

---

# Phase transition in PCA with missing data: Reduced signal-to-noise ratio, not sample size!

---

Niels Bruun Ipsen<sup>1</sup> Lars Kai Hansen<sup>1</sup>

## Abstract

How does missing data affect our ability to learn signal structures? It has been shown that learning signal structure in terms of principal components is dependent on the ratio of sample size and dimensionality and that a critical number of observations is needed before learning starts (Biehl and Mietzner, 1993). Here we generalize this analysis to include missing data. Probabilistic principal component analysis is regularly used for estimating signal structures in datasets with missing data. Our analytic result suggests that the effect of missing data is to effectively reduce signal-to-noise ratio rather than - as generally believed - to reduce sample size. The theory predicts a phase transition in the learning curves and this is indeed found both in simulation data and in real datasets.

## 1. Introduction

Principal component analysis (PCA) is a widely used tool for exploratory data analysis originating with Karl Pearson (1901) who described a method for obtaining the "best fit" of a plane or a line to a system of points, by minimizing the sum of squared orthogonal distances from the points to the line/plane. PCA has been rediscovered many times, for example by Hotelling (1933). In many research areas variants of PCA are known under specific names, viz., the Kosambi-Karhunen-Loève transform in signal processing, the Hotelling transform in multivariate quality control and Latent Semantic Analysis in natural language processing. A well known and illustrative application of PCA is the eigenfaces investigation by Turk and Pentland (1991). The probabilistic version of PCA (PPCA) was proposed independently by Roweis (1998) and Tipping and Bishop

---

<sup>1</sup>Department of Applied Mathematics and Computer Science, Technical University of Denmark, Denmark. Correspondence to: Niels Bruun Ipsen <nbp@dtu.dk>, Lars Kai Hansen <lkai@dtu.dk>.

(1999) and this formulation allows for estimating principal components in the presence of missing data via the EM algorithm (Dempster et al., 1977). Numerous modern PCA applications are reviewed in the recent volume (Sanguansat, 2012) and new variations on the classical scheme keep developing, see e.g., a recent ‘supervised’ principal component analysis for microbiology (Piironen & Vehtari, 2017).

Explorative data analysis is particularly relevant in high dimensional and correlated measurements such as images, times series, frequency spectra or text and where the data generating process yields relatively ‘clean’ measurements, i.e., the variance of the signal of interest  $\sigma_{\text{signal}}^2$  dominates the variance of the measurement noise  $\sigma_{\text{noise}}^2$  leaving a signal-to-noise variance ratio bigger than one:  $S = \sigma_{\text{signal}}^2 / \sigma_{\text{noise}}^2 > 1$ . The structure of such a clean signal can be learned from data by PCA: Let the measured signal be a  $D$ -dimensional vector  $\mathbf{x}$  and let the training sample consist of  $N$  such measurements  $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$ . Assuming that a non-informative mean has been subtracted from the data we can infer the pattern of covariance in measurement space  $\hat{\mathbf{a}}$  that maximally ‘explains’ variance, as  $\hat{\mathbf{a}} = \arg \max_{\|\mathbf{a}\|=1} \sum_{n=1}^N (\mathbf{a}^T \cdot \mathbf{x}_n)^2$ . Empirically, it is found that for sample sizes large enough, the estimate  $\hat{\mathbf{a}}$  stabilizes. In simulation studies where the ground truth direction  $\mathbf{a}_0$  is known we can use the quantity  $R^2 = (\hat{\mathbf{a}}^T \cdot \mathbf{a}_0)^2$  to measure the alignment between estimated and ground truth directions and there we find that as  $N \rightarrow \infty$  perfect alignment  $R^2 \rightarrow 1$  is obtained under weak conditions.

Principal component analysis has gained significant theoretical interest and a rather complete analytical understanding has been obtained in the big data limit  $D, N \rightarrow \infty$  (Biehl & Mietzner, 1993; Hoyle & Rattray, 2007). In fact in this limit a remarkable universality is found. If we define the ratio of samples to dimensionality  $\alpha = N/D$ , the analytical result

$$\langle R^2 \rangle_{\mathcal{X}} = \begin{cases} 0 & \alpha S^2 < 1, \\ \frac{\alpha S^2 - 1}{S + \alpha S^2} & \alpha S^2 \geq 1 \end{cases} \quad (1)$$

resembles a ‘phase transition’ at the critical value of the effective sample size:  $\alpha_{\text{critical}} = S^{-2}$ . The relation between signal-to-noise ratio and effective sample size is highly non-linear: If PCA is presented with too small an effective sample size, we essentially learn nothing  $\langle R^2 \rangle_{\mathcal{X}} = 0$ . Important

for applications PCA is quite robust to sample size in the large sample size limit:  $\langle R^2 \rangle_{\mathcal{X}} = 1 - \frac{S+1}{S^2} \frac{1}{\alpha} + o\left(\frac{1}{\alpha^2}\right)$ . The theoretical results are accurate already for modest sample sizes  $N \sim 10^2 - 10^3$ .

## 2. Learning from incomplete data

While conventional PCA can be said to be well-understood, an additional complication occurs often in practical applications. Namely, that the sampling process is imperfect leading to so-called ‘missing data’, see e.g., (Van Ginkel et al., 2014) for an example involving explorative analysis of questionnaires. How does missing data affect our ability to learn signal structures by PPCA?

Conventional wisdom is that missing data leads to effectively smaller sample size, viz., the following statement in the highly cited review by Schafer and Graham (2002): ‘*In missing data problems the sample may have to be larger than usual, because missing values effectively reduce the sample size*’. While it is true that the problems of missing data can be mitigated by increasing sample size, we here show that missing data is more accurately described as effectively reducing the signal-to-noise ratio.

A simple approach to handling missing data is to ‘impute’ the missing values, however, the specific imputation method may significantly impact the results (Dray & Josse, 2015). A probabilistic approach can formulate the inference problem based on the observed values only, assuming that the missing mechanism is ‘missing at random’, MAR, a weaker assumption than ‘missing completely at random’, MCAR (Little & Rubin, 2014). For example, so-called probabilistic PCA allows inference of principal components in the face of missing data, without imputation, when assuming the missing mechanism is MAR.

## 3. Extending replica analysis to handle missing data

While there are numerous empirical studies of imputation schemes and their performance, there has so far not been reported attempts to generalize Eq. (1) to missing data. We here present a theoretical analysis expanding the already mentioned analytical results for PCA learning to include the effects of missing data. While intuitively one would expect that missing data could take the form of an effective reduction of the sample size  $N$ , we find analytically that when data is ‘missing completely at random’ (Little & Rubin, 2014) at a rate of  $m$ , the role of missing data is instead to reduce the effective signal-to-noise ratio.

While the details are quite involved and deferred to the appendix, a surprisingly simple result is found, namely a straightforward generalization of Eq. (1) with  $S$  replaced by




Figure 1. Conceptually, missing data impacts the cosine similarity,  $R^2$ , between the true signal direction and the signal direction as estimated by probabilistic PCA. As the missing rate increases the estimated principal direction will deviate more and more from the true signal direction and  $R^2$  will decrease according to Eq. (2). In a high dimensional dataspace, the estimated principal direction will on average be orthogonal to the true signal direction when the missing rate is higher than some critical value. Conventional wisdom is that the impact of missing data is similar to a reduced sample size. We argue that a more accurate description is that the impact of missing data is to reduce the signal-to-noise ratio.

an effective signal-to-noise ratio  $S(m) = (1 - m)S$ ,

$$\langle R^2 \rangle_{\mathcal{X}} = \begin{cases} 0 & \alpha S(m)^2 < 1, \\ \frac{\alpha S(m)^2 - 1}{S(m) + \alpha S(m)^2} & \alpha S(m)^2 \geq 1. \end{cases} \quad (2)$$

We have applied probabilistic PCA to high dimensional datasets and simulated data, to investigate how missing data influences the principal components learned. As illustrated in Figure 1 an increasing missing rate implies a decrease in how well we can expect signal structures to be inferred using probabilistic PCA. A phase transition occurs in the learning curves at a critical missing rate,  $m_c = 1 - \frac{1}{\sqrt{\alpha S}}$ . The theoretical result Eq. (2) is remarkably accurate as seen in Figure 2a and 2b. While at the transition point  $m_c$ , scaling  $\alpha$  by  $(1 - m)^2$  and scaling  $S$  by  $(1 - m)$  are equivalent, only the latter leads to correct scaling for all  $m$ .

## 4. Design of simulation studies and relevance for real world data

We generate learning curves on simulated datasets and five high dimensional datasets: Faces in the Wild (Berg et al., 2005), MNIST (LeCun et al., 1998), NCI60 Cancer Microarray Project (Ross et al., 2000), a food pairing dataset (Ahn et al., 2011) and FashionMNIST (Xiao et al., 2017).

The generative process in probabilistic PCA assumes that data arise as a linear mapping of a  $k$ -dimensional, normally distributed latent variable  $\mathbf{z}$  into the  $D$ -dimensional

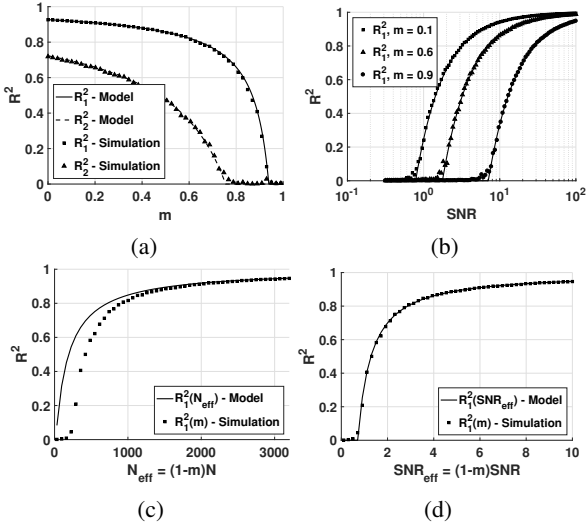


Figure 2.

(a) Learning curves for the first two principal components, varied missing rate. (b) Learning curves for the first principal component at three different missing rates, varied signal-to-noise ratio. (c) Learning curve for the first principal component, varied missing rate. The full line represents the learning curve that would result from an effective sample size as  $N(m) = (1 - m)N$  in Eq. (1). (d) Learning curve for the first principal component, varied missing rate. The full line represents the learning curve that would result from an effective signal-to-noise ratio as  $S(m) = (1 - m)S$ , i.e. the proposed relation in Eq. (2).

data space  $\mathbf{x}$  through the transformation matrix  $\mathbf{A}$ ,  $\mathbf{x} = \mathbf{A}\mathbf{z} + \epsilon$ ,  $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ ,  $\epsilon \sim \mathcal{N}(\mathbf{0}, \sigma^2\mathbf{I})$ . The signal is degraded by isotropic zero mean noise of variance  $\sigma^2$ . By assigning values to  $\mathbf{A}$  and  $\sigma^2$  we implicitly control the signal-to-noise ratio, the ratio between the signal variance and noise variance  $S_i = \frac{\|\mathbf{a}_i\|^2}{\sigma^2}$ , where  $\mathbf{a}_i$  is the  $i$ 'th column of the  $\mathbf{A}$  matrix, i.e., the  $i$ 'th signal direction. When a full dataset has been generated a missing data process is introduced, for this purpose missing completely at random. At a given missing rate  $m$ , each element of the data matrix is set to 'missing' with probability  $m$ .

For a range of missing rates the signal directions can be estimated using PPCA and subsequently compared to the true symmetry breaking directions used to create the data, in terms of the squared cosine similarity,  $R^2$ . This experiment can be repeated a number of times to produce an estimate of the expectation  $\langle R^2 \rangle_{\mathcal{X}}$ . To show the dependency on the signal-to-noise ratio, simulations have also been made where the noise  $\sigma^2$  is varied instead of the missing rate. For each repetition a new noise-free dataset is generated and data-matrix elements are set to 'missing' according to a fixed missing rate. For a range of noise variances, noise is

added to the data and signal directions are estimated using PPCA, to get an estimate of the expectation  $\langle R^2 \rangle_{\mathcal{X}}$ .

In order to apply the theory to the real datasets, we estimate and manipulate their signal-to-noise ratios. The signal-to-noise ratio can be estimated from the eigenvalues of a singular value decomposition. First the noise variance is estimated as the average of the eigenvalues in the non-signal dimensions. As the noise is assumed to be isotropic, the variance in the signal direction consists of the signal variance plus the noise variance. The signal variance is therefore obtained as the eigenvalue in the signal direction minus the noise variance. Finally the signal-to-noise ratio is the ratio between the signal variance and the noise variance  $\hat{\sigma}^2 = \frac{1}{D-k} \sum_{i=k+1}^D \lambda_i$ ,  $S_i = \frac{\lambda_i - \hat{\sigma}^2}{\hat{\sigma}^2}$ , where  $k$  is the number of signal directions. In order to manipulate the signal-to-noise ratio, zero-mean Gaussian noise with isotropic variance,  $\sigma_{\text{added}}^2$ , is added to the dataset and the resulting signal-to-noise ratio becomes  $S_i = \frac{\lambda_i - \hat{\sigma}^2}{\hat{\sigma}^2 + \sigma_{\text{added}}^2}$ .

## 5. Results

Figure 2a was generated using 50 missing rates,  $N = 2000$ ,  $D = 3000$  and  $\sigma^2 = 0.05$ . Furthermore we set  $\|\mathbf{a}_1\|, \|\mathbf{a}_2\| = 1.0, 0.5$  resulting in signal-to-noise ratios of  $S_1 = 20$  and  $S_2 = 5$ . The full and dashed lines represent our proposed theory, Eq. (2), and the markers are simulation results averaged over 5 repetitions, approximating the expectation  $\langle R^2 \rangle_{\mathcal{X}}$ .

Following the learning curves from right to left; when everything is missing nothing can be learned about the underlying signal direction. As we decrease the missing rate a critical point is reached where learning for the 1st signal direction starts, with a steep phase transition. As even lesser data is missing, learning of the 2nd signal direction also initiates with a sharp phase transition. After these steep increases, the learning curves enter a more flat plateau phase where learning is robust to the missing rate.

In figure 2b we vary the signal-to-noise ratio for three different fixed missing rates, all three learning curves are for the 1st signal direction. Here we see that with higher missing rates, we need a better signal-to-noise ratio before learning starts for the 1st signal direction.

The figure was generated using 100 signal-to-noise ratios,  $N = 3200$  and  $D = 1600$ . The signal-to-noise ratio was obtained by adding zero-mean gaussian noise of different magnitude in variance to the linear transformation, keeping  $\|\mathbf{a}_1\| = 1$ . We plot averages over 10 datasets (error bars are smaller than the marker size).

Following (Schafer & Graham, 2002), we could hypothesize that missing data changes the learning curves by decreasing the effective number of samples in the data, in the simplest case this would amount to  $N(m) = (1 - m)N$ . However,

the theory in Eq. (2) suggests that the missing rate affects the learning curves by reducing the effective signal-to-noise ratio as  $S(m) = (1 - m)S$ . This is illustrated in figure 2c and 2d. Figure 2c compares the theory in Eq. (1), assuming an effective sample size, to corresponding missing rate simulations. In figure 2d the theory in Eq. (2) is compared to corresponding missing rate simulations.

In five high dimensional datasets we manipulate the signal-to-noise ratio and missing rate in order to investigate the learning curves for the 1st signal direction. Results are found in figure 3, three learning curves for each dataset, each of them for the first signal direction at a fixed missing rate for varying signal-to-noise ratio.

### Faces in the Wild

This is a collection of face images captured under natural conditions and the principal components are eigenfaces capturing natural face appearance variation (Turk & Pentland, 1991). The first signal direction in the dataset is found using regular PCA in the no-missing case and used as a surrogate for the true signal direction. Now we set a fixed missing rate and the signal-to-noise ratio is gradually decreased by adding gaussian noise of increasing magnitude, to get the learning curves seen in figure 3a. Our theory, Eq. (2), is seen to match this dataset very well.

Plot symbols are averages over 10 datasets and solid lines are theoretical learning curves as described by Eq. (2). The estimate of the true signal direction deteriorates slowly as more and more noise is introduced, until a phase transition occurs and learning stops. At higher missing rates learning stops at a higher signal-to-noise ratio, i.e. missing data affects learning so that we need a better signal-to-noise ratio in order to learn with the same learning performance as in the non-missing case.

### MNIST

The corresponding analysis for the handwritten digit dataset is provided in figure 3b. For a missing rate less than 0.6 the learning curve quite accurately follows the theory, while at larger rates, there are some deviations between theory and experiment. We observe a pronounced phase transition and that an increased signal-to-noise ratio is needed for higher missing rates before learning starts. Even though the MNIST dataset and the Faces in the Wild dataset are both image databases, the differences in the general appearance of the images between the two databases are affecting compliance with the PPCA model assumptions. Probabilistic PCA assumes an isotropic noise variance and this assumption is violated in the MNIST dataset, where there are pixel locations along the image boundary where there is no variation at all throughout the dataset.

### NCI60

The microarray dataset is a collection of expression levels of genes in a number of different cell types. The signal

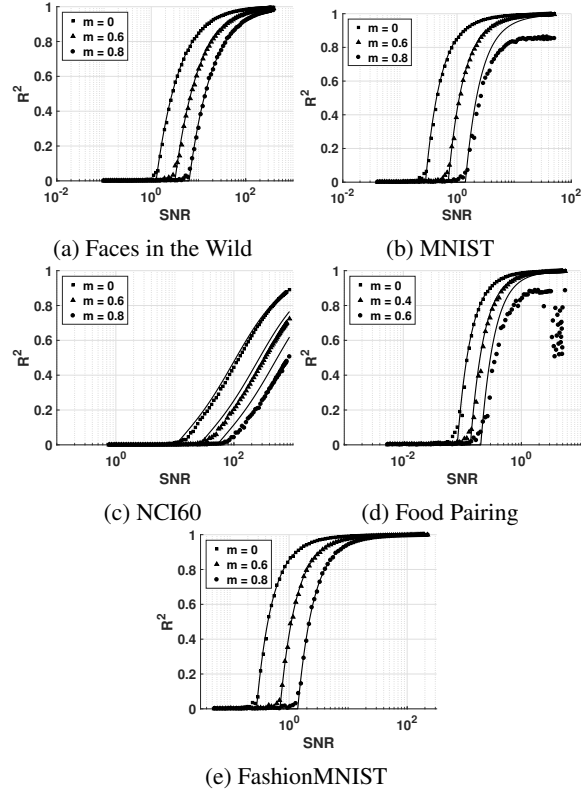


Figure 3. Learning curves for varying signal-to-noise ratio at different missing rates for five different high dimensional datasets. Plot symbols are averages over 10 different noise simulations and the full lines are learning curves as described by Eq. (2).

variance in this dataset is due to different genes having the same kind of expression across cell types and the principal components are referred to as eigengenes (Alter et al., 2000). Learning curves for the micro array dataset are shown in figure 3c. For a missing rate of 0.6 and 0.8 an increasing deviation between the theoretical learning curve and the actual learning curve is seen. This dataset only has  $N = 64$  samples and therefore we may see finite size effects.

### Food pairing

The food pairing dataset consists of food recipes, each of which containing one or more ingredients. The learning curves are found in figure 3d. The model assumptions for PPCA are violated in the sense that we are not looking at continuous data, but a binary indication of whether an ingredient was present in a recipe or not. Some ingredients are only in one recipe, leaving very little variance in this direction ( $\sim 1/N = 0.000017$ ). This may explain the poor adherence for this dataset to the model curves. We have changed the missing rates here from 0.6 and 0.8 to 0.4 and 0.6 instead. The learning curve at a missing rate of 0.8 does not learn.

### FashionMNIST

In figure 3e we present learning curves for the FashionMNIST dataset, which is designed to match the MNIST data in dimensions, sample size etc, but different in image contents. Here we have enough variation in all feature dimensions that this dataset adheres to the generative model of PPCA, and the learning curves accurately follow the model.

Overall, our theory fits simulation studies while the theory seems to match real world datasets with deviations. The Faces in the Wild dataset adheres very well to the theory and a corresponding  $\langle R^2 \rangle_{\mathcal{X}}$  experiment where only the missing rate is varied can be seen in figure 4. On the other hand the MNIST and Food Pairing datasets show some clear deviations between experimental and theoretical learning curves. These deviations may be due to violations of the PPCA model assumption that the noise variance is isotropic. This assumption is definitely violated in the MNIST dataset as there are several dimensions with no variance at all, for example corner pixels, and the covariance matrix becomes degenerate. Similar observations are made in the food pairing dataset. The PPCA model is limited in its noise variance modelling and a more flexible model like factor analysis could potentially better capture the noise in these datasets. Code for generating learning curves available at [https://github.com/nbip/ppca\\_ICML2019](https://github.com/nbip/ppca_ICML2019).




Figure 4. Missing rate experiment for the Faces in the Wild dataset. The images in the bottom row are example images illustrating the missing rate. The images in the top row are the corresponding first principal component as found by PPCA.

## 6. Discussion

We have discussed the effects of missing data on our ability to learn signal structure. We generalized the well-established theoretical results for principal components (Biehl and Mietzner, 1993) to include missing data, to find a result which is surprising in two ways. First, the generalization is remarkably simple modification of the known result for learning in PCA. Secondly, we find that the effect of missing data scales the effective signal-to-noise ratio, rather than scaling the sample size which has been earlier thought.

The theory further predicts a phase transition in the learning curves as the missing rate increases. We used probabilistic

principal component analysis for estimating signal structures in datasets with missing data and found consistent results in both in simulation data and in real datasets. Hence in conclusion: Missing data reduces signal-to-noise ratio and gives rise to a phase transition as the missing rate increases - from almost perfect learning to no learning at all.

In the Low Rank Matrix Completion (LRMC) literature empirical and theoretical results have been presented indicating phase transition like behavior as reported here (e.g (Gao & Ganguli, 2015)). In our presentation we have focused on methods that avoid imputation by marginalization. Our theoretical results are likewise based on the analysis of marginal distributions over the non-missing data, i.e., the generating function is a function of the measured data only in Eq. (30). This is in contrast to the LRMC literature’s focus on imputation.

### Appendix A: Learning curves in the limit of large data

In the simple PCA model we seek a single direction  $\mathbf{a} \in \mathbb{R}^d$  for which the mean projection is maximal

$$\hat{\mathbf{a}} = \arg \max_{\|\mathbf{v}\|=1} \sum_{n=1}^N (\mathbf{v}^T \cdot \mathbf{x}_n)^2. \quad (3)$$

The key question is how  $\hat{\mathbf{a}}(\mathcal{X})$  relates to  $\mathbf{a}_0$  as a function of  $N, D, \sigma^2$ ? This problem has been the subject of quite detailed analysis in the machine learning / statistical mechanics community (Hansel et al., 1992; Biehl & Mietzner, 1993; Watkin et al., 1993; Reimann & Van den Broeck, 1996; Buhot & Gordon, 1998). A particularly detailed account generalizing the transition to more general PCA was given in (Hoyle & Rattray, 2007) using the replica approach. The first use of the replica method for such unsupervised learning is often attributed to the Biehl and Mietzner preprint of the paper (Biehl & Mietzner, 1993). Earlier work on closely related supervised models include (Sompolinsky et al., 1990; Seung et al., 1992; Biehl & Mietzner, 1993).

We are interested in quantifying the typical values of  $(\mathbf{a}_0^T \cdot \mathbf{a})^2$  as function of the physical parameters, hence we follow (Hoyle & Rattray, 2007) and define the generating function (partition function),

$$Z(\mathbf{h}) = \int d\mathbf{a} \delta(\|\mathbf{a}\|^2 - 1) \exp \left( \beta \sum_{n=1}^N (\mathbf{a}^T \cdot \mathbf{x}_n)^2 + \mathbf{h}^T \cdot \mathbf{a} \right). \quad (4)$$

The quantity of interest for a specific data set can be obtained directly from  $Z$

$$R^2 = (\mathbf{a}_0^T \cdot \mathbf{a}(\mathcal{X}))^2 = \lim_{\mathbf{h} \rightarrow 0, \beta \rightarrow \infty} \left( \frac{\partial \log Z(\mathbf{h})}{\partial \mathbf{h}} \cdot \mathbf{a}_0 \right)^2. \quad (5)$$

It is not hard to see that

$$(\mathbf{a}_0^T \cdot \mathbf{a}(\mathcal{X}))^2 = (\mathbf{a}_0^T \cdot \mathbf{u}_1(\mathcal{X}))^2, \quad (6)$$

where  $\mathbf{u}_1(\mathcal{X})$  is eigenvector with the maximal eigenvalue of the empirical covariance matrix  $\widehat{\mathbf{C}} = \frac{1}{N} \sum_{n=1}^N \mathbf{x}_n \mathbf{x}_n^T$ .

More interesting is it to find the typical value of  $R^2$ , defined as

$$\langle (\mathbf{a}_0^T \cdot \mathbf{a}(\mathcal{X}))^2 \rangle_{\mathcal{X}} = \int (\mathbf{a}_0^T \cdot \mathbf{u}_1(\mathcal{X}))^2 \prod_{n=1}^N p(\mathbf{x}_n | \mathbf{C}) d\mathbf{x}_n, \quad (7)$$

where the data follows a zero mean multivariate normal law  $p(\mathbf{x} | \mathbf{C})$  with covariance matrix  $\mathbf{C} = \sigma^2 \mathbf{I} + \mathbf{a}_0 \mathbf{a}_0^T$ . Unfortunately, it is entirely non-trivial to compute the distribution of this eigenvector. A brute force calculation entails the averaging of the log-characteristic function with respect to the data distribution. In statistical mechanics this is approached using a series of tricks referred to as "the replica method". For completeness we here first reproduce the arguments for PCA and then make a generalization to the case with missing data. The first step is to invoke the limit

$$\log Z = \lim_{w \rightarrow 0} \frac{Z^w - 1}{w} \quad (8)$$

and use the fact that averaging powers of  $Z$  is more feasible than averaging the logarithm. Formally, the  $w$ 'th power of  $Z$  can be written

$$Z^w(\mathcal{X}) = \int \prod_{j=1}^w d\mathbf{a}_j \delta(\|\mathbf{a}_j\|^2 - 1) \exp \left( \beta \sum_{n=1}^N \sum_{j=1}^w (\mathbf{a}_j^T \cdot \mathbf{x}_n)^2 + \sum_{j=1}^w \mathbf{h}^T \cdot \mathbf{a}_j \right).$$

The average over  $\mathcal{X}$  then becomes

$$\langle Z^w(\mathcal{X}) \rangle_{\mathcal{X}} = \int \prod_{n=1}^N p(\mathbf{x}_n | \mathbf{C}) d\mathbf{x}_n \int \prod_{j=1}^w d\mathbf{a}_j \delta(\|\mathbf{a}_j\|^2 - 1) \exp \left( \beta \sum_{n=1}^N \sum_{j=1}^w (\mathbf{a}_j^T \cdot \mathbf{x}_n)^2 + \sum_{j=1}^w \mathbf{h}^T \cdot \mathbf{a}_j \right).$$

When exchanging the sequence of integrals the average appears as an integral over a product of  $N$  identical terms, i.e.,

$$\langle Z^w(\mathcal{X}) \rangle_{\mathcal{X}} = \int \prod_{j=1}^w d\mathbf{a}_j \delta(\|\mathbf{a}_j\|^2 - 1) \left( \int p(\mathbf{x} | \mathbf{C}) d\mathbf{x} \exp \left( \beta \sum_{j=1}^w (\mathbf{a}_j^T \cdot \mathbf{x})^2 + \sum_{j=1}^w \mathbf{h}^T \cdot \mathbf{a}_j \right) \right)^N.$$

For simplicity, we ignore the generating function variable  $\mathbf{h}$  and rewrite the quadratic terms

$$\begin{aligned} \exp(\beta(\mathbf{a}_j^T \cdot \mathbf{x})^2) &= \exp\left(\frac{1}{2} 2\beta(\mathbf{a}_j^T \cdot \mathbf{x})^2\right) \\ &= \int \prod_{j=1}^w \frac{du_j}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}u_j^2 + \sqrt{2\beta}u_j\mathbf{a}_j^T \cdot \mathbf{x}_n\right), \end{aligned}$$

this introduces a  $w$  dimensional array of normal integrals over replica parameters  $u_j$ . We next compute the data integral

$$\begin{aligned} &\int p(\mathbf{x} | \mathbf{C}) d\mathbf{x} \exp\left(\beta \sum_{j=1}^w (\mathbf{a}_j^T \cdot \mathbf{x})^2\right) \\ &= \int p(\mathbf{x} | \mathbf{C}) d\mathbf{x} \int \prod_{j=1}^w \frac{du_j}{\sqrt{2\pi}} \exp\left(-\frac{1}{2} \sum_{j=1}^w u_j^2 + \sum_{j=1}^w \sqrt{2\beta} u_j \mathbf{a}_j^T \cdot \mathbf{x}_n\right) \\ &= \int \prod_{j=1}^w \frac{du_j}{\sqrt{2\pi}} \exp\left(-\frac{1}{2} \sum_{j=1}^w u_j^2\right) \int p(\mathbf{x} | \mathbf{C}) d\mathbf{x} \exp\left(\sum_{j=1}^w \sqrt{2\beta} u_j \mathbf{a}_j^T \cdot \mathbf{x}_n\right). \end{aligned} \quad (9)$$

The Gaussian data integral can be performed using

$$\begin{aligned} &\int p(\mathbf{x} | \mathbf{C}) d\mathbf{x} \exp\left(\sum_{j=1}^w \sqrt{2\beta} u_j \mathbf{a}_j^T \cdot \mathbf{x}\right) \\ &= \exp\left(\beta \sum_{j,j'=1}^w u_j u_{j'} \mathbf{a}_j^T \mathbf{C}^{-1} \mathbf{a}_{j'}\right). \end{aligned} \quad (11)$$

Combining Eqs. (10) and (11) we get the final replica integral

$$\begin{aligned} \langle Z^w(\mathcal{X}) \rangle_{\mathcal{X}} &= \int \prod_{j=1}^w d\mathbf{a}_j \delta(\|\mathbf{a}_j\|^2 - 1) \\ &\left( \int \prod_{j=1}^w \frac{du_j}{\sqrt{2\pi}} \exp\left(-\frac{1}{2} \sum_{j=1}^w u_j^2 + \beta \sum_{j,j'=1}^w u_j u_{j'} \mathbf{a}_j^T \mathbf{C} \mathbf{a}_{j'}\right) \right)^N. \end{aligned}$$

To proceed, we use  $\mathbf{C} = \sigma^2 \mathbf{I} + \mathbf{a}_0 \mathbf{a}_0^T$ , to re-write the Gaus-

sian integral over replica variables

$$\langle Z^w(\mathcal{X}) \rangle_{\mathcal{X}} = \int \prod_{j=1}^w d\mathbf{a}_j \delta(\|\mathbf{a}_j\|^2 - 1) \left( \int \prod_{j=1}^w \frac{du_j}{\sqrt{2\pi}} \exp\left(-\frac{1}{2} \sum_{j,j'=1}^w u_j M_{j,j'} u_{j'}\right) \right)^N \quad (12)$$

with

$$M_{j,j'} = \delta_{j,j'} - 2\beta\sigma^2 (q_{j,j'} + S r_j r_{j'}), \quad (13)$$

and where we have defined the key quantities  $q_{j,j'} = \mathbf{a}_j^T \cdot \mathbf{a}_{j'}$  and  $r_j = \mathbf{a}_j^T \cdot \mathbf{a}_0$  relating the inferred and generating parameters in the replicated system. The Gaussian integrals yield

$$\langle Z^w(\mathcal{X}) \rangle_{\mathcal{X}} = (2\pi)^{-Nw} \int \prod_{j=1}^w d\mathbf{a}_j \delta(\|\mathbf{a}_j\|^2 - 1) |\mathbf{M}(\{\mathbf{a}_j\})|^{-N/2}.$$

The integrals over the  $w$  unit spheres are next parameterized as integrals over the  $q$  matrix and the  $r$  vector,

$$\langle Z^w(\mathcal{X}) \rangle_{\mathcal{X}} = (2\pi)^{-Nw} \int \prod_{j=1}^w d\mathbf{a}_j \delta(\|\mathbf{a}_j\|^2 - 1) \times \int \prod_{j>j'} dq_{j,j'} \delta(q_{j,j'} - \mathbf{a}_j^T \cdot \mathbf{a}_{j'}) \int \prod_j dr_j \delta(r_j - \mathbf{a}_j^T \cdot \mathbf{a}_0) |\mathbf{M}(\mathbf{Q}, \mathbf{r})|^{-N/2}. \quad (14)$$

As the delta functions relate  $\mathbf{Q}$  and the inner products of  $\mathbf{a}$ , we can replace  $\delta(\|\mathbf{a}_j\|^2 - 1)$  by  $\delta(q_{j,j} - 1)$ . The  $\mathbf{Q}$  and  $\mathbf{r}$  integrals can be combined by defining the matrix

$$\mathbf{Q}_+ = \begin{bmatrix} \mathbf{a}_0^T \cdot \mathbf{a}_0 & \mathbf{a}_0^T \cdot \mathbf{a}_1 & \dots & \mathbf{a}_0^T \cdot \mathbf{a}_w \\ \mathbf{a}_1^T \cdot \mathbf{a}_0 & \mathbf{a}_0^T \cdot \mathbf{a}_1 & \dots & \mathbf{a}_1^T \cdot \mathbf{a}_w \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{a}_w^T \cdot \mathbf{a}_0 & \mathbf{a}_w^T \cdot \mathbf{a}_1 & \dots & \mathbf{a}_w^T \cdot \mathbf{a}_w \end{bmatrix} = \begin{bmatrix} 1 & \mathbf{r} \\ \mathbf{r}^T & \mathbf{Q} \end{bmatrix}. \quad (15)$$

With this definition,

$$\langle Z^w(\mathcal{X}) \rangle_{\mathcal{X}} = (2\pi)^{-Nw} \int \prod_{j=1}^w d\mathbf{a}_j \times \int d\mathbf{Q}_+ \prod_{j>j'=0}^w \delta(\mathbf{Q}_{+,j,j'} - \mathbf{a}_j^T \cdot \mathbf{a}_{j'}) \prod_{j=0}^w \delta(\mathbf{Q}_{+,j,j} - 1) |\mathbf{M}(\mathbf{Q}, \mathbf{r})|^{-N/2}, \quad (16)$$

the  $\mathbf{a}$  integrals can be performed to leading order in  $D$  with a procedure outlined in Appendix A of (Ahr et al., 1999). First

we write the integral in matrix form with  $\mathbf{A}$  a  $D \times (w+1)$  matrix

$$\int d\mathbf{A} \delta(\mathbf{Q}_+ - \mathbf{A}^T \cdot \mathbf{A}) = |\mathbf{U}^T \mathbf{\Lambda}|^D \int d\tilde{\mathbf{A}} \delta(\mathbf{I} - \tilde{\mathbf{A}}^T \cdot \tilde{\mathbf{A}}). \quad (17)$$

Here we used the transformation  $\mathbf{A} \rightarrow \tilde{\mathbf{A}} = \mathbf{A}\mathbf{U}\mathbf{\Lambda}$  based on the spectral representation of  $(w+1) \times (w+1)$  matrix  $\mathbf{Q}_+ = \mathbf{U}^T \mathbf{\Lambda} \mathbf{U}$ . Here  $\mathbf{U}$  is an orthogonal matrix and  $\mathbf{\Lambda}$  is a diagonal matrix with the square roots of the eigenvalues of symmetric real matrix  $\mathbf{Q}_+$ . The integral  $\int d\tilde{\mathbf{A}} \delta(\mathbf{I} - \tilde{\mathbf{A}}^T \cdot \tilde{\mathbf{A}})$  is seen to be a constant with respect to the integration variables  $\mathbf{Q}_+$ . Hence, to leading order the dependent term is

$$\int d\mathbf{A} \delta(\mathbf{Q}_+ - \mathbf{A}^T \cdot \mathbf{A}) \propto |\mathbf{U}^T \mathbf{\Lambda}|^D \propto |\mathbf{Q}_+|^{\frac{D}{2}}. \quad (18)$$

The determinant can be computed using the identity

$$\begin{vmatrix} \mathbf{X}_{11} & \mathbf{X}_{12} \\ \mathbf{X}_{21} & \mathbf{X}_{22} \end{vmatrix} = |\mathbf{X}_{11}| |\mathbf{X}_{22} - \mathbf{X}_{21} \mathbf{X}_{11}^{-1} \mathbf{X}_{12}| \quad (19)$$

with  $\mathbf{X}_{11} = 1$ ,  $\mathbf{X}_{21} = \mathbf{r}$  and  $\mathbf{X}_{22} = \mathbf{Q}$ ,

$$\langle Z^w(\mathcal{X}) \rangle_{\mathcal{X}} \propto \int \prod_{j>j'} dq_{j,j'} \int \prod_j dr_j |\mathbf{Q} - \mathbf{r}\mathbf{r}^T|^{D/2} |\mathbf{M}(\mathbf{Q}, \mathbf{r})|^{-N/2}. \quad (20)$$

To bring the scaling with  $D, N \rightarrow \infty$  out more clearly, we rewrite

$$\langle Z^m(\mathcal{X}) \rangle_{\mathcal{X}} \propto \int d\mathbf{Q} d\mathbf{r} \exp \left[ D \left( \frac{1}{2} \log |\mathbf{Q} - \mathbf{r}\mathbf{r}^T| - \frac{\alpha}{2} \log |\mathbf{I} - 2\beta\sigma^2(\mathbf{Q} + S\mathbf{r}\mathbf{r}^T)| \right) \right], \quad (21)$$

with the definitions from the main text  $\alpha = \frac{N}{D}$ ,  $S = \frac{1}{\sigma^2}$ .

We are interested in the limit of large  $D$ , where the integral is dominated by the stationary points of the integrand. We further assume "replica symmetry", hence, the following simple structures of  $\mathbf{Q}, \mathbf{r}$

$$\mathbf{Q} = \begin{bmatrix} 1 & q & \dots & q \\ q & 1 & \dots & q \\ \vdots & \vdots & \ddots & \vdots \\ q & q & \dots & 1 \end{bmatrix}, \quad \mathbf{r} = \begin{bmatrix} r \\ r \\ \vdots \\ r \end{bmatrix} \quad (22)$$

The two  $w \times w$  determinants can then be computed using simple algebra

$$\begin{vmatrix} z & y & \dots & y \\ y & z & \dots & y \\ \vdots & \vdots & \ddots & \vdots \\ y & y & \dots & z \end{vmatrix} = \left(1 + \frac{wy}{z-y}\right) (z-y)^w. \quad (23)$$

Hence,

$$\log |\mathbf{Q} - \mathbf{r}\mathbf{r}^T| = \log \left( 1 + \frac{w(q - r^2)}{1 - q} \right) + w \log(1 - q) \quad (24)$$

$$\begin{aligned} & \log |\mathbf{I} - 2\beta\sigma^2(\mathbf{Q} + S\mathbf{r}\mathbf{r}^T)| = \\ & \log \left( 1 - w \frac{2\beta\sigma^2(q + Sr^2)}{1 - \beta\sigma^2(1 - Sr^2)} \right) \\ & + w \log(1 - 2\beta\sigma^2(1 - Sr^2)). \end{aligned}$$

Finally, we send  $w \rightarrow 0$  to get,

$$\lim_{w \rightarrow 0} \frac{1}{w} \log |\mathbf{Q} - \mathbf{r}\mathbf{r}^T| = \frac{q - r^2}{1 - q}, \quad (25)$$

and

$$\begin{aligned} & \lim_{m \rightarrow 0} \frac{1}{w} \log |\mathbf{I} - \beta\sigma^2(\mathbf{Q} + S\mathbf{r}\mathbf{r}^T)| \\ & = - \frac{2\beta\sigma^2(q + Sr^2)}{1 - 2\beta\sigma^2(1 - q)}. \end{aligned}$$

Following (Reimann & Van den Broeck, 1996) we locate the stationary points first wrt.  $q$  for fixed  $r$  and then find the resulting stationary point wrt.  $r$ . To leading order the averaged log partition function is given by

$$\langle \log Z(\mathcal{X}) \rangle_{\mathcal{X}} \propto \frac{1 + r^2}{v} + 2\alpha \frac{1 + Sr^2}{1 - 2v}, \quad (26)$$

where  $v = \beta\sigma^2(1 - q)$ . To simplify the location of the stationary points we further define the two functions of  $r$ :  $\phi_1 = 1 + r^2$  and  $\phi_2 = \alpha(1 + Sr^2)$ .

$$\langle \log Z(\mathcal{X}) \rangle_{\mathcal{X}} \propto \left( \frac{\phi_1}{v} + \frac{\phi_2}{1 - 2v} \right). \quad (27)$$

The stationary point is given by

$$\hat{v} = \frac{1}{2} \frac{\sqrt{\phi_1}}{\sqrt{\phi_1} \pm \sqrt{\phi_2}} \quad (28)$$

We choose the positive square root and substitute back into the log partition function to obtain

$$\begin{aligned} & \langle \log Z(\mathcal{X}) \rangle_{\mathcal{X}} \propto \left( \sqrt{\phi_1} \pm \sqrt{\phi_2} \right) \\ & = \left( \sqrt{1 + r^2} + \sqrt{\alpha(1 + Sr^2)} \right)^2. \end{aligned}$$

Finally, the extremal point wrt.  $r$  is given by

$$\langle (\mathbf{a}_0 \cdot \mathbf{a})^2 \rangle_{\mathcal{X}} = \hat{r}^2 = \begin{cases} \frac{\alpha S^2 - 1}{S(1 + \alpha S)} & \text{if } \alpha > \frac{1}{S^2} \\ 0 & \text{if } \alpha \leq \frac{1}{S^2} \end{cases} \quad (29)$$

producing the non-smooth learning curve result, i.e., relation between  $\mathbf{a}$  and  $\mathbf{a}_0$  as a function of the relative sample size  $\alpha$ . It is remarkable that in the large scale system, there is a range of sample sizes  $N < N_{\text{critical}} \equiv D\sigma^4$  for which learning is completely absent.

### Statistical mechanics of simple PCA with missing data

To account for missing data we introduce the missing indicator  $s_{n,d} \in [0, 1]$ , with  $s_{n,d} = 1$  when the feature is present. The averaging procedure increases a bit in bookkeeping as we only average over the actual present features in each sample, hence, we obtain

$$\begin{aligned} \langle Z^w(\mathcal{X}) \rangle_{\mathcal{X}} &= \int_{\|\mathbf{a}_j\|^2=1} d\mathbf{A} \prod_{n=1}^N \int p(\mathbf{x}_n | \mathbf{C}_n) d\mathbf{x}_n \\ & \exp \left( \beta \sum_{n=1}^N \sum_{j=1}^w (\mathbf{a}_j^T \cdot (\mathbf{s}_n \odot \mathbf{x}_n))^2 + \sum_{j=1}^w \mathbf{h}^T \cdot \mathbf{a}_j \right), \end{aligned} \quad (30)$$

where the element wise product is  $(\mathbf{s}_n \odot \mathbf{x}_n)_d = s_{n,d}x_{n,d}$ , and  $\mathbf{C}_n$  is the sub-matrix defined by the present features in the  $n$ 'th sample. The generating function in Equation (30) is equivalent to the marginalization in probabilistic PCA, i.e., we integrate only the present features. Thus, Equation (12) becomes

$$\begin{aligned} \langle Z^w(\mathcal{X}) \rangle_{\mathcal{X}} &= \int_{\|\mathbf{a}_j\|^2=1} d\mathbf{A} \prod_{n=1}^N \int \prod_{j=1}^w \frac{du_j^n}{\sqrt{2\pi}} \\ & \exp \left( -\frac{1}{2} \sum_{j,j'=1}^w u_j^n M_{j,j'}^n u_{j'}^n \right). \end{aligned} \quad (31)$$

Here,

$$M_{j,j'}^n = \delta_{j,j'} - 2\beta\sigma^2 (q_{j,j'}^n + Sr_j^n r_{j'}^n), \quad (32)$$

with  $q_{j,j'}^n = \sum_{d=1}^D s_{n,d} \mathbf{a}_{j,d} \mathbf{a}_{j',d}$  and  $r_j^n = \sum_{d=1}^D s_{n,d} \mathbf{a}_{j,d} \mathbf{a}_{0,j',d}$ .

Re-writing the product over samples as an exponentiated sum over  $M_{j,j'}^n$  terms, we may invoke a simple self-averaging assumption (Grant, 2004) (the sum can be replaced by its average),

$$\langle M_{j,j'}^n \rangle = \delta_{j,j'} - 2\beta\sigma^2 ((1 - m)q_{j,j'}^n + (1 - m)^2 Sr_j^n r_{j'}^n), \quad (33)$$

where  $m$  is the missing rate and we used that the  $q$  term involves a single  $(1 - m)$  factor while  $r_j^n r_{j'}^n$  term involves two independent  $(1 - m)$  factors.

By inspection we find that this is equivalent to the earlier expression if  $\beta \rightarrow (1 - m)\beta$  and  $S \rightarrow (1 - m)S$ . Hence, the asymptotic expression for  $R^2$  in presence of missing data at rate  $m$  is as already shown in Equation 2,

$$\langle R^2 \rangle_{\mathcal{X}} = \begin{cases} 0 & \alpha((1 - m)S)^2 < 1, \\ \frac{\alpha((1 - m)S)^2 - 1}{(1 - m)S + \alpha((1 - m)S)^2} & \alpha((1 - m)S)^2 \geq 1. \end{cases} \quad (34)$$



## Acknowledgements

The Danish Innovation Foundation supported this work through Danish Center for Big Data Analytics driven Innovation (DABAI).

## References

- Ahn, Y.-Y., Ahnert, S. E., Bagrow, J. P., and Barabási, A.-L. Flavor network and the principles of food pairing. *Scientific reports*, 1, 2011.
- Ahr, M., Biehl, M., and Urbanczik, R. Statistical physics and practical training of soft-committee machines. *The European Physical Journal B-Condensed Matter and Complex Systems*, 10(3):583–588, 1999.
- Alter, O., Brown, P. O., and Botstein, D. Singular value decomposition for genome-wide expression data processing and modeling. *Proceedings of the National Academy of Sciences*, 97(18):10101–10106, 2000.
- Berg, T. L., Berg, A. C., Edwards, J., and Forsyth, D. A. Who’s in the picture. In *Advances in neural information processing systems*, pp. 137–144, 2005.
- Biehl, M. and Mietzner, A. Statistical mechanics of unsupervised learning. *EPL (Europhysics Letters)*, 24(5):421, 1993.
- Buhot, A. and Gordon, M. B. Phase transitions in optimal unsupervised learning. *Physical Review E*, 57(3):3326, 1998.
- Dempster, A. P., Laird, N. M., and Rubin, D. B. Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society. Series B (methodological)*, pp. 1–38, 1977.
- Dray, S. and Josse, J. Principal component analysis with missing values: a comparative survey of methods. *Plant Ecology*, 216(5):657–667, 2015.
- Gao, P. and Ganguli, S. On simplicity and complexity in the brave new world of large-scale neuroscience. *Current opinion in neurobiology*, 32:148–155, 2015.
- Grant, M. Dirty tricks for statistical mechanics, physics department, mcgill university version 0.91, 2004.
- Hansel, D., Mato, G., and Meunier, C. Memorization without generalization in a multilayered neural network. *EPL (Europhysics Letters)*, 20(5):471, 1992.
- Hotelling, H. Analysis of a complex of statistical variables into principal components. *Journal of educational psychology*, 24(6):417, 1933.
- Hoyle, D. C. and Rattay, M. Statistical mechanics of learning multiple orthogonal signals: asymptotic theory and fluctuation effects. *Physical review E*, 75(1):016101, 2007.
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Little, R. J. and Rubin, D. B. *Statistical analysis with missing data*. John Wiley & Sons, 2014.
- Pearson, K. Liii. on lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11):559–572, 1901.
- Piironen, J. and Vehtari, A. Iterative Supervised Principal Components. *ArXiv e-prints*, October 2017.
- Reimann, P. and Van den Broeck, C. Learning by examples from a nonuniform distribution. *Physical Review E*, 53(4):3989, 1996.
- Ross, D. T., Scherf, U., Eisen, M. B., Perou, C. M., Rees, C., Spellman, P., Iyer, V., Jeffrey, S. S., Van de Rijn, M., Waltham, M., et al. Systematic variation in gene expression patterns in human cancer cell lines. *Nature genetics*, 24(3):227, 2000.
- Roweis, S. Em algorithms for pca and spca. *Advances in neural information processing systems*, pp. 626–632, 1998.
- Sanguansat, P. *Principal component analysis-multidisciplinary applications*. InTech, 2012.
- Schafer, J. L. and Graham, J. W. Missing data: our view of the state of the art. *Psychological methods*, 7(2):147, 2002.
- Seung, H., Sompolinsky, H., and Tishby, N. Statistical mechanics of learning from examples. *Physical Review A*, 45(8):6056, 1992.
- Sompolinsky, H., Tishby, N., and Seung, H. S. Learning from examples in large neural networks. *Physical Review Letters*, 65(13):1683, 1990.
- Tipping, M. E. and Bishop, C. M. Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 61(3):611–622, 1999.
- Turk, M. and Pentland, A. Eigenfaces for recognition. *Journal of cognitive neuroscience*, 3(1):71–86, 1991.

- Van Ginkel, J. R., Kroonenberg, P. M., and Kiers, H. A. Missing data in principal component analysis of questionnaire data: a comparison of methods. *Journal of Statistical Computation and Simulation*, 84(11):2298–2315, 2014.
- Watkin, T. L., Rau, A., and Biehl, M. The statistical mechanics of learning a rule. *Reviews of Modern Physics*, 65(2):499, 1993.
- Xiao, H., Rasul, K., and Vollgraf, R. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, 2017.