \mathbb{E}_{x,z} \log \mathbb{E}_{\tilde{x}} e^{-\frac{(\tilde{x} - (x + z \Sigma_b(\lambda, \mathcal{E}_k)))^2}{2 \Sigma_b^2(\lambda, \mathcal{E}_k)}} - \frac{1 + \alpha_b}{2}.\tag{15}$$

Proof sketch. The proof is a slight generalization of the Guerra construction in [24]. Using properties of the Shannon entropy, the mutual information can be written as

$$I(Y, R; X) = H(Y, R) - H(Y|X) - H(R|X) = \mathbb{E}_{\mathbf{y}, \mathbf{r}} \log \mathbb{E}_{\mathbf{x}} e^{-\frac{\|\mathbf{y} - \Phi \mathbf{x}\|_2^2}{2\Delta}} e^{-\frac{\|\mathbf{r} - \mathbf{x}\|_2^2}{2\lambda^{-1}}} - \frac{1 + \alpha_b}{2}.$$

Computation of this expectation is simplified by noticing that it appears as equation (41) in the Guerra construction of [24], where it was used as a proof method for the offline result by interpolating from a pure noisy Gaussian channel (at "time" $t=0$) to the actual linear channel (at "time" $t=1$). Authors of [24] denoted $\lambda(t)$ as the variance of the Gaussian channel and $\gamma(t)$ as the variance of the linear channel. Our computation corresponds instead to a "time" $0 \leq \tau \leq 1$ where both channels are used. Using Sec. V of [24], with the change of notation $\gamma(\tau) \rightarrow \Delta$ and $\lambda(\tau) \rightarrow \lambda$, we reach

$$\lim_{N \rightarrow \infty} \frac{I(Y, R; X)}{N} = i_{RS}^b(\mathcal{E}_k) - \int_0^\tau R_{\mathcal{E}_k}(t) dt + O(1),\tag{16}$$

where $0 \leq \tau \leq 1$ and $R_{\mathcal{E}_k}(t)$ a non-negative function called the reminder. The validity of the mutual information formula in the offline situation [24, 25] implies that the integral of the reminder in $[0, 1]$ is zero when $\mathcal{E}^* = \operatorname{argmin} i_{RS}^b(\mathcal{E}_b)$. Since $R(t)$ is non-negative, this implies that it is zero almost everywhere, thus $\int_0^\tau R_{\mathcal{E}_k}(t) dt = 0$. \square

Theorem 3 (MMSE for each mini-batch). *With the same hypothesis of Theorem 2, the MMSE when one has access to a noisy estimate with i.i.d. noise $\mathcal{N}(0, \lambda^{-1})$ and the data from the mini-batch at step k , is*

$$\text{MMSE} = \text{argmin } i_{\text{RS}}^b(\mathcal{E}_k). \quad (17)$$

Proof sketch. The proof follows again directly from generic results on the Guerra interpolation in [24] and the so-called I-MMSE formula $di(\Delta)/d\Delta^{-1} = \alpha_b y_{\text{MMSE}}/2$ and y-MMSE formula $y_{\text{MMSE}} = \text{MMSE}/(1 + \Delta \text{MMSE})$ linking the mutual information and the MMSE. \square

One can show through explicit computation that the extrema of $i_{\text{RS}}^b(\mathcal{E})$ correspond – just as in the offline case – to the fixed points of the state evolution.

Using these results, we can analyze, both algorithmically and information theoretically, the mini-batch program (2). Indeed, the new information on the parameters \mathbf{x} passed from mini-batch $k - 1$ to k contained in (11) is simply a (Gaussian) noisy version of \mathbf{x} with inverse variance λ_{k-1} . This is true for the AMP estimate (see (11)) and, in the large N limit, for the exact marginalized posterior distribution as well (see e.g. [26, 35, 36]). The optimal MSE at each mini-batch is thus given by the recursive application of Theorem 3, where at each mini-batch $k = 1, \dots, B$ we minimize (15) using $\lambda_{k-1} = \lambda_{k-2} + \alpha_b/(\Delta + \text{MMSE}_{k-1})$.

Now we can compare the MSE reached by the Mini-AMP algorithms to the MMSE. If, for each mini-batch, the MMSE is reached by the state evolution of the Mini-AMP algorithm starting from the previously reached MSE, then we have the remarkable result that Mini-AMP performs a *Bayes-optimal* and *computationally efficient* implementation of the mini-batch program (2). Otherwise, the Mini-AMP is suboptimal. It remains an open question whether in that case any polynomial algorithm can improve upon the MSE reached by Mini-AMP.

All our results can be directly generalized to the case of AMP for matrix or tensor factorization, as derived and proven in [20, 26, 37]. They can also be adapted to the case of GLMs with non-linear output channels. However, in this setting, the formula for the mutual information has not been yet proven rigorously.

5 Performance and phase transitions on GLMs

5.1 Optimality & efficiency trade-offs with Mini-AMP

We now illustrate the above results on some examples. In Figure 1, we consider the SLR model and the perceptron with binary ± 1 parameters, both with random matrices $\Phi \in \mathbb{R}^{M \times N}$, $\Phi_{\mu i} \sim \mathcal{N}(0, 1/N)$. Our analysis quantifies the loss coming from using mini-batches with respect to a fully offline implementation. In the limit of small mini-batch $\alpha_b \rightarrow 0$, we recover the results of the ADF algorithm which performs fully online learning, processing one sample at a time [9, 11]. This suggests that the state evolution accurately describes the behavior of Mini-AMP beyond the theoretical assumption of $\alpha_b = O(1)$, even for mini-batches as small as a single sample.

The effect of the mini-batch sizes varies greatly with the problem. For the perceptron with ± 1 weights, a zero error is eventually obtained after a sufficient number of mini-batches have been processed. Moreover, the dependence on the mini-batch size is mild: while the offline scheme achieves zero error at $\alpha \approx 1.5$ [27, 28], the fully online does it at $\alpha \approx 4.4$ [9], that is, going from offline to a fully online scheme costs only about three times more data points. The behavior of the Mini-AMP for SLR shows instead rather drastic changes with the mini-batch size. The MSE decays smoothly when the mini-batch size is small. However, as we increase it, a sudden decay occurs after a few mini-batches have been processed. For the noiseless case ($\Delta = 0$), the study of the state evolution shows that the asymptotic (in α) MSE is given by

$$\text{MSE}_x(\alpha) \sim e^{-\frac{1}{\alpha_b} \log(1 - \frac{\alpha_b}{\rho}) \alpha}, \quad (18)$$

if $\alpha_b \leq \rho$, and by 0 otherwise. These results provide a basis for an optimal choice of mini-batch size. Given the drastic change in behavior past a certain mini-batch size, one concludes that small investments in memory might be worthwhile, since they can lead to large gains in performance.

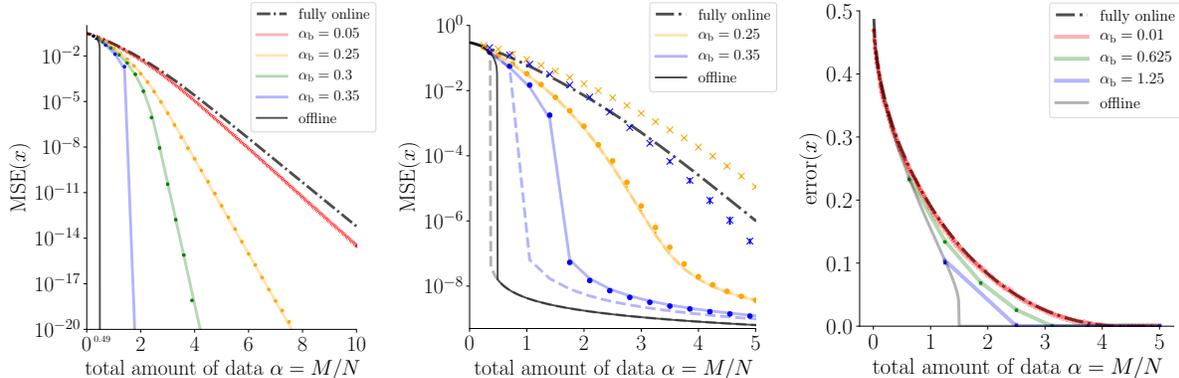


Figure 1: Accuracy of Mini-AMP inference as a function of the total amount of presented data. *Left*: SLR with $\Delta = 0$ and sparsity $\rho = 0.3$. *Center*: SLR with $\Delta = 10^{-8}$ and $\rho = 0.3$. *Right*: Perceptron learning with Rademacher (± 1) parameters/synapses. For different mini-batch sizes (*colors*), we show both the state evolution predictions for Mini-AMP (*solid lines*), the predicted MMSE (*dashed lines*, only center pannel), and empirical experiments for Mini-AMP (\bullet) and streaming VB [12, 38] (\times) averaged over 10 realizations of size $N = 2000$. We also show the results for both fully offline (*solid black line*) and fully online (*dash-dot black line*) algorithms. Even for moderate N the state evolution is found to almost perfectly describe Mini-AMP’s behavior. For the parameters of the center plot we observed that for $\alpha_b \lesssim 0.33$ Mini-AMP is asymptotically optimal.

Finally, we have compared the Mini-AMP scheme with the streaming VB approach [12] using the mean-field algorithm described in [38] for SLR. While the mean-field approach is found to give results comparable to AMP in the offline case in [38], we see here that the results are considerably worse in the streaming problem. In fact, as shown in Figure 1 (center), mean-field can give worse performance than the fully online ADF, even when processing rather large mini-batches.

5.2 Phase transitions

It turns out that, just as for the offline setting, there are phase transitions appearing for mini-batch learning, in terms of the learning error that is achievable information-theoretically (MMSE) *or* computationally efficiently (by Mini-AMP). These can be understood by an analysis of the function i_{RS}^b , since the minimum of i_{RS}^b gives the MMSE, and since AMP is effectively trying to minimize i_{RS}^b starting from the MSE reached at the previous mini-batch steps.

Let us illustrate the reason behind the sharp phenomenon in the behavior of AMP in Fig.1. We show, in Fig. 2 (left), an example of the function $i_{\text{RS}}^b(\mathcal{E})$ for the streaming SLR problem as a function of the MSE \mathcal{E} as each mini-batch is being processed. Initially, it presents a “good” and a “bad” minimum, at small and large MSEs respectively. In the very first batch, AMP reaches the bad minimum. As more batches are processed, the good minimum becomes global, but AMP is yet not able to reach it, and keeps returning the bad one instead. This indicates a computational phase transition, and we expect that other algorithms will, as AMP, fail to deliver the MMSE in polynomial time when this happens. Eventually, the good minimum becomes unique, at which point AMP is able to reach it, thus yielding the sudden decay observed in Figure 1.

Consider now the Bayes-optimal “streaming-MMSE” given by the global minimum of the mutual information at each step, regardless of whether AMP achieves it. In the offline noiseless case, the MMSE is achieved by AMP only if the processed batch has size $\alpha \geq \alpha_{\text{offline}}$ or $\alpha \leq \rho$ [18]. In the streaming case, we also observe that Mini-AMP reaches the streaming-MMSE if the mini-batch size is sufficiently small *or* sufficiently large. In Figure 2 (right) we compare MMSE to the MSE reached by Mini-AMP, with a region between the full and dashed line where the algorithm is sub-optimal.

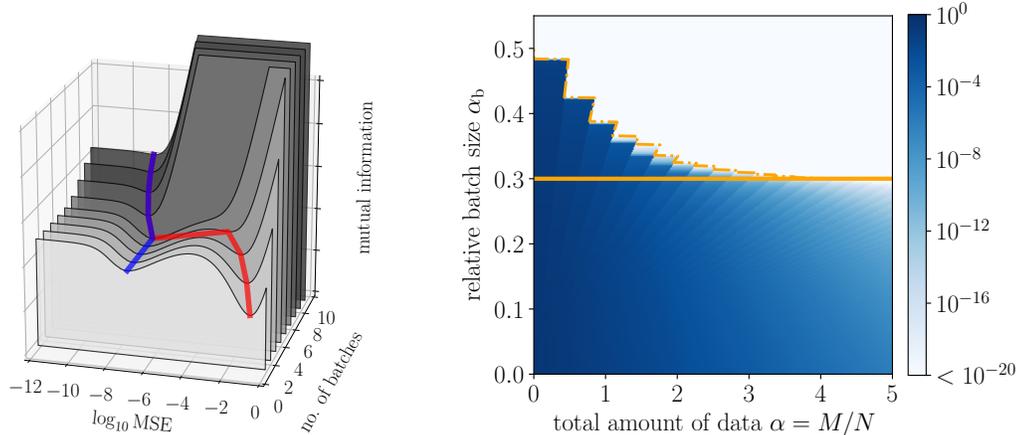


Figure 2: Phase transitions in streaming SLR. *Left*: evolution of the mutual information in the streaming SLR problem as each mini-batch is processed. Parameters are set to $\rho = 0.3$, $\Delta = 10^{-8}$ and $\alpha_b = 0.35$. *Right*: MSE of Mini-AMP for different mini-batch sizes. Mini-AMP achieves the MMSE for $\alpha_b < \rho$. For $\alpha_b > \rho$ the MMSE is zero after processing a single batch, while for batch sizes between $\rho < \alpha_b < 0.49$ the Mini-AMP is suboptimal unless a sufficient number of mini-batches is processed.

6 Mini-AMP for matrix factorization problems

We now consider the case of low-rank matrix factorization, and in particular clustering using the Gaussian mixture model (GMM) with R clusters. For such problems, the generative model reads

$$P(U, V, Y) = \prod_{ij} \mathcal{N}(Y_{ij}; \mathbf{U}_i \cdot \mathbf{V}_j, \Delta) \prod_{i=1}^N \mathcal{N}(\mathbf{U}_i; \mathbf{0}, \mathbb{I}_R) \prod_{j=1}^M \frac{1}{R} \sum_{k=1}^R \delta(\mathbf{V}_j - \mathbf{e}_k), \quad (19)$$

where \mathbf{U}_i and \mathbf{V}_j give the i -th row of U and j -th row of V respectively. Each of the R columns of U describe the mean of a N -variate i.i.d. Gaussian, and V has the role of picking one of these Gaussians. Finally, each column of Y is given by the chosen column of U plus Gaussian noise. In clustering, these are the data points, and the objective is to figure out the position of the centroids as well as the label assignment, given by the columns of U and the rows of V respectively. In the streaming setting the columns of the matrix Y are arriving in mini-batches. The offline AMP algorithm, its state evolution, and corresponding proofs are known for matrix factorization from [17, 19, 23, 26, 37, 40]. The Mini-AMP is obtained by adjusting the update of the estimators using (12).

In GMM clustering with prior on U having zero mean, there is an interesting "undetectability" phase transition for $R \leq 4$. If the number of samples is such that $\alpha = M/N < \alpha_c = R^2 \Delta^2$, then the Bayes optimal posterior asymptotically does not contain any information about the ground truth parameters [19]. This transition survives even when $R > 4$, in the sense that AMP and other tractable algorithms are unable to find any information on the ground truth parameters.

In the streaming problem, this undetectability implies that for mini-batches of relative size $\alpha_b < \alpha_c$, Mini-AMP does not improve the error of the random estimator, no matter the number of mini-batches presented. In particular, the fully online algorithm does not provide *any* useful output in this scenario. If $\alpha_b > \alpha_c$, on the other hand, an accurate reconstruction of the unknown values becomes possible. We illustrate the MSE as a function of the mini-batch size in Figure 3.

While we have presented Mini-AMP as a means for a theoretical analysis, it can be applied to real data, performing concrete learning tasks. To illustrate its efficacy, we have considered the classical problem of unsupervised clustering using the GMM. In Figure 3, Mini-AMP is shown to obtain better performance for real data clustering than mini-batch K-means, a state-of-the-art algorithm for streaming clustering [39].

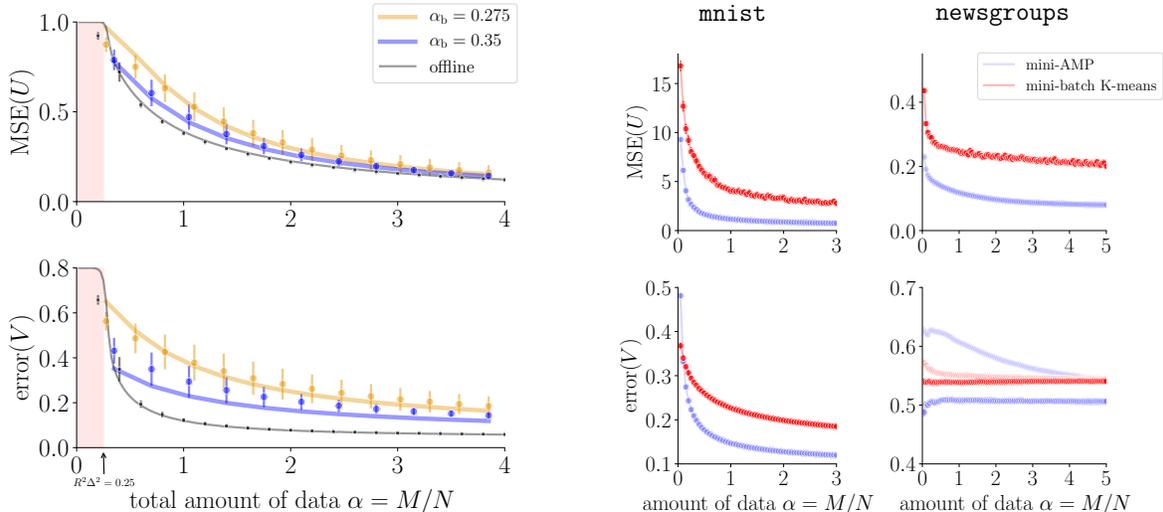


Figure 3: Clustering with Mini-AMP on synthetic Gaussian mixture data (left) and real-world data (right). *Left*: mean-squared error in U (centroids) and 0-1 loss in V (labels) using different batch sizes. Solid lines give state evolution and symbols give averages over 100 instances of size $N = 1000$. A transition at $\alpha_c = R^2 \Delta^2 = 0.25$ prevents Mini-AMP from giving non-zero overlap when $\alpha_b < \alpha_c$. Parameters are set to $R = 5$, $\Delta = 0.1$. *Right*: clustering on MNIST and the 20 newsgroups dataset using Mini-AMP for model (19), with the prior on U replaced with a non-negative Gaussian of mean zero and variance 0.1, and the noise variance Δ estimated from the data. On MNIST, digits of size $N = 784$ were clustered on $K = 3$ classes (0, 1 and 2), whereas for the 20 newsgroups dataset, frequency statistics of $N = 1000$ words and $K = 3$ top-level hierarchies (`comp`, `rec`, and `sci`) were used. Batch sizes were set so that $\alpha_b = 0.05$. Blue/red circles give the cumulative performance of Mini-AMP and mini-batch K-means [39], respectively, averaged over a 100 different orders of presentation. For the newsgroups dataset, a 2nd pass over the data was performed so that all labels could be recomputed with accurate estimation of the centroids; results of 1st and 2nd passes are shown in light/dark lines respectively.

7 Conclusion

Let us conclude by stating that the Mini-AMP algorithm can be applied to any problem for which the streaming can be defined and for which offline AMP exists. Therefore, we expect that this novel development will improve the usefulness of AMP algorithms in more practical situations.

Acknowledgments

This work has been supported by the ERC under the European Union’s FP7 Grant Agreement 307087-SPARCS. AM thanks Paulo V. Rossi and Thibault Lesieur for insightful discussions.

References

- [1] L. Bottou. Large-scale machine learning with stochastic gradient descent. In *Proc. COMPSTAT*, pages 177–186, 2010.
- [2] M. D. Hoffman, D. M. Blei, C. Wang, and J. Paisley. Stochastic variational inference. *J. Machine Learning Research*, 14(1):1303–1347, 2013.
- [3] M. Opper. A Bayesian approach to online learning. In D. Saad, editor, *On-line learning in Neural Networks*, pages 363–378. Cambridge University Press, 1998.

- [4] T. P. Minka. Expectation propagation for approximate Bayesian inference. In *Proc. Conf. on Uncertainty in Artificial Intelligence*, pages 362–369, 2001.
- [5] I. Mitliagkas, C. Caramanis, and P. Jain. Memory limited, streaming pca. In *Adv. in Neural Info. Proc. Systems*, 2013.
- [6] C. Wang and Y. M. Lu. Online learning for sparse PCA in high dimensions: Exact dynamics and phase transitions. In *Proc. IEEE Info. Theory Workshop*, pages 186–190, 2016.
- [7] O. Kinouchi and N. Caticha. Optimal generalization in perceptrons. *Journal of Physics A*, 25(23):6243, 1992.
- [8] M. Biehl and P. Riegler. On-line learning with a perceptron. *EPL*, 28(7):525, 1994.
- [9] S. Solla and O. Winther. Optimal perceptron learning: as online Bayesian approach. In D. Saad, editor, *On-line Learning in Neural Networks*, pages 379–398. Cambridge, 1998.
- [10] D. Saad. *On-line learning in neural networks*, volume 17. Cambridge, 1999.
- [11] P. V. Rossi, Y. Kabashima, and J. Inoue. Bayesian online compressed sensing. *Phys. Rev. E*, 94(2):022137, 2016.
- [12] T. Broderick, N. Boyd, A. Wibisono, A. C. Wilson, and M. I. Jordan. Streaming variational Bayes. In *Adv. in Neural Info. Proc. Systems*, pages 1727–1735, 2013.
- [13] D. L. Donoho, A. Maleki, and A. Montanari. Message-passing algorithms for compressed sensing. *Proc. Nat. Acad. Sci.*, 106(45):18914–18919, 2009.
- [14] S. Rangan. Generalized approximate message passing for estimation with random linear mixing. In *Proc. IEEE Int. Symp. on Info. Theory*, pages 2168–2172, 2011.
- [15] L. Zdeborová and F. Krzakala. Statistical physics of inference: Thresholds and algorithms. *Advances in Physics*, 65(5):453–552, 2016.
- [16] M. Mézard and A. Montanari. *Information, Physics, and Computation*. Oxford, 2009.
- [17] S. Rangan and A. K. Fletcher. Iterative estimation of constrained rank-one matrices in noise. In *IEEE Int. Symp. on Info. Theory*, pages 1246–1250, 2012.
- [18] F. Krzakala, M. Mézard, F. Sausset, Y. Sun, and L. Zdeborová. Probabilistic reconstruction in compressed sensing: algorithms, phase diagrams, and threshold achieving matrices. *J. of Statistical Mechanics: Theory and Experiment*, (08):P08009, 2012.
- [19] T. Lesieur, C. De Bacco, J. Banks, F. Krzakala, C. Moore, and L. Zdeborová. Phase transitions and optimal algorithms in high-dimensional Gaussian mixture clustering. *Proc. Allerton Conf. on Comm., Control, and Computing*, pages 601–608, 2016.
- [20] T. Lesieur, L. Miolane, M. Lelarge, F. Krzakala, and L. Zdeborová. Statistical and computational phase transitions in spiked tensor estimation. arXiv Preprint [math.ST]:1701.08010, 2017.
- [21] M. Bayati and A. Montanari. The Dynamics of Message Passing on Dense Graphs, with Applications to Compressed Sensing. *IEEE Trans. on Inf. Th.*, 57(2):764–785, February 2011.
- [22] M. Bayati, M. Lelarge, and A. Montanari. Universality in polytope phase transitions and message passing algorithms. *Annals of Applied Probability*, 25(2):753–822, 2015.
- [23] J. Barbier, M. Dia, N. Macris, F. Krzakala, T. Lesieur, and L. Zdeborová. Mutual information for symmetric rank-one matrix estimation: A proof of the replica formula. In *Adv. in Neural Info. Proc. Systems*, 2016.

- [24] J. Barbier, N. Macris, M. Dia, and F. Krzakala. Mutual information and optimality of approximate message-passing in random linear estimation. *arXiv Preprint [cs.IT]:1701.05823*, 2017.
- [25] G. Reeves and H. D. Pfister. The replica-symmetric prediction for compressed sensing with Gaussian matrices is exact. In *IEEE Int. Symp. on Info. Theory*, pages 665–669, 2016.
- [26] L. Miolane. Fundamental limits of low-rank matrix estimation. *arXiv Preprint [math.PR]:1702.00473*, 2017.
- [27] E. Gardner and B. Derrida. Three unfinished works on the optimal storage capacity of networks. *J. of Phys. A: Mathematical and General*, 22(12):1983, 1989.
- [28] G. Györfyi. First-order transition to perfect generalization in a neural network with binary synapses. *Phys. Rev. A*, 41:7097–7100, 1990.
- [29] E. J. Candès, J. K. Romberg, and T. Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59(8):1207–1223, 2006.
- [30] S. Rangan, P. Schniter, and A. Fletcher. Vector approximate message passing. *arXiv Preprint [cs.IT]:1601.03082*, 2016.
- [31] B. Çakmak, O. Winther, and B. H. Fleury. S-amp: Approximate message passing for general matrix ensembles. In *IEEE Info. Theory Workshop*, pages 192–196. IEEE, 2014.
- [32] E. Gardner and B. Derrida. Optimal storage properties of neural network models. *J. of Phys. A: Mathematical and General*, 21(1):271, 1988.
- [33] F. Krzakala, J. Xu, and L. Zdeborová. Mutual information in rank-one matrix estimation. In *IEEE Info. Theory Workshop*, pages 71–75, 2016.
- [34] Y. Kabashima, F. Krzakala, M. Mézard, A. Sakata, and L. Zdeborová. Phase transitions and sample complexity in bayes-optimal matrix factorization. *IEEE Trans. on Info. Theory*, 62:4228–4265, 2016.
- [35] M. Talagrand. *Spin glasses: a challenge for mathematicians: cavity and mean field models*, volume 46. Springer Science & Business Media, 2003.
- [36] M. Lelarge and L. Miolane. Fundamental limits of symmetric low-rank matrix estimation. *arXiv Preprint [math.PR]:1611.03888*, 2016.
- [37] T. Lesieur, F. Krzakala, and L. Zdeborová. Constrained low-rank matrix estimation: Phase transitions, approximate message passing and applications. *arXiv Preprint [math.ST]:1701.00858*, 2017.
- [38] F. Krzakala, A. Manoel, E. W. Tramel, and L. Zdeborová. Variational free energies for compressed sensing. In *IEEE Int. Symp. on Info. Theory*, pages 1499–1503, 2014.
- [39] D. Sculley. Web-scale k-means clustering. In *Proc. Int. Conf. on World Wide Web*, pages 1177–1178, 2010.
- [40] R. Matsushita and T. Tanaka. Low-rank matrix reconstruction and clustering via approximate message passing. In *Adv. in Neural Info. Proc. Systems*, 2013.
- [41] T. Lesieur, F. Krzakala, and L. Zdeborová. MMSE of probabilistic low-rank matrix estimation: Universality with respect to the output channel. In *Proc. Allerton Conf. on Communication, Control, and Computing*, pages 680–687, 2015.
- [42] D. Arthur and S. Vassilvitskii. k-means++: The advantages of careful seeding. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1027–1035. Society for Industrial and Applied Mathematics, 2007.

- [43] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

A AMP equations for different classes of models

We present here the AMP equations for different models. As before, adapting them to the streaming setting is done by introducing $\mathbf{\Lambda}$, $\mathbf{\Theta}$ variables and replacing the $\eta(A, B)$ function with (12).

A.1 Generalized linear models

Denote by $\mathbf{y} \in \mathbb{R}^M$ the response variable, by $\Phi \in \mathbb{R}^{M \times N}$ the design matrix, and by $\mathbf{x} \in \mathbb{R}^N$ the parameter vector that we want to estimate. Then our generative model reads

$$P(\mathbf{y}, \mathbf{x} | \Phi) = \prod_{\mu=1}^M P(y_\mu | z_\mu \equiv \Phi_\mu \cdot \mathbf{x}) \prod_{i=1}^N P_X(x_i). \quad (20)$$

The GAMP algorithm provides the following approximation to the marginals of \mathbf{x}

$$q_x(x_i | A_i, B_i) = \frac{1}{Z_x(A_i, B_i)} P_X(x_i) e^{-\frac{1}{2} A_i x_i^2 + B_i x_i}, \quad (21)$$

and, to the marginals of \mathbf{z}

$$q_z(z_\mu | y_\mu, \omega_\mu, V_\mu) = \frac{1}{Z_z(y_\mu, \omega_\mu, V_\mu)} P(y_\mu | z_\mu) \frac{e^{-\frac{(z_\mu - \omega_\mu)^2}{2V_\mu}}}{\sqrt{2\pi V_\mu}}. \quad (22)$$

The parameters \mathbf{A} , \mathbf{B} , $\boldsymbol{\omega}$ and \mathbf{V} are determined by iterating the GAMP equations. We denote the mean and variance of $q_x(A, B)$ by $\eta(A, B) = \frac{\partial}{\partial B} \log Z_x(A, B)$ and $\eta'(A, B) = \frac{\partial^2}{\partial B^2} \log Z_x(A, B)$ respectively; moreover, we define $g_{\text{out}}(y, \omega, V) = \frac{\partial}{\partial \omega} \log Z_z(y, \omega, V)$. The GAMP equations then read [14, 15]

$$\begin{aligned} \boldsymbol{\omega}^{(t)} &= \Phi \hat{\mathbf{x}}^{(t)} - \mathbf{V}^{(t)} \circ \mathbf{g}^{(t-1)}, & \mathbf{V}^{(t)} &= (\Phi \circ \Phi) \hat{\boldsymbol{\sigma}}^{(t)}, \\ g_\mu^{(t)} &= g_{\text{out}}(y_\mu, \omega_\mu^{(t)}, V_\mu^{(t)}) \forall \mu, & \partial_\omega g_\mu^{(t)} &= \partial_\omega g_{\text{out}}(y_\mu, \omega_\mu^{(t)}, V_\mu^{(t)}) \forall \mu, \\ \mathbf{B}^{(t)} &= \Phi^T \mathbf{g}^{(t)} + \mathbf{A}^{(t)} \circ \hat{\mathbf{x}}^{(t)}, & \mathbf{A}^{(t)} &= -(\Phi \circ \Phi)^T \partial_\omega \mathbf{g}^{(t)}, \\ \hat{x}_i^{(t+1)} &= \eta(A_i^{(t)}, B_i^{(t)}) \forall i, & \hat{\sigma}_i^{(t+1)} &= \eta'(A_i^{(t)}, B_i^{(t)}) \forall i. \end{aligned} \quad (23)$$

Note that, for a Gaussian likelihood $P(y|z) = \mathcal{N}(y; z, \Delta)$, $g_{\text{out}}(y, \omega, V) = \frac{y - \omega}{\Delta + V}$. Then, by defining $\tilde{\mathbf{z}}^{(t)} = \mathbf{y} - \boldsymbol{\omega}^{(t)}$ and replacing $\mathbf{V}^{(t)}$ and $\mathbf{A}^{(t)}$ by its averages, we get via the central limit theorem

$$\mathbf{V}^{(t)} \approx \frac{1}{N} \sum_{i=1}^N \sigma^{(t)}(A_i^{(t)}, B_i^{(t)}), \quad \mathbf{A}^{(t)} \approx \frac{\boldsymbol{\alpha}}{\Delta + V^{(t)}}, \quad (24)$$

where we have assumed the $\Phi_{\mu i}$ are i.i.d. and have zero mean and variance $1/N$. From these (5)-(7) follow through.

A.1.1 Variational Bayes

We compare AMP equations to the Variational Bayes (VB) ones, which we use with the Streaming Variational Bayes scheme. For simplicity we restrict ourselves to the Gaussian case. As usual, VB is derived by determining the $q_i(x_i)$ which minimize

$$KL[\prod_i q_i(x_i) \| P(\mathbf{x} | \Phi, \mathbf{y})] = -\mathbb{E}_{\{q_i\}} \log P(\mathbf{y} | \Phi, \mathbf{x}) + \sum_{i=1}^N KL[q_i(x_i) \| P_X(x_i)] - \log P(\mathbf{y} | \Phi) \quad (25)$$

If done by means of a fixed-point iteration, this minimization leads to [38]

$$\hat{x}_i^{(t+1)} = \eta \left(\frac{1}{\Delta} \sum_{\mu} \Phi_{\mu i}^2, \frac{1}{\Delta} \sum_{\mu} \Phi_{\mu i} (y_{\mu} - \sum_{j \neq i} \Phi_{\mu j} \hat{x}_j^{(t)}) \right) \forall i, \quad (26)$$

where η is defined as before. A closer inspection shows that the same equations are obtained by setting $\mathbf{V}^{(t)} = 0$ in (23).

As shown in [38], this iteration leads to good results when performed sequentially, *if* the noise Δ is not fixed but learned. We employ this same strategy in our experiments.

A.1.2 Assumed Density Filtering

The assumed density filtering (ADF) algorithm [3, 4] replaces the posterior at each step k , $P(\mathbf{x}|\Phi_k, y_k) \propto P(y_k|\Phi_k, \mathbf{x}) Q_{k-1}(\mathbf{x})$, by the distribution $Q_k(\mathbf{x})$ that minimizes

$$\text{KL}[P(\mathbf{x}|\Phi_k, y_k)||Q_k(\mathbf{x})]. \quad (27)$$

Note this is the *direct* KL divergence, and not the *reverse* one $\text{KL}[Q(\mathbf{x})||P(\mathbf{x}|\Phi, \mathbf{y})]$ that is minimized in Variational Bayes. In particular if $Q_k(\mathbf{x}) = \prod_{i=1}^N q_k(x_i)$, then minimizing this KL divergence leads to the following integral

$$q_k(x_i) = q_{k-1}(x_i) \int \left[\prod_{j \neq i} dx_j q_{k-1}(x_j) \right] P(y_k|\Phi_k, \mathbf{x}) \forall i, \quad (28)$$

which is tractable since we are processing a single sample y_k at a time, i.e. since the likelihood consists of a single factor. These are actually the exact marginals of $P(\mathbf{x}|\Phi_k, y_k)$, and also the equations given by the belief propagation (BP) algorithm in the single sample limit.

The ADF equations for GLMs can be derived by using (28) together with the central limit theorem [11]. Because BP gives ADF in the single sample limit, AMP (which is based on BP) gives the equations derived by [11] when $M = 1$, if one additionally neglects the correction term $-\mathbf{V}^{(t)} \mathbf{g}^{(t-1)}$ (analogously, if a single iteration is performed).

Algorithm 2 Assumed Density Filtering for GLMs [11]

- 1: initialize $\Lambda_{0,i} = 0, \Theta_{0,i} = 0 \forall i$
 - 2: **for** $k \leftarrow 1$ **to** M **do**
 - 3: compute $\omega_k = \Phi_k \cdot \hat{\mathbf{x}}_{k-1}, V_k = (\Phi_k \circ \Phi_k) \cdot \hat{\sigma}_{k-1}$
 - 4: compute $g_k, \partial_{\omega} g_k$ following (23)
 - 5: compute $\mathbf{A}_k, \mathbf{B}_k$ following (23)
 - 6: compute $\hat{\sigma}_k, \hat{\mathbf{x}}_k$ following (12)
 - 7: accumulate $\Lambda_k \leftarrow \Lambda_{k-1} + \mathbf{A}_k$
 - 8: accumulate $\Theta_k \leftarrow \Theta_{k-1} + \mathbf{B}_k$
 - 9: **end for**
-

A.2 Low-rank matrix factorization

Denote by $Y \in \mathbb{R}^{N \times M}$ the matrix we want to factorize, and by $U \in \mathbb{R}^{N \times R}, V \in \mathbb{R}^{M \times R}$ the matrices which product approximates Y . The generative model then reads

$$P(Y, U, V) = \prod_{ij} P(Y_{ij}|W_{ij} \equiv \mathbf{U}_i \cdot \mathbf{V}_j) \prod_{i=1}^N P_U(\mathbf{U}_i) \prod_{j=1}^M P_V(\mathbf{V}_j), \quad (29)$$

where \mathbf{U}_i and \mathbf{V}_j denote the i -th and j -th rows of U and V respectively. The algorithm provides the following approximation to the marginal of \mathbf{U}_i

$$q_U(\mathbf{U}_i|A_U, \mathbf{B}_{U,i}) = \frac{1}{Z_U(A_U, \mathbf{B}_{U,i})} P_U(\mathbf{U}_i) e^{-\frac{1}{2}\mathbf{U}_i^T A_U \mathbf{U}_i + \mathbf{B}_{U,i}^T \mathbf{U}_i}. \quad (30)$$

and q_V is analogously defined as the marginal of \mathbf{V}_j . As in the previous case, $A_U \in \mathbb{R}^{R \times R}$ and $\mathbf{B}_{U,i} \in \mathbb{R}^R$ are to be determined by iterating a set of equations. We denote by B_U the $N \times R$ matrix which rows are given by $\mathbf{B}_{U,i}$, $i = 1, \dots, N$. The functions $\eta_U(A, \mathbf{B}) = \mathbb{E}_{q_U} \mathbf{U} = \nabla_{\mathbf{B}} \log Z_U(A, \mathbf{B})$ and $\eta'_U(A, \mathbf{B}) = \nabla_{\mathbf{B}} \eta_U(A, \mathbf{B})$ give the mean and covariance of q_U , and η_V and η'_V , defined analogously, the mean and covariance of q_V .

In order to write the AMP equations, we first introduce

$$J_{ij} = \frac{1}{\sqrt{N}} \frac{\partial \ln P(Y_{ij}|w=0)}{\partial w}, \quad (31)$$

$$\beta = \frac{1}{N} \mathbb{E}_{P(y|w=0)} \left[\frac{\partial \ln P(y|w=0)}{\partial w} \right]^2.$$

so as to define an effective Gaussian channel [41]. The equations to be iterated are then

$$\begin{aligned} B_U^{(t)} &= J \hat{V}^{(t)} - \beta \Sigma_V^{(t)} \hat{U}^{(t-1)}, & A_U^{(t)} &= \beta \hat{V}^{(t)} \hat{V}^{(t)T}, \\ \hat{U}_i^{(t)} &= \eta(A_U^{(t)}, \mathbf{B}_{U,i}^{(t)}) \forall i, & \Sigma_U^{(t)} &= \sum_{i=1}^N \eta'(A_U^{(t)}, \mathbf{B}_{U,i}^{(t)}), \\ B_V^{(t)} &= J^T \hat{U}^{(t)} - \beta \Sigma_U^{(t)} \hat{V}^{(t)}, & A_V^{(t)} &= \beta \hat{U}^{(t)} \hat{U}^{(t)T}, \\ \hat{V}_j^{(t+1)} &= \eta(A_V^{(t)}, \mathbf{B}_{V,j}^{(t)}) \forall j, & \Sigma_V^{(t+1)} &= \sum_{j=1}^M \eta'(A_V^{(t)}, \mathbf{B}_{V,j}^{(t)}). \end{aligned} \quad (32)$$

In order to adapt this algorithm to the online setting, we repeat procedure (12) and, as the k -th batch is processed, replace calls to $\eta_U(A_U, \mathbf{B}_U)$ by

$$\eta_U \left(\underbrace{\sum_{\ell=1}^{k-1} A_{U,\ell}}_{\Lambda_{k-1}} + A_{U,k}^{(t)}, \underbrace{\sum_{\ell=1}^{k-1} \mathbf{B}_{U,\ell} + \mathbf{B}_{U,k}^{(t)}}_{\Theta_{k,i}} \right). \quad (33)$$

We assume that U is fixed and V_k changes for each batch k ; thus, the calls to η_V do not change.

B State evolution and asymptotic limits

Through the state evolution equations, we analyze the behaviour and performance of the algorithms described in the previous section. We restrict ourselves to the Bayes-optimal case (i.e. the Nishimori line), where the generative model is known. The strategy we use to go from the offline to the streaming setting is easily adapted to the non-optimal case.

B.1 Generalized linear models

The state evolution equations for a GLM with likelihood $P(y|z)$ and prior $P_X(x)$ are

$$\begin{cases} \hat{m}^{(t)} = -\alpha \mathbb{E}_{y,z,w} \partial_w g(y, w, \rho - m^{(t)}), \\ m^{(t+1)} = \mathbb{E}_{x,b} x \eta(\hat{m}^{(t)}, b), \end{cases} \quad (34)$$

where we denote $\rho = \mathbb{E}x^2$, and the averages are taken with respect to $P(x, b) = P_X(x)\mathcal{N}(b; \hat{m}x, \hat{m})$ and $P(y, \omega, z) = P(y|z)\mathcal{N}(z; \omega, \rho - m)\mathcal{N}(\omega; 0, m)$. The MSE at each step is obtained from $\mathcal{E}^{(t)} = \rho - m^{(t)}$. For a Gaussian likelihood $P(y|z) = \mathcal{N}(y; z, \Delta)$, $\hat{m}^{(t)} = \frac{\alpha}{\Delta + \mathcal{E}^{(t)}}$ and we recover (9).

The fixed points of the state evolution extremize the so-called replica free energy

$$\phi(m, \hat{m}) = \frac{1}{2}m\hat{m} - \mathbb{E}_{b,x} \log Z_x(\hat{m}, b) - \alpha \mathbb{E}_{y,\omega,z} \log Z_z(y, \omega, \rho - m) \quad (35)$$

which gives the large system limit of the Bethe free energy (extremized by AMP). The mutual information (10) differs from $\phi(m) = \text{extr}_{\hat{m}} \phi(m, \hat{m})$ by a constant – more specifically, by the entropy of $P(y|z)$, $i_{\text{RS}}(m) = \phi(m) - \alpha H[P(y|z)]$. For a Gaussian likelihood, $H[P(y|z)] = \frac{1}{2} \log(2\pi e \Delta)$.

In order to adapt this to the streaming setting, we introduce $\lambda_k^{(t)} = \sum_{\ell=1}^{k-1} \hat{m}_\ell^{(t_{\max})} + \hat{m}_k^{(t)}$ and iterate instead, for each mini-batch k

$$\begin{cases} \lambda_k^{(t)} = \lambda_{k-1} - \alpha_b \mathbb{E}_{y,z,w} \partial_w g(y, w, \rho - m_k^{(t)}), \\ m_k^{(t+1)} = \mathbb{E}_{x,b_k} x \eta(\lambda_k^{(t)}, b), \end{cases} \quad (36)$$

with the averages now computed over $P(x, b_k) = P_X(x)\mathcal{N}(b_k; \lambda_k x, \lambda_k)$. These equations should be iterated for $t = 1, \dots, t_{\max}$, at which point we assign $\lambda_k = \lambda_{k-1}^{(t_{\max})}$. The MSE on \mathbf{x} after mini-batch k is processed is then given by $\mathcal{E}_k = \rho - m_k^{(t_{\max})}$.

Note that in the small batch size limit ($\alpha_b \rightarrow 0$), the equation for λ becomes an ODE

$$\frac{d\lambda}{d\alpha} = -\mathbb{E}_{y,z,w} \partial_w g(y, w, \rho - m(\lambda)), \quad (37)$$

which describes the performance of the ADF algorithm [3, 11].

The free energy is also easily rewritten

$$\phi_k(m_k, \hat{m}_k; \lambda_{k-1}) = \frac{1}{2}m_k\hat{m}_k - \mathbb{E}_{b,x} \log Z_x(\lambda_{k-1} + \hat{m}_k, b) - \alpha_b \mathbb{E}_{y,\omega,z} \log Z_z(y, \omega, \rho - m_k), \quad (38)$$

or analogously, by working with $\lambda_k = \lambda_{k-1} + \hat{m}_k$ instead of \hat{m}_k

$$\phi_k(m_k, \lambda_k; \lambda_{k-1}) = \frac{1}{2}m_k(\lambda_k - \lambda_{k-1}) - \mathbb{E}_{b,x} \log Z_x(\lambda_k, b) - \alpha_b \mathbb{E}_{y,\omega,z} \log Z_z(y, \omega, \rho - m_k). \quad (39)$$

from which it is clear that the extrema of ϕ_k are given by the fixed points of (36).

B.1.1 Asymptotic behavior

Equations (36) can be put in the following form

$$\begin{cases} \mathcal{E}_k = \varepsilon(\lambda_k), \\ \lambda_k = \lambda_{k-1} + \alpha_b \delta(\mathcal{E}_k), \end{cases} \quad (40)$$

where $\varepsilon(\lambda_k)$ and $\delta(\mathcal{E}_k)$ are functions that depend on the prior/channel respectively. Assuming ε is invertible, we rewrite this system of equations as a function of \mathcal{E}_k only

$$\varepsilon^{-1}(\mathcal{E}_k) = \varepsilon^{-1}(\mathcal{E}_{k-1}) + \alpha_b \delta(\mathcal{E}_k), \quad (41)$$

and then solve this equation for \mathcal{E}_k ; that gives us a recurrence relation which is unsolvable in most cases. We use instead asymptotic forms for ε and δ , obtained in the $\lambda \rightarrow \infty$, $\mathcal{E} \rightarrow 0$ limit. For the Bernoulli-Gaussian prior $P_0(x_i) = \rho \mathcal{N}(x_i; 0, 1) + (1 - \rho) \delta(x_i)$, we have

$$\varepsilon(\lambda) \sim \frac{\rho}{\lambda}, \quad (42)$$

while a Gaussian likelihood gives, in the $\Delta \rightarrow 0$ limit, $\delta(\mathcal{E}) = \frac{1}{\mathcal{E}}$. Thus, for SLR

$$\frac{\rho}{\mathcal{E}_k} \approx \frac{\rho}{\mathcal{E}_{k-1}} + \frac{\alpha_b}{\mathcal{E}_k} \Rightarrow \mathcal{E}_k \approx \left(1 - \frac{\alpha_b}{\rho}\right)^k \mathcal{E}_0, \quad (43)$$

leading to (18). In the $\alpha_b \rightarrow 0$ limit we recover the expression obtained by [11], $\text{MSE}(\alpha) \sim e^{-\frac{\alpha}{\rho}}$.

B.2 Low-rank matrix factorization

Also for low-rank models the large N limit can be analyzed by taking into account that $A_{U,V}^{(t)}$ and $B_{U,V}^{(t)}$ converge in distribution to [41]

$$\begin{aligned} A_U^{(t)} &= \beta M_V^{(t)}, & \mathbf{B}_U &\sim \mathcal{N}(\beta M_V^{(t)} \mathbf{U}, \beta M_V^{(t)}), \\ A_V^{(t)} &= \beta M_U^{(t)}, & \mathbf{B}_V &\sim \mathcal{N}(\beta M_U^{(t)} \mathbf{V}, \beta M_U^{(t)}), \end{aligned} \quad (44)$$

where $M_U, M_V \in \mathbb{R}^{R \times R}$ are the overlap matrices between the ground truth and the estimate at time t , that is

$$\begin{aligned} M_U^{(t)} &= \mathbb{E}_{\mathbf{U}, \mathbf{B}_U} \mathbf{U}^T \eta_U(\beta M_V^{(t)}, \mathbf{B}_U), \\ M_V^{(t+1)} &= \alpha \mathbb{E}_{\mathbf{V}, \mathbf{B}_V} \mathbf{V}^T \eta_V(\beta M_U^{(t)}, \mathbf{B}_V). \end{aligned} \quad (45)$$

While computing the expectations might become unfeasible for $R > 1$, an ansatz of the following form can often be used [19, 41]

$$M_{U,V} = a_{U,V} \mathbb{I}_R + b_{U,V} \mathbb{J}_R, \quad (46)$$

with \mathbb{J}_R denoting the $R \times R$ matrix of ones. This significantly simplifies the iteration above.

Again we adapt this to the online case by incrementing the matrices obtained as each batch is processed, that is, we replace M_V by

$$\lambda_{V,k}^{(t+1)} = \lambda_{V,k-1} + \alpha_b \mathbb{E}_{\mathbf{V}, \mathbf{B}_V} \mathbf{V}^T \eta_V(\beta M_{U,k}^{(t)}, \mathbf{B}_V). \quad (47)$$

Note that since U is fixed, the equation for M_U does not change.

The replica free energy reads, in the offline case

$$\phi(M_U, M_V) = \frac{\beta}{2} \text{Tr} M_U M_V^T - \mathbb{E}_{\mathbf{U}, \mathbf{B}_U} \log Z_U(\beta M_V, \mathbf{B}_U) - \mathbb{E}_{\mathbf{V}, \mathbf{B}_V} \log Z_V(\beta M_U, \mathbf{B}_V), \quad (48)$$

and we adapt it to the online case by taking into account that λ_V is being incremented

$$\begin{aligned} \phi(M_{U,k}, \lambda_{V,k}; \lambda_{V,k-1}) &= \frac{\beta}{2} \text{Tr} M_{U,k} (\lambda_{V,k} - \lambda_{V,k-1}) - \\ &\mathbb{E}_{\mathbf{U}, \mathbf{B}_U} \log Z_U(\beta \lambda_{V,k}, \mathbf{B}_U) - \mathbb{E}_{\mathbf{V}, \mathbf{B}_V} \log Z_V(\beta M_{U,k}, \mathbf{B}_V). \end{aligned} \quad (49)$$

C Performance for different number of iterations

For our experiments in Figure 1, Mini-AMP has been iterated until, for each block, convergence is achieved – that is, until $\frac{1}{N} \|\hat{\mathbf{x}}^{(t)} - \hat{\mathbf{x}}^{(t-1)}\|_1 < 10^{-13}$. Remarkably, our framework allows us to study the performance of the algorithm even if we do not iterate it until convergence, but only for a few steps t_{\max} instead. In Figure 4, we investigate the performance of Mini-AMP under the same settings of Figure 1 (center), for different values of t_{\max} . We observe that the performance deteriorates if convergence is not reached.

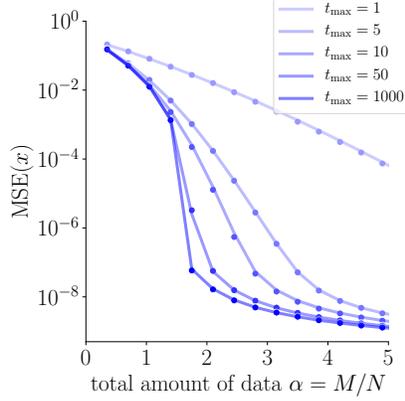


Figure 4: Performance of the Mini-AMP algorithm for different values of t_{\max} , under the same settings of Figure 1 (center), and $\alpha_b = 0.35$. Solid lines give state evolution, and symbols results of empirical experiments averaged over 10 realizations of size $N = 2000$. The performance deteriorates if the algorithm is not iterated until convergence.

D Experiments on real-world data

For the experiments with real data, we have used the following model

$$P(U, V, Y) = \prod_{ij} \mathcal{N}(Y_{ij}; \mathbf{U}_i \cdot \mathbf{V}_j, \Delta) \prod_{i=1}^N \mathcal{N}_{\geq 0}(\mathbf{U}_i; \mathbf{0}, \sigma^2 \mathbb{I}_R) \prod_{j=1}^M \frac{1}{R} \sum_{k=1}^R \delta(\mathbf{V}_j - \mathbf{e}_k), \quad (50)$$

where

$$\mathcal{N}_{\geq 0}(\mathbf{x}; \boldsymbol{\mu}, \sigma^2 \mathbb{I}_R) = \frac{1}{\mathcal{Z}(\boldsymbol{\mu}, \sigma^2)} \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \sigma^2 \mathbb{I}_R) \prod_{k=1}^R \theta(x_k) \quad (51)$$

is a truncated normal distribution supported on the positive quadrant of a R -dimensional space, and $\mathcal{Z}(\boldsymbol{\mu}, \sigma^2)$ ensures proper normalization.

Note that evaluating the $\eta(A, \mathbf{B})$ function in this case is not trivial, since it depends on the following integral

$$\begin{aligned} Z_U(A, \mathbf{B}) &= \frac{1}{\mathcal{Z}(\boldsymbol{\mu}, \sigma^2)} \int d\mathbf{x} e^{-\frac{1}{2} \mathbf{x}^T A \mathbf{x} + \mathbf{B}^T \mathbf{x}} \prod_{k=1}^R P_0(x_k) \\ &\propto \int \prod_{k=1}^R dx_k P_0(x_k) e^{-\frac{1}{2} A_{kk} x_k^2 + (B_k + \sum_{\ell \neq k} A_{k\ell} x_\ell) x_k} \end{aligned} \quad (52)$$

where, in this case, $P_0(x_k) = \mathcal{N}(x_k; \mu_k, \sigma^2) \theta(x_k)$. We proceed by performing a mean-field approximation. We first define

$$\tilde{\eta}(A, B) = \frac{\partial}{\partial B} \log \int dx P_0(x) e^{-\frac{1}{2} A x^2 + B x}. \quad (53)$$

for scalar A and B . Then, for each $i = 1, \dots, N$, we iterate

$$\hat{U}_{ik} = \tilde{\eta}(A_{kk}, B_{ik} - \frac{1}{2} \sum_{\ell \neq k} A_{k\ell} \hat{U}_{i\ell}) \quad (54)$$

sequentially in $k = 1, \dots, R$ until convergence is reached, at which point we use the values obtained for

assigning $\hat{U}^{(t)}$ in AMP. The variances are computed from

$$\tilde{\sigma}_{ik} = \tilde{\eta}' \left(A_{kk}, B_{ik} - \frac{1}{2} \sum_{\ell \neq k} A_{k\ell} \hat{U}_{i\ell} \right) \quad (55)$$

and the covariance matrix used in AMP is obtained as a function of these variances

$$\sigma_{k\ell} = \begin{cases} \sum_{i=1}^N \tilde{\sigma}_{ik}, & \text{if } k = \ell, \\ -\frac{1}{2} A_{k\ell} \sum_{i=1}^N \tilde{\sigma}_{ik} \tilde{\sigma}_{i\ell}, & \text{otherwise,} \end{cases} \quad (56)$$

where in order to assign the off-diagonal terms we have used a linear response approximation, $\sigma_{k\ell} = \sum_{i=1}^N \frac{\partial \hat{U}_{ik}}{\partial B_{i\ell}}$. We proceed by detailing other aspects of the experiments

Initialization At the first few mini-batches (usually the first five), we reinitialize the position of the centroids. We use the same strategy as the k-means++ algorithm [42]: the first centroid is picked at random from the data points, and the next ones are sampled so as to have them far apart from each other. The labels are initialized according to the closest centroid.

Stopping criterion For each batch, Mini-AMP was iterated either for 50 steps or until $\frac{1}{NR} \sum_{ik} |\hat{U}_{ik}^{(t)} - \hat{U}_{ik}^{(t-1)}| + \frac{1}{MR} \sum_{jk} |\hat{V}_{jk}^{(t)} - \hat{V}_{jk}^{(t-1)}| < 10^{-7}$.

Noise learning We do not assign a fixed value for Δ , but instead update it after each mini-batch is processed using a simple learning rule

$$\hat{\Delta}_k = \frac{1}{NM} \sum_{ij} (Y_{k,ij} - \hat{U}_{k,i} \cdot \mathbf{V}_{k,j})^2. \quad (57)$$

Preprocessing For MNIST, we work with all samples of digits 0, 1 and 2. They are rescaled so that the pixel intensities are between 0 and 1. For the 20 newsgroups dataset, we build Term Frequency Inverse Document Frequency (TF-IDF) features for 3 top-level hierarchies (`comp`, `rec` and `sci`), and use the 1000 most frequent words; we rescale each feature vector so that its maximum is equal to 1.

Mini-batch K-means We use the mini-batch K-means [39] implementation available on scikit-learn [43]. Default parameters were used, apart from the centroids initialization, which was set to random normal variables of zero mean and variance 10^{-3} – this seemed to improve the algorithm performance with respect to the standard choices.