



PAPER • OPEN ACCESS

Saddle hierarchy in dense associative memory

To cite this article: Robin Thériault and Daniele Tantari 2026 *Mach. Learn.: Sci. Technol.* **7** 015001

View the [article online](#) for updates and enhancements.

You may also like

- [Tunable magnetic properties in compositionally graded L₁-FePt thin films](#)
Isaak G Vasileiadis, Dimitrios Karfaridis, Nikolaos Maniotis et al.
- [ICRH modelling of DTT in full power and reduced-field plasma scenarios using full wave codes](#)
A Cardinali, C Castaldo, F Napoli et al.
- [Self-affine structure of glacial surfaces in Lombardy](#)
Marina Carpineti, Gabriele Tasca, Dorian Brogioli et al.



PAPER

Saddle hierarchy in dense associative memory

OPEN ACCESS

RECEIVED
27 August 2025REVISED
1 December 2025ACCEPTED FOR PUBLICATION
22 December 2025PUBLISHED
6 January 2026

Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](#).

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Robin Thériault^{1,*} and Daniele Tantari² ¹ Scuola Normale Superiore di Pisa, Piazza dei Cavalieri 7, 56126 Pisa (PI), Italy² Department of Mathematics, University of Bologna, Piazza di Porta San Donato 5, 40126 Bologna (BO), Italy

* Author to whom any correspondence should be addressed.

E-mail: robin.theriault@sns.it**Keywords:** fast optimization, statistical mechanics, machine learning theory, regularization, phase transitions**Abstract**

Dense associative memory (DAM) models have been attracting renewed attention since they were shown to be robust to adversarial examples and closely related to cutting edge machine learning paradigms, such as the attention mechanism and generative diffusion. We study a DAM built upon a three-layer Boltzmann machine with Potts hidden units, which represent data clusters and classes. Through a statistical mechanics analysis, we derive saddle-point equations that characterize both the stationary points of DAMs trained on real data and the fixed points of DAMs trained on synthetic data within a teacher-student framework. Based on these results, we propose a novel regularization scheme that makes training significantly more stable. Moreover, we show empirically that our DAM learns interpretable solutions to both supervised and unsupervised classification problems. Pushing our theoretical analysis further, we find that the weights learned by relatively small DAMs correspond to unstable saddle points in larger DAMs. We implement a network-growing algorithm that leverages this saddle-point hierarchy to drastically reduce the computational cost of training DAM.

1. Introduction

Studying the stationary points of machine learning algorithms is crucial to understand how they work. For example, [1, 2] demonstrated that local minima in the loss landscape of large artificial neural networks (NNs) are relatively close to the global minimum, explaining why they generalize well in practice. Moreover, [3, 4] showed that saddle points are much more numerous than local minima in large NNs. These breakthroughs were made by establishing deep connections between the loss landscape of NNs and the energy landscape of disordered systems studied in statistical mechanics. Beyond the broad insights provided by these studies, and despite the progress made by [5–8], the classification of stationary points in machine learning algorithms remains an open problem.

Energy-based models have been playing a central role in studies of NNs and their theoretical properties. The Hopfield network, one of the most historically important energy-based models, was originally introduced as a paradigmatic model of biological associative memory [9]. Generalized Hopfield networks [10, 11] were then developed to improve upon the limited storage capacity of the original [9, 12]. The scale of the improvement was determined rigorously in following studies [13–15]. A few years ago, these generalized networks, commonly referred to as dense associative memory (DAM) or modern Hopfield networks, were made into trainable machine learning models capable of accurate pattern classification by Krotov and Hopfield (K & H) [16]. In a nutshell, K & H's DAM learns prototypes of patterns in a trainable weight matrix. Each prototype casts a vote for a class, and the patterns awaiting classification are assigned based on the votes of the prototypes that most closely resemble them. The resulting classification scheme is considerably more adversarially robust [17, 18] and interpretable [16] than that of feedforward NNs with ReLU activation functions. Since their debut as trainable machine learning architectures, deep connections have been made between modern Hopfield networks and other well-known and established machine learning paradigms, such as attention [19] and generative diffusion [20, 21]. In particular, modern Hopfield networks were used to implement the attention mechanism of transformers

[19], which has been attracting a lot of interest in fundamental and applied research. See [22, 23] for a review of recent advances on the topic. Recently, it was observed that the trainable weights of K & H's DAM are channeled toward minima by a low-dimensional network of valleys in the loss landscape [24]. Moreover, the points where valleys branch out from one another were identified as saddles in the simple case where the DAM has two patterns to learn. In general, it is not straightforward to classify the stationary points of machine learning algorithms [8]. However, the results of [24] and the interpretability of DAMs suggest that their stationary points are both fundamental to their learning dynamics and easier to characterize than that of generic NNs. With this goal in mind, we revisit DAM for pattern classification [16] using the framework of Boltzmann machines (BMs) [25–28] and statistical mechanics.

The body of this paper is structured as follows. In section 2, we introduce the DAM model that we study and the analytical tools that we use throughout our work. In particular, in section 2.1, we derive the DAM in question from a BM template, and, in section 2.2, we explain the setting that we use to analyze DAM stationary points.

Section 3 presents our theoretical contributions. In section 3.1, we formulate saddle-point equations that characterize DAM stationary points, and, in section 3.2, we leverage these equations to establish a saddle-point hierarchy principle stating that weights learned by DAMs of a given width are embedded as saddle points in wider DAMs.

Section 4 explains how the theoretical insights of section 3 can be used to enhance training and explores how DAMs represent data once trained. In section 4.1, we introduce a regularization method that facilitates supervised learning. Next, in section 4.2, we show that our DAM, although designed for supervised tasks, can discover interpretable solutions in both supervised and unsupervised classification settings. Finally, in section 4.3, we link our findings to the learning dynamics shaped by valleys and saddles studied in [24], and we implement a network-growing algorithm [29] that exploits the saddle-point hierarchy to significantly reduce DAM training costs.

The code and hyperparameter configurations used in our experiments are available in the following public repository [30].

2. Model

A BM is a canonical graphical model of correlations in discrete data [25]. It is customary to partition BMs into a visible layer $\mathbf{v} = \{v_i\}_{i=1}^N \in \mathbb{R}^N$ and a hidden layer $\mathbf{h} = \{h_\mu\}_{\mu=1}^P \in \mathbb{R}^P$ such that connections between the two layers are allowed, but connections within them are prohibited [26]. In this case, the visible layer represents concrete features of the data, whose mutual correlations are encoded in connections with the hidden layer. The restricted BM (RBM) obtained using this partition is much easier to train than a generic BM [27, 28] and still has considerable generating power [27, 31], making it more practical in machine learning applications [32–34]. The visible and hidden units of an RBM follow the Gibbs distribution

$$P_\beta(\mathbf{v}, \mathbf{h} | \mathbf{J}) = Z_\beta(\mathbf{J})^{-1} P_0(\mathbf{v}) P_0(\mathbf{h}) \exp(-\beta H[\mathbf{v}, \mathbf{h}; \mathbf{J}]),$$

where $\beta \geq 0$ is known as the inverse temperature, $P_0(\mathbf{v})$ and $P_0(\mathbf{h})$ are priors on \mathbf{v} and \mathbf{h} , $H[\mathbf{v}, \mathbf{h}; \mathbf{J}] = -\sum_{i=1}^N \sum_{\mu=1}^P J_i^\mu v_i h_\mu$ is called the energy function or Hamiltonian, $\mathbf{J} = \{J_i^\mu\}_{1 \leq i \leq N, 1 \leq \mu \leq P}$ are trainable weights, and $Z_\beta(\mathbf{J})$ is a normalization constant called the partition function. The inverse temperature β represents the absolute strength of the RBM connections, or equivalently controls the amount of noise $T = 1/\beta$ in the RBM. In this regard, $P_0(\mathbf{v})$ and $P_0(\mathbf{h})$, which restrict the form of the Gibbs distribution to help the RBM represent the data, are the marginal laws of \mathbf{v} and \mathbf{h} when there are no connections, i.e. $\beta = 0$. Their contribution to the Gibbs distribution can be tuned with β , which can therefore be interpreted as a regularization parameter.

2.1. A DAM model

As mentioned in the Introduction, we will now derive a DAM model for classification from a BM template. We will explain why our model is a DAM at the end of this section, once we have clearly defined it. We make three basic assumptions on the distribution of data to be classified:

1. the data is scale invariant, i.e. for all positive scalars c , data points \mathbf{x} and $c\mathbf{x}$ are equivalent;
2. the data can be partitioned in disjoint clusters;
3. the clusters can be grouped into mutually exclusive classes.

In order to exploit these three assumptions, we study a BM partitioned into three layers with different roles: the data layer \mathbf{x} , the hidden layer \mathbf{h} and the class layer \mathbf{q} , which represent data, cluster membership and class membership, respectively. The corresponding energy is

$$-H[\mathbf{x}, \mathbf{q}, \mathbf{h}; \mathbf{J}] = \sum_{i=1}^N \sum_{\mu=1}^P w_i^\mu x_i h_\mu + \sum_{y=1}^C \sum_{\mu=1}^P u_y^\mu q_y h_\mu + \sum_{\mu=1}^P h_\mu b^\mu, \quad (1)$$

where $\mathbf{J} = \{\mathbf{w}, \mathbf{u}, \mathbf{b}\}$ is the set of the trainable weights $\mathbf{w} = \{w_i^\mu\}_{1 \leq i \leq N, 1 \leq \mu \leq P}$, $\mathbf{u} = \{u_y^\mu\}_{1 \leq y \leq C, 1 \leq \mu \leq P}$ and $\mathbf{b} = \{b^\mu\}_{\mu=1}^P$. There are no direct interactions between the visible layer and the class layer. In other words, conditional on the cluster layer, the visible layer and the class layer are independent. Therefore, this BM is a deep BM (DBM) with three layers [35], which can also be thought of as an RBM whose visible layer \mathbf{v} is further divided into \mathbf{x} and \mathbf{q} .

Since the data is scale invariant (assumption 1), the scale of individual data points contains no information about the classification task, so we normalize them by their (Euclidean) norm in the data layer. In other terms, we take the data units x_i to be continuous variables with unit norm $\sqrt{\sum_{i=1}^N (x_i)^2} = 1$. We assume no further knowledge about \mathbf{x} , so we take the prior $P_0(\mathbf{x})$ to be the uniform distribution on the $N - 1$ dimensional unit hypersphere S^{N-1} , i.e. is the set of all \mathbf{x} with $\sqrt{\sum_{i=1}^N (x_i)^2} = 1$. Data normalization is a very common practice in machine learning. For example, normalization by the Euclidean norm has been popular in text document clustering even since its introduction in the 1980 s [36]. Various types of normalization also occur in the brain and retina [37].

Since the hidden layer and the class layer aim to represent disjoint clusters and classes, respectively (assumptions 2 and 3), we take their respective units to be mutually exclusive binary variables, i.e. $\mathbf{h} \in \{0, 1\}^P$ with $\sum_{\mu=1}^P h_\mu \in \{0, 1\}$ and $\mathbf{q} \in \{0, 1\}^C$ with $\sum_{y=1}^C q_y \in \{0, 1\}$. At any given time, at most one unit per layer can take the value 1, representing the fact that the clusters and classes are disjoint. In other words, we take each of the two layers to be the vector representation of a single categorical (or Potts [38, 39]) variable with $P + 1$ and $C + 1$ categories, respectively. As such, $P_0(\mathbf{h})$ and $P_0(\mathbf{q})$ simplify to probability mass functions $P_0(\mathbf{h} = \mathbf{e}_\gamma)$ and $P_0(\mathbf{q} = \mathbf{e}_y)$, where we introduce $\mathbf{e}_\gamma = \{\delta_{\gamma\mu}\}_{\mu=1}^P$ for $\gamma \in \{0, \dots, P\}$ and define $\mathbf{e}_y \in \{0, 1\}^C$ analogously for $y \in \{0, \dots, C\}$. In particular $\mathbf{e}_0 = \mathbf{0}$ represents a state outside the P clusters or the C classes.

These priors on the hidden layer and the class layer can also be obtained by introducing fixed inhibitory connections within the hidden layer and the class layer, respectively [40, 41]. Since at most one hidden unit h_μ can be activated at once, the hidden layer is a very sparse representation of the visible layer. In machine learning, sparsity can improve interpretability [42], generalization, computational efficiency [43], and adversarial robustness [44]. The sparsity of the brain suggests that it is also beneficial for biological NNs [45].

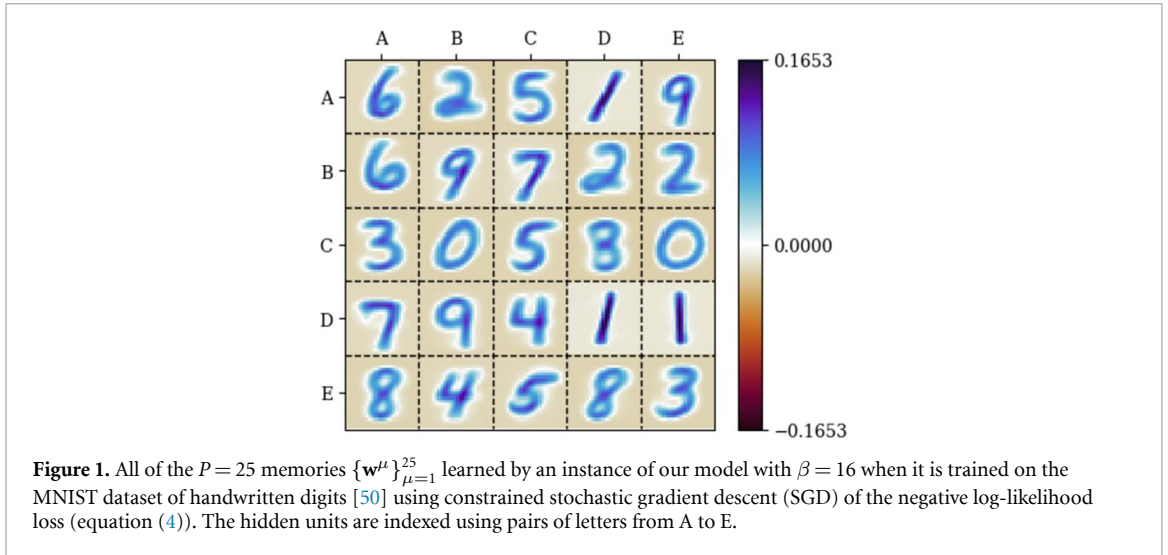
Given these priors $P_0(\mathbf{x})$, $P_0(\mathbf{h})$ and $P_0(\mathbf{q})$, we derive the marginal distribution of the visible layer (\mathbf{x}, \mathbf{q}) (see appendix A for details). We start our derivation by showing that the conditional distribution of the data layer given the hidden layer has the form

$$P_\beta(\mathbf{x}|\mu, \mathbf{J}) := P_\beta(\mathbf{x}|\mathbf{h} = \mathbf{e}_\mu, \mathbf{J}) \propto \exp\left(\beta \sum_{i=1}^N w_i^\mu x_i\right) \quad \forall \mathbf{x} \in S^{N-1} \text{ and } \mu \in \{1, \dots, P\}. \quad (2)$$

In other words, the probability density $P_\beta(\mathbf{x}|\mu, \mathbf{J})$ corresponding to each cluster $\mu > 0$ is a von Mises–Fisher (vMF) distribution centered on the direction $\mathbf{w}^\mu = \{w_i^\mu\}_{i=1}^N$ (see appendix B). In order to interpret the \mathbf{w}^μ as centroids for their respective clusters, we assume that they belong to S^{N-1} like the data layer. Under this assumption, we find the normalization constant of $P_\beta(\mathbf{x}|\mu, \mathbf{J})$ to be $\Omega_N(\beta) = \frac{(2\pi)^{N/2} I_{N/2-1}(\beta)}{\beta^{N/2-1}}$, where $I_n(x)$ is the modified Bessel function of the first kind of order n (see appendix B).

After slightly more work, we find the marginal distribution of the visible layer to be

$$P_\beta(\mathbf{x}, y|\mathbf{w}, \mathbf{p}) := P_\beta(\mathbf{x}, \mathbf{q} = \mathbf{e}_y|\mathbf{J}) = \sum_{\mu=1}^P p_y^\mu \frac{\exp\left(\beta \sum_{i=1}^N w_i^\mu x_i\right)}{\Omega_N(\beta)} + p_y^0 \frac{1}{\Omega_N(0)} \quad \forall \mathbf{x} \in S^{N-1} \text{ and } y \in \{0, \dots, C\}, \quad (3)$$



where $\mathbf{p} = \{p_\gamma^\gamma\}_{0 \leq \gamma \leq P} = \{P_\beta(\mathbf{q} = \mathbf{e}_y, \mathbf{h} = \mathbf{e}_y | \mathbf{J})\}_{0 \leq \gamma \leq P}$. $\Omega_N(0) = \frac{2\pi^{N/2}}{\Gamma(N/2)}$ is the surface area of S^{N-1} , where $\Gamma(x)$ is the Gamma function, so $\frac{1}{\Omega_N(0)}$ is the uniform distribution on S^{N-1} . The uniform distribution term of equation (3) encourages the model to ignore very noisy data during training, which may or may not be desirable depending on the application.

The detailed derivation of equation (3) (in appendix A) is inspired by the derivation of Gaussian mixtures from RBMs presented in [46] based on the framework of [40, 41]. In our case, the marginal distribution $P_\beta(\mathbf{x} | \mathbf{w}, \mathbf{p}) = \sum_{y=0}^C P_\beta(\mathbf{x}, y | \mathbf{w}, \mathbf{p})$ and conditional distributions $P_\beta(\mathbf{x} | y; \mathbf{w}, \mathbf{p}) = \frac{P_\beta(\mathbf{x}, y | \mathbf{w}, \mathbf{p})}{\sum_{\gamma=0}^C P_\beta(y, \gamma | \mathbf{w}, \mathbf{p})}$ are convex mixtures between the uniform distribution on S^{N-1} and the vMF distributions $P_\beta(\mathbf{x} | \mu, \mathbf{J})$ (see equation (2)). Probabilistic models similar to $P_\beta(\mathbf{x} | \mathbf{w}, \mathbf{p})$ are notably used in text document clustering [47]. Mixture distributions that have class-dependent weights like $P_\beta(\mathbf{x} | y; \mathbf{w}, \mathbf{p})$ are also used in Gaussian mixture discriminant analysis [48].

The class weights \mathbf{p} depend on the trainable parameters \mathbf{u} and \mathbf{b} of equation (1) (see appendix A). Without loss of generality, we choose to directly study (and train) \mathbf{p} instead of \mathbf{u} and \mathbf{b} . Recall that p_γ^γ is a probability distribution, and in particular $p_\gamma^\gamma \geq 0$ (see equation (3)). We constrain the marginal $\sum_{y=0}^C p_\gamma^\gamma$, i.e. the fraction of data in each cluster, to a fixed distribution $p_h(\gamma)$ and the marginal $\sum_{\gamma=0}^P p_\gamma^\gamma$, i.e. the proportion of data in each class y , to another fixed distribution $p_q(y)$. In sum, we end up with the constraints $p_\gamma^\gamma \geq 0$, $\sum_{\gamma=0}^P p_\gamma^\gamma = p_q(y)$ and $\sum_{y=0}^C p_\gamma^\gamma = p_h(\gamma)$. Since each cluster belongs to a single class (assumption 3), we expect the p_γ^γ of a trained model to be close to $\sum_{y'=0}^C p_{y'}^\gamma$ for a given y and close to 0 otherwise.

Given a dataset of P^* patterns $\{\mathbf{x}^{*\mu}\}_{\mu=1}^{P^*}$ with soft labels $q_y^{*\mu}$ [49], we can train the weights \mathbf{w} and \mathbf{p} by minimizing the negative log-likelihood loss

$$L(\mathbf{w}, \mathbf{p}) = -\frac{1}{P^*} \sum_{\mu=1}^{P^*} \sum_{y=0}^C q_y^{*\mu} \log P_\beta(\mathbf{x}^{*\mu}, y | \mathbf{w}, \mathbf{p}), \tag{4}$$

which is a form of maximum likelihood estimation. We do so using constrained Stochastic Gradient Descent with momentum (simply called SGD in this paper). We explain how the constraints on \mathbf{w} and \mathbf{p} are enforced in appendix C, and we briefly discuss the initial conditions and the learning rate in appendix D.

The weights \mathbf{w}^μ learned by SGD of the loss (equation (4)) [50] are interpretable prototypes, or memories, of the data (see figure 1 in the case of the MNIST dataset of handwritten digits [50]), which is consistent with their role as cluster centroids. Once the model is trained, we can use its conditionals $P_\beta(y | \mathbf{x}; \mathbf{w}, \mathbf{p})$ and $P_\beta(\mathbf{x} | y; \mathbf{w}, \mathbf{p})$ to efficiently reconstruct y from \mathbf{x} or \mathbf{x} from y , respectively. In particular, we can classify unseen patterns using the Bayes classification rule $y = \operatorname{argmax}_{y'} \{\log P_\beta(y' | \mathbf{x}; \mathbf{w}, \mathbf{p})\}$ and reconstruct patterns of a given class by finding the local minima of the effective energy $-\log P_\beta(\mathbf{x} | y; \mathbf{w}, \mathbf{p})$ as a function of $\mathbf{x} \in S^{N-1}$.

High-dimensional probabilistic models that store or learn prototypes, such as our model, can typically reconstruct a limited number of patterns with nontrivial accuracy. In other words, they have limited storage capacity. For instance, Hopfield's model of associative memory, the Hopfield network [9], has a capacity of up to $\mathcal{O}(N)$ patterns [9, 12]. DAM models, which are inspired by the Hopfield network, are a class of models with an asymptotically much higher capacity [14, 16]. The effective energy $-\log P_\beta(\mathbf{x}|y; \mathbf{w}, \mathbf{p})$ that we can minimize to reconstruct patterns with our model (see equation (3)) is very similar to that of the DAM with exponential capacity introduced in [19], which uses energy minimization to implement the attention mechanism of transformers. In fact, according to [51], our model also belongs to the class of DAMs with exponential capacity.

2.2. Teacher-student setting

Among the extensive research on the properties of artificial NNs from the perspective of statistical mechanics (see [52] for a review), there have been many studies of simple RBMs trained by maximum likelihood estimation (a line of work pioneered in [53–55]) or by averaging samples from the posterior distribution of the weights \mathbf{J} given some observed data \mathcal{D} . The latter approach has notably been used to characterize the fundamental limits of learning for many types of RBMs [56–58], with different priors [59–61] and architectures [62, 63], in the *teacher-student setting* [64–66] where the data used to train \mathbf{J} is sampled from another RBM with planted weights \mathbf{J}^* . This teacher-student setting can also be used to study our DAM. In this scenario, a *teacher* DAM with weights \mathbf{w}^* and \mathbf{p}^* generates a large amount $M = \alpha N$ of noisy data $\mathcal{D} = \{\mathbf{x}^c, y^c\}_{c=1}^M$ and feeds them to a *student* DAM, which then trains its weights \mathbf{w} and \mathbf{p} by averaging samples from the posterior distribution

$$P_\beta(\mathbf{w}, \mathbf{p}|\mathcal{D}) = \mathcal{Z}_\beta(\mathcal{D})^{-1} P_0(\mathbf{w}, \mathbf{p}) \prod_{c=1}^M P_\beta(\mathbf{x}^c, y^c|\mathbf{w}, \mathbf{p}), \quad (5)$$

where $\mathcal{Z}_\beta(\mathcal{D}) = \mathbb{E}_{\mathbf{w}, \mathbf{p}} \left[\prod_{c=1}^M P_\beta(\mathbf{x}^c, y^c|\mathbf{w}, \mathbf{p}) \right]$ is the posterior partition function and $P_0(\mathbf{w}, \mathbf{p})$ is the prior on \mathbf{w} and \mathbf{p} , which for simplicity we choose as uniform over the sets in which \mathbf{w} and \mathbf{p} are constrained.

We use a statistical mechanics approach to derive a single set of *saddle-point* equations that simultaneously characterize the weights that are stationary points of maximum likelihood estimation (equation (4)) for generic data and the typical weight configurations obtained by averaging samples from equation (5) in the teacher-student setting. We assume that the student does not know the number of hidden units P^* and the inverse temperature β^* of the teacher, so it cannot match them with its own. In particular, we consider the case where the noise injected by the teacher in the data is relatively small, i.e. $\beta^*/N > 0$ as $N \rightarrow \infty$, while the student chooses a conservative inverse temperature $\beta \ll N$ to avoid overfitting. Moreover, we fix $\sum_{y=0}^C p_y^{*0} = p_h^*(0) = 0$ and $\sum_{y=0}^C p_y^{*\mu} = p_h^*(\mu) = 1/P^*$ for all $\mu > 0$ so that each $\mathbf{g}^{*\mu} := P_\beta(y|\mu; \mathbf{w}, \mathbf{p}) = \mathbf{p}^{*\mu}/p_h^*(\mu) = P^* \mathbf{p}^{*\mu}$ is a soft label for the corresponding $\mathbf{w}^{*\mu}$. On the contrary, we do not give $\sum_{\gamma=0}^{P^*} p_y^{*\gamma} = p_q^*(y)$ a restrictive form. In other words, $p_q^*(y)$ is free to be any given probability mass function.

Table 2 of appendix E summarizes, for quick reference, the most important symbols introduced in this section and their meaning.

3. Theoretical results

3.1. Saddle-point equations

In this section, we introduce a set of equations for the stationary points of the loss (equation (4)), which we then relate to the saddle-point equations emerging from the statistical mechanics analysis of posterior sampling (see equation (5)) in the teacher-student setting.

Let us first establish a few definitions that we will use frequently throughout the section. Given two matrices $\mathbf{w}^* \in \mathbb{R}^{P^* \times N}$ and $\mathbf{w} \in \mathbb{R}^{P \times N}$, we define the overlap matrix $\mathbf{m}(\mathbf{w}^*, \mathbf{w}) = \mathbf{w}^* \mathbf{w}^T \in \mathbb{R}^{P^* \times P}$. We write its entries as $m^{\mu_* \mu}(\mathbf{w}^*, \mathbf{w}) = \sum_{i=1}^N w_i^{*\mu_*} w_i^\mu$ and its row vectors as $m^{\mu_*}(\mathbf{w}^*, \mathbf{w})$, where $1 \leq \mu_* \leq P^*$ and $1 \leq \mu \leq P$. Moreover, for any matrix $\mathbf{m} \in \mathbb{R}^{P^* \times (P+1)}$ with entries $m^{\mu_* \gamma}$ and row vectors m^{μ_*} , we use

$$\sigma_\gamma(m^{\mu_*}) = \frac{\exp(m^{\mu_* \gamma})}{\sum_{\nu=0}^P \exp(m^{\mu_* \nu})}$$

to represent the entry number $\gamma \in \{0, 1, \dots, P\}$ of the softmax function applied to the row vector m^{μ_*} . In this context, m^{μ_*} has a zeroth component $m^{\mu_* 0}$, and so does its softmax.

In appendix F, we show that the stationary points of the negative log-likelihood loss (equation (4)) satisfy

$$\begin{aligned}
 w_i^\mu &= \frac{\bar{w}_i^\mu}{\sqrt{\sum_{j=1}^N [\bar{w}_j^\mu]^2}} \\
 p_y^\gamma &= \frac{\bar{p}_y^\gamma}{\zeta_\gamma^\gamma(\bar{\mathbf{p}}; p_{\mathbf{h}})} \tag{6} \\
 \text{with } \bar{w}_i^\mu &= \sum_{\mu_*=1}^{P^*} x_i^{*\mu_*} \sum_{y=0}^C q_y^{*\mu_*} \sigma_\mu \left(\beta m^{\mu_*}(\mathbf{x}^*, \mathbf{w}) + \log [\mathbf{p}_y] \right) \\
 \bar{p}_y^\gamma &= \sum_{\mu_*=1}^{P^*} q_y^{*\mu_*} \sigma_\gamma \left(\beta m^{\mu_*}(\mathbf{x}^*, \mathbf{w}) + \log [\mathbf{p}_y] \right)
 \end{aligned}$$

for all $1 \leq \mu \leq P$ and $0 \leq \gamma \leq P$, where $m^{\mu_*0}(\mathbf{x}^*, \mathbf{w}) = \frac{1}{\beta} \log [\Omega_N(\beta) / \Omega_N(0)]$ and the normalization constant $\zeta_\gamma^\gamma(\bar{\mathbf{p}}; p_{\mathbf{h}})$ is defined in appendix C.

Similar equations arise naturally in our statistical mechanics analysis, which amounts to using the replica method [67, 68] to compute the limiting free entropy

$$f(\varrho, v, \beta, P^*, P) = \lim_{\beta^*, M, N \rightarrow \infty} \frac{1}{N} \mathbb{E}_{\mathbf{w}^*, \mathbf{p}^*, \mathcal{D}} \log [\mathcal{Z}_\beta(\mathcal{D})], \tag{7}$$

where $v = \beta^* / N$, $\varrho = \frac{M}{P^* N}$ and $\mathbb{E}_{\mathbf{w}^*, \mathbf{p}^*, \mathcal{D}}$ is the joint expectation over the distribution of examples $\mathcal{D} = \{\mathbf{x}^c, \mathbf{y}^c\}_{c=1}^M$ generated by the teacher and the priors on the teacher weights. To be more precise, we show that the free entropy can be computed with a variational principle of the form

$$f(\varrho, v, \beta, P^*, P) = \text{Extr}_{\mathbf{m}, \hat{\mathbf{m}}, \mathbf{p}} f(\mathbf{m}, \hat{\mathbf{m}}, \mathbf{p}), \tag{8}$$

whose extremizer $\mathbf{m} = \{m^{\mu_*\mu}\}_{1 \leq \mu_* \leq P^*}$ can be interpreted as the $N \rightarrow \infty$ limit of the expected value of the teacher-student overlaps $\mathbf{m}(\mathbf{w}^*, \mathbf{w})$. We show the derivation of the variational principle in appendix G together with an explicit expression for the trial function $f(\mathbf{m}, \hat{\mathbf{m}}, \mathbf{p})$. If we assume that there are $P^* \ll N$ teacher memories and that the priors on $\mathbf{w}^{*\mu_*}$ and \mathbf{g}^* are uniform like those of the student, we find that the expected teacher-student overlaps must satisfy the saddle-point equations

$$\begin{aligned}
 m^{\mu_*\mu} &= \varsigma \left(2\beta_{\text{eff}}\varrho \sqrt{\sum_{\nu_*=1}^{P^*} [\hat{m}^{\nu_*\mu}]^2} \right) \frac{\hat{m}^{\mu_*\mu}}{\sqrt{\sum_{\nu_*=1}^{P^*} [\hat{m}^{\nu_*\mu}]^2}} \\
 p_y^\gamma &= \frac{\bar{p}_y^\gamma}{\zeta_\gamma^\gamma(\bar{\mathbf{p}}; p_{\mathbf{h}})} \\
 \text{with } \hat{m}^{\mu_*\mu} &= \sum_{y=0}^C p_{\mathbf{q}}^*(y) \sigma_\mu \left(\beta_{\text{eff}} m^{\mu_*} + \log [\mathbf{p}_y] \right) \\
 \bar{p}_y^\gamma &= p_{\mathbf{q}}^*(y) \sum_{\mu_*=1}^{P^*} \sigma_\gamma \left(\beta_{\text{eff}} m^{\mu_*} + \log [\mathbf{p}_y] \right) \tag{9}
 \end{aligned}$$

for all $1 \leq \mu \leq P$ and $0 \leq \gamma \leq P$, where $\varsigma(x) = \frac{x}{\sqrt{x^2+1+1}}$, $\beta_{\text{eff}} = \varsigma(2v)\beta$ and $m^{\mu_*0} = \frac{1}{\beta_{\text{eff}}} \log [\Omega_N(\beta) / \Omega_N(0)]$.

If we instead clamp the teacher weights $\mathbf{w}^{*\mu_*}$ and $\mathbf{g}^{*\mu_*}$ to fixed patterns $\bar{\mathbf{x}}^{\mu_*}$ and their corresponding soft labels $\mathbf{q}^{*\mu_*}$ [49], respectively, to mimic a more general distribution for the data, then equations (9) become (see appendix H)

$$\begin{aligned}
 m^{\mu_*\mu} &= \varsigma \left(2\beta_{\text{eff}}\varrho \sqrt{\sum_{i=1}^N [\bar{x}_i^\mu]^2} \right) \frac{\sum_{i=1}^N x_i^{*\mu_*} \bar{x}_i^\mu}{\sqrt{\sum_{i=1}^N [\bar{x}_i^\mu]^2}} \\
 p_y^\gamma &= \frac{\bar{p}_y^\gamma}{\zeta_\gamma^\gamma(\bar{\mathbf{p}}; p_{\mathbf{h}})}
 \end{aligned}$$

$$\begin{aligned} \text{with } \bar{x}_i^\mu &= \sum_{\mu_*=1}^{P^*} x_i^{*\mu_*} \sum_{y=0}^C q_y^{*\mu_*} \sigma_\mu \left(\beta_{\text{eff}} m^{\mu_*} + \log [\mathbf{p}_y] \right) \\ \bar{p}_y^\gamma &= \sum_{\mu_*=1}^{P^*} q_y^{*\mu_*} \sigma_\gamma \left(\beta_{\text{eff}} m^{\mu_*} + \log [\mathbf{p}_y] \right). \end{aligned} \tag{10}$$

for all $1 \leq \mu \leq P$ and $0 \leq \gamma \leq P$, where we recall that $m^{\mu_*0} = \frac{1}{\beta_{\text{eff}}} \log [\Omega_N(\beta) / \Omega_N(0)]$.

In the limit of $\varrho, v \rightarrow \infty$, which indicates a large number of examples and a low level of teacher noise, equation (10) become equivalent to equation (6) if we make the identification $\bar{x}_i^\mu = \bar{w}_i^\mu$, i.e.

$$w_i^\mu = \frac{\bar{x}_i^\mu}{\sqrt{\sum_{j=1}^N [\bar{x}_j^\mu]^2}}. \tag{11}$$

Let us explain this limit step by step. For any v and ϱ , the M examples $\{\mathbf{x}^c\}_{c=1}^M$ given to the student are corrupted versions of the teacher patterns \mathbf{x}^* with a noise level of $1/v$. Therefore, in the limit of $v \rightarrow \infty$ with finite ϱ , each example \mathbf{x}^c is one of the original patterns \mathbf{x}^* , as in equation (6). However, even in this case, the empirical distribution of the examples deviates from that of \mathbf{x}^* , of which it is merely a bootstrap sample. At fixed P^* and P , this mismatch disturbs the accurate learning of \mathbf{x}^* when $\varrho = \frac{M}{P^*N}$ is finite, or equivalently the expected number of repetitions M/P^* of each pattern is not sufficiently large compared to N (see equation (10)). This influence progressively weakens as ϱ grows and the ς function approaches 1, reflecting the convergence of the empirical distribution of the examples to that of the teacher patterns.

In the limit of $\varrho \rightarrow \infty$ with finite v , equation (10) is still very similar to equation (6). In that case, the only difference between them is that equation (10) has β_{eff} instead of β in the argument of σ and in the denominator of the definition of m^{μ_*0} . Using ς as a shorthand for $\varsigma(2v)$, we find that the fixed points of equations (10) with $\varrho \rightarrow \infty$ are related, through the identification made in equation (11), to the stationary points of the effective loss

$$\begin{aligned} \mathcal{L}(\mathbf{w}, \mathbf{p}) &= -\frac{1}{P^*} \sum_{\mu=1}^{P^*} \sum_{y=0}^C q_y^{*\mu} \log \mathcal{P}_{\beta, \varsigma}(\mathbf{x}^{*\mu}, y | \mathbf{w}, \mathbf{p}), \\ \text{where } \mathcal{P}_{\beta, \varsigma}(\mathbf{x}, y | \mathbf{w}, \mathbf{p}) &= \sum_{\mu=1}^P p_y^\mu \frac{\exp\left(\varsigma \beta \sum_{i=1}^N w_i^\mu x_i\right)}{\Omega_N(\beta)} + p_y^0 \frac{1}{\Omega_N(0)}. \end{aligned} \tag{12}$$

What distinguishes this equation from the standard loss (equation (4)) is that $\mathcal{P}_{\beta, \varsigma}(\mathbf{x}, y | \mathbf{w}, \mathbf{p})$ has $\beta_{\text{eff}} = \varsigma \beta$ in the argument of the exponential function instead of β . As a consequence, $\mathcal{P}_{\beta, \varsigma}(\mathbf{x}, y | \mathbf{w}, \mathbf{p})$ is not a probability distribution unless $\varsigma = 1$, in the limit of $v \rightarrow \infty$ (see appendix B). A value of ς less than one in the effective loss (equation (12)) is reminiscent of the presence of noise in the data generation process, so we propose to use it as a regularizer for the weights. We discuss this point in more detail in section 4.1.

3.2. Saddle-point hierarchy

As shown in [8], the loss landscape of any NN with unconstrained weights contains the stationary points of narrower NNs with the same architecture. In appendix I, we show that this result also applies to the teacher-student setting with $\varrho \rightarrow \infty$ and any nonzero v . To be more precise, we show that, if the parameters $\bar{x}_i^{\text{fixed}, \mu}, \bar{p}_y^{\text{fixed}, \gamma}, m^{\text{fixed}, \mu_* \gamma}, p_y^{\text{fixed}, \gamma}$ with hidden unit prior $p_h^{\text{given}}(\gamma)$ are a fixed point of equation (10) with P hidden units, then the duplicated parameters

$$\begin{aligned} \bar{x}_i^{\text{dupli}, \mu} &= \begin{cases} \bar{x}_i^{\text{fixed}, \mu} & 0 < \mu \leq P \\ \bar{x}_i^{\text{fixed}, \mu-P} & P < \mu \leq P+R \end{cases} \\ \bar{p}_y^{\text{dupli}, \gamma} &= \begin{cases} \bar{p}_y^{\text{fixed}, 0} & \gamma = 0 \\ \frac{1}{2} \bar{p}_y^{\text{fixed}, \gamma} & 0 < \gamma \leq R \\ \bar{p}_y^{\text{fixed}, \gamma} & R < \gamma \leq P \\ \frac{1}{2} \bar{p}_y^{\text{fixed}, \gamma-P} & P < \gamma \leq P+R \end{cases} \end{aligned} \tag{12}$$

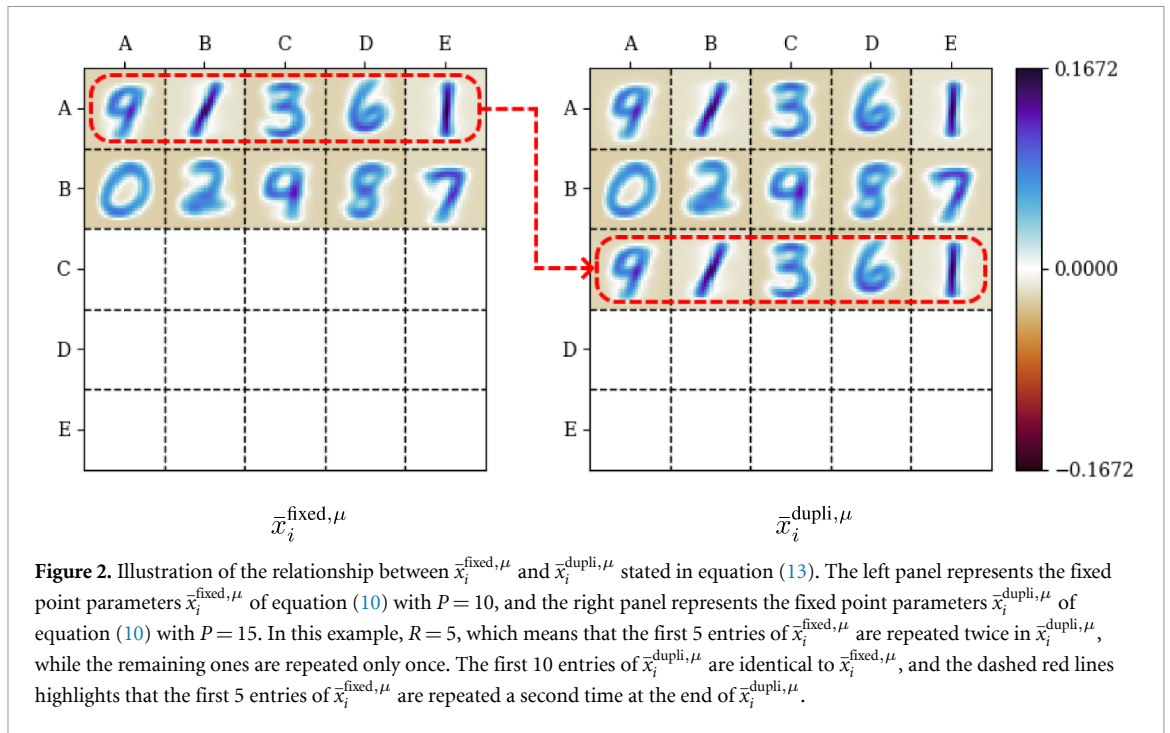


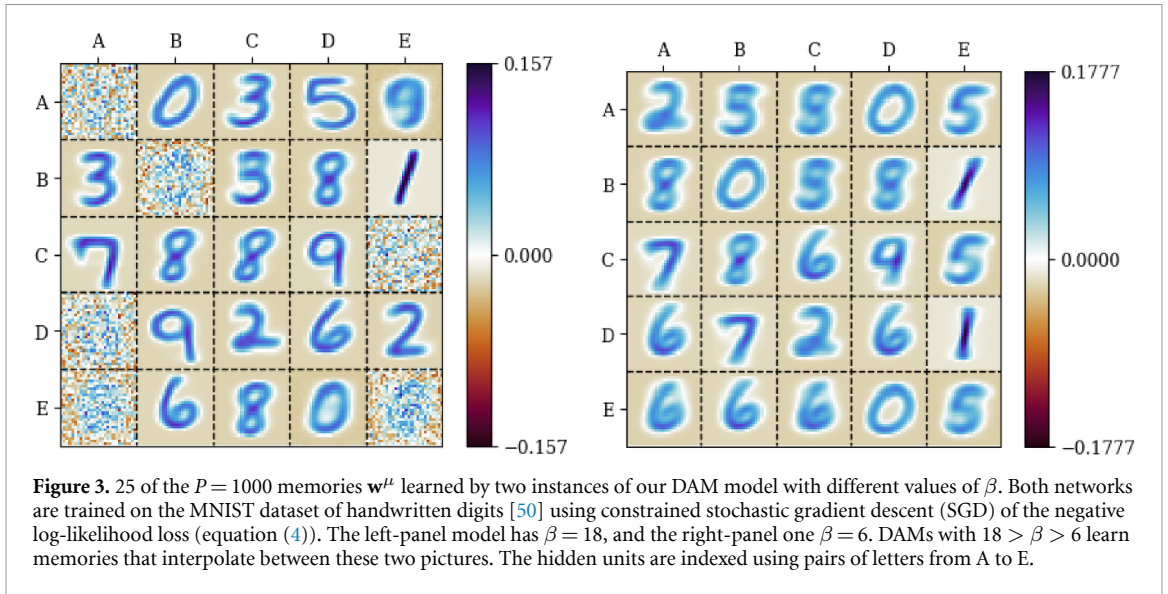
Figure 2. Illustration of the relationship between $\bar{x}_i^{\text{fixed}, \mu}$ and $\bar{x}_i^{\text{dupli}, \mu}$ stated in equation (13). The left panel represents the fixed point parameters $\bar{x}_i^{\text{fixed}, \mu}$ of equation (10) with $P = 10$, and the right panel represents the fixed point parameters $\bar{x}_i^{\text{dupli}, \mu}$ of equation (10) with $P = 15$. In this example, $R = 5$, which means that the first 5 entries of $\bar{x}_i^{\text{fixed}, \mu}$ are repeated twice in $\bar{x}_i^{\text{dupli}, \mu}$, while the remaining ones are repeated only once. The first 10 entries of $\bar{x}_i^{\text{dupli}, \mu}$ are identical to $\bar{x}_i^{\text{fixed}, \mu}$, and the dashed red lines highlights that the first 5 entries of $\bar{x}_i^{\text{fixed}, \mu}$ are repeated a second time at the end of $\bar{x}_i^{\text{dupli}, \mu}$.

$$\begin{aligned}
 m^{\text{dupli}, \mu * \gamma} &= \begin{cases} m^{\text{fixed}, \mu * 0} & \gamma = 0 \\ m^{\text{fixed}, \mu * \gamma} & 0 < \gamma \leq P \\ m^{\text{fixed}, \mu * \gamma - P} & P < \gamma \leq P + R \end{cases} \\
 p_y^{\text{dupli}, \gamma} &= \begin{cases} p_y^{\text{fixed}, 0} & \gamma = 0 \\ \frac{1}{2} p_y^{\text{fixed}, \gamma} & 0 < \gamma \leq R \\ p_y^{\text{fixed}, \gamma} & R < \gamma \leq P \\ \frac{1}{2} p_y^{\text{fixed}, \gamma - P} & P < \gamma \leq P + R \end{cases} \tag{13}
 \end{aligned}$$

$$\text{along with } p_h(\gamma) = \begin{cases} p_h^{\text{given}}(0) & \gamma = 0 \\ \frac{1}{2} p_h^{\text{given}}(\gamma) & 0 < \gamma \leq R \\ p_h^{\text{given}}(\gamma) & R < \gamma \leq P \\ \frac{1}{2} p_h^{\text{given}}(\gamma - P) & P < \gamma \leq P + R, \end{cases}$$

are a fixed point of the same saddle-point equations with $P + R \in \{P, \dots, 2P\}$ hidden units (See appendix 1 for a detailed derivation and figure 2 for an illustrative example). In other words, we duplicate some of the weights solving equation (10) to construct a fixed point for a wider network. In that sense, wide DAMs contain the fixed points of narrower DAMs. In particular (see section 3.1), this property also holds for the stationary points of both the standard loss (equation (4)) and the effective loss (equation (12)).

The saddle-point equations (equation (10)) with duplicated order parameters (equation (13)) are invariant to the permutation of any hidden unit $\gamma \in \{1, \dots, R\}$ and its duplicate $\gamma + P$. This kind of symmetry can be spontaneously broken if it leads to a higher free entropy (see equations (7) and (26)). However, symmetry-breaking transitions can be prohibitively slow when the symmetric state is stable to local perturbations [69]. Interestingly, the DAM introduced in [16] quickly undergoes many successive permutation symmetry-breaking bifurcations during training [24]. This observation and additional empirical evidence suggest that the symmetric states are unstable, or in other words that they are saddles, which was verified analytically for DAMs with only two data points to memorize [24]. In appendix 1, we prove that, if β is large enough, equation (13) is a saddle, which is a major step toward explaining why permutation symmetry breaking is relatively fast in DAMs trained with a large number of data points. We call this result *the saddle-point hierarchy principle*.



4. Empirical results

In this section, we use our theoretical results to improve training, and we show empirically that our DAM learns interpretable solutions to both supervised and unsupervised classification problems. Unless explicitly stated otherwise, we perform our numerical experiments on the MNIST dataset of handwritten digits [50]. The code and hyperparameter values used in our numerical experiments are available in the following public repository [30].

4.1. Learning by minimizing the effective loss

As the number of hidden units P increases, standard maximum likelihood estimation with fixed inverse temperature β becomes progressively less apt to train our DAM to its full potential. At high P and β , many memories stay stuck in noisy states that do not contribute to classification, which is wasteful (see figure 3). Using a lower β helps the memories converge, but also reduces their resolution and diversity, possibly because it also lowers the DAM’s capacity [51]. This change is far from being only cosmetic. In fact, it comes with a gradual reduction in classification accuracy (see table 1).

Optimization algorithms with a parameter analogous to β often converge to better solutions when this parameter is increased from a small value during optimization. For example, annealing schedules [70, 71] can be incorporated into minimization algorithms to help them find deeper local minima than they could otherwise. This observation suggests that our DAM would learn better memories if β increased during training. It is tempting to do so by SGD of equation (4) with respect to β , but it makes β increase so quickly that many memories still stay stuck in noisy states.

We find that optimizing the effective loss (equation (12)) with a suitable ς slows down the evolution of β enough to let the DAM learn much cleaner memories (see figure 4, top panel) and achieve much better classification accuracy (see table 1) than with standard training (equation (4)). Looking back at section 3.1, we propose to interpret ς as a regularization parameter that helps the DAM take noise from the data into account during training. There is no obvious theoretically motivated way to find the best value of ς for a generic dataset, so we choose it by hand. Despite this limitation, we believe that the simplicity and interpretability of our method still make it an interesting alternative to annealing schedules.

When we use this training method, we compute the DAM classification accuracy from the effective predictions

$$\mathcal{P}_{\beta,\varsigma}(y|\mathbf{x};\mathbf{w},\mathbf{p}) = \frac{\mathcal{P}_{\beta,\varsigma}(\mathbf{x},y|\mathbf{w},\mathbf{p})}{\mathcal{P}_{\beta,\varsigma}(\mathbf{x}|\mathbf{w},\mathbf{p})} \tag{14}$$

$$\text{where } \mathcal{P}_{\beta,\varsigma}(\mathbf{x}|\mathbf{w},\mathbf{p}) = \sum_{y=0}^C \mathcal{P}_{\beta,\varsigma}(\mathbf{x},y|\mathbf{w},\mathbf{p}),$$

instead of the true predictions $P_\beta(y|\mathbf{x};\mathbf{w},\mathbf{p}) = P_\beta(\mathbf{x},y|\mathbf{w},\mathbf{p}) / P_\beta(\mathbf{x}|\mathbf{w},\mathbf{p})$. This approach allows us to calculate both the accuracy and the loss through a single evaluation of $\mathcal{P}_{\beta,\varsigma}(y|\mathbf{x};\mathbf{w},\mathbf{p})$, which is more

Table 1. DAM classification accuracy (rounded down to two significant figures) for $P = 1000$ and various values of β . The last line is for β trained by SGD of the effective loss (equation (12)) with $\varsigma = 0.25$.

Inverse temperature β	Classification accuracy
6	79%
10	85%
14	89%
18	91%
Trained with $\varsigma = 0.25$	96%

efficient than computing $\mathcal{P}_{\beta,\varsigma}(\mathbf{x}, y | \mathbf{w}, \mathbf{p})$ and $\mathbb{P}_{\beta}(\mathbf{x}, y | \mathbf{w}, \mathbf{p})$ separately when monitoring the progress of training.

4.2. DAM is interpretable, even in unsupervised classification

Now that we understand how to train our DAM reasonably well, we investigate one of the most interesting properties of the solutions that it learns: their interpretability. As advertised in the Introduction, we will explain how to further improve training in the following section. We already mentioned that the regularization parameter ς of the effective loss is interpretable (see section 4.1 and equation (12)). Here, we point out that the learned weights are as well. In fact, each learned $\mathbf{p}^{\mu}/p_{\mathbf{h}}(\mu)$ can be interpreted as a soft label for the corresponding \mathbf{w}^{μ} (see figure 4). This property was also observed by [24] in K & H's DAM for pattern classification [16].

At test time, we observe that our DAM classifies approximately 98% of the test data points into the class $y = \operatorname{argmax}_{y'} \{p_{y'}^{\mu}\}$ of the memory \mathbf{w}^{μ} to which they are the most similar. In other words, a 1-nearest neighbor classifier [72] conditioned on the memories \mathbf{w}^{μ} and their soft labels $\mathbf{p}^{\mu}/p_{\mathbf{h}}(\mu)$ approximates the classification of our model with 98% fidelity. This behavior is reminiscent of K & H's DAM [16], where only a few memories participate in the classification of each data point. The 1-nearest neighbor classifier approximates DAM classification more faithfully for correctly classified data points (approximately 99% fidelity) than incorrectly classified data points (approximately 70% fidelity).

We now show that the notion of interpretability described in this section extends beyond supervised learning. Given a dataset of P^* unlabeled patterns $\{\mathbf{x}^{*\mu}\}_{\mu=1}^{P^*}$, we find that we can train our DAM for unsupervised classification by replacing the soft labels $q_y^{*\mu}$ in the effective loss (equation (12)) with the softened DAM predictions $(1 - \varepsilon)\mathcal{P}_{\beta,\varsigma}(y | \mathbf{x}^{*\mu}; \mathbf{w}, \mathbf{p}) + \varepsilon \frac{1}{C+1}$, where $\varepsilon \in [0, 1]$ and $\mathcal{P}_{\beta,\varsigma}(y | \mathbf{x}^{*\mu}; \mathbf{w}, \mathbf{p})$ is defined in equation (14). In this scenario, we are minimizing

$$\mathcal{L}_{\text{unsup}}(\mathbf{w}, \mathbf{p}) = -\frac{1}{P^*} \sum_{\mu=1}^{P^*} \sum_{y=0}^C \left[(1 - \varepsilon)\mathcal{P}_{\beta,\varsigma}(y | \mathbf{x}^{*\mu}; \mathbf{w}, \mathbf{p}) + \varepsilon \frac{1}{C+1} \right] \log \mathcal{P}_{\beta,\varsigma}(\mathbf{x}^{*\mu}, y | \mathbf{w}, \mathbf{p}). \quad (15)$$

As before, we do so using constrained SGD (see appendix C) with the initialization and learning rate described in appendix D. Equivalently, we can also view this algorithm as minimizing the combined loss

$$\mathcal{L}_{\text{total}}(\mathbf{w}, \mathbf{p}) = \mathcal{L}_{\text{margin}}(\mathbf{w}, \mathbf{p}) + \lambda \mathcal{L}_{\text{cond}}(\mathbf{w}, \mathbf{p}), \quad (16)$$

$$\text{where } \mathcal{L}_{\text{margin}}(\mathbf{w}, \mathbf{p}) = -\frac{1}{P^*} \sum_{\mu=1}^{P^*} \log \mathcal{P}_{\beta,\varsigma}(\mathbf{x}^{*\mu} | \mathbf{w}, \mathbf{p})$$

$$\text{and } \mathcal{L}_{\text{cond}}(\mathbf{w}, \mathbf{p}) = -\frac{1}{P^*} \sum_{\mu=1}^{P^*} \sum_{y=0}^C \mathcal{P}_{\beta,\varsigma}(y | \mathbf{x}^{*\mu}; \mathbf{w}, \mathbf{p}) \log \mathcal{P}_{\beta,\varsigma}(y | \mathbf{x}^{*\mu}; \mathbf{w}, \mathbf{p}),$$

where $\mathcal{P}_{\beta,\varsigma}(\mathbf{x}^{*\mu} | \mathbf{w}, \mathbf{p}) = \sum_{y=0}^C \mathcal{P}_{\beta,\varsigma}(\mathbf{x}^{*\mu}, y | \mathbf{w}, \mathbf{p})$ (see equation (14)). $\mathcal{L}_{\text{cond}}(\mathbf{w}, \mathbf{p})$ is called the minimum entropy regularization term in unsupervised machine learning [73]. Intuitively, it encourages the DAM to learn different class weights p_y^{μ} for each class y , which is not possible by minimizing only $\mathcal{L}_{\text{margin}}(\mathbf{w}, \mathbf{p})$ because $\mathcal{P}_{\beta,\varsigma}(\mathbf{x}^{*\mu} | \mathbf{w}, \mathbf{p})$ does not depend on y . Exploiting this characteristic, we train our DAM on patches of MNIST digits [50] and find that it learns reasonable latent classes $y = \operatorname{argmax}_{y'} \{p_{y'}^{\mu}\}$ for the memories \mathbf{w}^{μ} (see figure 5 and figures 8–11 of appendix J). This approach could potentially be useful for feature extraction [74].

4.3. Fast training with splitting steepest descent

We now explain how to further improve the training method of section 4.1. More specifically, we use the saddle-point hierarchy derived in section 3.2 to accelerate training.

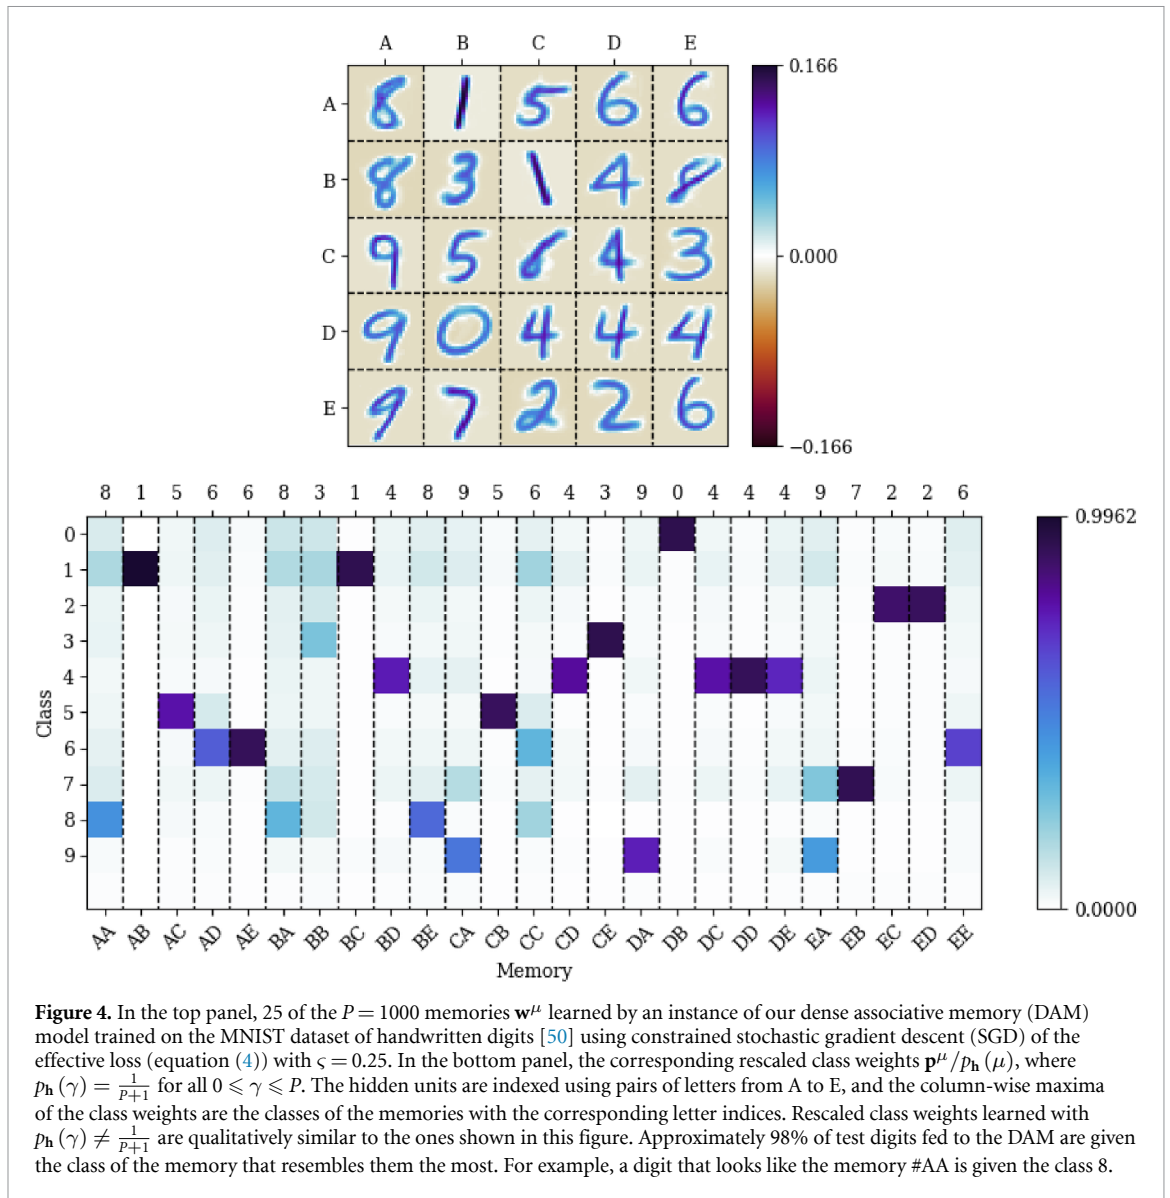


Figure 4. In the top panel, 25 of the $P = 1000$ memories \mathbf{w}^μ learned by an instance of our dense associative memory (DAM) model trained on the MNIST dataset of handwritten digits [50] using constrained stochastic gradient descent (SGD) of the effective loss (equation (4)) with $\zeta = 0.25$. In the bottom panel, the corresponding rescaled class weights $p^\mu / p_h(\mu)$, where $p_h(\gamma) = \frac{1}{P+1}$ for all $0 \leq \gamma \leq P$. The hidden units are indexed using pairs of letters from A to E, and the column-wise maxima of the class weights are the classes of the memories with the corresponding letter indices. Rescaled class weights learned with $p_h(\gamma) \neq \frac{1}{P+1}$ are qualitatively similar to the ones shown in this figure. Approximately 98% of test digits fed to the DAM are given the class of the memory that resembles them the most. For example, a digit that looks like the memory #AA is given the class 8.

Machine learning models that experience permutation symmetry breaking as described in section 3.2—such as K & H’s DAM [24], RBMs with binary units [62, 63] and Gaussian mixtures [40, 41, 71]—have characteristic tree-shaped learning curves. Intuitively, permutation symmetry-breaking transitions are points where model parameters differentiate from each other, so they correspond to the bifurcations of the tree. In the left panel of figure 6, we show that the learning dynamics of our DAM also follows a tree of permutation symmetry-breaking transitions.

The saddle-point hierarchy principle introduced in section 3.2 suggests the following idea to accelerate learning: train a relatively narrow DAM for cheap, then repeatedly duplicate (or ‘split’) some of the hidden units, escape the corresponding saddle point and continue training. Intuitively, we expect to save a lot of computing resources by using relatively few hidden units at the start of training, when the learning dynamics follows a few branches close to the root of the learning dynamics tree. The splitting steepest descent algorithm introduced in [29] formalizes this idea in an efficient and theoretically motivated way. We implement a variant of this algorithm that takes the constraint $\mathbf{w}^\mu \in S^{N-1}$ into account, but is otherwise very similar to the original version (see algorithm 1), with fast splitting implemented as in [75]. The constraint $\mathbf{w}^\mu \in S^{N-1}$ only matters in steps 5 and 11, which are also arguably the most conceptually difficult parts of algorithm 1, so we explain them in detail in appendix K. Figure (6) shows that the learning dynamics tree of splitting steepest descent has a more sparsely populated trunk than that of SGD without splitting (see figure 6). In other words, by using a relatively small number P of memories at the beginning of training, we reduce the total number of values $P \times T$ that they take over a period of time T , which is consistent with our intuition that it can save computational resources. The situation could have been different if using fewer memories slowed down training.

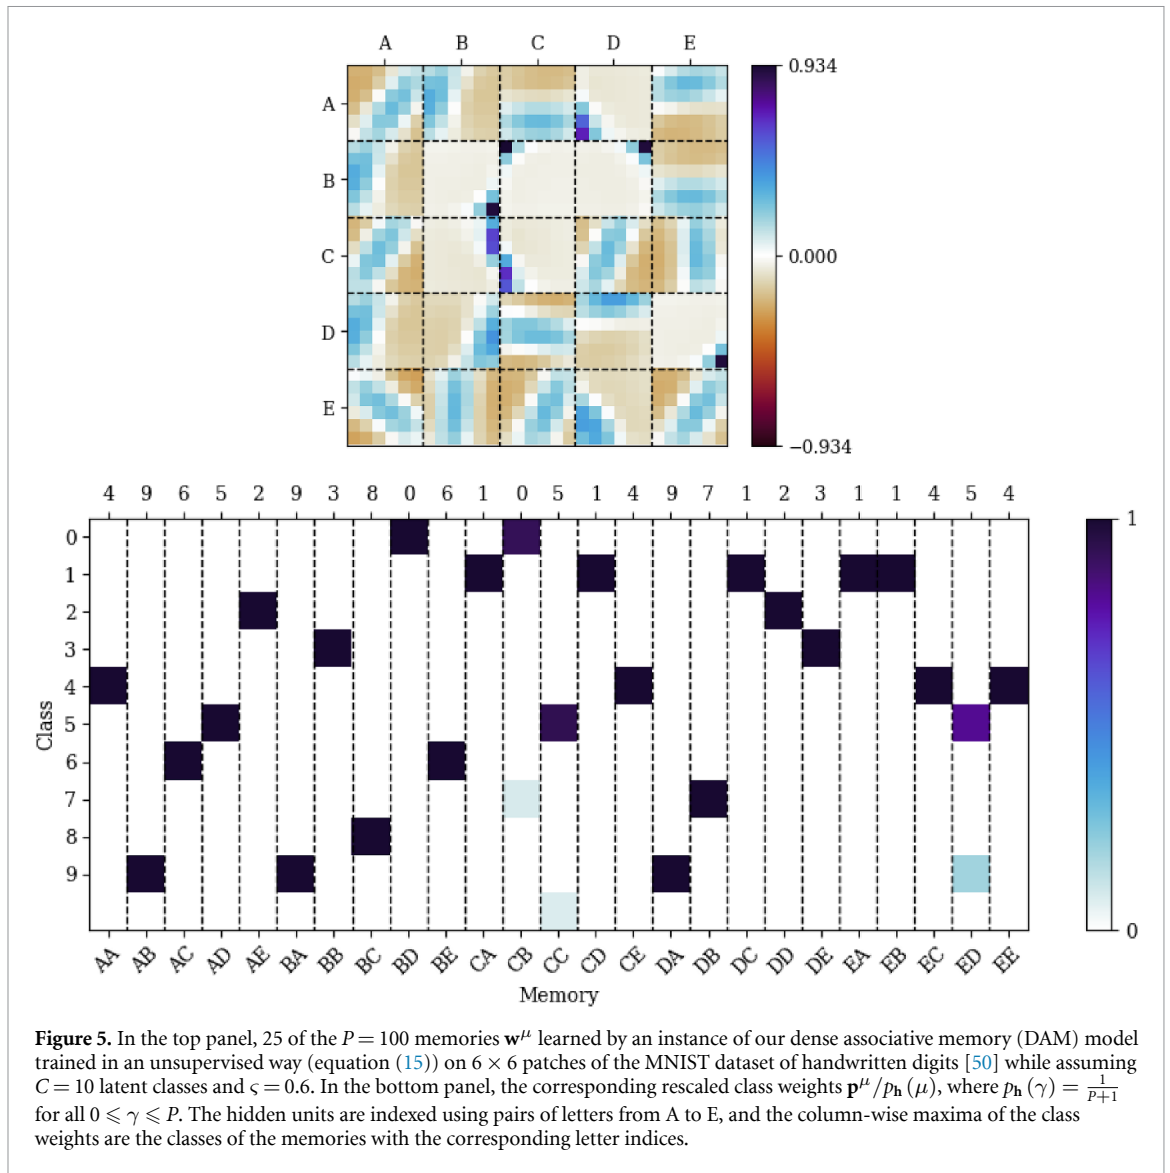
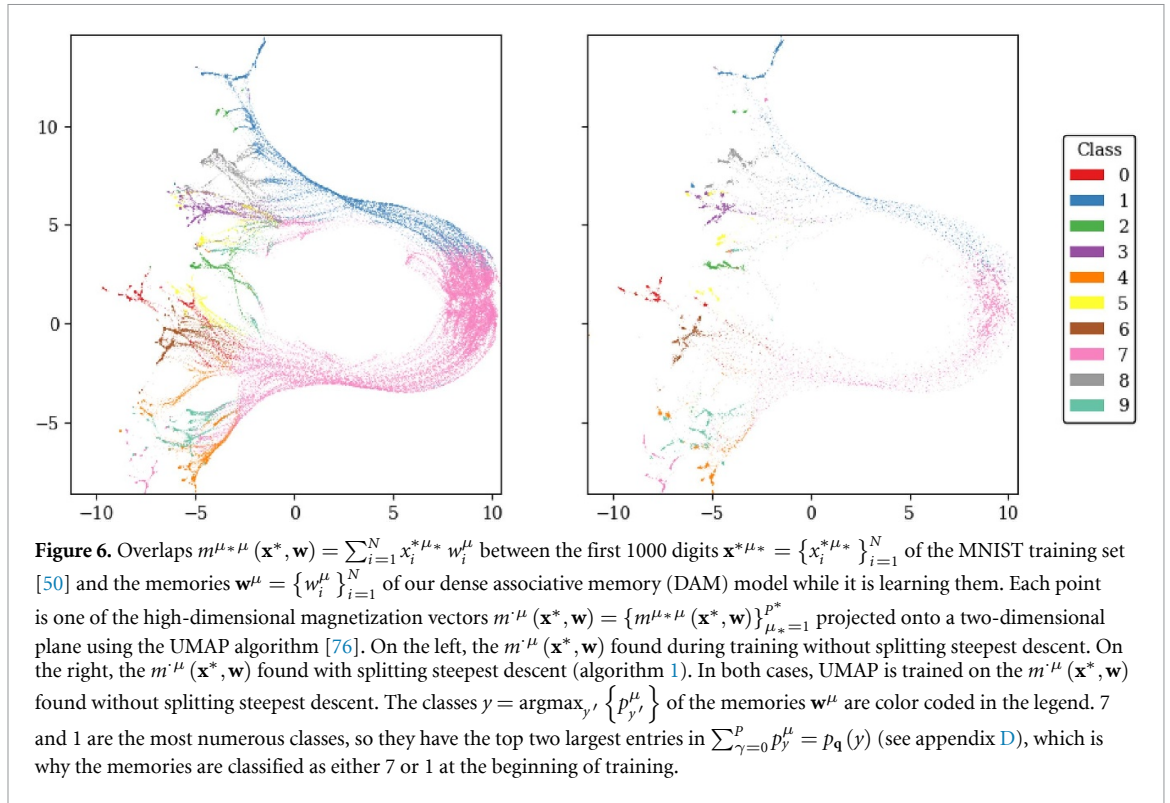


Figure 5. In the top panel, 25 of the $P = 100$ memories \mathbf{w}^μ learned by an instance of our dense associative memory (DAM) model trained in an unsupervised way (equation (15)) on 6×6 patches of the MNIST dataset of handwritten digits [50] while assuming $C = 10$ latent classes and $\varsigma = 0.6$. In the bottom panel, the corresponding rescaled class weights $\mathbf{p}^\mu / p_h(\mu)$, where $p_h(\gamma) = \frac{1}{P+1}$ for all $0 \leq \gamma \leq P$. The hidden units are indexed using pairs of letters from A to E, and the column-wise maxima of the class weights are the classes of the memories with the corresponding letter indices.

We now establish a methodology that we will use to compare the training times of DAMs trained on MNIST [50] and Fashion-MNIST [77] using the effective loss (equation (12)) with and without splitting steepest descent. It is more interesting and meaningful to compare the training times of NNs with similar performance. Therefore, we pick hyperparameters such that DAMs trained with and without splitting have similar classification accuracy. For that purpose, we find the general region of hyperparameter space where DAMs trained without splitting steepest descent have the best classification accuracy. The accuracy does not change much in this region, so we pick generic hyperparameters inside. Next, we manually tune the hyperparameters of splitting steepest descent so that the resulting DAMs have comparable accuracy. The accuracy obtained for Fashion-MNIST is generically slightly higher with splitting steepest descent than without it. In contrast, splitting steepest descent does not significantly affect the accuracy obtained for MNIST.

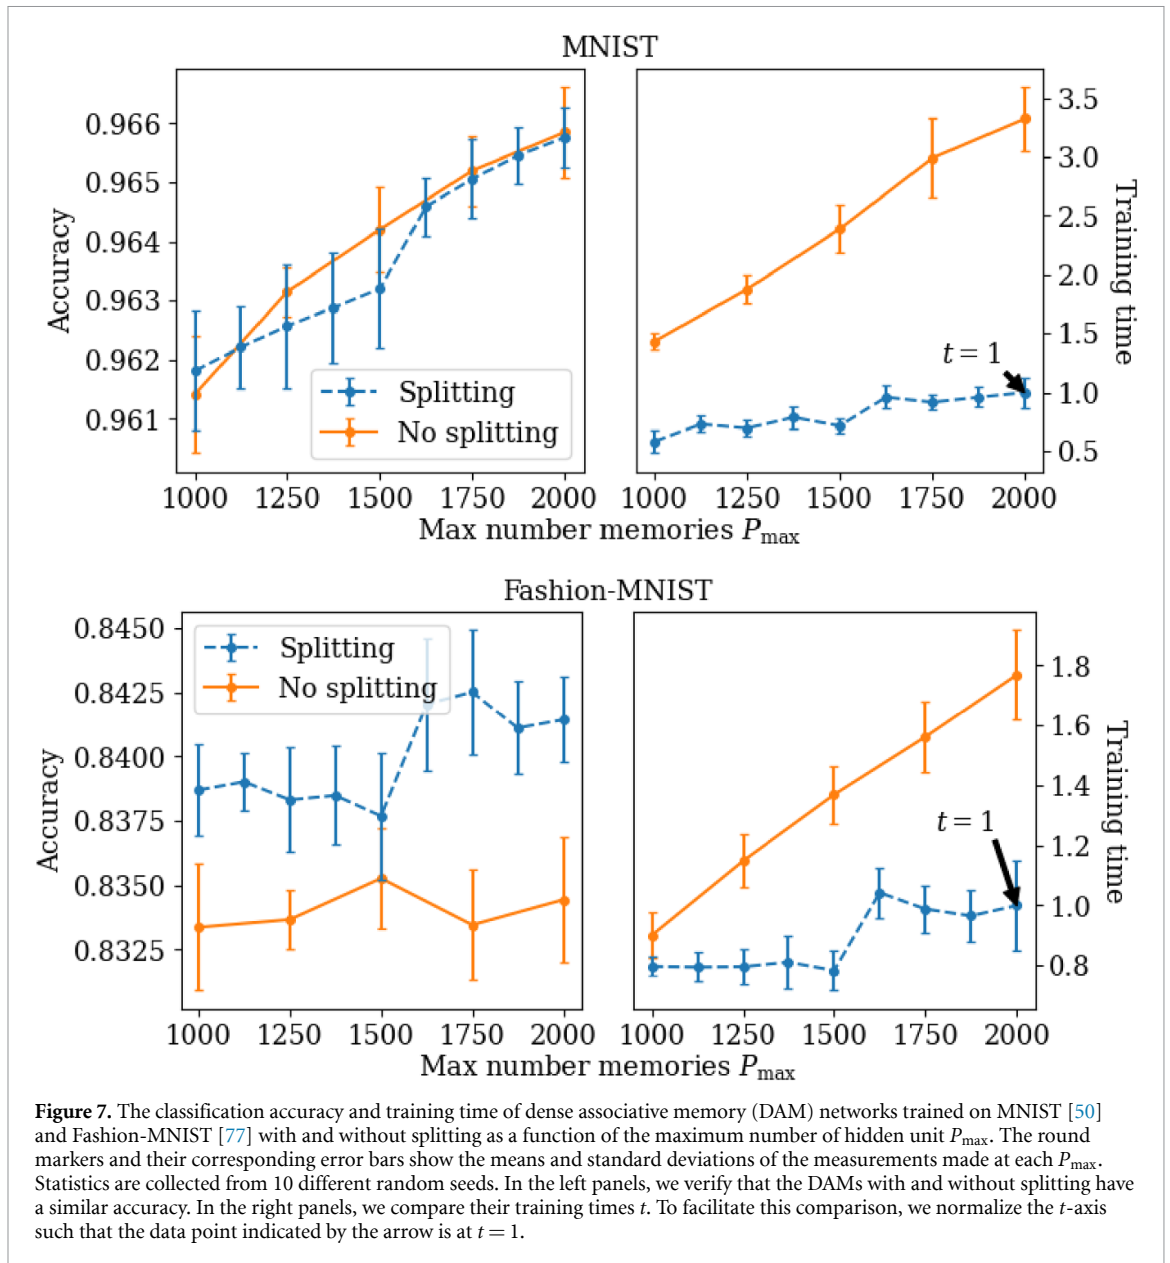
Once the hyperparameters are set, we collect statistics of the accuracy and training time. We run our experiment on a CPU and manually set the seed of pseudorandom number generation to make training deterministic and reproducible. Two DAMs with the same hyperparameters and the same seed are guaranteed to be trained using the same number of computer operations and to have the same classification accuracy after training. However, background processes and other external factors can change between two runs with the same seed, which adds noise to the bare training time that we want to measure. As such, we approximate the bare training time as the minimum over many runs with the same seed. We then calculate the average and standard deviation of the classification accuracy and bare training time over multiple seeds.



Algorithm 1. Splitting steepest descent [29, 75].

- 1: Preallocate space for a DAM with P_{\max} hidden units and the corresponding weights \mathbf{w} and \mathbf{p}
 - 2: Initialize the weights \mathbf{w}^μ and \mathbf{p}^μ connected to the P_{cur} first hidden units $\mu \in \{1, \dots, P_{\text{cur}}\}$, as well as \mathbf{p}^0
 - 3: $\min L(\mathbf{w}, \mathbf{p})$ with SGD
 - 4: **while** $P_{\text{cur}} < P_{\max}$ **do**
 - 5: Identify a subset $\mu_{\text{copy}} \subseteq \{1, \dots, P_{\text{cur}}\}$ of $R \leq P_{\max} - P_{\text{cur}}$ hidden units to split, **return** if empty
 - 6: Let $\mu_{\text{paste}} = \{P_{\text{cur}} + 1, \dots, P_{\text{cur}} + R\}$
 - 7: Build weights $\mathbf{w}^{\mu_{\text{paste}}} = \mathbf{w}^{\mu_{\text{copy}}}$ for μ_{paste}
 - 8: Rescale $\mathbf{p}^{\mu_{\text{copy}}} \leftarrow \mathbf{p}^{\mu_{\text{copy}}} / 2$ and $p_h(\mu_{\text{split}}) \leftarrow p_h(\mu_{\text{split}}) / 2$
 - 9: Build weights $\mathbf{p}^{\mu_{\text{paste}}} = \mathbf{p}^{\mu_{\text{copy}}}$ and $p_h(\mu_{\text{paste}}) = p_h(\mu_{\text{copy}})$ for μ_{paste}
 - 10: Update $P_{\text{cur}} \leftarrow P_{\text{cur}} + R$
 - 11: Escape the saddle point by 2nd order descent of $L(\mathbf{w}, \mathbf{p})$ w.r.t. \mathbf{w}
 - 12: $\min L(\mathbf{w}, \mathbf{p})$ with SGD
 - 13: **end while**
 - 14: **return** ▷ See appendix K and [29, 75] for details about steps 5 and 11
-

In figure 7, we compare DAM training time and accuracy with and without splitting steepest descent, using round markers and error bars to show their respective means and standard deviations as a function of the maximum number of hidden units P_{\max} . Splitting stops at P_{\max} hidden units, and DAMs without splitting have P_{\max} hidden units from the start. All hidden units are split in each while-loop iteration of algorithm 1, except for the last iteration, where exactly $P_{\max} - P_{\text{cur}}$ hidden units are split. As wanted, DAMs have similar accuracy with and without splitting (left panels). Moreover, the reasonably small error bars of the training time indicate that the residual noise of the background processes is controlled and that the seed has a limited impact on the training speed (right panels). Splitting steepest descent has a significant speed advantage that scales very advantageously with P_{\max} (right panels). Without splitting, the training time is proportional to P_{\max} . With splitting, it appears to be piecewise constant. The clearest sign that it increases with P_{\max} is the jump between $P_{\max} = 1500$ and $P_{\max} = 1625$, where the number of while-loop iterations in algorithm 1 increases by 1. In other words, the run time of algorithm 1 is better explained as proportional to $\lceil \log P_{\max} \rceil$ than P_{\max} in the range of P_{\max} of our numerical experiment, which is a dramatic improvement over the run time without splitting. The benefit of the algorithm becomes evident primarily when the accuracy improves consistently across a wide range of hidden unit counts, i.e. when training a substantially large network is beneficial. This pattern, for



example, does not apply to Fashion MNIST (see figure 7, bottom plot), where the accuracy gain from training a larger network is marginal. As a result, the need for strategies that accelerate training becomes less critical, although such acceleration still occurs.

5. Conclusion

In this work, we study a DAM model based on the framework of BMs that is capable of learning interpretable solutions to both supervised and unsupervised problems.

We derive two sets of equations that respectively characterize the stationary points of DAMs trained on real data using maximum likelihood estimation and the fixed points of DAMs trained on synthetic data in the teacher-student setting. Guided by their similarity, we then establish that the maximum likelihood equations are a special case of the teacher-student equations. Building on this equivalence, we introduce an effective loss function that makes training significantly more stable thanks to a regularization parameter that mimics the noise present in the data generation process.

We show that the stationary points of the effective loss of DAMs with a given number of hidden units are saddle points of larger DAMs, and we considerably accelerate training using a custom implementation of the splitting steepest descent network-growing algorithm [29] inspired by this *saddle-point hierarchy principle*. In our numerical experiments, the training time of splitting steepest descent is

approximately logarithmic in the number of hidden units, which is a huge improvement over the linear training time that we observe without splitting steepest descent. Further research is needed to clearly determine the asymptotic functional form of the training time.

In section 4.1, we briefly discuss how our DAM is interpretable. The ingredients needed to make an interpretable NN are not well understood, and we believe that our DAM offers a promising opportunity to explore that area. Studying the teacher-student setting with teacher inverse temperature β^* less than order N is another interesting research direction that we leave to future work. Beyond the framework of our model, [78, 79] have designed a technique that trains Gaussian mixture models by replicating mixture components. In contrast with our work, it uses a handpicked cooling schedule and breaks permutation symmetries with random noise. By design, splitting steepest descent finds the directions in the loss landscape that allow NNs to escape saddle points with permutation symmetries as quickly as possible, but finding these directions (as described in appendix K and [29, 75]) is more expensive than generating random noise. It would be interesting to study this tradeoff and identify the regimes where each method is preferable over the other. More generally, one could investigate how permutation symmetry breaking takes place in energy models related to DAMs, such as attentional BMs [80], and study its relationship with the permutation symmetry-breaking [62, 63] and dynamical transitions [53, 54, 81] found in other types of RBMs, whose careful analysis was also used to improve training [82]. Recent works [24, 83] have made connections between DAM learning dynamics and cellular differentiation. Similarly, it could be possible to investigate whether there are similarities between cellular division and splitting steepest descent in our model. Finally, it would be interesting to see if splitting steepest descent and our proposed regularization technique can be used to improve the training of transformers and generative diffusion.

Data availability statement

The data that support the findings of this study are openly available at the following URL/DOI: <https://doi.org/10.5281/zenodo.17632013> [30].

Acknowledgments

This work was partially supported by project SERICS (PE00000014) under the MUR National Recovery and Resilience Plan funded by the European Union—NextGenerationEU. The work was also supported by the project PRIN22TANTARI ‘Statistical Mechanics of Learning Machines: from algorithmic and information-theoretical limits to new biologically inspired paradigms’ 20229T9EAT—CUP J53D23003640001. DT also acknowledges GNFM-Indam.

RT acknowledges the computational resources of the Center for High Performance Computing (CHPC) at the Scuola Normale Superiore of Pisa, as well as Carlo Lucibello and Matteo Negri for insightful discussions.

Conflicts of interest

The authors have no conflicts of interest to declare.

Author contributions

Robin Thériault  0009-0004-9555-4088

Conceptualization (lead), Data curation (lead), Formal analysis (lead), Investigation (lead), Methodology (lead), Software (lead), Visualization (lead), Writing – original draft (lead)

Daniele Tantari  0000-0001-9982-5720

Conceptualization (supporting), Funding acquisition (lead), Project administration (lead), Supervision (lead), Validation (lead), Writing – original draft (supporting), Writing – review & editing (lead)

Appendix A. Derivation of the model

Consider the RBM Hamiltonian

$$-\beta H[\mathbf{x}, \mathbf{q}, \mathbf{h}; \mathbf{J}] = \beta \sum_{\mu=1}^P h_{\mu} \sum_{i=1}^N w_i^{\mu} x_i + \beta \sum_{\mu=1}^P h_{\mu} \sum_{y=1}^C u_y^{\mu} q_y + \beta \sum_{\mu=1}^P h_{\mu} b^{\mu}.$$

where $\mathbf{J} = \{\mathbf{w}, \mathbf{u}, \mathbf{b}\}$ is the set of all RBM parameters, i.e. $\mathbf{w} = \{w_i^{\mu}\}_{1 \leq \mu \leq P, 1 \leq i \leq N}$, $\mathbf{u} = \{u_y^{\mu}\}_{1 \leq \mu \leq P, 1 \leq y \leq C}$ and $\mathbf{b} = \{b^{\mu}\}_{\mu=1}^P$. The prior densities $P_0(\mathbf{h})$ and $P_0(\mathbf{q})$ are nonzero only when $\mathbf{h} \in \{\mathbf{e}_{\gamma}\}_{\gamma=0}^P$ and $\mathbf{q} \in \{\mathbf{e}_{\gamma}\}_{\gamma=0}^C$, in which case we have

$$-\beta H[\mathbf{x}, \mathbf{e}_{\gamma}, \mathbf{e}_{\gamma}; \mathbf{J}] = \begin{cases} 0 & \gamma = 0 \\ \beta \sum_{i=1}^N w_i^{\gamma} x_i + a_{\gamma}^{\gamma} & 0 < \gamma \leq P \end{cases}$$

where $a_y^{\mu} = \begin{cases} \beta b^{\mu} & y = 0 \\ \beta u_y^{\mu} + \beta b^{\mu} & 0 < y \leq C. \end{cases}$

For convenience, we write the corresponding Gibbs distribution $P_{\beta}(\mathbf{x}, \mathbf{q} = \mathbf{e}_{\gamma}, \mathbf{h} = \mathbf{e}_{\gamma} | \mathbf{J})$ as $P_{\beta}(\mathbf{x}, y, \mu | \mathbf{J})$, and the priors $P_0(\mathbf{q} = \mathbf{e}_{\gamma})$ and $P_0(\mathbf{h} = \mathbf{e}_{\gamma})$ as $\pi_{\mathbf{q}}(y)$ and $\pi_{\mathbf{h}}(\gamma)$, respectively. Given a uniform prior on \mathbf{x} , we find

$$P_{\beta}(\mathbf{x}, y, \mu | \mathbf{J}) = \frac{1}{Z_{\beta}(\mathbf{J})} \pi_{\mathbf{q}}(y) \begin{cases} \pi_{\mathbf{h}}(0) & \gamma = 0 \\ \pi_{\mathbf{h}}(\gamma) \exp(a_{\gamma}^{\gamma}) \exp\left(\beta \sum_{i=1}^N w_i^{\gamma} x_i\right) & 0 < \gamma \leq P, \end{cases}$$

where $Z_{\beta}(\mathbf{J})$ is an unknown normalization constant. The data \mathbf{x} belonging to each cluster $\mu > 0$ follows a vMF distribution $P_{\beta}(\mathbf{x} | \mu, \mathbf{J}) \propto \exp\left(\beta \sum_{i=1}^N w_i^{\mu} x_i\right)$ centered on $\mathbf{w}^{\mu} = \{w_i^{\mu}\}_{i=1}^N$ (see appendix B). We assume that $\sqrt{\sum_{i=1}^N (w_i^{\mu})^2} = 1$ so that the corresponding cluster centroids \mathbf{w}^{μ} lie on the hypersphere S^{N-1} like the data $\mathbf{x} = \{x_i\}_{i=1}^N$. We marginalize $P_{\beta}(\mathbf{x}, y, \mu | \mathbf{J})$ over the hidden units μ and get

$$P_{\beta}(\mathbf{x}, y | \mathbf{J}) = \frac{1}{Z_{\beta}(\mathbf{J})} \pi_{\mathbf{q}}(y) \left[\sum_{\mu=1}^P \pi_{\mathbf{h}}(\mu) \exp(a_{\gamma}^{\mu}) \exp\left(\beta \sum_{i=1}^N w_i^{\mu} x_i\right) + \pi_{\mathbf{h}}(0) \right].$$

We will now find the normalization constant $Z_{\beta}(\mathbf{J})$. Marginalizing $P_{\beta}(\mathbf{x}, y | \mathbf{J})$ over \mathbf{x} , we obtain

$$\begin{aligned} P_{\beta}(y | \mathbf{J}) &= \int_{S^{N-1}} d\mathbf{x} P_{\beta}(\mathbf{x}, y | \mathbf{J}) \\ &= \frac{1}{Z_{\beta}(\mathbf{J})} \pi_{\mathbf{q}}(y) \int_{S^{N-1}} d\mathbf{x} \left[\sum_{\mu=1}^P \pi_{\mathbf{h}}(\mu) \exp(a_{\gamma}^{\mu}) \exp\left(\beta \sum_{i=1}^N w_i^{\mu} x_i\right) + \pi_{\mathbf{h}}(0) \right] \\ &= \frac{1}{Z_{\beta}(\mathbf{J})} \pi_{\mathbf{q}}(y) \Omega_N(\beta) \sum_{\mu=1}^P \pi_{\mathbf{h}}(\mu) \exp(a_{\gamma}^{\mu}) + \Omega_N(0) \pi_{\mathbf{h}}(0) \\ &= \frac{1}{Z_{\beta}(\mathbf{J})} \pi_{\mathbf{q}}(y) \Omega_N(\beta) \left[\sum_{\mu=1}^P \pi_{\mathbf{h}}(\mu) \exp(a_{\gamma}^{\mu}) + \pi_{\mathbf{h}}(0) \exp(a^0) \right], \end{aligned}$$

where $a^0 = \log[\Omega_N(0) / \Omega_N(\beta)]$. Normalization of $P_{\beta}(y | \mathbf{J})$ requires that

$$P_{\beta}(y | \mathbf{J}) = \frac{\pi_{\mathbf{q}}(y) \left[\sum_{\mu=1}^P \pi_{\mathbf{h}}(\mu) \exp(a_{\gamma}^{\mu}) + \pi_{\mathbf{h}}(0) \exp(a^0) \right]}{\sum_{y'=0}^C \pi_{\mathbf{q}}(y') \left[\sum_{\nu=1}^P \pi_{\mathbf{h}}(\nu) \exp(a_{\gamma'}^{\nu}) + \pi_{\mathbf{h}}(0) \exp(a^0) \right]},$$

so we deduce that $Z_{\beta}(\mathbf{J}) = \Omega_N(\beta) \sum_{y=0}^C \pi_{\mathbf{q}}(y) \left[\sum_{\mu=1}^P \pi_{\mathbf{h}}(\mu) \exp(a_{\gamma}^{\mu}) + \pi_{\mathbf{h}}(0) \exp(a^0) \right]$. We define

$$p_y^{\gamma} = \begin{cases} \frac{\pi_{\mathbf{q}}(y) \pi_{\mathbf{h}}(0) \exp(a^0)}{\sum_{y'=0}^C \pi_{\mathbf{q}}(y') \left[\sum_{\nu=1}^P \pi_{\mathbf{h}}(\nu) \exp(a_{\gamma'}^{\nu}) + \pi_{\mathbf{h}}(0) \exp(a^0) \right]} & \gamma = 0 \\ \frac{\pi_{\mathbf{q}}(y) \pi_{\mathbf{h}}(\mu) \exp(a_{\gamma}^{\mu})}{\sum_{y'=0}^C \pi_{\mathbf{q}}(y') \left[\sum_{\nu=1}^P \pi_{\mathbf{h}}(\nu) \exp(a_{\gamma'}^{\nu}) + \pi_{\mathbf{h}}(0) \exp(a^0) \right]} & 0 < \gamma \leq P, \end{cases}$$

from which we obtain

$$P_\beta(\mathbf{x}, \mathbf{y} | \mathbf{J}) = \sum_{\mu=1}^P p_y^\mu \frac{\exp\left(\beta \sum_{i=1}^N w_i^\mu x_i\right)}{\Omega_N(\beta)} + p_y^0 \frac{1}{\Omega_N(0)}.$$

The conditional distributions $P_\beta(\mathbf{x} | \mathbf{y}, \mathbf{J})$ and their marginal $P_\beta(\mathbf{x} | \mathbf{J}) = \sum_{y=0}^C P_\beta(\mathbf{x} | \mathbf{y}, \mathbf{J}) P_\beta(y | \mathbf{J})$ are vMF mixtures [47] with weights $\frac{p_y^\gamma}{\sum_{\nu=0}^C p_y^\nu}$ and $\sum_{y=0}^C p_y^\gamma$, respectively. As explained in section 2, we constrain the marginals $\sum_{y=0}^C p_y^\gamma = P_\beta(\gamma | \mathbf{J})$ and $\sum_{\gamma=0}^P p_y^\gamma = P_\beta(y | \mathbf{J})$ to be equal to fixed distributions $p_h(\gamma)$ and $p_q(y)$, respectively. Formally, this means that $\mathbf{p} = \left\{ p_y^\gamma \right\}_{0 \leq \gamma \leq C, 0 \leq y \leq C}$ belongs to the transportation polytope with sum constraints $\sum_{\gamma=0}^P p_y^\gamma = p_q(y)$ and $\sum_{y=0}^C p_y^\gamma = p_h(\gamma)$ [84].

Appendix B. Integration of the von Mises–Fisher (vMF) density

The vMF distribution [85] is an isotropic Gaussian distribution restricted to the $N - 1$ dimensional unit sphere S^{N-1} . It takes the form

$$p(\mathbf{x}) = \Omega(r)^{-1} \exp\left(\sum_{i=1}^N r_i x_i\right),$$

where $r = \sqrt{\sum_i [r_i]^2}$ is called the concentration parameter. When $r = 0$, it reduces to the uniform distribution on the unit sphere, whose surface area $\frac{2\pi^{N/2}}{\Gamma(N/2)}$ is thus also the normalization constant $\Omega_N(0)$. When $r > 0$, we define the mean direction $\hat{\mathbf{r}} = \{r_i/r\}_{i=1}^N$ and find the normalization constant to be

$$\begin{aligned} \Omega_N(r) &= \int_{S^{N-1}} d\mathbf{x} \exp\left(\sum_{i=1}^N r_i x_i\right) \\ &= \int_{S^{N-1}} d\mathbf{x} \exp\left(r \sum_{i=1}^N \hat{r}_i x_i\right) \\ &= \Omega_{N-1}(0) \int_{-1}^1 du (1-u^2)^{(N-3)/2} \exp(ru) \\ &= \Omega_N(0) \left(\frac{r}{2}\right)^{1-N/2} \Gamma\left(\frac{N}{2}\right) I_{N/2-1}(r), \end{aligned} \tag{17}$$

where $I_n(x)$ is the modified Bessel function of the first kind of order n . The third line of (17) comes from the change of variables $u = \sum_{i=1}^N \hat{r}_i x_i$, and the fourth line is a consequence of Poisson’s Bessel function integral [86]. In the limit of large N with $\rho = \frac{r}{N-2} \approx \frac{r}{N}$, we find

$$\begin{aligned} \int_{S^{N-1}} d\mathbf{x} \exp\left(\sum_{i=1}^N r_i x_i\right) &\approx \frac{\Omega_N(0)}{\sqrt{2\pi(N/2-1)}} \left(\frac{N}{2}-1\right)^{1-N/2} (1+(2\rho)^2)^{-1/4} \\ &\quad \Gamma\left(\frac{N}{2}\right) \exp\left[\left(\frac{N}{2}-1\right)(\eta(2\rho)+1)\right] \\ \text{where } \eta(x) &= (1+x^2)^{1/2} - 1 - \log\left[1+(1+x^2)^{1/2}\right] + \log 2, \end{aligned}$$

by exploiting the large N asymptotic expansion of $I_{N/2-1}([N/2-1] \cdot 2\rho)$ found in [87, 88]. We use Stirling’s approximation

$$\Gamma\left(\frac{N}{2}\right) \approx \sqrt{2\pi(N/2-1)} \left(\frac{N}{2}-1\right)^{N/2-1} \exp\left(-\frac{N}{2}+1\right)$$

to simplify it to

$$\begin{aligned} \int_{S^{N-1}} d\mathbf{x} \exp\left(\sum_{i=1}^N r_i x_i\right) &\approx \Omega_N(0) (1+(2\rho)^2)^{-1/4} \exp\left[\left(\frac{N}{2}-1\right)\eta(2\rho)\right] \\ \text{where } \eta(x) &= (1+x^2)^{1/2} - 1 - \log\left[1+(1+x^2)^{1/2}\right] + \log 2. \end{aligned} \tag{18}$$

Appendix C. Normalization of the weights

In order to enforce $\sqrt{\sum_{i=1}^N (w_i^\mu)^2} = 1$ at each SGD step, we project \mathbf{w}^μ and the gradient of the loss with respect to \mathbf{w}^μ onto the unit sphere S^{N-1} and the tangent space of \mathbf{w}^μ , respectively. We divide w_i^μ by its norm to project it onto S^{N-1} , and we multiply the gradient by $\delta_{jk} - w_j^\mu w_k^\mu$ to project it onto the tangent space. Projecting the gradient onto the tangent space is a mathematically sound way to obtain the gradient of a function restricted to a manifold embedded in \mathbb{R}^N [89].

The tasks of deriving the saddle-point equations for \bar{p}_y^γ (equation (10)), finding the stationarity condition of the loss with respect to the class weights p_y^γ (equation (6)) and efficiently training p_y^γ all amount to solving the extremization problem

$$\text{Extr}_{\mathbf{p}, \omega, \lambda} \left\{ f(\mathbf{p}) + \sum_{y=0}^C \lambda_y \left(\sum_{\gamma=0}^P p_y^\gamma - p_{\mathbf{q}}(y) \right) + \sum_{\gamma=0}^P \omega^\gamma \left(\sum_{y=0}^C p_y^\gamma - p_{\mathbf{h}}(\gamma) \right) \right\}, \quad (19)$$

where λ_y and ω^γ are Lagrange multipliers that enforce the constraints $\sum_{\gamma} p_y^\gamma = p_{\mathbf{q}}(y)$ and $\sum_y p_y^\gamma = p_{\mathbf{h}}(\gamma)$. In the former task, $f(\mathbf{p})$ is the free entropy (equation (26)) at fixed \mathbf{m} and $\hat{\mathbf{m}}$. In the latter two, it is the negative log-likelihood loss (equation (4)) at a given \mathbf{w} . To find the saddle-point equations and the stationarity conditions of the loss, we derive an implicit solution of equation (19) that is useful in analytical calculations. On the other hand, to train p_y^γ , we design an algorithm to quickly compute a numerical solution of equation (19). We start by noting that, since the extrema of equation (19) are the points where the gradient vanishes, they take the form

$$\begin{aligned} \partial_{p_y^\gamma} f(\mathbf{p}) &= \lambda_y + \omega^\gamma \\ \sum_{\gamma=0}^P p_y^\gamma &= p_{\mathbf{q}}(y) \\ \sum_{y=0}^C p_y^\gamma &= p_{\mathbf{h}}(\gamma). \end{aligned} \quad (20)$$

Define $\bar{p}_y^\gamma = \partial_{p_y^\gamma} \exp(f(\mathbf{p})) = p_y^\gamma \partial_{p_y^\gamma} f(\mathbf{p})$, then

$$\frac{\bar{p}_y^\gamma}{p_y^\gamma} = \lambda_y + \omega^\gamma,$$

along with the previously established row and column sum constraints on \mathbf{p} . Rearranging terms, we get

$$p_y^\gamma = \frac{1}{\lambda_y + \omega^\gamma} \bar{p}_y^\gamma.$$

Using the row and column constraints $\sum_{\gamma=0}^P p_y^\gamma = p_{\mathbf{q}}(y)$ and $\sum_{y=0}^C p_y^\gamma = p_{\mathbf{h}}(\gamma)$, we find that the Lagrange multipliers λ_y and ω^γ solve the nonlinear equations

$$\begin{aligned} \lambda_y &= \frac{1}{p_{\mathbf{q}}(y)} \sum_{\gamma=0}^P \frac{\lambda_y}{\lambda_y + \omega^\gamma} \bar{p}_y^\gamma \\ \omega^\gamma &= \frac{1}{p_{\mathbf{h}}(\gamma)} \sum_{y=0}^C \frac{\omega^\gamma}{\lambda_y + \omega^\gamma} \bar{p}_y^\gamma, \end{aligned} \quad (21)$$

For conciseness, we define $\zeta_y^\gamma(\bar{\mathbf{p}}; p_{\mathbf{h}}) = \lambda_y(\bar{\mathbf{p}}; p_{\mathbf{h}}) + \omega^\gamma(\bar{\mathbf{p}}; p_{\mathbf{h}})$, where $\lambda_y(\bar{\mathbf{p}}; p_{\mathbf{h}})$ and $\omega^\gamma(\bar{\mathbf{p}}; p_{\mathbf{h}})$ are the λ_y and ω^γ solving equation (21) at given P and $p_{\mathbf{h}}$. Using these definitions, we find the implicit solution $p_y^\gamma = \bar{p}_y^\gamma / \zeta_y^\gamma(\bar{\mathbf{p}}; p_{\mathbf{h}})$. As wanted, this equation is useful in analytical calculations involving the saddle-point equations and the stationarity conditions of the loss. However, it is also quite slow to solve numerically for a given $\bar{\mathbf{p}}$, as reported in multiple studies [90–92]. Therefore, although we can train p_y^γ by iterating $p_y^\gamma = \bar{p}_y^\gamma / \zeta_y^\gamma(\bar{\mathbf{p}}; p_{\mathbf{h}})$, it is not a very efficient method.

In order to efficiently train p_y^γ , we devise a faster way to solve equation (20) than through equation (21). Exponentiating both sides of the first line of equation (20), we find

$$\exp\left(\eta \partial_{p_y^\gamma} f(\mathbf{p})\right) \exp(-\eta \lambda_y) \exp(-\eta \omega^\gamma) = 1$$

$$\exp(-\eta\omega^\gamma) p_y^\gamma \exp\left(\eta\partial_{p_y^\gamma} f(\mathbf{p})\right) \exp(-\eta\lambda_y) = p_y^\gamma,$$

where η is an arbitrary scalar that will play the role of a learning rate. Fix $k_y^\gamma = p_y^\gamma \exp\left(\eta\partial_{p_y^\gamma} f(\mathbf{p})\right)$. By the Sinkhorn–Knopp theorem [93, 94], there is a rescaled matrix of the form $p_y^{\prime\gamma} = D_L^\gamma(\mathbf{k}) k_y^\gamma D_R^\gamma(\mathbf{k})$ that satisfies the same constraints $p_y^{\prime\gamma} \geq 0$, $\sum_{\gamma=0}^P p_y^{\prime\gamma} = p_{\mathbf{q}}(y)$ and $\sum_{y=0}^C p_y^{\prime\gamma} = p_{\mathbf{h}}(\gamma)$ as p_y^γ if some technical conditions are satisfied [95]. Moreover, if $p_y^{\prime\gamma}$ exists, then it is unique [96, 97], and we can quickly compute suitable scaling factors $D_L^\gamma(\mathbf{k})$ and $D_R^\gamma(\mathbf{k})$ for $\mathbf{k} = \left\{k_y^\gamma\right\}_{0 \leq y \leq C}^{0 \leq \gamma \leq P}$ using the Sinkhorn–Knopp algorithm first proposed in [98] and theoretically justified in subsequent works [93, 94]. See [95] for a concise review of all these points. $p_y^{\prime\gamma}$ is generally not equal to p_y^γ . However, we observe that the iteration

$$\begin{aligned} k_y^\gamma(t) &= p_y^\gamma(t) \exp\left(\eta\partial_{p_y^\gamma(t)} f(\mathbf{p}(t))\right) \\ p_y^\gamma(t+1) &= D_L^\gamma(\mathbf{k}(t)) k_y^\gamma(t) D_R^\gamma(\mathbf{k}(t)) \end{aligned} \quad (22)$$

converges at small η (for example $\sim 0.1/P$), which means that it must converge to a solution of equation (20). Since the Sinkhorn–Knopp algorithm is fast, it is significantly more efficient to iterate equation (22) than $p_y^\gamma = \bar{p}_y^\gamma / \zeta_y^\gamma(\bar{\mathbf{p}}; p_{\mathbf{h}})$ to solve equation (20).

In practice, we train p_y^γ using a stochastic variant of equation (22) where the gradient is estimated over small batches of data and smoothed with a momentum hyperparameter. Furthermore, we multiply the gradient by $\delta_{yy'} - p_y^\gamma / p_{\mathbf{h}}(\gamma)$ to reduce the size of the components of $\exp\left(\eta\partial_{p_y^\gamma(t)} f(\mathbf{p}(t))\right)$ that move p_y^γ away from the constraints $\sum_{\gamma=0}^P p_y^\gamma = p_{\mathbf{q}}(y)$ and $\sum_{y=0}^C p_y^\gamma = p_{\mathbf{h}}(\gamma)$. This step is a simple approximation of the gradient projection proposed in [99].

Appendix D. Initialization and learning rate

To make all our figures, except figure 6, we initialize the memories \mathbf{w}^μ uniformly at random on the unit hypersphere S^{N-1} [100]. To make figure 6, we instead use the algorithm of [101] to sample the initial memories \mathbf{w}^μ from a vMF distribution (see appendix B) with mean direction $\tilde{\mathbf{x}}^* = \mathbf{x}^* / \|\mathbf{x}^*\|$, where $\mathbf{x}^* = \frac{1}{N} \sum_{\mu^*=1}^{P^*} \mathbf{x}^{*\mu^*}$ is the mean of the patterns $\mathbf{x}^{*\mu^*}$ and $\|\mathbf{x}^*\| = \sqrt{\sum_{i=1}^N [\bar{x}_i^*]^2}$. By construction, the mean direction $\tilde{\mathbf{x}}^*$ is the ordered solution of equation (10) with the most permutation symmetries, or in other words the root of the learning dynamics tree shown in figure 6, so initializing the memories around it helps reveal the tree structure.

We use a learning rate of $\eta \sim 0.1$ to train the memories \mathbf{w}^μ . Based on our experience, \mathbf{p} trains well when its own learning rate is approximately $\left(1 + \frac{p_{\mathbf{h}}(0)}{1-p_{\mathbf{h}}(0)}\right)P$ times smaller. Without this rescaling, the multiplicative update factor $\exp\left(\eta\partial_{p_y^\gamma} L(\mathbf{w}, \mathbf{p})\right)$ described in appendix C can be relatively large even when η is relatively small, making it ill behaved. Equivalently, we can also train $\mathbf{g}_y^\gamma = \left(1 + \frac{p_{\mathbf{h}}(0)}{1-p_{\mathbf{h}}(0)}\right)P\mathbf{p}$ with the same learning rate η as \mathbf{w} . We adopt this approach in our code available at [30].

Let $P_0 = \frac{p_{\mathbf{h}}(0)}{1-p_{\mathbf{h}}(0)}P$ so that $\mathbf{g}_y^\gamma = (P+P_0)p_y^\gamma$. In terms of $\mathbf{g}_y^{\text{fixed},\gamma} = (P+P_0)p_y^{\text{fixed},\gamma}$, the duplicated parameters (equation (13)) of the saddle-point hierarchy principle are

$$\begin{aligned} \bar{x}_i^{\text{dupli},\mu} &= \begin{cases} \bar{x}_i^{\text{fixed},\mu} & 0 < \mu \leq P \\ \bar{x}_i^{\text{fixed},\mu-P} & P < \mu \leq P+R \end{cases} \\ \bar{g}_y^{\text{dupli},\gamma} &= \frac{P+R}{P} \begin{cases} \bar{g}_y^{\text{fixed},0} & \gamma = 0 \\ \frac{1}{2}\bar{g}_y^{\text{fixed},\gamma} & 0 < \gamma \leq R \\ \bar{g}_y^{\text{fixed},\gamma} & R < \gamma \leq P \\ \frac{1}{2}\bar{g}_y^{\text{fixed},\gamma-P} & P < \gamma \leq P+R \end{cases} \\ m^{\text{dupli},\mu*\gamma} &= \begin{cases} m^{\text{fixed},\mu*0} & \gamma = 0 \\ m^{\text{fixed},\mu*\gamma} & 0 < \gamma \leq P \\ m^{\text{fixed},\mu*,\gamma-P} & P < \gamma \leq P+R \end{cases} \\ g_y^{\text{dupli},\gamma} &= \frac{P+R}{P} \begin{cases} g_y^{\text{fixed},0} & \gamma = 0 \\ \frac{1}{2}g_y^{\text{fixed},\gamma} & 0 < \gamma \leq R \\ g_y^{\text{fixed},\gamma} & R < \gamma \leq P \\ \frac{1}{2}g_y^{\text{fixed},\gamma-P} & P < \gamma \leq P+R \end{cases} \end{aligned}$$

Algorithm 2. Rescaled splitting steepest descent [29, 75].

- 1: Preallocate space for a DAM with P_{\max} hidden units and the corresponding weights \mathbf{w} and \mathbf{g}
- 2: Initialize the weights \mathbf{w}^μ and \mathbf{g}^μ connected to the P_{cur} first hidden units $\mu \in \{1, \dots, P_{\text{cur}}\}$, as well as \mathbf{g}^0
- 3: $\min L(\mathbf{w}, \mathbf{g})$ with SGD
- 4: **while** $P_{\text{cur}} < P_{\max}$ **do**
- 5: Identify a subset $\mu_{\text{copy}} \subseteq \{1, \dots, P_{\text{cur}}\}$ of $R \leq P_{\max} - P_{\text{cur}}$ hidden units to split, **return** if empty
- 6: Let $\mu_{\text{paste}} = \{P_{\text{cur}} + 1, \dots, P_{\text{cur}} + R\}$ and $\mu_{\text{dupli}} = \{1, \dots, P_{\text{cur}} + R\}$
- 7: Build weights $\mathbf{w}^{\mu_{\text{paste}}} = \mathbf{w}^{\mu_{\text{copy}}}$ for μ_{paste}
- 8: Rescale $\mathbf{g}^{\mu_{\text{copy}}} \leftarrow \mathbf{g}^{\mu_{\text{copy}}} / 2$ and $g_{\text{h}}(\mu_{\text{split}}) \leftarrow g_{\text{h}}(\mu_{\text{split}}) / 2$
- 9: Build weights $\mathbf{g}^{\mu_{\text{paste}}} = \mathbf{g}^{\mu_{\text{copy}}}$ and $g(\mu_{\text{paste}}) = g(\mu_{\text{copy}})$ for μ_{paste}
- 10: Rescale $\mathbf{g}^{\mu_{\text{dupli}}} \leftarrow \frac{P_{\text{cur}} + R}{P_{\text{cur}}} \mathbf{g}^{\mu_{\text{dupli}}}$ and $g_{\text{h}}(\mu_{\text{dupli}}) \leftarrow \frac{P_{\text{cur}} + R}{P_{\text{cur}}} g_{\text{h}}(\mu_{\text{dupli}})$
- 11: Update $P_{\text{cur}} \leftarrow P_{\text{cur}} + R$
- 12: Escape the saddle point by 2nd order descent of $L(\mathbf{w}, \mathbf{g})$ w.r.t. \mathbf{w}
- 13: $\min L(\mathbf{w}, \mathbf{g})$ with SGD
- 14: **end while**
- 15: **return**

▷ See appendix K and [29, 75] for details about steps 5 and 12

$$\text{along with } g_{\text{h}}(\gamma) = \frac{P+R}{P} \begin{cases} g_{\text{h}}^{\text{given}}(0) & \gamma = 0 \\ \frac{1}{2} g_{\text{h}}^{\text{given}}(\gamma) & 0 < \gamma \leq R \\ g_{\text{h}}^{\text{given}}(\gamma) & R < \gamma \leq P \\ \frac{1}{2} g_{\text{h}}^{\text{given}}(\gamma - P) & P < \gamma \leq P + R, \end{cases}$$

where $g_{\text{h}}^{\text{given}}(\gamma) = (P + P_0) p_{\text{h}}^{\text{given}}(\gamma)$ and $g_{\text{h}}(\gamma) = (P + P_0) p_{\text{h}}(\gamma)$. The newly introduced scaling factor of $\frac{P+R}{P}$ must be taken into account in our implementation of splitting steepest descent (algorithm 1), which gives algorithm 2, shown below.

We initialize the weights \mathbf{g}^γ and the hidden unit distribution $g_{\text{h}}(\gamma)$ according to

$$g_y^\gamma = \begin{cases} P_0 p_{\text{q}}(y) & \gamma = 0 \\ p_{\text{q}}(y) & \gamma > 0 \end{cases}$$

$$\text{and } g_{\text{h}}(\gamma) = \begin{cases} \frac{P_0}{P+P_0} & \gamma = 0 \\ \frac{1}{P+P_0} & \gamma > 0. \end{cases}$$

When we train our DAM without splitting steepest descent, we set $P_0 = 1$. As discussed in [102, 103], the resulting weights are, in some sense, the most ‘typical’ for the constraints $g_y^\gamma \geq 0$, $\sum_{\gamma=0}^P g_y^\gamma = (P+1) p_{\text{q}}(y)$ and $\sum_{y=0}^C g_y^\gamma = (P+1) p_{\text{h}}(\gamma) = 1$ (see section 2). With splitting steepest descent, we set $P_0 = P_{\text{cur}}/P_{\text{final}}$ to ensure that $g_{\text{h}}(0) = \sum_{y=0}^C g_y^0$ is approximately the same size as the other entries of $g_{\text{h}}(\gamma)$ after training. For simplicity, we set $p_{\text{q}}(y)$ to the proportions of classes in the data during supervised training. For unsupervised training (see section 4.2), we break the permutation symmetry between the different classes by using different values for all the entries of $p_{\text{q}}(y)$. Otherwise, class weights with unbroken permutation symmetries stay stuck at their initial conditions.

Appendix E. Important symbols introduced in the model section

The following table summarizes the important symbols introduced in section 2 and their meaning.

Appendix F. Stationarity conditions of the loss

Since we constrain the memories \mathbf{w}^μ to have unit norm (see section 2), the method of Lagrange multipliers tells us any set of memories \mathbf{w} that minimizes equation (4) must solve the extremization problem

$$\text{Extr}_{\mathbf{w}, \varphi} \left\{ L(\mathbf{w}, \mathbf{p}) + \frac{1}{2} \sum_{\mu=1}^P \varphi^\mu \left(\sum_{i=1}^N [w_i^\mu]^2 - 1 \right) \right\}.$$

Table 2. Summary of the important symbols introduced in section 2 and their meaning.

Symbol	Meaning
$\mathbf{x} = \{x_i\}_{i=1}^N$	Data layer
$\mathbf{h} = \{h_\mu\}_{\mu=1}^P$	Hidden layer
$\mathbf{q} = \{q_y\}_{y=1}^C$	Class layer
N	# data units
P	# hidden units (and clusters)
C	# class units (and classes)
\mathbf{J}	Set of all trainable weights
$\mathbf{w} = \{w_i^\mu\}_{1 \leq \mu \leq P, 1 \leq i \leq N}$	Weight matrix of the data layer
$\mathbf{p} = \{p_y^\gamma\}_{0 \leq \gamma \leq P, 0 \leq y \leq C}$	Weight matrix of the class layer
$p_q(y)$	Marginal distribution $\sum_{\gamma=0}^P p_y^\gamma$
$p_h(\gamma)$	Marginal distribution $\sum_{y=0}^C p_y^\gamma$
β	Inverse temperature
$\Omega_N(\beta)$	Normalization constant of the vMF distribution
$\mathbf{x}^{*\mu} = \{x_i^{*\mu}\}_{i=1}^N$	Pattern from a given dataset
$\mathbf{q}^{*\mu} = \{q_y^{*\mu}\}_{y=0}^C$	Soft label of the pattern
P^*	# hidden units (and clusters) of the teacher
$\mathbf{w}^* = \{w_i^{*\mu}\}_{1 \leq \mu \leq P^*, 1 \leq i \leq N}$	Weight matrix of the data layer of the teacher DAM
$\mathbf{p}^* = \{p_y^{*\gamma}\}_{0 \leq \gamma \leq P^*, 0 \leq y \leq C}$	Weight matrix of the class layer of the teacher DAM
\mathbf{g}^*	Rescaled weight matrix $\mathbf{g}^* = P^* \mathbf{p}^*$ of the class layer of the teacher DAM
$p_q^*(y)$	Marginal distribution $\sum_{\gamma=0}^{P^*} p_y^{*\gamma}$
$p_h^*(\gamma)$	Marginal distribution $\sum_{y=0}^C p_y^{*\gamma}$
β^*	Inverse temperature of the teacher
$D = \{\mathbf{x}^c, \mathbf{y}^c\}_{c=1}^M$	Set of data points \mathbf{x}^c and labels \mathbf{y}^c generated by the teacher
M	Number of data points generated by the teacher
$\alpha = M/N$	Load of data generated by the teacher

The extrema are the points where the gradient vanishes, so they take the form

$$\partial_{w_i^\mu} L(\mathbf{w}, \mathbf{p}) + \varphi^\mu w_i^\mu = 0 \sum_{i=1}^N [w_i^\mu]^2 = 1$$

for all $1 \leq \mu \leq P$. We calculate $\partial_{w_i^\mu} L(\mathbf{w}, \mathbf{p})$ and write the solution in the more explicit form

$$\begin{aligned} \bar{w}_i^\mu &= \sum_{\mu^*=1}^{P^*} x_i^{*\mu^*} \sum_{y=0}^C q_y^{*\mu^*} \sigma_\mu \left(\beta m^{\mu^*}(\mathbf{x}^*, \mathbf{w}) + \log [\mathbf{p}_y] \right) \\ w_i^\mu &= \frac{\bar{w}_i^\mu}{\sqrt{\sum_{j=1}^N [\bar{w}_j^\mu]^2}}, \end{aligned}$$

for all $1 \leq \mu \leq P$, where $\sigma_\gamma(x^{\gamma*}) = \frac{\exp(x^{\gamma*})}{\sum_{\nu=0}^P \exp(x^{\gamma*\nu})}$, $m^{\mu^*}(\mathbf{x}^*, \mathbf{w}) = \sum_{i=1}^N x_i^{*\mu^*} w_i^\mu$ for $1 \leq \mu \leq P$ and $m^{\mu^*0}(\mathbf{x}, \mathbf{w}) = \frac{1}{\beta} \log [\Omega_N(\beta) / \Omega_N(0)]$.

Similarly, any set of class weights \mathbf{p} that minimizes equation (4) must solve the extremization problem

$$\text{Extr}_{\mathbf{p}, \omega, \lambda} \left\{ L(\mathbf{w}, \mathbf{p}) + \sum_{y=0}^C \lambda_y \left(\sum_{\gamma=0}^P p_y^\gamma - p_q(y) \right) + \sum_{\gamma=0}^P \omega^\gamma \left(\sum_{y=0}^C p_y^\gamma - p_h(\gamma) \right) \right\}.$$

In appendix C, we show that its solution is

$$\begin{aligned} \bar{p}_y^\gamma &= \sum_{\mu^*=1}^{P^*} q_y^{*\mu^*} \sigma_\gamma \left(\beta m^{\mu^*}(\mathbf{x}, \mathbf{w}) + \log [\mathbf{p}_y] \right) \\ p_y^\gamma &= \frac{\bar{p}_y^\gamma}{\zeta_y^\gamma(\bar{\mathbf{p}}; p_h)} \end{aligned}$$

for all $0 \leq \gamma \leq P$, where the normalization constant $\zeta_y^\gamma(\bar{\mathbf{p}}; p_h)$ is defined using equation (21) of appendix C. Combining these equations with the ones for \mathbf{w} , we find the stationarity conditions

$$\begin{aligned}\bar{w}_i^\mu &= \sum_{\mu^*=1}^{P^*} x_i^{*\mu^*} \sum_{y=0}^C q_y^{*\mu^*} \sigma_\mu \left(\beta m^{\mu^*}(\mathbf{x}^*, \mathbf{w}) + \log[\mathbf{p}_y] \right) \\ \bar{p}_y^\gamma &= \sum_{\mu^*=1}^{P^*} q_y^{*\mu^*} \sigma_\gamma \left(\beta m^{\mu^*}(\mathbf{x}^*, \mathbf{w}) + \log[\mathbf{p}_y] \right) \\ w_i^\mu &= \frac{\bar{w}_i^\mu}{\sqrt{\sum_{j=1}^N [\bar{w}_j^\mu]^2}} \\ p_y^\gamma &= \frac{\bar{p}_y^\gamma}{\zeta_y^\gamma(\bar{\mathbf{p}}; p_h)}\end{aligned}$$

for all $1 \leq \mu \leq P$ and $0 \leq \gamma \leq P$.

Appendix G. Replicated partition function and free entropy

Suppose a student DAM model is trained using a dataset $\mathcal{D} = \{\mathbf{x}^c, y^c\}_{c=1}^M$ of M i.i.d. examples \mathbf{x}^c with labels y^c . By Bayes' theorem, the student weights \mathbf{w} and \mathbf{p} follow the distribution

$$P_\beta(\mathbf{w}, \mathbf{p} | \mathcal{D}) = \mathcal{Z}(\mathcal{D})^{-1} P(\mathbf{w}, \mathbf{p}) \prod_{c=1}^M P_\beta(\mathbf{x}^c, y^c | \mathbf{w}, \mathbf{p}),$$

where $\mathcal{Z}(\mathcal{D}) = \mathbb{E}_{\mathbf{w}, \mathbf{p}} \left[\prod_{c=1}^M P_\beta(\mathbf{x}^c, y^c | \mathbf{w}, \mathbf{p}) \right]$. Assuming the examples are sampled from a teacher DAM with weights \mathbf{w}^* and \mathbf{p}^* , the average replicated partition takes the form

$$\begin{aligned}\langle \mathcal{Z}^L \rangle &= \sum_{\{y^c\}_{c=1}^M} \int \left[\prod_{c=1}^M d\mathbf{x}^c \right] \mathbb{E}_{\mathbf{w}^*, \mathbf{p}^*} \left[\prod_{c=1}^M P_\beta(\mathbf{x}^c, y^c | \mathbf{w}^*, \mathbf{p}^*) \right] \mathcal{Z}(\mathcal{D})^L \\ &= \sum_{\{y^c\}_{c=1}^M} \int \left[\prod_{c=1}^M d\mathbf{x}^c \right] \mathbb{E}_{\mathbf{w}^*, \mathbf{p}^*} \left[\prod_{c=1}^M P_\beta(\mathbf{x}^c, y^c | \mathbf{w}^*, \mathbf{p}^*) \right] \prod_{a=1}^L \mathbb{E}_{\mathbf{w}^a, \mathbf{p}^a} \left[\prod_{c=1}^M P_\beta(\mathbf{x}^c, y^c | \mathbf{w}^a, \mathbf{p}^a) \right] \\ &= \mathbb{E}_{\mathbf{w}^*, \mathbf{w}} \mathbb{E}_{\mathbf{p}^*, \mathbf{p}} \left[\prod_{c=1}^M \sum_{y^c=0}^C \int_{S^{N-1}} d\mathbf{x}^c P_\beta(\mathbf{x}^c, y^c | \mathbf{w}^*, \mathbf{p}^*) \prod_{a=1}^L P_\beta(\mathbf{x}^c, y^c | \mathbf{w}^a, \mathbf{p}^a) \right] \\ &= \mathbb{E}_{\mathbf{w}^*, \mathbf{w}} \mathbb{E}_{\mathbf{p}^*, \mathbf{p}} \left[\left(\sum_{y=0}^C \int_{S^{N-1}} d\mathbf{x} P_\beta(\mathbf{x}, y | \mathbf{w}^*, \mathbf{p}^*) \prod_{a=1}^L P_\beta(\mathbf{x}, y | \mathbf{w}^a, \mathbf{p}^a) \right)^M \right],\end{aligned}$$

where we redefined $\mathbf{w} = \{\mathbf{w}^a\}_{a=1}^L$ and $\mathbf{p} = \{\mathbf{p}^a\}_{a=1}^L$. In order to simplify the argument of $(\cdot)^M$, it is convenient to write equation (3) in the form

$$P_\beta(\mathbf{x}, y | \mathbf{w}^a, \mathbf{p}^a) = \sum_{\gamma=0}^P \Omega_N(\beta[1 - \delta_{\gamma 0}])^{-1} p_y^{a\gamma} \exp\left(\beta[1 - \delta_{\gamma 0}] \sum_{i=1}^N w_i^{a\gamma} x_i\right), \quad (23)$$

where the value of \mathbf{w}^{a0} is arbitrary. Until the end of this section, all sums over the hidden units will include the index 0 unless explicitly indicated otherwise. Define $I_y(\mathbf{x}) = P_\beta(\mathbf{x}, y | \mathbf{w}^*, \mathbf{p}^*) \prod_{a=1}^L P_\beta(\mathbf{x}, y | \mathbf{w}^a, \mathbf{p}^a)$, then

$$\begin{aligned}I_y(\mathbf{x}) &= \left[\sum_{\gamma^*=0}^{P^*} \Omega_N(\beta^*[1 - \delta_{\gamma^* 0}])^{-1} p_y^{*\gamma^*} \exp\left(\beta^*[1 - \delta_{\gamma^* 0}] \sum_{i=1}^N w_i^{*\gamma^*} x_i\right) \right] \\ &\quad \times \prod_{a=1}^L \left[\sum_{\gamma=0}^P \Omega_N(\beta[1 - \delta_{\gamma 0}])^{-1} p_y^{a\gamma} \exp\left(\beta[1 - \delta_{\gamma 0}] \sum_{i=1}^N w_i^{a\gamma} x_i\right) \right] \\ &= \sum_{\gamma^* \gamma_1 \dots \gamma_L} \Omega_N(\beta^*[1 - \delta_{\gamma^* 0}])^{-1} \left[\prod_a \Omega_N(\beta[1 - \delta_{\gamma_a 0}]) \right]^{-1}\end{aligned}$$

$$\times p_y^{*\gamma_*} \left[\prod_a p_y^{a\gamma_a} \right] \exp \left(\beta^* [1 - \delta_{\gamma_*0}] \sum_i w_i^{*\gamma_*} x_i + \beta \sum_a [1 - \delta_{\gamma_a0}] \sum_i w_i^{a\gamma_a} x_i \right).$$

We will now evaluate the integral $\int_{S^{N-1}} d\mathbf{x} I_y(\mathbf{x})$. Using equation (18) of appendix B with $\rho = \rho_{\gamma_*\gamma} := \frac{1}{N} \sqrt{\sum_i [\beta^* [1 - \delta_{\gamma_*0}] w_i^{*\gamma_*} + \beta \sum_a [1 - \delta_{\gamma_a0}] w_i^{a\gamma_a}]^2}$, we get

$$\int_{S^{N-1}} d\mathbf{x} \exp \left(\beta^* [1 - \delta_{\gamma_*0}] \sum_i w_i^{*\gamma_*} x_i + \beta \sum_a [1 - \delta_{\gamma_a0}] \sum_i w_i^{a\gamma_a} x_i \right) \approx \Omega_N(0) \left(1 + (2\rho_{\gamma_*\gamma})^2 \right)^{-1/4} \exp \left[\left(\frac{N}{2} - 1 \right) \eta(2\rho_{\gamma_*\gamma}) \right],$$

in the limit of large N . Since $\sqrt{\sum_i (w_i^{*\gamma_*})^2} = 1$, the square of $\rho_{\gamma_*\gamma}$ simplifies to

$$\begin{aligned} (\rho_{\gamma_*\gamma})^2 &= \left(\frac{1}{N} \sqrt{\sum_i \left[\beta^* [1 - \delta_{\gamma_*0}] w_i^{*\gamma_*} + \beta \sum_a [1 - \delta_{\gamma_a0}] w_i^{a\gamma_a} \right]^2} \right)^2 \\ &= v^2 [1 - \delta_{\gamma_*0}] + 2 \frac{\beta}{N} v [1 - \delta_{\gamma_*0}] \sum_a [1 - \delta_{\gamma_a0}] \sum_i w_i^{*\gamma_*} w_i^{a\gamma_a} \\ &\quad + \frac{\beta^2}{N^2} \sum_{a,b} [1 - \delta_{\gamma_a0}] [1 - \delta_{\gamma_b0}] \sum_i w_i^{a\gamma_a} w_i^{b\gamma_b}, \end{aligned}$$

where $v = \frac{\beta^*}{N}$. To leading order in $\frac{\beta}{N}$ we find

$$\begin{aligned} \left(1 + (2\rho_{\gamma_*\gamma})^2 \right)^{-1/4} &\approx \left(1 + (2v [1 - \delta_{\gamma_*0}])^2 \right)^{-1/4} \\ \text{and } \left(\frac{N}{2} - 1 \right) \eta(2\rho_{\gamma_*\gamma}) &\approx \left(\frac{N}{2} - 1 \right) \eta(2v [1 - \delta_{\gamma_*0}]) + \beta_{\text{eff}} [1 - \delta_{\gamma_*0}] \sum_a [1 - \delta_{\gamma_a0}] \sum_i w_i^{*\gamma_*} w_i^{a\gamma_a} \\ &\quad + \frac{[\beta \xi_{\gamma_*}]^2}{2N} \sum_{a,b} [1 - \delta_{\gamma_a0}] [1 - \delta_{\gamma_b0}] \sum_i w_i^{a\gamma_a} w_i^{b\gamma_b}, \end{aligned}$$

where $\beta_{\text{eff}} = \frac{2v}{\sqrt{[2v]^2 + 1 + 1}} \beta$ and $\xi_{\gamma_*} = \delta_{\gamma_*0} + \sqrt{\frac{2}{\sqrt{[2v]^2 + 1 + 1}}} [1 - \delta_{\gamma_*0}]$. Assuming that $v \gg 1/N$ and $\beta \ll N$, we drop the last term. Finally, we use equation (18) backwards with $\rho = v [1 - \delta_{\gamma_*0}]$ and obtain

$$\begin{aligned} \Omega_N(\beta^* [1 - \delta_{\gamma_*0}])^{-1} \int_{S^{N-1}} d\mathbf{x} \exp \left(\beta^* [1 - \delta_{\gamma_*0}] \sum_i w_i^{*\gamma_*} x_i + \beta \sum_a [1 - \delta_{\gamma_a0}] \sum_i w_i^{a\gamma_a} x_i \right) \\ \approx \exp \left(\beta_{\text{eff}} [1 - \delta_{\gamma_*0}] \sum_a [1 - \delta_{\gamma_a0}] \sum_i w_i^{*\gamma_*} w_i^{a\gamma_a} \right), \end{aligned}$$

from which we conclude that

$$\begin{aligned} \sum_{y=0}^C \int_{S^{N-1}} d\mathbf{x} I_y(\mathbf{x}) &\approx \sum_{\gamma_* \gamma_1 \dots \gamma_L} \sum_y p_y^{*\gamma_*} \left[\prod_a p_y^{a\gamma_a} \right] \left[\prod_a \Omega_N(\beta [1 - \delta_{\gamma_a0}]) \right]^{-1} \\ &\quad \times \exp \left(\beta_{\text{eff}} [1 - \delta_{\gamma_*0}] \sum_a [1 - \delta_{\gamma_a0}] \sum_i w_i^{*\gamma_*} w_i^{a\gamma_a} \right). \end{aligned}$$

We define $\alpha = M/N$, introduce the order parameter

$$m^{a\gamma_*\nu} \text{ for } \sum_i w_i^{*\gamma_*} w_i^{a\nu}, \tag{24}$$

and insert into $\langle Z^L \rangle$ the identity operator

$$1 = \int_{\mathbb{R}} \prod_{\gamma_*, \nu; a} dm^{a\gamma_*\nu} \delta \left(m^{a\gamma_*\nu} - \sum_i w_i^{*\gamma_*} w_i^{a\nu} \right)$$

$$= \int_{\mathbb{R}} \prod_{\gamma_*, \nu; a} dm^{a\gamma_*\nu} \int_{i\mathbb{R}} \prod_{\gamma_*, \nu; a} d\hat{m}^{a\gamma_*\nu} \exp \left\{ \frac{\beta_{\text{eff}}\alpha}{P^*} \sum_{\gamma_*, \nu; a} \hat{m}^{a\gamma_*\nu} \left(\sum_i w_i^{*\gamma_*} w_i^{a\nu} - m^{a\gamma_*\nu} \right) \right\}$$

so that we can rewrite it as

$$\begin{aligned} \langle \mathcal{Z}^L \rangle &= \int \prod_{\gamma_*, \nu; a} d\hat{m}^{a\gamma_*\nu} dm^{a\gamma_*\nu} \mathbb{E}_{\mathbf{w}^*, \mathbf{w}} \exp \{ NH_S(\mathbf{w}, \mathbf{w}^*; \hat{\mathbf{m}}) \} \exp \{ -NH_Q(\mathbf{m}, \hat{\mathbf{m}}) \} \\ &\times \mathbb{E}_{\mathbf{p}^*, \mathbf{p}} \exp \left\{ \alpha N \log \left[\sum_{\gamma_* \gamma_1 \dots \gamma_L} \sum_y p_y^{*\gamma_*} \left[\prod_a p_y^{a\gamma_a} \right] \exp \{ H_E(\gamma, \gamma_*; \mathbf{m}) \} \right] \right\}, \end{aligned} \tag{25}$$

$$\begin{aligned} \text{where } H_S(\mathbf{w}, \mathbf{w}^*; \hat{\mathbf{m}}) &= \frac{\beta_{\text{eff}}\alpha}{P^*} \sum_{\gamma_*, \nu; a} \hat{m}^{a\gamma_*\nu} \sum_i w_i^{*\gamma_*} w_i^{a\nu} \\ H_Q(\mathbf{m}, \hat{\mathbf{m}}) &= \frac{\beta_{\text{eff}}\alpha}{P^*} \sum_{\gamma_*, \nu; a} \hat{m}^{a\gamma_*\nu} m^{a\gamma_*\nu}, \\ H_E(\gamma, \gamma_*; \mathbf{m}) &= - \sum_a \log [\Omega_N(\beta [1 - \delta_{\gamma_a 0})]] + \beta_{\text{eff}} [1 - \delta_{\gamma_* 0}] \sum_a [1 - \delta_{\gamma_a 0}] m^{a\gamma_*\gamma_a}. \end{aligned}$$

Using equation (18) of appendix B, the exponential of H_S integrates to

$$\begin{aligned} &\mathbb{E}_{\mathbf{w}} \exp \{ NH_S(\mathbf{w}, \mathbf{w}^*; \hat{\mathbf{m}}) \} \\ &= \Omega_N(0)^{-PL} \int_{S^{N-1}} d\mathbf{w} \exp \left(\frac{\beta_{\text{eff}}\alpha}{P^*} \sum_{\gamma_*, \nu; a} \hat{m}^{a\gamma_*\nu} \sum_i w_i^{*\gamma_*} w_i^{a\nu} \right) \\ &\approx \prod_{\nu; a} \exp \left[-\frac{1}{4} \log \left(1 + \left[\frac{2\beta_{\text{eff}}\alpha}{P^*} \right]^2 \sum_i \left[\sum_{\gamma_*} \hat{m}^{a\gamma_*\nu} w_i^{*\gamma_*} \right]^2 \right) \right] \\ &\times \exp \left[\left(\frac{N}{2} - 1 \right) \eta \left(\frac{2\beta_{\text{eff}}\alpha}{P^*} \sqrt{\sum_i \left[\sum_{\gamma_*} \hat{m}^{a\gamma_*\nu} w_i^{*\gamma_*} \right]^2} \right) \right]. \end{aligned}$$

In order to continue the calculations, we make the replica-symmetric ansatz

$$\begin{aligned} m^{a\gamma_*\nu} &= m^{\gamma_*\nu} \text{ and } \hat{m}^{a\gamma_*\nu} = \hat{m}^{\gamma_*\nu} \text{ for all } a; \gamma_*, \nu \\ p_y^{a\nu} &= p_y^\nu \text{ for all } a; \nu \end{aligned}$$

so that we can use the replica trick to simplify

$$\begin{aligned} &\frac{1}{L} \log \left[\sum_{\gamma_* \gamma_1 \dots \gamma_L} \sum_y p_y^{*\gamma_*} \left[\prod_a p_y^{a\gamma_a} \right] \exp \{ H_E(\gamma, \gamma_*; \mathbf{m}) \} \right] \\ &= \frac{1}{L} \log \left[\sum_{\gamma_*} \sum_y p_y^{*\gamma_*} \sum_{\gamma_1 \dots \gamma_L} \left[\prod_a p_y^{a\gamma_a} \right] \exp \{ H_E(\gamma, \gamma_*; \mathbf{m}) \} \right] \\ &\approx \sum_{\gamma_*} \sum_y p_y^{*\gamma_*} \log \left[\sum_{\gamma} p_y^\gamma \Omega_N(\beta [1 - \delta_{\gamma 0}])^{-1} \exp(\beta_{\text{eff}} [1 - \delta_{\gamma_* 0}] [1 - \delta_{\gamma 0}] m^{\gamma_*\gamma}) \right], \end{aligned}$$

as we take L to zero. Similarly, we get

$$\frac{1}{L} \log [\mathbb{E}_{\mathbf{w}} \exp \{NH_S(\mathbf{w}, \mathbf{w}^*; \hat{\mathbf{m}})\}] = -\frac{1}{4} \log \left(1 + \left[\frac{2\beta_{\text{eff}}\alpha}{P^*} \right]^2 \sum_i \left[\sum_{\gamma_*} \hat{m}^{\gamma_*\nu} w_i^{*\gamma_*} \right]^2 \right) + \left(\frac{N}{2} - 1 \right) \eta \left(\frac{2\beta_{\text{eff}}\alpha}{P^*} \sqrt{\sum_i \left[\sum_{\gamma_*} \hat{m}^{\gamma_*\nu} w_i^{*\gamma_*} \right]^2} \right).$$

From now on, we assume for simplicity that $\sum_y p_y^{*0} = p_{\mathbf{h}}^*(0) = 0$ and $\sum_y p_y^{*\mu_*} = p_{\mathbf{h}}^*(\mu_*) = 1/P^*$ for all $\mu_* > 0$. Defining $g_y^{*\mu_*} = P^* p_y^{*\mu_*} = P_{\beta}(y|\mu_*, \mathbf{J})$ for all $\mu_* > 0$, $\mathbf{g}^* = \left\{ g_y^{*\gamma_*} \right\}_{0 \leq y \leq C}^{1 \leq \mu_* \leq P^*}$ and $\varrho = \frac{\alpha}{P^*} = \frac{M}{P^*N}$, the free entropy f then takes the form

$$f \approx \text{Extr}_{\mathbf{m}, \hat{\mathbf{m}}, \mathbf{p}} \{f(\mathbf{m}, \hat{\mathbf{m}}, \mathbf{p})\} \quad \text{such that} \quad \sum_{\gamma=0}^P p_y^\gamma = p_{\mathbf{q}}(y) \quad \text{and} \quad \sum_{y=0}^C p_y^\gamma = p_{\mathbf{h}}(\gamma) \quad (26)$$

$$\text{with} \quad f(\mathbf{m}, \hat{\mathbf{m}}, \mathbf{p}) = -\beta_{\text{eff}}\varrho \sum_{\gamma_*, \gamma=0}^{P^*, P} \hat{m}^{\gamma_*\gamma} m^{\gamma_*\gamma} + \frac{1}{2} \mathbb{E}_{\mathbf{w}^*} \left[\sum_{\gamma=0}^P \eta \left(2\beta_{\text{eff}}\varrho \sqrt{\sum_{i=1}^N \left[\sum_{\gamma_*=0}^{P^*} \hat{m}^{\gamma_*\gamma} w_i^{*\gamma_*} \right]^2} \right) \right] + \varrho \sum_{\mu_*=1}^{P^*} \sum_{y=0}^C \mathbb{E}_{\mathbf{g}^*} [g_y^{*\mu_*}] \log \left[\sum_{\gamma=0}^P p_y^\gamma \Omega_N(\beta[1 - \delta_{\gamma 0}])^{-1} \exp(\beta_{\text{eff}}[1 - \delta_{\gamma 0}] m^{\mu_*\gamma}) \right]$$

where we approximated the expectation over \mathbf{p} as a (constrained) extremization problem using Laplace’s method. By inspection, the order parameters $m^{0\gamma}$ and $\hat{m}^{0\gamma}$ always vanish, so we ignore them in the incoming derivation of the saddle-point equations.

Appendix H. Saddle-point equations

In this appendix, we adopt the convention $1 \leq \mu_* \leq P^*$, $0 \leq \gamma \leq P$ and $1 \leq \mu \leq P$. By the Lagrange multiplier theorem, any set of class weights \mathbf{p} that extremizes equation (26) must solve the extremization problem

$$\text{Extr}_{\mathbf{p}, \omega, \lambda} \left\{ f(\mathbf{m}, \hat{\mathbf{m}}, \mathbf{p}) + \sum_{y=0}^C \lambda_y \left(\sum_{\gamma=0}^P p_y^\gamma - p_{\mathbf{q}}(y) \right) + \sum_{\gamma=0}^P \omega^\gamma \left(\sum_{y=0}^C p_y^\gamma - p_{\mathbf{h}}(\gamma) \right) \right\}.$$

In appendix C, we show that its solution is

$$\bar{p}_y^\gamma = \sum_{\mu_*=1}^{P^*} \mathbb{E}_{\mathbf{g}^*} [g_y^{*\mu_*}] \sigma_\gamma \left(\beta_{\text{eff}}[1 - \delta_{\gamma 0}] m^{\mu_*} - \log[\Omega_N(\beta[1 - \delta_{\gamma 0})] + \log[\mathbf{p}_y] \right) \\ p_y^\gamma = \frac{\bar{p}_y^\gamma}{\zeta_y^\gamma(\bar{\mathbf{p}}; p_{\mathbf{h}})},$$

for all $0 \leq \gamma \leq P$, where $\zeta_y^\gamma(\bar{\mathbf{p}}; p_{\mathbf{h}})$ is defined using equations (21) of appendix C and $\sigma_\gamma(x^{\mu_*}) = \frac{\exp(x^{\mu_*\gamma})}{\sum_{\kappa=0}^P \exp(x^{\mu_*\kappa})}$ is the softmax function. We extremize equation (26) with respect to the remaining parameters by solving for the gradient equal to zero. $\partial_{m^{\mu_*\mu}} f(\mathbf{m}, \hat{\mathbf{m}}, \mathbf{p}) = 0$ immediately yields

$$\hat{m}^{\mu_*\gamma} = [1 - \delta_{\gamma 0}] \sum_{y=0}^C \mathbb{E}_{\mathbf{g}^*} [g_y^{*\mu_*}] \frac{p_y^\gamma \Omega_N(\beta[1 - \delta_{\gamma 0}])^{-1} \exp(\beta_{\text{eff}}[1 - \delta_{\gamma 0}] m^{\mu_*\gamma})}{\sum_{\kappa=0}^P p_y^\kappa \Omega_N(\beta[1 - \delta_{\kappa 0}])^{-1} \exp(\beta_{\text{eff}}[1 - \delta_{\kappa 0}] m^{\mu_*\kappa})}. \quad (27)$$

On the other hand, $\partial_{\hat{m}^{\mu_*\gamma}} f(\mathbf{m}, \hat{\mathbf{m}}, \mathbf{p}) = 0$ gives an equation that depends on the choice of prior for the teacher memories $\mathbf{w}^{*\mu_*}$. We first investigate the case where the teacher memories $\mathbf{w}^{*\mu_*}$ and the class weights \mathbf{g}^* are distributed uniformly at random over the sets to which they are constrained (see section 2.2). In this scenario, we have $\mathbb{E}_{\mathbf{g}^*} [g_y^{*\mu_*}] = p_{\mathbf{q}}^*(y)$. Moreover, assuming that $P^* \ll N$, random

vectors on the unit sphere are orthonormal with high probability, which means that

$$\begin{aligned} \sum_{i=1}^N \left[\sum_{\mu_*=1}^{P^*} \hat{m}^{\mu_*\gamma} w_i^{*\mu_*} \right]^2 &= \sum_{i=1}^N \sum_{\mu_*, \nu_*=1}^{P^*} \hat{m}^{\mu_*\gamma} \hat{m}^{\nu_*\gamma} w_i^{*\mu_*} w_i^{*\nu_*} \\ &= \sum_{\mu_*=1}^{P^*} [\hat{m}^{\mu_*\gamma}]^2 + \sum_{\mu_* \neq \nu_*} \hat{m}^{\mu_*\gamma} \hat{m}^{\nu_*\gamma} \sum_{i=1}^N w_i^{*\mu_*} w_i^{*\nu_*} \\ &\approx \sum_{\mu_*=1}^{P^*} [\hat{m}^{\mu_*\gamma}]^2. \end{aligned}$$

Therefore, $\partial_{\hat{m}^{\mu_*\gamma}} f(\mathbf{m}, \hat{\mathbf{m}}, \mathbf{p}) = 0$ gives

$$\begin{aligned} 0 &= \partial_{\hat{m}^{\mu_*\gamma}} f(\mathbf{m}, \hat{\mathbf{m}}, \mathbf{p}) \\ 0 &= -\beta_{\text{eff}} \varrho m^{\mu_*\gamma} + \frac{1}{2} \partial_{\hat{m}^{\mu_*\gamma}} \sum_{\kappa=0}^P \eta \left(2\beta_{\text{eff}} \varrho \sqrt{\sum_{\nu_*=1}^{P^*} [\hat{m}^{\nu_*\kappa}]^2} \right) \\ m^{\mu_*\gamma} &= \varsigma \left(2\beta_{\text{eff}} \varrho \sqrt{\sum_{\nu_*=1}^{P^*} [\hat{m}^{\nu_*\gamma}]^2} \right) \frac{\hat{m}^{\mu_*\gamma}}{\sqrt{\sum_{\nu_*=1}^{P^*} [\hat{m}^{\nu_*\gamma}]^2}}, \end{aligned}$$

where $\varsigma(x) = \frac{x}{\sqrt{x^2+1}+1}$. $\hat{m}^{\mu_*\gamma}$ vanishes (see equation (27)) and $m^{\mu_*\gamma}$ is arbitrary (see equations (23) and (24)) when $\gamma=0$. Therefore, we may update only $\hat{m}^{\mu_*\mu}$ and $m^{\mu_*\mu}$ with $1 \leq \mu \leq P$ when solving for $\nabla f = 0$ by fixed-point iteration. Defining $m^{\mu_*0} = \frac{1}{\beta_{\text{eff}}} \log[\Omega_N(\beta)/\Omega_N(0)]$, we then obtain the saddle-point equations

$$\begin{aligned} \hat{m}^{\mu_*\mu} &= \sum_{y=0}^C p_{\mathbf{q}}^*(y) \sigma_{\mu} \left(\beta_{\text{eff}} m^{\mu_*} + \log[\mathbf{p}_y] \right) \\ m^{\mu_*\mu} &= \varsigma \left(2\beta_{\text{eff}} \varrho \sqrt{\sum_{\nu_*=1}^{P^*} [\hat{m}^{\nu_*\mu}]^2} \right) \frac{\hat{m}^{\mu_*\mu}}{\sqrt{\sum_{\nu_*=1}^{P^*} [\hat{m}^{\nu_*\mu}]^2}}, \end{aligned}$$

for all $1 \leq \mu \leq P$, where we simplified the argument of the softmax function $\sigma_{\mu}(x^{\mu_*}) = \frac{\exp(x^{\mu_*\mu})}{\sum_{\kappa=0}^P \exp(x^{\mu_*\kappa})}$ using $m^{\mu_*0} = \frac{1}{\beta_{\text{eff}}} \log[\Omega_N(\beta)/\Omega_N(0)]$. Putting the equations for $m^{\mu_*\mu}$ and p_y^{γ} together, we find

$$\begin{aligned} \hat{m}^{\mu_*\mu} &= \sum_{y=0}^C p_{\mathbf{q}}^*(y) \sigma_{\mu} \left(\beta_{\text{eff}} m^{\mu_*} + \log[\mathbf{p}_y] \right) \\ \bar{p}_y^{\gamma} &= p_{\mathbf{q}}^*(y) \sum_{\mu=1}^{P^*} \sigma_{\gamma} \left(\beta_{\text{eff}} m^{\mu_*} + \log[\mathbf{p}_y] \right) \\ m^{\mu_*\mu} &= \varsigma \left(2\beta_{\text{eff}} \varrho \sqrt{\sum_{\nu_*=1}^{P^*} [\hat{m}^{\nu_*\mu}]^2} \right) \frac{\hat{m}^{\mu_*\mu}}{\sqrt{\sum_{\nu_*=1}^{P^*} [\hat{m}^{\nu_*\mu}]^2}} \\ p_y^{\gamma} &= \frac{\bar{p}_y^{\gamma}}{\zeta_y^{\gamma}(\bar{\mathbf{p}}; \mathbf{p}_{\mathbf{h}})} \end{aligned}$$

for all $1 \leq \mu \leq P$ and $0 \leq \gamma \leq P$. If we instead clamp $\mathbf{w}^{*\mu_*}$ and $\mathbf{g}^{*\mu_*}$ to some fixed patterns $\mathbf{x}^{*\mu_*}$ and their soft labels $\mathbf{q}^{*\mu_*}$, respectively, then the solutions of equation (26) take the form

$$\begin{aligned} \hat{m}^{\mu_*\mu} &= \sum_{y=0}^C q_y^{*\mu_*} \sigma_{\mu} \left(\beta_{\text{eff}} m^{\mu_*} + \log[\mathbf{p}_y] \right) \\ \bar{p}_y^{\gamma} &= \sum_{\mu_*=1}^{P^*} q_y^{*\mu_*} \sigma_{\gamma} \left(\beta_{\text{eff}} m^{\mu_*} + \log[\mathbf{p}_y] \right) \end{aligned}$$

$$m^{\mu^* \mu} = \varsigma \left(2\beta_{\text{eff}} \varrho \sqrt{\sum_{i=1}^N \left[\sum_{\nu^*=1}^{P^*} \hat{m}^{\nu^* \mu} x_i^{*\nu^*} \right]^2} \right) \frac{\sum_{i=1}^N x_i^{*\nu^*} \sum_{\nu^*=1}^{P^*} \hat{m}^{\nu^* \mu} x_i^{*\nu^*}}{\sqrt{\sum_{i=1}^N \left[\sum_{\nu^*=1}^{P^*} \hat{m}^{\nu^* \mu} x_i^{*\nu^*} \right]^2}}$$

$$p_y^\gamma = \frac{\bar{p}_y^\gamma}{\zeta_y^\gamma(\bar{\mathbf{p}}; p_{\mathbf{h}})}.$$

Defining $\bar{x}_i^\mu = \sum_{\mu^*} \hat{m}^{\mu^* \mu} x_i^{*\mu^*}$, we thus find

$$\bar{x}_i^\mu = \sum_{\mu^*=1}^{P^*} x_i^{*\mu^*} \sum_{y=0}^C q_y^{*\mu^*} \sigma_\mu \left(\beta_{\text{eff}} m^{\mu^*} + \log [\mathbf{p}_y] \right)$$

$$\bar{p}_y^\gamma = \sum_{\mu^*=1}^{P^*} q_y^{*\mu^*} \sigma_\gamma \left(\beta_{\text{eff}} m^{\mu^*} + \log [\mathbf{p}_y] \right)$$

$$m^{\mu^* \mu} = \varsigma \left(2\beta_{\text{eff}} \varrho \sqrt{\sum_{i=1}^N [\bar{x}_i^\mu]^2} \right) \frac{\sum_{i=1}^N x_i^{*\mu^*} \bar{x}_i^\mu}{\sqrt{\sum_{i=1}^N [\bar{x}_i^\mu]^2}}$$

$$p_y^\gamma = \frac{\bar{p}_y^\gamma}{\zeta_y^\gamma(\bar{\mathbf{p}}; p_{\mathbf{h}})},$$

for all $1 \leq \mu \leq P$ and $0 \leq \gamma \leq P$.

Appendix I. Saddle-point hierarchy

Suppose that the set of parameters $\bar{x}_i^{\text{fixed}, \mu}$, $\bar{p}_y^{\text{fixed}, \gamma}$, $m^{\text{fixed}, \mu^* \gamma}$, $p_y^{\text{fixed}, \gamma}$ with hidden unit prior $p_{\mathbf{h}}^{\text{given}}(\gamma)$ is a fixed point of equations (10) with P hidden units. Substitute into the same saddle-point equations with $P + R \in \{P, \dots, 2P\}$ hidden units the duplicated order parameters

$$\bar{x}_i^{\text{dupli}, \mu} = \begin{cases} \bar{x}_i^{\text{fixed}, \mu} & 0 < \mu \leq P \\ \bar{x}_i^{\text{fixed}, \mu - P} & P < \mu \leq P + R \end{cases}$$

$$\bar{p}_y^{\text{dupli}, \gamma} = \begin{cases} \bar{p}_y^{\text{fixed}, 0} & \gamma = 0 \\ \frac{1}{2} \bar{p}_y^{\text{fixed}, \gamma} & 0 < \gamma \leq R \\ \bar{p}_y^{\text{fixed}, \gamma} & R < \gamma \leq P \\ \frac{1}{2} \bar{p}_y^{\text{fixed}, \gamma - P} & P < \gamma \leq P + R \end{cases}$$

$$m^{\text{dupli}, \mu^* \gamma} = \begin{cases} m^{\text{fixed}, \mu^* 0} & \gamma = 0 \\ m^{\text{fixed}, \mu^* \gamma} & 0 < \gamma \leq P \\ m^{\text{fixed}, \mu^*, \gamma - P} & P < \gamma \leq P + R \end{cases}$$

$$p_y^{\text{dupli}, \gamma} = \begin{cases} p_y^{\text{fixed}, 0} & \gamma = 0 \\ \frac{1}{2} p_y^{\text{fixed}, \gamma} & 0 < \gamma \leq R \\ p_y^{\text{fixed}, \gamma} & R < \gamma \leq P \\ \frac{1}{2} p_y^{\text{fixed}, \gamma - P} & P < \gamma \leq P + R \end{cases}$$

along with $p_{\mathbf{h}}(\gamma) = \begin{cases} p_{\mathbf{h}}^{\text{given}}(0) & \gamma = 0 \\ \frac{1}{2} p_{\mathbf{h}}^{\text{given}}(\gamma) & 0 < \gamma \leq R \\ p_{\mathbf{h}}^{\text{given}}(\gamma) & R < \gamma \leq P \\ \frac{1}{2} p_{\mathbf{h}}^{\text{given}}(\gamma - P) & P < \gamma \leq P + R, \end{cases}$

where the hidden units $\gamma \in \{P + 1, \dots, P + R\}$ and their corresponding order parameters are duplicates, or copies, of $\gamma \in \{1, \dots, R\}$. $\gamma = 0$ can also be duplicated, but the result is less interesting. By definition (see equation (21)), $\zeta_y^\gamma(\bar{\mathbf{p}}^{\text{dupli}}; p_{\mathbf{h}}) = \lambda_y(\bar{\mathbf{p}}^{\text{dupli}}; p_{\mathbf{h}}) + \omega^\gamma(\bar{\mathbf{p}}^{\text{dupli}}; p_{\mathbf{h}})$, where $\lambda_y(\bar{\mathbf{p}}^{\text{dupli}}; p_{\mathbf{h}})$ and $\omega^\gamma(\bar{\mathbf{p}}^{\text{dupli}}; p_{\mathbf{h}})$ are the λ_y and ω^γ solving

$$\omega^\gamma = \frac{1}{p_{\mathbf{h}}(\gamma)} \sum_{y=0}^C \frac{\omega^\gamma}{\lambda_y + \omega^\gamma} \bar{p}_y^{\text{dupli}, \gamma}$$

$$\begin{aligned}
 &= \frac{1}{p_h^{\text{given}}(\gamma)} \sum_{y=0}^C \frac{\omega^\gamma}{\lambda_y + \omega^\gamma} \bar{p}_y^{\text{fixed},\gamma} \\
 \lambda_y &= \frac{1}{p_q(y)} \sum_{\gamma=0}^{P+R} \frac{\lambda_y}{\lambda_y + \omega^\gamma} \bar{p}_y^{\text{dupli},\gamma} \\
 &= \frac{1}{p_q(y)} \sum_{\gamma=0}^P \frac{\lambda_y}{\lambda_y + \omega^\gamma} \bar{p}_y^{\text{fixed},\gamma}.
 \end{aligned}$$

Therefore, $\zeta_y^\gamma(\bar{\mathbf{p}}^{\text{dupli}}; p_h) = \zeta_y^\gamma(\bar{\mathbf{p}}^{\text{fixed}}; p_h^{\text{given}})$, and the saddle-point equations simplify to

$$\begin{aligned}
 \bar{x}_i^{\text{fixed},\mu} &= \frac{1}{2} (1 + \mathbb{I}(\mu > R)) \sum_{\mu_*=1}^{P^*} x_i^{*\mu_*} \sum_{y=0}^C q_y^{*\mu_*} \sigma_\mu \left(\beta_{\text{eff}} m^{\text{fixed},\mu_*} + \log [\mathbf{p}_y^{\text{fixed}}] \right) \\
 \bar{p}_y^{\text{fixed},\gamma} &= \sum_{\mu_*=1}^{P^*} q_y^{*\mu_*} \sigma_\gamma \left(\beta_{\text{eff}} m^{\text{fixed},\mu_*} + \log [\mathbf{p}_y^{\text{fixed}}] \right) \\
 m^{\text{fixed},\mu_*\mu} &= \varsigma \left(2\beta_{\text{eff}} \varrho \sqrt{\sum_{i=1}^N [\bar{\mathbf{x}}_i^{\text{fixed},\mu}]^2} \right) \frac{\sum_{i=1}^N x_i^{*\mu_*} \bar{\mathbf{x}}_i^{\text{fixed},\mu}}{\sqrt{\sum_{i=1}^N [\bar{\mathbf{x}}_i^{\text{fixed},\mu}]^2}} \\
 p_y^{\text{fixed},\gamma} &= \frac{\bar{p}_y^{\text{fixed},\gamma}}{\zeta_y^\gamma(\bar{\mathbf{p}}^{\text{fixed}}; p_h^{\text{given}})},
 \end{aligned} \tag{28}$$

where $\mathbb{I}(\mu > R)$ is the indicator function equal to 1 when $\mu > R$ and 0 otherwise. Assume that $\varrho \rightarrow \infty$ so that $\varsigma \left(2\beta_{\text{eff}} \varrho \sqrt{\sum_{i=1}^N [\bar{x}_i^\mu]^2} \right) \rightarrow \mathbb{I} \left(\sqrt{\sum_{i=1}^N [\bar{x}_i^\mu]^2} > 0 \right)$, then the saddle-point equations are the same no matter how the prefactor of $\frac{1}{2} (1 + \mathbb{I}(\mu > R))$ affects the norm of \bar{x}_i^μ , and equations (13) are a fixed point of the saddle-point equations with $P + R$ hidden units. In particular, equation (13) is a stationary point of the loss (equation (4)) when $\beta_{\text{eff}} = \beta$.

For the rest of this appendix, all sums are understood as having the same bounds as equations (28) and (10). The stability of any fixed point of the form $\mathbf{x} = \mathbf{F}(\mathbf{x})$, such as those of equation (10), can be evaluated using the Jacobian matrix \mathbf{J} of \mathbf{F} . If all eigenvalues λ of the Jacobian satisfy $|\lambda| < 1$, then the fixed point is stable. Conversely, if the Jacobian has an eigenvalue λ with $|\lambda| > 1$, then the fixed point is unstable. In particular, if the quadratic form $\mathbf{v}^T \mathbf{J} \mathbf{v}$ is larger than 1 for some \mathbf{v} with $\|\mathbf{v}\| = 1$, then the fixed point is unstable. For the rest of this appendix, we evaluate the stability of equation (10) with duplicated order parameters. Keeping $\bar{\mathbf{p}}$ and \mathbf{p} fixed, the Jacobian of the saddle-point equation for $\bar{\mathbf{x}}$ is

$$\begin{aligned}
 \partial_{\bar{x}_k^\nu} \bar{x}_j^\mu &= \sum_{\mu_*} x_j^{*\mu_*} \partial_{\bar{x}_k^\nu} m^{\mu_*\nu} \sum_y q_y^{*\mu_*} \partial_{m^{\mu_*\nu}} \sigma_\mu \left(\beta_{\text{eff}} m^{\mu_*} + \log [\mathbf{p}_y] \right) \\
 &= \beta_{\text{eff}} \sum_{\mu_*} x_j^{*\mu_*} \partial_{\bar{x}_k^\nu} m^{\mu_*\nu} \sum_y q_y^{*\mu_*} \sigma_\mu \left(\beta_{\text{eff}} m^{\mu_*} + \log [\mathbf{p}_y] \right) \left(\delta_{\mu\nu} - \sigma_\nu \left(\beta_{\text{eff}} m^{\mu_*} + \log [\mathbf{p}_y] \right) \right) \\
 &= \beta_{\text{eff}} \sum_{\mu_*} x_j^{*\mu_*} \frac{1}{\sqrt{\sum_i [\bar{x}_i^\nu]^2}} \left(x_k^{*\mu_*} - \left[\sum_i x_i^{*\mu_*} \bar{x}_i^\nu \right] \bar{x}_k^\nu \right) \\
 &\quad \times \sum_y q_y^{*\mu_*} \sigma_\mu \left(\beta_{\text{eff}} m^{\mu_*} + \log [\mathbf{p}_y] \right) \left(\delta_{\mu\nu} - \sigma_\nu \left(\beta_{\text{eff}} m^{\mu_*} + \log [\mathbf{p}_y] \right) \right),
 \end{aligned}$$

where $\bar{x}_j^\nu = \frac{\bar{x}_j^\nu}{\sqrt{\sum_i [\bar{x}_i^\nu]^2}}$. Suppose that \bar{x}_i^μ , \bar{p}_y^γ , $m^{\mu_*\gamma}$ and p_y^γ are the duplicated parameters of equation (13) with $0 < \mu, \nu \leq R$ or $P < \mu, \nu \leq P + R$, then

$$\begin{aligned}
 \partial_{\bar{x}_k^\nu} \bar{x}_j^\mu &= \beta_{\text{eff}} \sum_{\mu_*} x_j^{*\mu_*} \frac{1}{\sqrt{\sum_i [\bar{x}_i^\nu]^2}} \left(x_k^{*\mu_*} - \left[\sum_i x_i^{*\mu_*} \bar{x}_i^\nu \right] \bar{x}_k^\nu \right) \\
 &\quad \times \sum_y q_y^{*\mu_*} \frac{1}{2} \sigma_{\theta(\mu)} \left(\beta_{\text{eff}} m^{\text{fixed},\mu_*} + \log [\mathbf{p}_y^{\text{fixed}}] \right) \left(\delta_{\mu\nu} - \frac{1}{2} \sigma_{\theta(\nu)} \left(\beta_{\text{eff}} m^{\text{fixed},\mu_*} + \log [\mathbf{p}_y^{\text{fixed}}] \right) \right),
 \end{aligned}$$

$$\text{where } \theta(\mu) = \begin{cases} \mu & 0 < \mu \leq R \\ \mu - P & 0 < \mu \leq P + R. \end{cases}$$

If β_{eff} is relatively large (for instance of order \sqrt{N}), then the softmax $\sigma_{\theta(\mu)}$ splits the indices μ_* into a cover \mathcal{S} of sets $\mathcal{S}(\mu)$ such that $\sigma_{\theta(\mu)}(\beta_{\text{eff}} m^{\text{fixed}, \mu_*} + \log[\mathbf{p}_y]) \approx \mathbb{I}(\mu_* \in \mathcal{S}(\mu))$. Therefore, we have

$$\begin{aligned} \partial_{\tilde{x}_k^\nu} \tilde{x}_j^\mu &\approx \beta_{\text{eff}} \sum_{\mu_*} x_j^{*\mu_*} \frac{1}{\sqrt{\sum_i [\tilde{x}_i^\nu]^2}} \left(x_k^{*\mu_*} - \left[\sum_i x_i^{*\mu_*} \tilde{x}_i^\nu \right] \tilde{x}_k^\nu \right) \\ &\times \frac{1}{2} \mathbb{I}(\mu_* \in \mathcal{S}(\mu)) \left(\delta_{\mu\nu} - \frac{1}{2} \mathbb{I}(\mu_* \in \mathcal{S}(\nu)) \right). \end{aligned}$$

The subcovers $\mathcal{S}(0 < \mu \leq P)$ and $\mathcal{S}(R < \mu \leq P + R)$ are both partitions of the indices μ_* . As such, $\left[\partial_{\tilde{x}_k^\nu} \tilde{x}_j^\mu \right]_{\mu, \nu=1}^{P+R}$ is block diagonal with 2×2 blocks coupling the indices $0 < \mu \leq R$ and $\nu = \mu + P$. Without loss of generality, we investigate the stability of the block

$$\begin{aligned} &\begin{bmatrix} \partial_{\tilde{x}_k^1} \tilde{x}_j^1 & \partial_{\tilde{x}_k^{P+1}} \tilde{x}_j^1 \\ \partial_{\tilde{x}_k^1} \tilde{x}_j^{P+1} & \partial_{\tilde{x}_k^{P+1}} \tilde{x}_j^{P+1} \end{bmatrix} \\ &= \frac{1}{2} \frac{\beta_{\text{eff}}}{\sqrt{\sum_i [x_i^{*1}]^2}} \sum_{\mu_* \in \mathcal{S}(1)} \begin{bmatrix} x_j^{*\mu_*} (x_k^{*\mu_*} - [\sum_i x_i^{*\mu_*} \tilde{x}_i^1] \tilde{x}_k^1) & -x_j^{*\mu_*} (x_k^{*\mu_*} - [\sum_i x_i^{*\mu_*} \tilde{x}_i^1] \tilde{x}_k^1) \\ -x_j^{*\mu_*} (x_k^{*\mu_*} - [\sum_i x_i^{*\mu_*} \tilde{x}_i^1] \tilde{x}_k^1) & x_j^{*\mu_*} (x_k^{*\mu_*} - [\sum_i x_i^{*\mu_*} \tilde{x}_i^1] \tilde{x}_k^1) \end{bmatrix}. \end{aligned}$$

where $x_i^{*1} = \sum_{\mu_* \in \mathcal{S}(1)} x_i^{*\mu_*}$. Let \mathbf{u}^1 be an arbitrary vector orthogonal to $\tilde{\mathbf{x}}^1$ and define

$$v_j^\mu = \frac{1}{\sqrt{2}} \begin{cases} u_j^1 & \mu = 1 \\ -u_j^1 & \mu = P + 1, \end{cases}$$

then we have the quadratic form

$$\begin{aligned} \sum_{\nu \in \{1, P+1\}} \sum_{k=1}^N v_k^\nu \partial_{\tilde{x}_k^\nu} \tilde{x}_j^\mu &= \frac{\beta_{\text{eff}}}{\sqrt{\sum_i [x_i^{*1}]^2}} \sum_{\mu_* \in \mathcal{S}(1)} \frac{1}{\sqrt{2}} \begin{bmatrix} x_j^{*\mu_*} \sum_k x_k^{*\mu_*} u_k^1 \\ -x_j^{*\mu_*} \sum_k x_k^{*\mu_*} u_k^1 \end{bmatrix} \\ \sum_{\mu, \nu \in \{1, P+1\}} \sum_{j, k=1}^N v_j^\mu v_k^\nu \partial_{\tilde{x}_k^\nu} \tilde{x}_j^\mu &= \frac{\beta_{\text{eff}}}{\sqrt{\sum_i [x_i^{*1}]^2}} \sum_{\mu_* \in \mathcal{S}(1)} \left[\sum_j x_j^{*\mu_*} u_j^1 \right]^2. \end{aligned}$$

If \mathbf{u}^1 is orthogonal to all the $\mathbf{x}^{*\mu_*}$ such that $\mu_* \in \mathcal{S}(1)$, and in particular if $\mathcal{S}(1)$ contains a single pattern $\mathbf{x}^{*\mu_*} = \mathbf{x}^1$, then $\sum_{\mu_* \in \mathcal{S}(1)} \left[\sum_j x_j^{*\mu_*} u_j^1 \right]^2 = 0$, so the quadratic form vanishes. In this case, \mathbf{u}^1 is a stable direction of the saddle-point equations. Otherwise, the quadratic form does not vanish, so there is a β_{split} such that the direction \mathbf{u}^1 is unstable when $\beta_{\text{eff}} > \beta_{\text{split}}$.

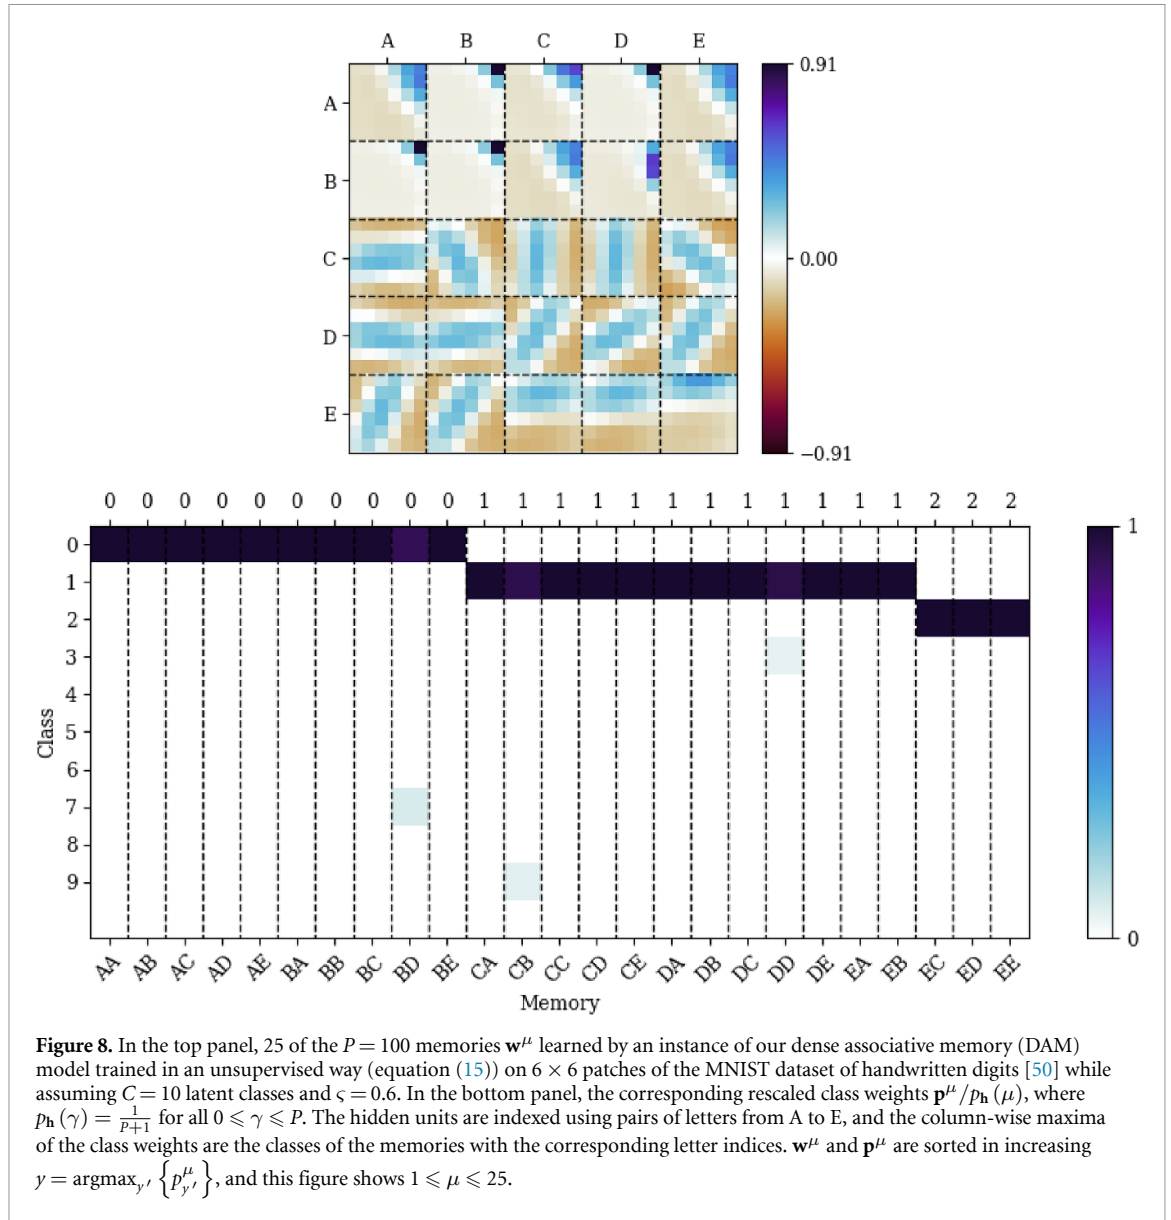
Appendix J. Weights learned with unsupervised training

This appendix contains plots of DAM weights learned in an unsupervised way (see section 4.2) and sorted in increasing $y = \text{argmax}_{y'} \{p_{y'}^\mu\}$. They are not in the main text because they take a lot of space.

Appendix K. Splitting steepest descent

Steps 5 and 11 of splitting steepest descent (algorithm 1) involve the splitting matrices $\mathcal{S}_\mu(\mathbf{w}, \mathbf{p})$ derived in [29]. In this appendix, we explain their role in the algorithm. Consult [29] for a detailed explanation of their interpretation and theoretical underpinnings.

Define the thresholds $\tau_{\text{thres}} \in (0, 1]$ and $\lambda_{\text{thres}} \leq 0$. At step 5 of algorithm 1, we duplicate the hidden units corresponding to the $R \leq \min\{\tau_{\text{thres}} P_{\text{cur}}, P_{\text{max}} - P_{\text{cur}}\}$ most negative minimum eigenvalues λ_{min}^μ

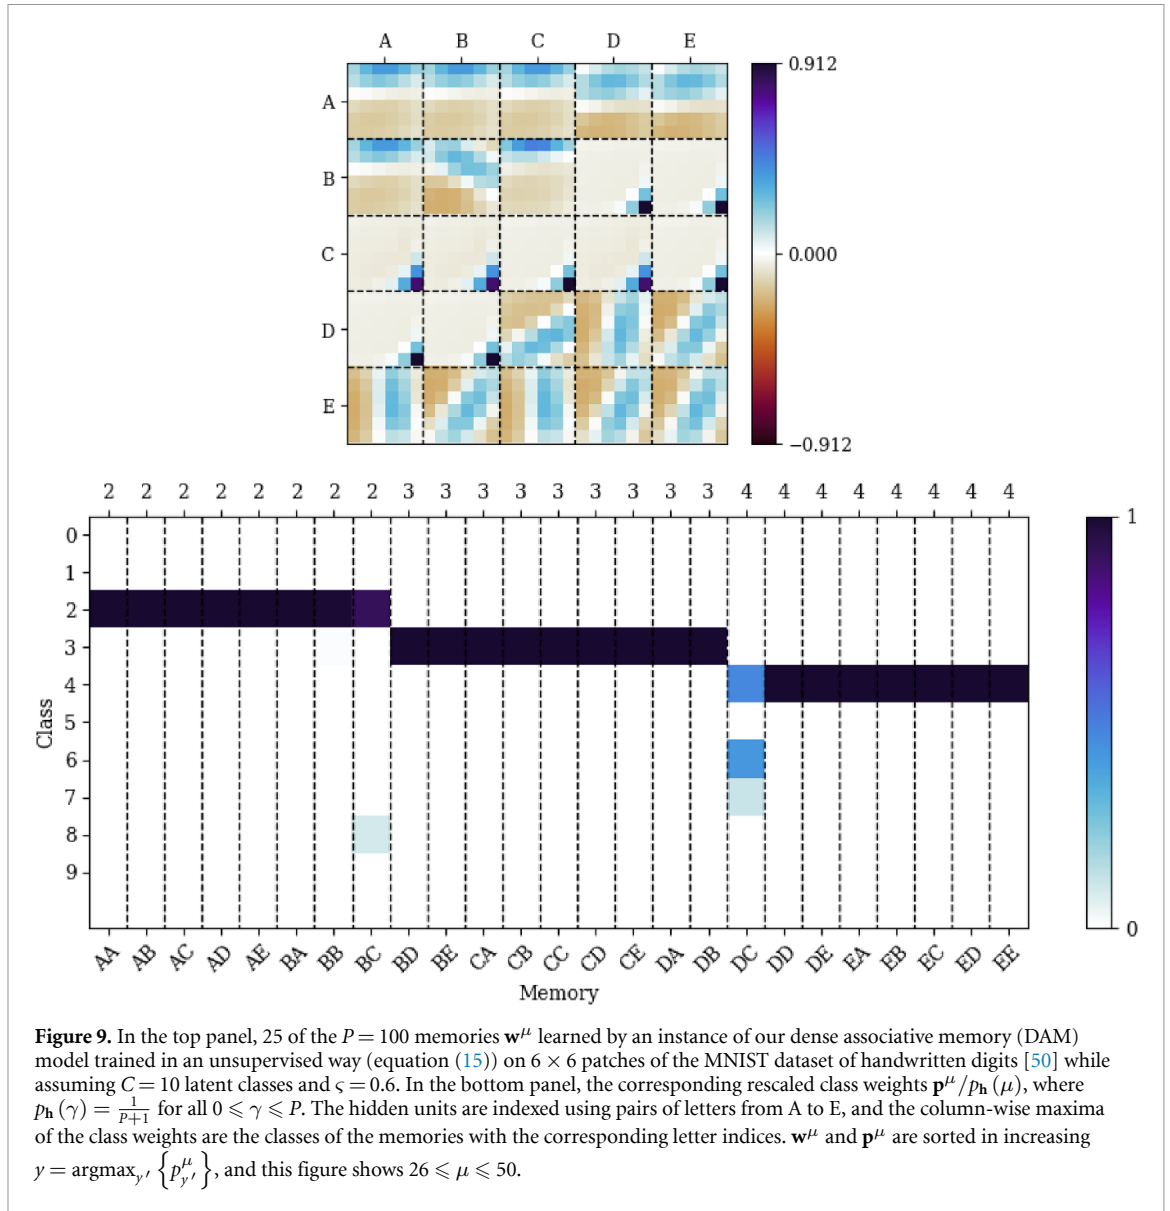


of the splitting matrices $S_\mu(\mathbf{w}, \mathbf{p})$ such that $\lambda_{\min}^\mu \leq \lambda_{\text{thres}}$ [29]. To make figure 7, we pick $\tau_{\text{thres}} = 1$ and $\lambda_{\text{thres}} \approx 0$. Our splitting matrices are

$$S_\mu(\mathbf{w}, \mathbf{p}) = -\mathbb{E}_{\mathbf{x}^*, y^*} \left[\frac{p_y^\mu \bar{\nabla}_{\mathbf{w}^\mu} \bar{\nabla}_{\mathbf{w}^\mu}^T \exp\left(\beta_{\text{eff}} \sum_{i=1}^N w_i^\mu x_i\right)}{\sum_{\nu=1}^P p_y^\nu \exp\left(\beta_{\text{eff}} \sum_{i=1}^N w_i^\nu x_i\right) + p_y^0 \frac{\Omega_N(\beta)}{\Omega_N(0)}} \right],$$

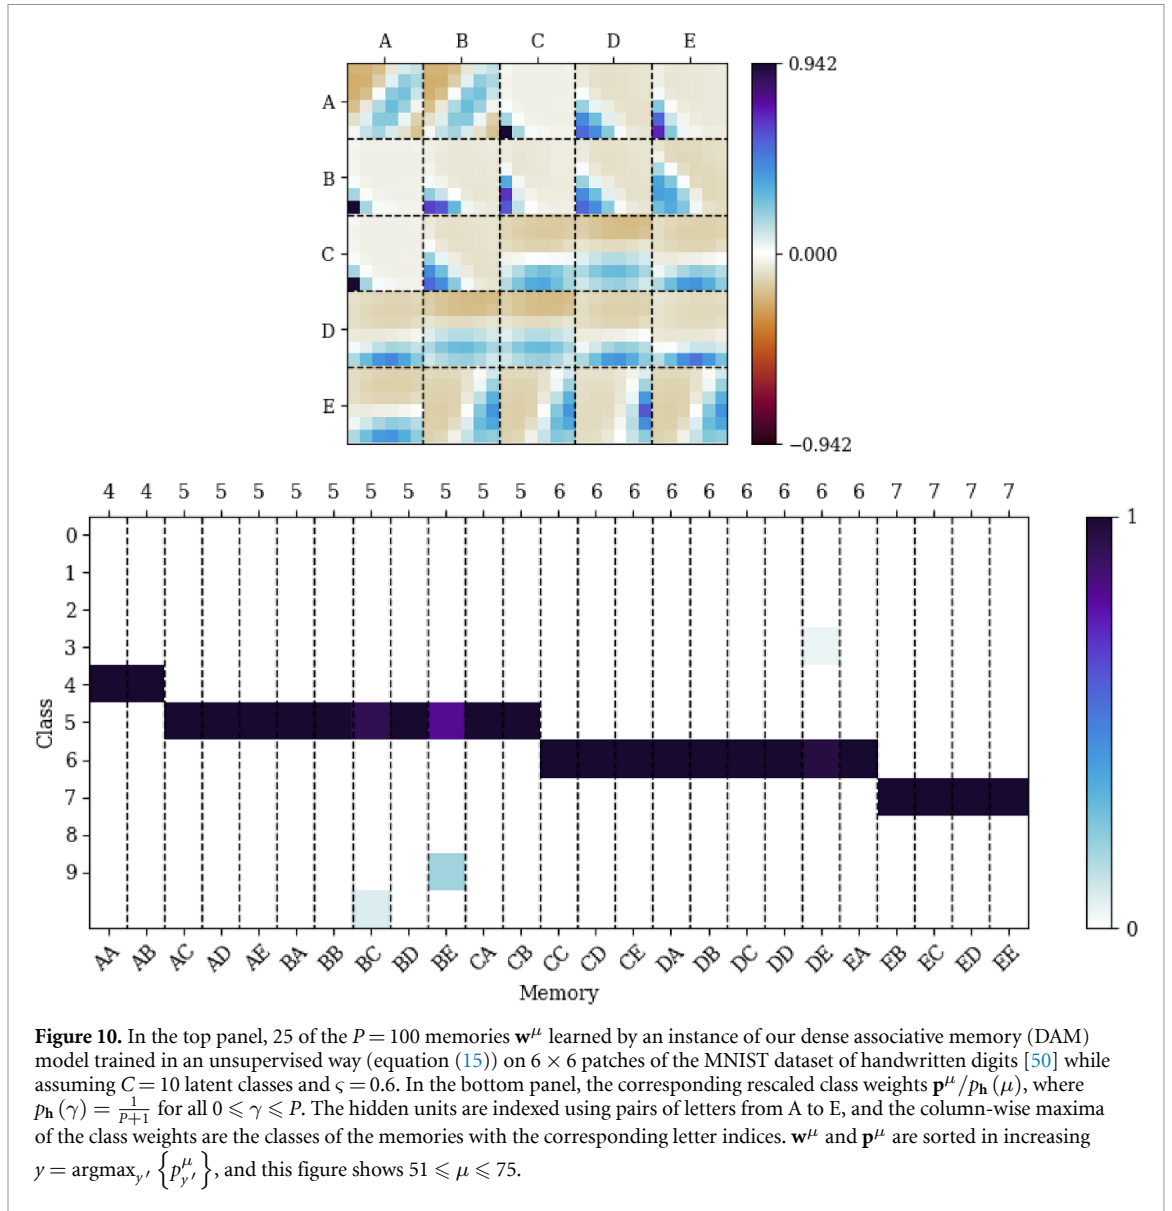
where $\bar{\nabla}_\theta$ is the gradient constrained to the unit hypersphere S^{N-1} (see appendix C) and β_{eff} can be written explicitly in terms of β as $\beta_{\text{eff}} = \varsigma\beta$ (see equations (12) and (10)). At step 11 of algorithm 1, we break permutation symmetries in the memories \mathbf{w}^μ by descending along the eigenvectors $\mathbf{u}^\mu \in S^{N-1}$ corresponding to the eigenvalues λ_{\min}^μ . To be more precise, we update the memories \mathbf{w}^μ and their duplicates $\mathbf{w}^{\text{dupli}, \mu}$ according to $\mathbf{w}^\mu \leftarrow \mathbf{w}^\mu + \delta \mathbf{u}^\mu$ and $\mathbf{w}^{\text{dupli}, \mu} \leftarrow \mathbf{w}^{\text{dupli}, \mu} - \delta \mathbf{u}^\mu$, respectively, where δ is a relatively small learning rate [29]. In physics terminology, the eigenvectors \mathbf{u}^μ are excitation modes that break the permutation symmetries of the memories. N and P are generally large, so it is prohibitively expensive to store $S_\mu(\mathbf{w}, \mathbf{p})$ explicitly for all hidden units $1 \leq \mu \leq P$. Therefore, as proposed in [75], we find the eigenvectors \mathbf{u}^μ and their eigenvalues λ_{\min}^μ by minimizing the Rayleigh quotients $Q_\mu[\mathbf{w}, \mathbf{p}] : S^{N-1} \ni \mathbf{u}^\mu \mapsto [\mathbf{u}^\mu]^T S_\mu(\mathbf{w}, \mathbf{p}) \mathbf{u}^\mu$, which can be evaluated without constructing $S_\mu(\mathbf{w}, \mathbf{p})$ explicitly. To derive $Q_\mu[\mathbf{w}, \mathbf{p}]$, we first compute $\bar{\nabla}_{\mathbf{w}^\mu} \bar{\nabla}_{\mathbf{w}^\mu}^T \exp(\beta_{\text{eff}} \sum_i w_i^\mu x_i)$. Given $f(\boldsymbol{\theta}) = \sum_i \theta_i x_i$, we find

$$\bar{\nabla}_\theta \bar{\nabla}_\theta^T \exp(\beta_{\text{eff}} f(\boldsymbol{\theta})) = \exp(\beta_{\text{eff}} f(\boldsymbol{\theta})) (\beta_{\text{eff}} \bar{\nabla}_\theta \bar{\nabla}_\theta^T f(\boldsymbol{\theta}) + \beta_{\text{eff}}^2 \bar{\nabla}_\theta f(\boldsymbol{\theta}) \bar{\nabla}_\theta^T f(\boldsymbol{\theta})).$$



As mentioned at the beginning of appendix C, $\bar{\nabla}_\theta f(\theta)$ is the unrestricted gradient of $f(\theta)$ projected onto the tangent space of θ . It remains to calculate $\bar{\nabla}_\theta \bar{\nabla}_\theta^T f(\theta)$. Following [89], we find it to be the iterated projected gradient

$$\begin{aligned}
 [\bar{\nabla}_\theta \bar{\nabla}_\theta^T f(\theta)]_{i\ell} &= \sum_k (\delta_{k\ell} - \theta_k \theta_\ell) \partial_{\theta_i} \left(\sum_j (\delta_{jk} - \theta_j \theta_k) \partial_{\theta_j} \left[\sum_h \theta_h x_h \right] \right) \\
 &= \sum_k (\delta_{k\ell} - \theta_k \theta_\ell) \partial_{\theta_i} \left(x_k - \left[\sum_j \theta_j x_j \right] \theta_k \right) \\
 &= - \sum_k (\delta_{k\ell} - \theta_k \theta_\ell) \left(x_i \theta_k + \left[\sum_j \theta_j x_j \right] \delta_{ik} \right) \\
 &= - \left[\sum_j \theta_j x_j \right] (\delta_{i\ell} - \theta_i \theta_\ell),
 \end{aligned}$$



so we obtain

$$Q_\mu[\mathbf{w}, \mathbf{p}](\mathbf{u}^\mu) = -\mathbb{E}_{\mathbf{x}^*, y^*} \left[\frac{p_{y^*}^\mu \exp\left(\beta_{\text{eff}} \sum_{i=1}^N w_i^\mu x_i\right)}{\sum_{\nu=1}^P p_{y^*}^\nu \exp\left(\beta_{\text{eff}} \sum_{i=1}^N w_i^\nu x_i\right) + p_{y^*}^0 \frac{\Omega_N(\beta)}{\Omega_N(0)}} F(\mathbf{u}^\mu; \mathbf{w}^\mu, \mathbf{x}) \right]$$

where $F(\boldsymbol{\varphi}; \boldsymbol{\theta}, \mathbf{x}) = \boldsymbol{\varphi}^T [\beta_{\text{eff}} \bar{\nabla}_{\boldsymbol{\theta}} \bar{\nabla}_{\boldsymbol{\theta}}^T f(\boldsymbol{\theta}) + \beta_{\text{eff}}^2 \bar{\nabla}_{\boldsymbol{\theta}} f(\boldsymbol{\theta}) \bar{\nabla}_{\boldsymbol{\theta}}^T f(\boldsymbol{\theta})] \boldsymbol{\varphi}$

$$= \beta_{\text{eff}}^2 \left(\sum_k \varphi_k x_k - \left[\sum_k \varphi_k \theta_k \right] \left[\sum_j \theta_j x_j \right] \right)^2$$

$$+ \beta_{\text{eff}} \left[\sum_j \theta_j x_j \right] \left[\sum_i \varphi_i \theta_i - 1 \right] \left[\sum_\ell \varphi_\ell \theta_\ell + 1 \right].$$

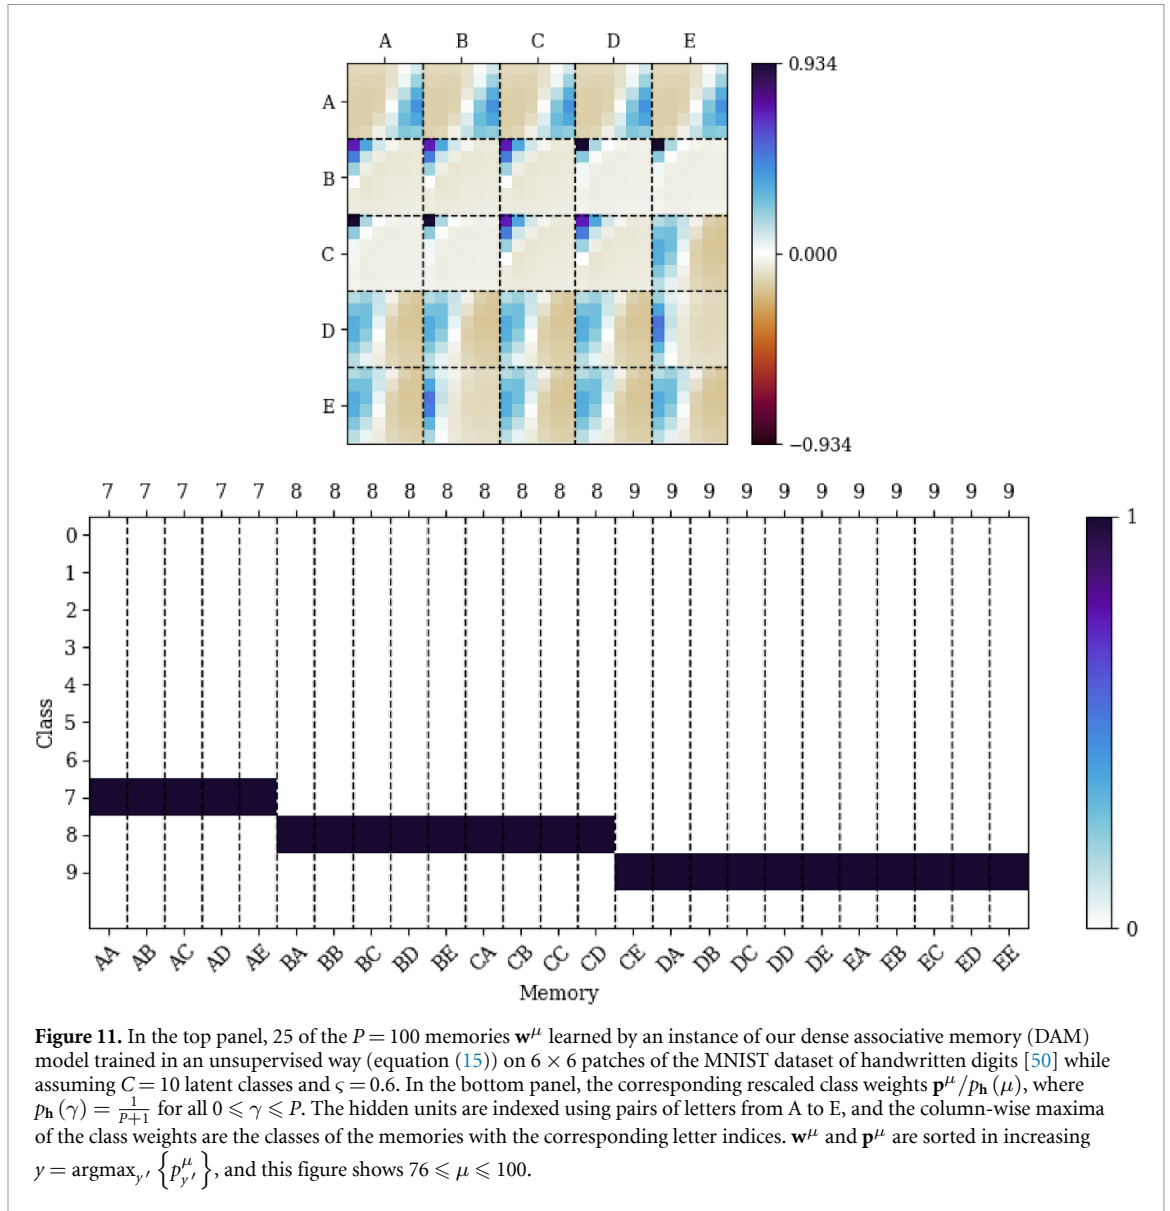


Figure 11. In the top panel, 25 of the $P = 100$ memories \mathbf{w}^μ learned by an instance of our dense associative memory (DAM) model trained in an unsupervised way (equation (15)) on 6×6 patches of the MNIST dataset of handwritten digits [50] while assuming $C = 10$ latent classes and $\varsigma = 0.6$. In the bottom panel, the corresponding rescaled class weights $\mathbf{p}^\mu / p_h(\mu)$, where $p_h(\gamma) = \frac{1}{p+1}$ for all $0 \leq \gamma \leq P$. The hidden units are indexed using pairs of letters from A to E, and the column-wise maxima of the class weights are the classes of the memories with the corresponding letter indices. \mathbf{w}^μ and \mathbf{p}^μ are sorted in increasing $y = \operatorname{argmax}_{y'} \{p_{y'}^\mu\}$, and this figure shows $76 \leq \mu \leq 100$.

We can directly minimize $Q[\mathbf{w}, \mathbf{p}](\mathbf{u}) = \sum_{\mu} Q_{\mu}[\mathbf{w}, \mathbf{p}](\mathbf{u}^{\mu})$ to find the set of all eigenvectors \mathbf{u}^{μ} simultaneously. However, it is more convenient to integrate the eigenvector calculations into the DAM architecture. As such, we define the modified loss

$$\begin{aligned} \mathcal{L}_{\epsilon}(\mathbf{w}, \mathbf{p}, \mathbf{u}) &= -\log \left[\sum_{\mu=1}^P p_y^{\mu} \exp \left(\varsigma \beta \sum_{i=1}^N w_i^{\mu} x_i \right) [1 + \epsilon F(\mathbf{u}^{\mu}; \mathbf{w}^{\mu}, \mathbf{x})] + p_y^0 \frac{\Omega_N(\beta)}{\Omega_N(0)} \right] \\ &= -\log \left[\sum_{\mu=1}^P p_y^{\mu} \exp \left(\varsigma \beta \sum_{i=1}^N w_i^{\mu} x_i + \log [1 + \epsilon F(\mathbf{u}^{\mu}; \mathbf{w}^{\mu}, \mathbf{x})] \right) + p_y^0 \frac{\Omega_N(\beta)}{\Omega_N(0)} \right], \end{aligned}$$

so that the four gradients used to trained the DAM can be calculated using the equations

$$\begin{aligned} \nabla_{\mathbf{u}} Q[\mathbf{w}, \mathbf{p}](\mathbf{u}) &= \lim_{\epsilon \rightarrow 0} \left\{ \frac{1}{\epsilon} \nabla_{\mathbf{u}} \mathcal{L}_{\epsilon}(\mathbf{w}, \mathbf{p}, \mathbf{u}) \right\}, \\ \nabla_{\mathbf{w}} \mathcal{L}(\mathbf{w}, \mathbf{p}) &= \nabla_{\mathbf{w}} \mathcal{L}_0(\mathbf{w}, \mathbf{p}, \mathbf{u}) \\ \nabla_{\mathbf{p}} \mathcal{L}(\mathbf{w}, \mathbf{p}) &= \nabla_{\mathbf{p}} \mathcal{L}_0(\mathbf{w}, \mathbf{p}, \mathbf{u}) \\ \text{and } \nabla_{\beta} \mathcal{L}(\mathbf{w}, \mathbf{p}) &= \nabla_{\beta} \mathcal{L}_0(\mathbf{w}, \mathbf{p}, \mathbf{u}). \end{aligned}$$

This technique is based on the automatic differentiation trick proposed in [75]. We numerically implement the limit in the first equation by setting $\epsilon = 0$ during loss evaluation and $\epsilon = 1$ during gradient computation. When we optimize $\mathcal{Q}[\mathbf{w}, \mathbf{p}](\mathbf{u})$, the minima of the different Rayleigh quotients $Q_\mu[\mathbf{w}, \mathbf{p}](\mathbf{u}^\mu)$ can span multiple orders of magnitude, so we normalize the gradients $\nabla_{\mathbf{u}^\mu} Q_\mu[\mathbf{w}, \mathbf{p}](\mathbf{u}^\mu)$ by a running average of their magnitudes to facilitate convergence. This is inspired by the use of RMSProp in [75]. Moreover, we constrain the eigenvectors \mathbf{u}^μ and the gradients $\nabla_{\mathbf{u}^\mu} Q_\mu[\mathbf{w}, \mathbf{p}](\mathbf{u}^\mu)$ to the unit hypersphere as we do for \mathbf{w}^μ (see appendix C). To make figure 7, we train the eigenvectors for 1 epoch, during which we monitor $\min_\mu \left\{ \frac{1}{2} [\mathbf{u}^\mu]^T \nabla_{\mathbf{u}^\mu} \mathcal{Q}[\mathbf{w}, \mathbf{p}](\mathbf{u}^\mu) \right\} = \min_\mu \{ \mathcal{Q}[\mathbf{w}, \mathbf{p}](\mathbf{u}^\mu) \} \sim \min_\mu \{ \lambda_{\min}^\mu \}$ as a metric. The Rayleigh quotients converge very quickly, and further training of the eigenvectors is not necessary.

References

- [1] Auffinger A, Arous G B and Černý J 2013 Random matrices and complexity of spin glasses *Commun. Pure Appl. Math.* **66** 165–201
- [2] Choromanska A, Henaff M, Mathieu M, Arous G B and LeCun Y 2015 The loss surfaces of multilayer networks *Proc. 18th Int. Conf. on Artificial Intelligence and Statistics (Proc. Machine Learning Research vol 38)*, ed G Lebanon and S V N Vishwanathan (PMLR) pp 192–204 (available at: <https://proceedings.mlr.press/v38/choromanska15.html>)
- [3] Pascanu R, Dauphin Y N, Ganguli S and Bengio Y 2014 On the saddle point problem for non-convex optimization (arXiv:1405.4604)
- [4] Dauphin Y N, Pascanu R, Gulcehre C, Cho K, Ganguli S and Bengio Y 2014 Identifying and attacking the saddle point problem in high-dimensional non-convex optimization *Advances in Neural Information Processing Systems vol 27*, ed Z Ghahramani, M Welling, C Cortes, N Lawrence and K Weinberger (Curran Associates, Inc.) (available at: https://proceedings.neurips.cc/paper_files/paper/2014/file/17e23e50bedc63b4095e3d8204ce063b-paper.pdf)
- [5] Fukumizu K and Amari S 2000 Local minima and plateaus in hierarchical structures of multilayer perceptrons *Neural Netw.* **13** 317–27
- [6] Fukumizu K, Yamaguchi S, Mototake Y-i and Tanaka M 2019 Semi-flat minima and saddle points by embedding neural networks to overparameterization *Advances in Neural Information Processing Systems vol 32*, ed H Wallach, H Larochelle, A Beygelzimer, F d' Alché-Buc, E Fox and R Garnett (Curran Associates, Inc.) (available at: https://proceedings.neurips.cc/paper_files/paper/2019/file/a4ee59dd868ba016ed2de90d330acb6a-paper.pdf)
- [7] Simsek B, Ged F, Jacot A, Spadaro F, Hongler C, Gerstner W and Brea J 2021 Geometry of the loss landscape in overparameterized neural networks: symmetries and invariances *Proc. 38th Int. Conf. on Machine Learning (Proc. Machine Learning Research) vol 139*, ed M Meila and T Zhang (PMLR) pp 9722–32 (available at: <https://proceedings.mlr.press/v139/simsek21a.html>)
- [8] Zhang Y, Zhang Z, Luo T and Xu Z J 2021 Embedding principle of loss landscape of deep neural networks *Advances in Neural Information Processing Systems vol 34*, ed M Ranzato, A Beygelzimer, Y Dauphin, P Liang and J W Vaughan (Curran Associates, Inc.) pp 14848–59 (available at: https://proceedings.neurips.cc/paper_files/paper/2021/file/7cc532d783a7461f227a5da8ea80bfe1-paper.pdf)
- [9] Hopfield J J 1982 Neural networks and physical systems with emergent collective computational abilities *Proc. Natl Acad. Sci.* **79** 2554–8
- [10] Chen H H, Lee Y C, Sun G Z, Lee H Y, Maxwell T and Giles C L 1986 High order correlation model for associative memory *AIP Conf. Proc.* **151** 86–99
- [11] Psaltis D and Park C H 1986 Nonlinear discriminant functions and associative memories *AIP Conf. Proc.* **151** 370–5
- [12] Amit D J, Gutfreund H and Sompolinsky H 1985 Storing infinite numbers of patterns in a spin-glass model of neural networks *Phys. Rev. Lett.* **55** 1530–3
- [13] Baldi P and Venkatesh S S 1987 Number of stable points for spin-glasses and neural networks of higher orders *Phys. Rev. Lett.* **58** 913–6
- [14] Gardner E 1987 Multiconnected neural network models *J. Phys. A: Math. Gen.* **20** 3453
- [15] Abbott L F and Arian Y 1987 Storage capacity of generalized networks *Phys. Rev. A* **36** 5091–4
- [16] Krotov D and Hopfield J J 2016 Dense associative memory for pattern recognition *Advances in Neural Information Processing Systems vol 29*, ed D Lee, M Sugiyama, U Luxburg, I Guyon and R Garnett (Curran Associates, Inc.) (available at: https://proceedings.neurips.cc/paper_files/paper/2016/file/eaee339c4d89fc102edd9bdb6a28915-paper.pdf)
- [17] Krotov D and Hopfield J 2018 Dense associative memory is robust to adversarial inputs *Neural Comput.* **30** 3151–67
- [18] Thériault R and Tantari D 2024 Dense Hopfield networks in the teacher-student setting *SciPost Phys.* **17** 040
- [19] Ramsauer H *et al* 2021 Hopfield networks is all you need *Int. Conf. on Learning Representations* (available at: <https://openreview.net/forum?id=tL89RnzIiCd>)
- [20] Hoover B, Strobel H, Krotov D, Hoffman J, Kira Z and Chau D H 2024 Memory in plain sight: Surveying the uncanny resemblances of associative memories and diffusion models (arXiv:2309.16750)
- [21] Ambrogioni L 2024 In search of dispersed memories: generative diffusion models are associative memory networks *Entropy* **26** 5
- [22] Krotov D 2023 A new frontier for Hopfield networks *Nat. Rev. Phys.* **5** 366–7
- [23] Krotov D, Hoover B, Ram P and Pham B 2025 Modern methods in associative memory (arXiv:2507.06211)
- [24] Boukacem N E, Leary A, Thériault R, Gottlieb F, Mani M and François P 2024 Waddington landscape for prototype learning in generalized Hopfield networks *Phys. Rev. Res.* **6** 033098
- [25] Ackley D H, Hinton G E and Sejnowski T J 1985 A learning algorithm for Boltzmann machines *Cogn. Sci.* **9** 147–69
- [26] Smolensky P 1986 Information processing in dynamical systems: foundations of harmony theory *Parallel Distributed Processing, Volume 1: Explorations in the Microstructure of Cognition: Foundations* (The MIT Press)
- [27] Freund Y and Haussler D 1991 Unsupervised learning of distributions on binary vectors using two layer networks *Advances in Neural Information Processing Systems vol 4*, ed J Moody, S Hanson and R Lippmann (Morgan-Kaufmann) (available at: https://proceedings.neurips.cc/paper_files/paper/1991/file/33e8075e9970de0cfea955afd4644bb2-paper.pdf)
- [28] Hinton G E 2002 Training products of experts by minimizing contrastive divergence *Neural Comput.* **14** 1771–800

- [29] Wu L, Wang D and Liu Q 2019 Splitting steepest descent for growing neural architectures *Advances in Neural Information Processing Systems* vol 32, ed H Wallach, H Larochelle, A Beygelzimer, F d' Alché-Buc, E Fox and R Garnett (Curran Associates, Inc.) (available at: https://proceedings.neurips.cc/paper_files/paper/2019/file/3a01fc0853ebeb94fde4d1cc6fb842a-paper.pdf)
- [30] Thériault R 2025 Robinther/dense_associative_network_vmf: saddle hierarchy in dense associative memory *Zenodo* (<https://doi.org/10.5281/zenodo.17632013>)
- [31] Roux N L and Bengio Y 2008 Representational power of restricted Boltzmann machines and deep belief networks *Neural Comput.* **20** 1631–49
- [32] Salakhutdinov R, Mnih A and Hinton G 2007 Restricted Boltzmann machines for collaborative filtering *Proc. 24th Int. Conf. on Machine Learning (ICML '07)* (Association for Computing Machinery) pp 791–8
- [33] Kivinen J and Williams C 2012 Multiple texture Boltzmann machines *Proc. 15th Int. Conf. on Artificial Intelligence and Statistics (Proc. Machine Learning Research* vol 22) ed N D Lawrence and M Girolami (PMLR) pp 638–46 (available at: <https://proceedings.mlr.press/v22/kivinen12.html>)
- [34] Srivastava N, Salakhutdinov R and Hinton G 2013 Modeling documents with a deep Boltzmann machine *Proc. 29th Conf. on Uncertainty in Artificial Intelligence (UAI'13)* (AUAI Press) pp 616–24
- [35] Salakhutdinov R and Hinton G 2009 Deep Boltzmann machines *Proc. 12th Int. Conf. on Artificial Intelligence and Statistics (Proc. Machine Learning Research* vol 5) ed D van Dyk and M Welling (PMLR) pp 448–55 (available at: <https://proceedings.mlr.press/v5/salakhutdinov09a.html>)
- [36] Salton G and McGill M J 1983 *Introduction to Modern Information Retrieval* (McGraw-Hill, Inc.)
- [37] Carandini M and Heeger D J 2012 Normalization as a canonical neural computation *Nat. Rev. Neurosci.* **13** 51–62
- [38] Potts R B 1952 Some generalized order-disorder transformations *Mathematical Proceedings of the Cambridge Philosophical Society* vol 48 pp 106–9
- [39] Wu F Y 1982 The Potts model *Rev. Mod. Phys.* **54** 235
- [40] Kappen B 1993 Using Boltzmann machines for probability estimation *ICANN '93*, ed S Gielen and B Kappen (Springer) pp 521–6
- [41] Kappen H J 1995 Deterministic learning rules for Boltzmann machines *Neural Netw.* **8** 537–48
- [42] Mozer M C and Smolensky P 1988 Skeletonization: a technique for trimming the fat from a network via relevance assessment *Advances in Neural Information Processing Systems* vol 1, ed D Touretzky (Morgan-Kaufmann) (available at: https://proceedings.neurips.cc/paper_files/paper/1988/file/07e1cd7dca89a1678042477183b7ac3f-paper.pdf)
- [43] Hoefler T, Alistarh D, Ben-Nun T, Dryden N and Peste A 2021 Sparsity in deep learning: pruning and growth for efficient inference and training in neural networks *J. Mach.: Learn. Res.* **22** 1–124
- [44] Guo Y, Zhang C, Zhang C and Chen Y 2018 Sparse dnns with improved adversarial robustness *Advances in Neural Information Processing Systems* vol 31, ed S Bengio, H Wallach, H Larochelle, K Grauman, N Cesa-Bianchi and R Garnett (Curran Associates, Inc.) (available at: https://proceedings.neurips.cc/paper_files/paper/2018/file/4c5bde74a8f110656874902f07378009-paper.pdf)
- [45] Friston K 2008 Hierarchical models in the brain *PLoS Comput. Biol.* **4** e1000211
- [46] Decelle A and Furtlehner C 2021 Restricted Boltzmann machine: Recent advances and mean-field theory* *Chin. Phys. B* **30** 040202
- [47] Banerjee A, Dhillon I S, Ghosh J and Sra S 2005 Clustering on the unit hypersphere using von Mises-Fisher distributions *J. Mach.: Learn. Res.* **6** 1345–82
- [48] Hastie T and Tibshirani R 1996 Discriminant analysis by Gaussian mixtures *J. R. Stat. Soc. B* **58** 155–76
- [49] Christian S, Vincent V, Sergey I, Jon S and Zbigniew W 2016 Rethinking the inception architecture for computer vision *2016 IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)* vol 2016 pp 2818–26 (available at: <https://cir.nii.ac.jp/crid/1361418519657311872>)
- [50] Lecun Y, Bottou L, Bengio Y and Haffner P 1998 Gradient-based learning applied to document recognition *Proc. IEEE* **86** 2278–324
- [51] Lucibello C and Mézard M 2024 Exponential capacity of dense associative memories *Phys. Rev. Lett.* **132** 077301
- [52] Charbonneau P, Marinari E, Mézard M, Parisi G, Ricci-Tersenghi F, Sicuro G and Zamponi F 2023 *Spin Glass Theory and Far Beyond* (World Scientific)
- [53] Decelle A, Fissore G and Furtlehner C 2017 Spectral dynamics of learning in restricted Boltzmann machines *Europhys. Lett.* **119** 60001
- [54] Decelle A, Fissore G and Furtlehner C 2018 Thermodynamics of restricted Boltzmann machines and related learning dynamics *J. Stat. Phys.* **172** 1576–608
- [55] Decelle A and Furtlehner C 2020 Gaussian-spherical restricted Boltzmann machines *J. Phys. A: Math. Theor.* **53** 184002
- [56] Huang H and Toyozumi T 2016 Unsupervised feature learning from finite data by message passing: Discontinuous versus continuous phase transition *Phys. Rev. E* **94** 062310
- [57] Huang H 2017 Statistical mechanics of unsupervised feature learning in a restricted Boltzmann machine with binary synapses *J. Stat. Mech.* **053302**
- [58] Barra A, Genovese G, Sollich P and Tantari D 2017 Phase transitions in restricted Boltzmann machines with generic priors *Phys. Rev. E* **96** 042156
- [59] Barra A, Genovese G, Sollich P and Tantari D 2018 Phase diagram of restricted Boltzmann machines and generalized Hopfield networks with arbitrary priors *Phys. Rev. E* **97** 022310
- [60] Huang H 2018 Role of zero synapses in unsupervised feature learning *J. Phys. A: Math. Theor.* **51** 08LT01
- [61] Manzan G and Tantari D 2025 The effect of priors on learning with restricted Boltzmann machines *Physica A* **674** 130766
- [62] Hou T, Wong K Y M and Huang H 2019 Minimal model of permutation symmetry in unsupervised learning *J. Phys. A: Math. Theor.* **52** 414001
- [63] Thériault R, Tosello F and Tantari D 2025 Modeling structured data learning with restricted Boltzmann machines in the teacher-student setting *Neural Netw.* **189** 107542
- [64] Gardner E and Derrida B 1989 Three unfinished works on the optimal storage capacity of networks *J. Phys. A: Math. Gen.* **22** 1983
- [65] Decelle A, Hwang S, Rocchi J and Tantari D 2021 Inverse problems for structured datasets using parallel TAP equations and restricted Boltzmann machines *Sci. Rep.* **11** 19990
- [66] Alemanno F, Camanzi L, Manzan G and Tantari D 2023 Hopfield model with planted patterns: a teacher-student self-supervised learning model *Appl. Math. Comput.* **458** 128253
- [67] Nishimori H 2001 *Statistical Physics of Spin Glasses and Information Processing: An Introduction* (Oxford University Press)

- [68] Charbonneau P 2022 From the replica trick to the replica symmetry breaking technique (arXiv:2211.01802)
- [69] Zdeborová L and Krzakala F 2016 Statistical physics of inference: thresholds and algorithms *Adv. Phys.* **65** 453–552
- [70] Kirkpatrick S, Gelatt C D and Vecchi M P 1983 Optimization by simulated annealing *Science* **220** 671–80
- [71] Rose K, Gurewitz E and Fox G C 1990 Statistical mechanics and phase transitions in clustering *Phys. Rev. Lett.* **65** 945–8
- [72] Fix E and Hodges J L 1989 Discriminatory analysis. Nonparametric discrimination: consistency properties *Int. Stat. Rev.* **57** 238–47
- [73] Grandvalet Y and Bengio Y 2004 Semi-supervised learning by entropy minimization *Proc. 18th Int. Conf. on Neural Information Processing Systems (NIPS'04)* (MIT Press) pp 529–36
- [74] Khalid S, Khalil T and Nasreen S 2014 A survey of feature selection and feature extraction techniques in machine learning 2014 *Science and Information Conf.* pp 372–8
- [75] Wang D, Li M, Wu L, Chandra V and Liu Q 2019 Energy-aware neural architecture optimization with fast splitting steepest descent (arXiv:2020.1910.03103)
- [76] McInnes L, Healy J, Saul N and Großberger L 2018 Umap: Uniform manifold approximation and projection *J. Open Source Softw.* **3** 861
- [77] Xiao H, Rasul K and Vollgraf R 2017 Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms (arXiv:1708.07747)
- [78] ROSE K, GUREWITZ E and FOX G C 1993 Constrained clustering as an optimization method *IEEE Trans. Pattern Anal. Mach. Intell.* **15** 785–94
- [79] ROSE K 1998 Deterministic annealing for clustering, compression, classification, regression and related optimization problems *Proc. IEEE* **86** 2210–39
- [80] Ota T and Karakida R 2023 Attention in a family of Boltzmann machines emerging from modern Hopfield networks *Neural Comput.* **35** 1463–80
- [81] Bachtis D, Biroli G, Decelle A and Seoane B 2025 Cascade of phase transitions in the training of energy-based models *J. Stat. Mech.* **074004**
- [82] Bereux N, Decelle A, Furtlehner C, Rosset L and Seoane B 2025 Fast training and sampling of Restricted Boltzmann Machines 15 figures, 31
- [83] Grishechkin A, Mukherjee A and Karin O 2025 Mathematical framework for hierarchical cell identity control by collective enhancer dynamics *PRX Life* **3** 043007
- [84] Loera J A D and Kim E D 2014 *Combinatorics and Geometry of Transportation Polytopes: An Update* vol 625 (American Mathematical Society) pp 37–76
- [85] Mardia K and Jupp P 2000 *Directional Statistics* (Wiley)
- [86] NIST Digital Library of Mathematical Functions ed, Olver F W J, Daalhuis A B O, Lozier D W, Schneider B I, Boisvert R F, Clark C W, Miller B R, Saunders B V, Cohl H S, and McClain M A, eds. (available at: <https://dlmf.nist.gov/10.9.E4>)
- [87] Fröhlich J and Spencer T 1981 The Kosterlitz-Thouless transition in two-dimensional Abelian spin systems and the coulomb gas *Commun. Math. Phys.* **81** 527–602
- [88] NIST Digital Library of Mathematical Functions *Release 1.2.2 of 2024-09-15*, Olver F W J, Daalhuis A B O, Lozier D W, Schneider B I, Boisvert R F, Clark C W, Miller B R, Saunders B V, Cohl H S, and McClain M A (available at: <https://dlmf.nist.gov/10.41.E3>)
- [89] Barilari D 2025 Lecture notes on differential geometry pp 137 (available at: <https://www.math.unipd.it/~barilari/notesv25.pdf>) (Accessed 26 December 2025)
- [90] Uribe P, de Leeuw C G and Theil H 1966 The information approach to the prediction of interregional trade flows1 *Rev. Econ. Stud.* **33** 209–20
- [91] Hewings G J D and Janson B N 1980 Exchanging regional input-output coefficients: a reply and further comments *Environ. Plan. A* **12** 843–54
- [92] McNeil S and Hendrickson C 1985 A note on alternative matrix entry estimation techniques *Transp. Res. B* **19** 509–19
- [93] Sinkhorn R 1964 A relationship between arbitrary positive matrices and doubly stochastic matrices *Ann. Math. Stat.* **35** 876–9
- [94] Knopp P and Sinkhorn R 1967 Concerning nonnegative matrices and doubly stochastic matrices *Pac. J. Math.* **21** 343–8
- [95] Idel M 2016 A review of matrix scaling and sinkhorn's normal form for matrices and positive maps (arXiv:1609.06349)
- [96] Menon M and Schneider H 1969 The spectrum of a nonlinear operator associated with a matrix *Linear Algebr. Appl.* **2** 321–34
- [97] Hershkowitz D, Rothblum U G and Schneider H 1988 Classifications of nonnegative matrices using diagonal equivalence *SIAM J. Matrix Anal. Appl.* **9** 455–60
- [98] Deming W E and Stephan F F 1940 On a least squares adjustment of a sampled frequency table when the expected marginal totals are known *Ann. Math. Stat.* **11** 427–44
- [99] Douik A and Hassibi B 2019 Manifold optimization over the set of doubly stochastic matrices: a second-order geometry *IEEE Trans. Signal Process.* **67** 5761–74
- [100] Muller M E 1959 A note on a method for generating points uniformly on n-dimensional spheres *Commun. ACM* **2** 19–20
- [101] Pinzón C and Jung K 2023 Fast Python sampler for the von Mises Fisher distribution (available at: <https://hal.science/hal-04004568>)
- [102] Barvinok A 2010 What does a random contingency table look like? *Comb. Probab. Comput.* **19** 517–39
- [103] Good I J 1963 Maximum entropy for hypothesis formulation, especially for multidimensional contingency tables *Ann. Math. Stat.* **34** 911–34