

Learning from Incomplete Features by Simultaneous Training of Neural Networks and Sparse Coding

Cesar F. Caiafa*

IAR - CONICET, ARGENTINA
Tensor Learning Team - RIKEN AIP, JAPAN
ccaiafa@gmail.com

Jordi Solé-Casals

University of Vic (UVic-UCC), SPAIN
jordi.sole@uvic.cat

Ziyao Wang

School of Automation - SEU, CHINA
Tensor Learning Team - RIKEN AIP, JAPAN
zy_wang@seu.edu.cn

Qibin Zhao

Tensor Learning Team - RIKEN AIP, JAPAN
qibin.zhao@riken.jp

Abstract

In this paper, the problem of training a classifier on a dataset with incomplete features is addressed. We assume that different subsets of features (random or structured) are available at each data instance. This situation typically occurs in the applications when not all the features are collected for every data sample. A new supervised learning method is developed to train a general classifier, such as a logistic regression or a deep neural network, using only a subset of features per sample, while assuming sparse representations of data vectors on an unknown dictionary. Sufficient conditions are identified, such that, if it is possible to train a classifier on incomplete observations so that their reconstructions are well separated by a hyperplane, then the same classifier also correctly separates the original (unobserved) data samples. Extensive simulation results on synthetic and well-known datasets are presented that validate our theoretical findings and demonstrate the effectiveness of the proposed method compared to traditional data imputation approaches and one state-of-the-art algorithm.

1. Introduction

Learning methods from limited or imperfect data has attracted great attention in the literature recently. Datasets with limited, weak, noisy labels or incomplete features represent an important and still open problem. In this paper, we address the problem of training a classifier on a dataset with incomplete features, which arises in many machine learning applications where sometimes the measurements are incomplete, noisy or affected by artifacts. Examples of this situation include: recommendation systems built upon

the information gathered by different users where not all the users have fully completed their forms; medical datasets where typically not all tests can be performed on every patient; or a self-driving vehicle or robot where objects in the view field can be partially occluded.

Handling correctly the incomplete-features problem is a classical challenge in machine learning. Skipping missing features by setting them to zero values damages the classification accuracy [4]. Most previous studies addressed this problem by using an imputation approach, which consists of performing data completion followed by training the classifier with those reconstructions (referred here as the sequential method). However, this strategy cannot ensure the statistical consistence of the classifier, as data completion is usually fully unsupervised or label information is partially or inefficiently exploited.

In this work, a new supervised learning method is developed to train a general classifier, such as a logistic regression or a deep neural network, using only a subset of features per sample, while assuming sparse representations of data vectors on an unknown dictionary. The proposed method simultaneously learns the classifier, the dictionary and the corresponding sparse representation of each input data sample. In this way, we combine the approximation power and simplicity of sparse coding with the extraordinary ability of neural networks (NNs) to model complex decision functions (classifiers) with the goal to successfully train a classifier based on incomplete features.

We analyze the limitations of the sequential approach (section 1.3), *i.e.* imputation followed by training, and introduce the simultaneous classification and coding approach in section 2. Our method consists of incorporating a sparse data representation model into a single cost function that is optimized for training the classifier and, at the same time, finding the best representation of the observed data. A learn-

*Corresponding author

ing algorithm is presented in section 2.1 to train a classifier on incomplete features and sufficient conditions under which such a classifier performs as good as the ideal classifier, *i.e.* the one that can be obtained from complete observations, is identified (section 2.2). Extensive experimental results are presented in section 3, using synthetic and well known benchmark datasets that validate our theoretical findings and illustrate the effectiveness of the proposed method.

1.1. Related work

Practical sequential methods based on statistical imputation, such as computing the “mean”, “regression” and “multiple” imputation techniques are common practice [27]. Remarkably, it was shown that imputing with a constant, *e.g.* the mean, is Bayes-risk consistent only when missing features are not informative [20]. More elaborated completion methods were also explored, such as K -nearest neighbor estimators, multilayer or recurrent NNs and others, see [8] and references therein.

However, sequential methods do not fully exploit label information. Data labels can provide valuable information about missing features that could potentially improve the classifier learning process. Recent advances on probabilistic generative models have allowed for a formulation of supervised learning with incomplete features as a statistical inference problem arriving at algorithms that significantly outperformed sequential methods. In the seminal work [9], a framework for maximum likelihood density estimation based on mixtures models was proposed and successfully applied to small incomplete-features problems. In particular, a Gaussian Mixture Model (GMM) was fitted to incomplete-data through an Expectation-Maximization (EM) algorithm. Building upon this generative model strategy, some approaches have considered integrating out the missing values based on a simple logistic regression function [40, 2]. Other versions of this approach proposed an explicit simultaneous learning of the model and the decision function [25, 7]. While probabilistic generative models provided a nice and elegant approach to the incomplete-features problem showing good results on small datasets, they are not suitable for many modern machine learning applications because: (1) despite some acceleration techniques were explored, *e.g.* [26, 6], those algorithms are computationally expensive becoming prohibitive for moderate to large datasets; (2) GMM is impractical for modeling high-dimensional datasets because the number of parameters to achieve good approximations becomes unmanageable; and (3) they do not consider complex classification functions as the ones provided by deep NN architectures.

Recently, some approaches based on the low-rank property of the features data matrix were investigated and algorithms for data completion were proposed incorporating the label information [10, 14, 17]. Since the rank estimation of

a matrix is a computationally expensive task, usually based on the Singular Value Decomposition (SVD), the obtained algorithms are prohibitive to solve modern machine learning problems with large datasets. Additionally, as in the case of the probabilistic generative models, none of these methods considered complex classification functions. To overcome this drawback, more recently, a framework based on various NN architectures such as autoencoders, multilayer perceptrons and Radial Basis Function Networks (RBFNs), was proposed for handling missing input data by setting a probabilistic model, *e.g.* a GMM, for every missing feature, which is trained together with the NN weights [35]. This method combined the great capability of NNs to approximate complex decision functions with the nice formulation of the GMM to model missing data. However, it inherited the drawbacks of GMMs, *i.e.* they are not well suited to higher-dimensional datasets.

On the other hand, during the last few years in the signal processing community, there has been a rapid development of theory and algorithms for sparse coding approximations which, by exploiting the redundancy of natural signals, are able to provide simple and accurate models of complex data distributions, see [24, 28, 11, 30, 5] and references therein. Sparse coding is nearly ubiquitous in Nature, for example, it is found in the way that neurons encode sensory information [32, 31]. Sparse representations of data showed to be useful also in classification problems. In [16], a Linear Discriminant Analysis (LDA) classifier was trained on corrupted data providing a robust classification method. In [29], algorithms for learning *discriminative* sparse models, instead of purely *reconstructive* ones, were proposed based on simple linear and bilinear classifiers. Similar methods were also studied by either using class-specific dictionaries [33, 36] or using a single one for all classes [37]. However, these proposed methods neither were applied to the incomplete-features problem nor considered deep NN classifiers.

1.2. Problem formulation

We assume a supervised learning scenario with vector samples and labels $\{\mathbf{x}_i, y_i\}$, $i = 1, 2, \dots, I$, $\mathbf{x}_i \in \mathbb{R}^N$ and $y_i \in \{0, 1, \dots, C - 1\}$ (C classes). However, we are constrained to observe only subsets of features and their labels: $\{\mathbf{x}_i^o, y_i\}$, $\mathbf{x}_i^o \in \mathbb{R}^{M_i}$ with $M_i < N$. Unobserved (missing) features are denoted by $\mathbf{x}_i^m \in \mathbb{R}^{N-M_i}$. We consider arbitrary patterns of missing features, which are allowed to be different for each data instance i . The set of indices of missing features at sample i is denoted by \mathcal{M}_i , *i.e.* $\mathbf{x}_i^m = \mathbf{x}_i(\mathcal{M}_i)$ and $\mathbf{x}_i^o = \mathbf{x}_i(\overline{\mathcal{M}}_i)$.

We define the set of all K -sparse vectors $\Sigma_K^P = \{\mathbf{s} \in \mathbb{R}^P \text{ s.t. } \|\mathbf{s}\|_0 \leq K\}$ (containing at most K non-zero entries) and assume that data vectors \mathbf{x}_i admit K -sparse representations over an unknown dictionary $\mathbf{D} \in \mathbb{R}^{N \times P}$ ($P \geq N$):

$$\mathbf{x}_i = \mathbf{D}\mathbf{s}_i, \text{ with } \mathbf{s}_i \in \Sigma_K^P. \quad (1)$$

The columns of a dictionary are called “atoms” because every data vector can be written as a linear combination of at most K elementary components. Sometimes dictionaries are orthogonal such as the ones derived from the Discrete Cosine or Wavelet [30] transforms. However, overcomplete ($P \geq N$) nonorthogonal dictionaries have demonstrated to play an important role in image processing tasks such as denoising, inpainting, etc [28, 19].

By partitioning \mathbf{D} according to the pattern of missing features at sample i , we obtain $\mathbf{D}_i^o = \mathbf{D}(\mathcal{M}_i, :) \in \mathbb{R}^{M_i \times P}$ and $\mathbf{D}_i^m = \mathbf{D}(\mathcal{M}_i^c, :) \in \mathbb{R}^{(N-M_i) \times P}$, which according to equation (1) implies:

$$\mathbf{x}_i^o = \mathbf{D}_i^o \mathbf{s}_i, \text{ and } \mathbf{x}_i^m = \mathbf{D}_i^m \mathbf{s}_i. \quad (2)$$

Let us assume that a perfect classifier, *e.g.* a logistic regression or deep NN, that assigns probability $p_\Theta(\hat{y}|\mathbf{x})$ to predicted label \hat{y} given data \mathbf{x} can be trained on the complete dataset $\{\mathbf{x}_i, y_i\}$, such that, in a two-classes scenario ($C = 2$), $p_\Theta(\hat{y} = y_i|\mathbf{x}_i) > p_\Theta(\hat{y} \neq y_i|\mathbf{x}_i), \forall i = 1, 2, \dots, I$, where Θ is the set of trained parameters. Our goal is to develop a method to obtain an estimate $\hat{\Theta}$ of parameters using only the incomplete dataset $\{\mathbf{x}_i^o, y_i\}$ and to identify conditions under which such a classifier is compatible with the ideal one.

1.3. Why training after imputation is difficult?

If the K -sparse representations of the observations \mathbf{x}_i^o were unique, then \mathbf{x}_i can be perfectly reconstructed from the incomplete observations and the classifier can be successfully trained using these reconstructions. In the particular case where the dictionary is known in advance, there exist conditions on the sampling patterns based on the *coherence*, *spark* or *RIP (Restricted Isometry Property)* of matrix \mathbf{D}_i^o that can guarantee uniqueness [5]. However, these conditions are difficult to meet in practice and determining RIP/Spark properties are NP-hard in general [39]. Moreover, in the general case where the dictionary \mathbf{D} is unknown and needs to be learned from data, it is even more difficult to obtain well separated reconstructions which certainly leads to suboptimal or wrong classifiers.

Next, we provide some intuition about the limitation of the sequential approach through a toy example. Let us consider the classification of hand-written digit images belonging to two classes: “3s” and “8s” and assume that they admit 2-sparse representations over a dictionary. Fig. 1 (a-b) shows the representations of two example vectors \mathbf{x}_i and \mathbf{x}_j belonging to classes “3” and “8”, respectively. If only the right halves of the images are observed and no label information is provided, we are clearly faced with a problem because our observed samples from two different classes are identical, *i.e.* $\mathbf{x}_i^o = \mathbf{x}_j^o$. It is obvious that at least two possible 2-sparse representations for the observed data exist as illustrated in Fig. 1(c). When the sparse solution is not unique, we may end up reconstructing wrong vectors that

could not be even well separated as illustrated in Fig. 1 (d-e). In general, sequential methods using only the information of observed features are prone to fail because the non-uniqueness of solutions can make the training of a good classifier an impossible task. However, we could solve this problem by incorporating the labelling information from the very beginning as it is proposed in the following section.

2. Simultaneous learning and coding approach

We propose to train the classifier and find the proper representation, not only as sparse as possible but also providing the best separation of classes. We want to combine the training of the classifier together with the learning of a dictionary and optimal sparse representations such that the reconstructed data vectors are compatible with observations and well separated. To do that we propose to minimize the following global cost function:

$$J(\Theta, \mathbf{D}, \mathbf{s}_i) = \underbrace{\frac{1}{I} \sum_{i=1}^I \{J_0(\Theta, \hat{\mathbf{x}}_i, y_i) + \lambda_1 J_1(\mathbf{D}, \mathbf{s}_i)\}}_{F(\Theta, \mathbf{D}, \mathbf{s}_i)} + \underbrace{\frac{1}{I} \sum_{i=1}^I \{\lambda_2 J_2(\mathbf{s}_i)\}}_{G(\mathbf{s}_i)}, \quad (3)$$

with respect to Θ , \mathbf{D} and \mathbf{s}_i ($i = 1, 2, \dots, I$), where Θ contains the classifier parameters, *i.e.* the vector of weights in a deep NN classifier architecture; $\mathbf{D} \in \mathbb{R}^{N \times P}$ ($P \geq N$) is a dictionary and $\mathbf{s}_i \in \Sigma_K^P$ are the representation coefficients such that the reconstructed data vectors are $\hat{\mathbf{x}}_i = \mathbf{D}\mathbf{s}_i$.

$J_0(\Theta, \hat{\mathbf{x}}_i, y_i)$ is a measure of the classification error for the reconstructed sample vector $\hat{\mathbf{x}}_i$. Typically, we use the crossentropy measure, *i.e.* $J_0(\Theta, \hat{\mathbf{x}}_i, y_i) = -\log[p_\Theta(y_i|\hat{\mathbf{x}}_i)]$, where $p_\Theta(y_i|\hat{\mathbf{x}}_i)$ is the probability assigned by the classifier to sample $\hat{\mathbf{x}}_i$ as belonging to class y_i . $J_1(\mathbf{D}, \mathbf{s}_i)$ is a measure of the approximation error of the reconstruction when it is restricted to observed features, which is defined as follows: $J_1(\mathbf{D}, \mathbf{s}_i) = \frac{M_i}{N} \|\mathbf{m}_i \odot (\mathbf{x}_i - \mathbf{D}\mathbf{s}_i)\|^2$, where \odot stands for the entry-wise product, $\mathbf{m}_i \in \mathbb{R}^N$ is the observation mask for sample i , *i.e.* $m_i(n) = 0$ (1) if data entry $\mathbf{x}_i(n)$ is missing (available); and $J_2(\mathbf{s}_i) = \frac{1}{N} \|\mathbf{s}_i\|_1$ is proportional to the ℓ_1 -norm whose minimization promotes the sparsity of the representation since ℓ_1 -norm is a convenient proxy for ℓ_0 -norm [3]. Finally, the hyper-parameters λ_1 and λ_2 allow us to give more or less importance to the representation accuracy and its sparsity, with respect to the classification error. Intuitively, minimizing equation (2) favors solutions that not only have sparse representations compatible with observed features, but also providing reconstructions that are best separated in the given classes.

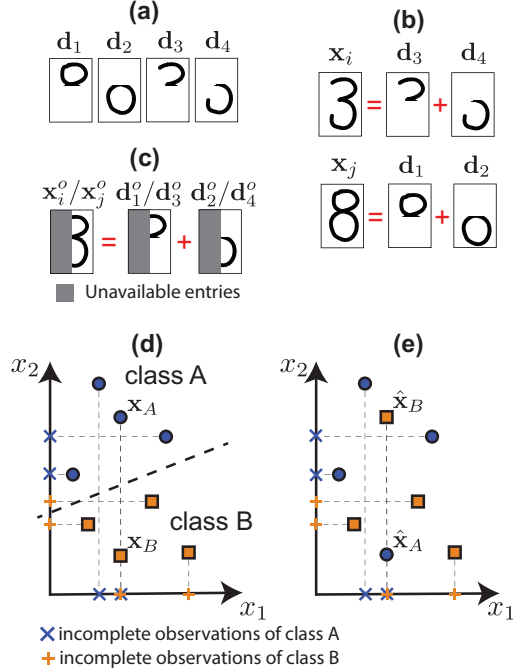


Figure 1. Toy example: (a) 4 out P dictionary elements d_i (atoms). (b) Digits “3” and “8” can be represented by combining only two atoms in the dictionary (2-sparse representations). (c) A left-half occluded digit “3” or “8” admits more than one 2-sparse representation (sum of d_1^o and d_2^o , or d_3^o and d_4^o). (d) Linearly separable samples from two classes (A and B) having two features: $\mathbf{x} = [x_1, x_2] \in \mathbb{R}^2$ where incomplete observations are taken by observing only one feature. Note that \mathbf{x}_A and \mathbf{x}_B belong to different classes but their observations are identical. (e) Without using label information, the sequential method could lead to wrong reconstructions of data vectors, *i.e.* $\hat{\mathbf{x}}_A \neq \mathbf{x}_A$ and $\hat{\mathbf{x}}_B \neq \mathbf{x}_B$ making the set of reconstructed vectors not linearly separable.

2.1. A sparsity-promoting sub-gradient optimization algorithm

To minimize the cost function in equation (2) we propose to alternate between the optimization over \mathbf{s}_i ($i = 1, 2, \dots, I$) and $\{\Theta, \mathbf{D}\}$ using the training dataset (incomplete).

For fixed $\{\Theta, \mathbf{D}\}$, the optimization with respect to \mathbf{s}_i is a non-smooth separable minimization sub-problem, which was extensively studied in the literature [38, 12]. In this sub-problem, the objective function is written as the sum of $F(\Theta, \mathbf{D}, \mathbf{s}_i)$ and a non-smooth separable function $G(\mathbf{s}_i)$, for which highly specialized, efficient and provable convergent solvers, namely the Coordinate Gradient Descent (CGD), already exists. However, the following key differences in our setting makes it not suitable for the CGD approach: first, our function $F(\Theta, \mathbf{D}, \mathbf{s}_i)$ involves evaluation of a multi-layer NN classifier, which can be non-smooth due to involved activation functions like ReLU or others; second, and more importantly, the computation of its second derivatives (Hessian) becomes prohibitive. Therefore, we choose a simpler and standard first order (stochastic sub-gradient based) search

of local minima with back-propagation. We take the strategy similar to the heuristics used in [34]. To update \mathbf{s}_i , we need to subtract $\sigma_s \frac{\partial J}{\partial \mathbf{s}_i}(j)$ from each coordinate j provided that we do not cross zero in the process in order to avoid escaping from a region where $G(\mathbf{s}_i)$ is differentiable. In such a case, we let the new value of $\mathbf{s}_i(j)$ be exactly zero. More specifically, we define $\Delta_i(j) = -\sigma_s \frac{\partial J}{\partial \mathbf{s}_i}(j)$ and, if $\mathbf{s}_i(j)[\mathbf{s}_i(j) + \Delta_i(j)] < 0$ (zero crossing condition), we re-define $\Delta_i(j) = -\mathbf{s}_i(j)$; finally we update $\mathbf{s}_i \leftarrow \mathbf{s}_i + \Delta_i$. It is noted that, once a coefficient $\mathbf{s}_i(j)$ reaches zero at a coordinate j , it becomes fixed, in other words, sparsity of solution \mathbf{s}_i is monotonically increasing with iterations.

When \mathbf{s}_i is fixed, our problem is reduced to minimize $F(\Theta, \mathbf{D}, \mathbf{s}_i)$ with respect to Θ and \mathbf{D} , which is easily done by standard first order (stochastic gradient based) search of local minima. The algorithm proposed for the training phase is presented as Algorithm 1.

In addition, for the testing phase, if the test dataset is incomplete, we need to find first the sparsest representation for the given observations, compute the reconstructions $\hat{\mathbf{x}}_i = \mathbf{D}\mathbf{s}_i$ and then apply the previously learned classifier to them as presented in Supp. material, Algorithm 2.

Algorithm 1 : Simultaneous classification and coding

Require: $\{\mathbf{x}_i^o, y_i\}$, $i = 1, 2, \dots, I$, hyper-parameters λ_1 and λ_2 , N_{iter} and update rates σ_Θ , $\sigma_{\mathbf{D}}$ and σ_s
Ensure: Weights Θ and reconstructions $\hat{\mathbf{x}}_i = \mathbf{D}\mathbf{s}_i, \forall i$

- 1: Randomly initialize $\Theta, \mathbf{D}, \mathbf{s}_i, \forall i$
- 2: **for** $n \leq N_{iter}$ **do**
- 3: Fix \mathbf{s}_i , update Θ and \mathbf{D} :
- 4: $\Theta = \Theta - \sigma_\Theta \frac{\partial J}{\partial \Theta}$
- 5: $\mathbf{D} = \mathbf{D} - \sigma_{\mathbf{D}} \frac{\partial J}{\partial \mathbf{D}}$
- 6: Normalize columns of matrix \mathbf{D}
- 7: Fix Θ and \mathbf{D} , update $\mathbf{s}_i, \forall i$:
- 8: $\Delta_i = -\sigma_s \frac{\partial J}{\partial \mathbf{s}_i}, \forall i$
- 9: **if** $\mathbf{s}_i(j)[\mathbf{s}_i(j) + \Delta_i(j)] < 0$ **then**
- 10: $\Delta_i(j) = -\mathbf{s}_i(j), \forall (i, j)$;
- 11: **end if**
- 12: $\mathbf{s}_i = \mathbf{s}_i + \Delta_i, \forall i$
- 13: **end for**
- 14: **return** $\Theta, \mathbf{D}, \mathbf{s}_i, \hat{\mathbf{x}}_i = \mathbf{D}\mathbf{s}_i, \forall i$

2.2. Theoretical analysis

Here, we investigate about conditions under which a perfect classifier of the complete data can be obtained from incomplete data samples.

2.2.1 Logistic regression

Let us first consider a logistic regression classifier [13] where the set of parameters $\Theta = \{\mathbf{w}, b\}$ are a vector $\mathbf{w} \in \mathbb{R}^N$ and a scalar b (bias). A perfect classifier exists if there is a hyper-plane that separates both classes, *i.e.*, for each data vector

\mathbf{x}_i : $f(\mathbf{x}_i) = \langle \mathbf{w}, \mathbf{x}_i \rangle + b > 0$ if $y_i = 1$, and $f(\mathbf{x}_i) \leq 0$ if $y_i = 0$. We consider data samples admitting a K -sparse representations $\mathbf{x}_i = \mathbf{D}\mathbf{s}_i$ with dictionary $\mathbf{D} \in \mathbb{R}^{N \times P}$ having unit-norm columns. We also assume an arbitrary pattern of missing features \mathcal{M}_i such that, data samples and dictionary are partitioned as $\{\mathbf{x}_i^m, \mathbf{x}_i^o\}$ and $\{\mathbf{D}_i^m, \mathbf{D}_i^o\}$, respectively. The following lemma identifies a sufficient condition under which, if we are able to train a classifier on incomplete observations such that the reconstructed data points are well separated by a hyperplane, then the same classifier correctly separates the original (unobserved) data vectors.

Lemma 2.1 (Sufficient condition type I). *Suppose that we have obtained an alternative dictionary $\mathbf{D}' \neq \mathbf{D} \in \mathbb{R}^{N \times P}$ such that, for the incomplete observations $\mathbf{x}_i^o \in \mathbb{R}^{M_i}$, the K -sparse representation solutions are non-unique, i.e. $\exists \mathbf{s}_i, \mathbf{s}'_i \in \Sigma_K^P$ such that $\mathbf{x}_i^o = \mathbf{D}'^o \mathbf{s}_i = \mathbf{D}^o \mathbf{s}'_i$, where $\mathbf{s}_i \in \mathbb{R}^P$ are the vectors of coefficients of the true data and \mathbf{s}'_i provides reconstructions $\hat{\mathbf{x}}_i = \mathbf{D}' \mathbf{s}'_i$. If a perfect classifier $\{\mathbf{w}, b\}$ of the reconstructions $\hat{\mathbf{x}}_i$ exists s.t. $|f(\hat{\mathbf{x}}_i)| > \epsilon_i > 0$ and*

$$\epsilon_i > |\langle \mathbf{w}_i^m, \mathbf{e}_i^m \rangle| \quad (4)$$

with $\mathbf{e}_i^m = \mathbf{x}_i^m - \hat{\mathbf{x}}_i^m$, then the full data vectors \mathbf{x}_i are also perfectly separated with this classifier, in other words: $f(\mathbf{x}_i) = \langle \mathbf{w}_i, \mathbf{x}_i \rangle + b > 0$ (≤ 0) if $y_i = 1$ ($y_i = 0$).

Proof. By using the missing/observed partition and omitting the sample index i , we can write: $f(\mathbf{x}) = \langle \mathbf{w}, \mathbf{x} \rangle + b = \langle \mathbf{w}^o, \mathbf{x}^o \rangle + \langle \mathbf{w}^m, \mathbf{x}^m \rangle + b$. If we add and subtract the term $\langle \mathbf{w}^m, \hat{\mathbf{x}}^m \rangle$ on the right left hand, arrange terms and use the fact that $\hat{\mathbf{x}}^o = \mathbf{D}'^o \mathbf{s}' = \mathbf{x}^o$, we get:

$$f(\mathbf{x}) = f(\hat{\mathbf{x}}) + \langle \mathbf{w}^m, \mathbf{e}^m \rangle. \quad (5)$$

Since we assumed that $f(\hat{\mathbf{x}}) > \epsilon > 0$ (for $y_i = 1$) and $|\langle \mathbf{w}^m, \mathbf{e}^m \rangle| < \epsilon$, it implies that $f(\mathbf{x}) > 0$. \square

Basically, condition (4) means requiring that reconstruction vector $\hat{\mathbf{x}}$ has a distance ϵ to the separating hyperplane larger than the absolute dot product between \mathbf{w} and the residual $\mathbf{e} = \mathbf{x} - \hat{\mathbf{x}}$, which of course is true when the reconstruction is accurate, i.e. $\mathbf{x} \approx \hat{\mathbf{x}}$. However, in practice, reconstructions are not accurate so we are interested in conditions under which Lemma 2.1 can still hold. Below, we derive a more restrictive but useful sufficient condition:

Proposition 2.1 (Sufficient condition type II). *Under the same hypothesis of Lemma 2.1, the following condition is enough to guarantee a proper classifier trained on incomplete data:*

$$\epsilon > |\langle \mathbf{w}^m, \mathbf{x}^m \rangle| + |\langle \mathbf{w}^m, \hat{\mathbf{x}}^m \rangle|. \quad (6)$$

Proof. By using the fact that $|\langle \mathbf{w}^m, \mathbf{e}^m \rangle| = |\langle \mathbf{w}^m, \mathbf{x}^m - \hat{\mathbf{x}}^m \rangle| \leq |\langle \mathbf{w}^m, \mathbf{x}^m \rangle| + |\langle \mathbf{w}^m, \hat{\mathbf{x}}^m \rangle|$, and applying Lemma 2.1 the proof is completed. \square

We highlight that, in our experiments, we were able to verify that Sufficient Condition type II is met in practice (see section 3, Fig. 3).

In Supp. material section 2, we derive an additional sufficient condition based on the Restricted Isometry Property (RIP) of the dictionary \mathbf{D} and sparsity level K , showing that sufficient condition (6) is easier to hold for datasets admitting highly sparse representations on dictionaries as close to orthogonal ones as possible.

2.2.2 Multilayer-perceptron

Lemma 2.1 can be straightforwardly generalized to multilayer-perceptron NNs where, if a softmax function is used at the output of the last layer then, as before, the prediction is based on the sign of the linear function:

$$f(\mathbf{x}) = \langle \mathbf{w}, \mathbf{x}^{(L)} \rangle + b, \text{ with } \mathbf{x}^{(l)} = h\left(\mathbf{W}_l^T \mathbf{x}^{(l-1)} + \mathbf{B}_l\right), \quad (7)$$

$l = 1, 2, \dots, L$, where $L + 1$ is the total number of layers, N_l is the number of neurons in layer l , $\mathbf{w} \in \mathbb{R}^{N_{L+1}}$ contains the weights in the last layer, $h(\cdot)$ is an activation function, e.g. ReLU, $\mathbf{W}_l \in \mathbb{R}^{N_{l-1} \times N_l}$ and $\mathbf{B}_l \in \mathbb{R}_l^N$ contain the weights and biases associated to neurons at layer l ; and $\mathbf{x}^{(0)} = \mathbf{x}$ is the input data vector. In this case, the first layer matrix $\mathbf{W}_1 \in \mathbb{R}^{N \times N_1}$ can be partitioned into submatrices $\mathbf{W}_{1i}^o \in \mathbb{R}^{M \times N_1}$ and $\mathbf{W}_{1i}^m \in \mathbb{R}^{(N-M) \times N_1}$ according to the observed and missing input features, respectively.

Proposition 2.2. *Under the same conditions of Lemma 2.1, if a NN-classifier $\{\mathbf{W}_l, \mathbf{B}_l (l = 1, 2, \dots, L), \mathbf{w}, b, h = \text{ReLU}\}$ of the reconstruction $\hat{\mathbf{x}}_i$ exists such that*

$$\epsilon_i > A \max_j |\langle \mathbf{W}_{1i}^m(:, j), \mathbf{e}_i^m \rangle|, \quad (8)$$

where $A = \|\mathbf{w}\| \prod_{l=2}^L \|\mathbf{W}_l\|_2$ and $\mathbf{e}_i^m = \mathbf{x}_i^m - \hat{\mathbf{x}}_i^m$, then the full data vector \mathbf{x}_i is also perfectly separated, in other words: $f(\mathbf{x}_i) > 0$ (< 0) if $y_i = 1$ ($y_i = 0$).

In the proof of Lemma 2.1, we were interested in finding a bound of the output error when the input \mathbf{x} of a classifier is perturbed, i.e. we found conditions such that $|f(\mathbf{x}) - f(\mathbf{x} + \delta)| < \epsilon$. By generalizing the classifier to the case of a multilayer perceptron we can derive the proof as follows:

Proof. Given a perturbation $\delta^{(l-1)} \in \mathbb{R}^{N_l}$ at the input of layer $l - 1$, i.e. $\hat{\mathbf{x}}^{(l-1)} = \mathbf{x}^{(l-1)} + \delta^{(l-1)}$, it is propagated to the output of layer l . By writing the error at the output we obtain:

$$\delta^{(l)} = h(\mathbf{W}_l^T \mathbf{x}^{(l-1)} + \mathbf{B}_l + \mathbf{W}_l^T \delta^{(l-1)}) - h(\mathbf{W}_l^T \mathbf{x}^{(l-1)} + \mathbf{B}_l), \quad (9)$$

and, by using the sub-additivity of ReLU function $h(\cdot)$, i.e. $h(a + b) \leq h(a) + h(b)$, we derive the following entry-wise inequality:

$$\delta^{(l)} \leq h\left(\mathbf{W}_l^T \delta^{(l-1)}\right), \quad (10)$$

and, by considering the property of ReLU activation function $\|h(\mathbf{x})\| \leq \|\mathbf{x}\|$, it turns out:

$$\|\delta^{(l)}\| \leq \|\mathbf{W}_l^T \delta^{(l-1)}\|, \quad (11)$$

Since the last layer of the NN is a linear classifier as in the case of Lemma 2.1, we can ask that $\langle \mathbf{w}, \delta^{(L)} \rangle < \epsilon$. Thus, by recursively using equation (11), we write

$$\langle \mathbf{w}, \delta^{(L)} \rangle \leq \|\mathbf{w}\| \|\delta^{(L)}\| \leq \|\mathbf{w}\| \|\mathbf{W}_L\| \|\mathbf{W}_{L-1}\| \dots \|\mathbf{W}_2\| \|\delta^{(1)}\|. \quad (12)$$

By defining $A = \|\mathbf{w}\| \prod_{l=2}^L \|\mathbf{W}_l\|$, evaluating equation (11) with $l = 1$ and taking into account that perturbation at the input of first layer is $\delta^{(0)} = \mathbf{e}$ with $\mathbf{e}^o = \mathbf{0}$, we arrive at:

$$\langle \mathbf{w}, \delta^{(L)} \rangle \leq A \|\mathbf{W}_1^{mT} \mathbf{e}^m\| \leq A \max_j |\langle \mathbf{W}_1^m(:, j), \mathbf{e}^m \rangle| < \epsilon, \quad (13)$$

which completes the proof. \square

It is interesting to note that $A = 1$ is attained when unit-norm filters (columns of \mathbf{W}_l) are orthogonal, which can be imposed by using orthogonality regularization [1].

3. Experimental results

We implemented all the algorithms in Pytorch 1.0.0 on a single GPU. Implementation details are reported in Supplemental material, sections 3.1 and 3.2. The code is available at ¹.

Synthetic datasets: We synthetically generated $I = 11,000$ (10,000 training + 1,000 test) K -sparse data vectors $\mathbf{x}_i \in \mathbb{R}^{100}$ using a dictionary $\mathbf{D} \in \mathbb{R}^{100 \times 200}$ obtained from a Gaussian distribution with normalized atoms, *i.e.* $\|\mathbf{D}(:, j)\| = 1, \forall j$. A random hyperplane $\{\mathbf{w}, b\}$ with $\mathbf{w} \in \mathbb{R}^N, b \in \mathbb{R}$ was randomly chosen dividing data vectors into two classes according to the sign of the expression $\langle \mathbf{w}, \mathbf{x}_i \rangle + b$, which defined the label y_i . We also controlled the degree of separation between classes by discarding all data vectors with distances to the hyperplane lower than a pre-specified threshold, *i.e.* $|\langle \mathbf{w}, \mathbf{x}_i \rangle + b| < d$. We used $n = 10$ repetitions of each experiment with different masks and input data in order to compute statistics.

We applied our simultaneous method (**Simult.**) with hyperparameters λ_1 and λ_2 in the cost function (2) tuned via cross-validation to train a logistic regression classifier on incomplete datasets with randomly distributed missing features. Then, we computed the classification accuracy on the complete test dataset and compared the results against the following standard sequential methods:

Sequential Sparsity based (Seq. Sp.): reconstructions are obtained by finding the sparsest representation compatible with the observations solving a LASSO problem. We used Algorithm 3 as shown in the Supp. material;

Zero Fill (ZF): missing features are filled with zeros, which

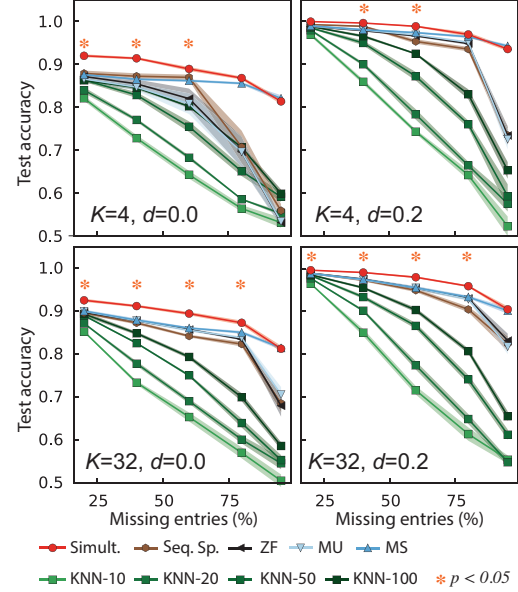


Figure 2. Experimental results on synthetic dataset with random masks using our algorithm (red) and compared to various sequential methods. Test accuracy (mean \pm s.e.m with $n = 10$) is shown as a function of the percentage of missing features for separation of classes $d = 0.0, 0.2$ and levels of sparsity $K = 4, 32$. Statistical significance for the difference between Simult. and MS is shown ($p < 0.05$).

is equivalent to ignore unknown values;

Mean Unsupervised (MU): missing features are filled with the mean computed on the available values;

Mean Supervised (MS): as in the previous case but the mean is computed on the same class vectors only;

K-Nearest Neighbor (KNN): as in the previous case but the mean is computed on the K -nearest neighbors of the same class vectors only.

To compare the performance of classifiers, we computed the mean accuracy \pm standard error of the mean (s.e.m.), with $n = 10$, on complete test datasets using all the methods for two levels of separation between classes ($d = 0.0, 0.2$), two levels of sparsity ($K = 4, 32$) and missing features in the training dataset ranging from 25% to 95% as shown in Fig. 2. Our results show that the simultaneous algorithm clearly outperforms all the sequential methods. A t-test was performed to evaluate the statistical significance with $p < 0.05$ of the difference between our algorithm and MS. It is interesting to note that, when classes has some degree of separation ($d = 0.2$), using the simple MS method, can give good results but not better than our algorithm.

In the second experiment, we generated $I = 10,000$ K -sparse data vectors $\mathbf{x}_i \in \mathbb{R}^{100}$ using $\mathbf{D} \in \mathbb{R}^{100 \times 100}$ and we evaluated the sufficient condition of equation (6) on $n = 10$ repetitions of the experiment with 95% missing features and separation $d = 0.0$. Fig. 3 clearly shows that the sufficient condition is mostly met in practice, especially for highly sparse representations of input data (small K).

¹<https://github.com/ccaiafa/SimultCodClass>




Figure 3. Verification of the sufficient condition (6) for various levels of sparsity K : 2D-histogram of ϵ versus $g = |\langle \mathbf{w}^m, \mathbf{x}^m \rangle| + |\langle \mathbf{w}^m, \hat{\mathbf{x}}^m \rangle|$. Mean + s.e.m ($n = 10$) percentage of correctly classified data samples are shown for $\epsilon > g$ and $\epsilon < g$.

This means that in practice it is not necessary to accurately reconstruct the input vectors, it is enough to capture the intrinsic characteristics of the classes such that the distances of reconstructions to the separating hyperplane satisfy the sufficient condition (3).

Benchmark datasets: We also considered three popular computer vision datasets: MNIST [22] and Fashion [41] consisting of 70,000 images (60,000 train + 10,000 test) each; and CIFAR10 [21] having 60,000 images (50,000 train + 10,000 test). MNIST/Fashion datasets contains 28×28 gray scale images while CIFAR10 dataset is built upon $32 \times 32 \times 3$ color images of different objects. The corresponding data sample size is $N = 28 \times 28 = 784$ for MNIST/Fashion and $N = 32 \times 32 \times 3 = 3,072$ for CIFAR10. We considered a dictionary of size 784×784 (MNIST/Fashion) and $1,024 \times 1,024$ (CIFAR10) and applied our simultaneous algorithm to learn the classifier on incomplete data using uniform random missing masks with several levels of missing data (25%, 50% and 75%) and 50% for random partial occlusions with MNIST/Fashion.

We used a logistic regression classifier (single layer NN) and a 4-layer convolutional neural network [23] (CNN4) for the MNIST/Fashion dataset using batch normalization (BN) [18] in the Fashion dataset. For CIFAR10 dataset, an 18-layer residual neural network, Resnet-18 [15] was implemented. We did not use any data augmentation strategy. The hyper-parameters λ_1 and λ_2 in cost function (2) were adjusted by cross-validation through a grid-search, as shown in Supp. material (Table S2 and Fig. S1). We compared our proposed algorithm with the following standard sequential methods: ZF, MS, KNN-10, KNN-20, KNN-50 and KNN-100; and against the recently proposed method from [35], referred here as NN-GMM, which uses the same NN classi-

fier as in our method and models missing features through GMM². We trained the classifiers on incomplete data with random masks and tested them on complete data for MNIST and CIFAR10 datasets. The obtained mean Test Accuracy \pm s.e.m ($n = 10$) are reported in Table 1. It is noted that NN-GMM provided good results with MNIST dataset compared to sequential methods, however, our simultaneous method outperformed all the methods. Interestingly, NN-GMM performed worst than any other method with CIFAR10 dataset. It seems that NN-GMM is not robust to large amount of missing data because, when we reduced the missing entries to 10%, the test accuracy sensibly increased to 52.57%. Additionally, our method showed to have little variability (small s.e.m) compared to the second best method (NN-GMM for MNIST and ZF for CIFAR10).

In Table 2, test accuracies obtained when the learned model is applied to incomplete and complete test data, are shown. The right-most column shows the baseline results obtained by training the model on complete datasets using a CNN4 [23] and a Resnet-18 [15], whose implementations can be found at ³ and ⁴. It is interesting to note that for the logistic regression classifier, we obtained better results when training with incomplete data rather than using complete data. Also, it is highlighted that training on incomplete data with 50% or fewer random missing features, provides similar test accuracy as training on complete data for MNIST dataset. This could be explained by noting two facts: (1) random missing features is similar to applying dropout, with the exception that missing data do not change during training; and (2) our model has more parameters (Dictionary + sparse coefficients + Linear layer) compared to the baseline logistic regression classifier. To provide a deeper understanding of this effect, we ran the baseline with Dropout at the input and we obtained: 91.95%, 91.97% and 92.01% for $p = 0.0, 0.1$ and 0.25, respectively, which shows that the improvement we obtained with our method is not solely caused by a dropout alike behavior.

In Fig. 4, we present some randomly selected visual examples comparing the original images in the MNIST/Fashion test dataset, their observations using random masks and partial occlusions, and the reconstructions using the dictionary learned from the incomplete training data. It is clear that, despite the reconstructions may be not very similar to the original images (see “5” digit example), they clearly own the properties of the class to which they belong to. Additional examples are provided in Supp. material (Figs. S2 and S3).

4. Discussion

It is well known that sparse coding has the ability to accurately model complex distributions of data, such as natural

²<https://github.com/lstruski/Processing-of-missing-data-by-neural-networks>

³<https://github.com/pytorch/examples/tree/master/mnist>

⁴<https://github.com/kuangliu/pytorch-cifar>

Table 1. Test accuracy (mean \pm s.e.m with $n = 10$) of various methods trained on incomplete data and tested on complete ones for MNIST and CIFAR10.

MNIST (CNN4)								
Miss.	ZF	MS	KNN10	KNN20	KNN50	KNN100	NN-GMM	Simult.
75%	84.86 \pm 0.02	83.79 \pm 0.01	88.16 \pm 0.01	87.94 \pm 0.01	87.03 \pm 0.002	86.52 \pm 0.01	96.36 \pm 0.12	98.09 \pm 0.04
50%	90.13 \pm 0.06	88.55 \pm 0.01	91.36 \pm 0.02	91.11 \pm 0.02	90.87 \pm 0.01	90.82 \pm 0.01	97.57 \pm 0.37	98.23 \pm 0.10
CIFAR10 (Resnet18)								
Miss.	ZF	MS	KNN10	KNN20	KNN50	KNN100	NN-GMM	Simult.
75%	32.22 \pm 2.09	21.30 \pm 0.40	22.84 \pm 0.87	25.67 \pm 0.80	26.52 \pm 0.70	26.01 \pm 0.52	12.10 \pm 0.61	54.81 \pm 0.47
50%	46.37 \pm 1.93	17.90 \pm 0.94	30.94 \pm 0.54	29.68 \pm 0.46	30.01 \pm 0.51	26.23 \pm 1.01	14.02 \pm 0.75	62.50 \pm 0.95

Table 2. Test Accuracies obtained with our method on MNIST, Fashion and CIFAR10 datasets training with incomplete data and testing on incomplete/complete data. Baseline results obtained by training the models on complete data are shown for reference in the right-most column.

Dataset	Classifier	Random missing features						Occlusion		Baseline
		%Train / %Test						% Train / %Test		% Train / %Test
		75/75	50/50	25/25	75/0	50/0	25/0	50/50	50/0	0/0
MNIST	Log. Reg.	90.45	93.68	94.14	91.94	93.44	94.43	-	-	91.95
	CNN4	94.62	98.34	98.94	98.09	98.23	98.95	88.55	91.37	98.95
Fashion	CNN4+BN	83.71	86.09	86.38	86.39	87.11	87.04	81.73	82.47	90.76
CIFAR10	Resnet18	53.82	61.08	63.73	54.81	62.50	63.87	-	-	80.13

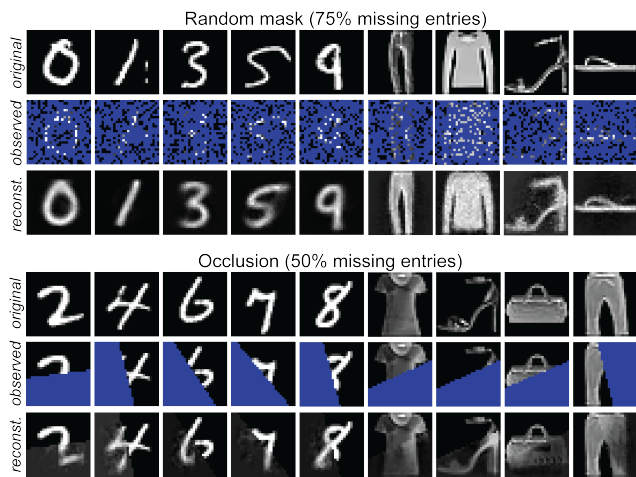


Figure 4. Original (top), observed (middle) and reconstructed (bottom) MNIST and Fashion test images.

signals (images, audio, EEG, etc). In this work, we demonstrated that assuming a sparse representation for input data allows for the successful training of a general NN when incomplete data is given outperforming traditional sequential approaches and other start-of-the-art methods. It is highlighted that our method can be used with potentially any deep NN architecture, thus relying on their extraordinary capability to accommodate complex decision boundaries as usually needed in modern machine learning.

Our method overcomes well known issues of previous approaches: (1) compared to imputation methods, our algorithm successfully incorporates the labelling information into the modeling of missing features; (2) sparse coding allows for a simple way to train dictionaries through linear methods such as stochastic gradient descent with back-propagation compared to the very expensive EM estimators for GMM used in probabilistic generative models, or SVD based algorithms for matrix rank minimization in matrix completion; (3) sparse coding can be more accurate modeling missing values in natural signals compared to GMM, especially for high

dimensional data where GMM may require a huge number of parameters making it computationally prohibitive.

We analyzed the limitations of the classical imputation approach and demonstrated through experiments with synthetic and real-world datasets that our simultaneous algorithm always outperforms them for various cases such as LASSO, zero-filling, supervised/unsupervised mean and KNN based methods as well as the *state-of-the-art* method based on NNs and GGM recently proposed in [35]. Nevertheless, our experimental results on synthetic and real-world dataset showed that, even though we only constrained dictionaries to have unit-norm columns but not enforcing any other kind of constraint like maximum coherence, the obtained results seem to be satisfactory enough. However, further analysis on the required properties of dictionaries could provide deeper insights and alternative ways to improve the algorithm, which we aim to address in a future work.

While current simple sub-gradient based optimization approach provided satisfactory results in terms of performance, it is remarked that observed convergence is slow requiring a thousand of iterations sometimes. We believe, it could be improved by trying to incorporate some second-order derivatives information for computing the updates. Although, full Hessian computation becomes prohibitive with multi-layer NNs a diagonal approximation approach could be explored. Also, a rigorous convergence analysis in the line of the analysis in [38, 12] and taking special properties of multi-layer NN classifier functions can be conducted in a future work.

Finally, we provided theoretical insights of the problem by providing sufficient conditions under which, if it is possible to train a classifier on incomplete observations so that its reconstructions are well separated by a hyperplane, then the same classifier also correctly separates the original (un-observed) data samples.

Acknowledgments: We are thankful for the RAIDEN computing system and its support team at RIKEN AIP, Tokyo. This work was supported by the JSPS KAKENHI (Grant No. 20H04249, 20H04208).

References

- [1] Nitin Bansal, Xiaohan Chen, and Zhangyang Wang. Can We Gain More from Orthogonality Regularizations in Training Deep CNNs? In *NeurIPS*, 2018.
- [2] Chiranjib Bhattacharyya, Pannagadatta K Shivaswamy, and Alexander J Smola. A Second Order Cone programming Formulation for Classifying Missing Data. In *NIPS*, 2004.
- [3] Emmanuel J. Candès and Terrence Tao. Decoding by Linear Programming. *IEEE Transactions on Information Theory*, 51(12):4203–4215, 2005.
- [4] Gal Chechik, Jeremy Heitz, Gal Elidan, Pieter Abbeel, and Daphne Koller. Max-margin Classification of Data with Absent Features. *Journal of Machine Learning Research (JMLR)*, 9:1–21, 2008.
- [5] M Davenport, Marco F Duarte, Yonina C Eldar, and G Kutyniok. *Introduction to compressed sensing*. Theory and Applications. Cambridge University Press, Cambridge, 2012.
- [6] Olivier Delalleau, Aaron C Courville, and Yoshua Bengio. Efficient EM Training of Gaussian Mixtures with Missing Data. *arXiv*, cs.LG:1209.0521, 2012.
- [7] Uwe Dick, Peter Haider, and Tobias Scheffer. Learning from incomplete data with infinite imputations. In *ICML*, 2008.
- [8] Pedro J García-Laencina, José-Luis Sancho-Gómez, and Aníbal R Figueiras-Vidal. Pattern classification with missing data: a review. *Neural Computing and Applications*, 19(2):263–282, 2009.
- [9] Zoubin Ghahramani and Michael I Jordan. Supervised learning from incomplete data via an EM approach. In *NIPS*, 1993.
- [10] Andrew B Goldberg, Xiaojin Zhu, Ben Recht, Jun-Ming Xu, and Robert D Nowak. Transduction with Matrix Completion - Three Birds with One Stone. In *NIPS*, 2010.
- [11] Karol Gregor and Yann LeCun. Learning Fast Approximations of Sparse Coding. In *ICML*, 2010.
- [12] Elaine T Hale, Wotao Yin, and Yin Zhang. Fixed-Point Continuation for ℓ_1 -Minimization: Methodology and Convergence. *SIAM J. OPTIM.*, 19(3):1107–1130, 2008.
- [13] Trevor Hastie, Robert Tibshirani, and Jerome H Friedman. *The elements of statistical learning - data mining, inference, and prediction, 2nd Edition*. Springer Series in Statistics. Springer New York, New York, NY, 2009.
- [14] Elad Hazan, Roi Livni, and Yishay Mansour. Classification with Low Rank and Missing Data. In *ICML*, 2015.
- [15] Kaifeng He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778. IEEE, 2016.
- [16] Ke Huang and Selin Aiyente. Sparse Representation for Signal Classification. In *NIPS*, 2006.
- [17] Sheng-Jun Huang, Miao Xu, Ming-Kun Xie, Masashi Sugiyama, Gang Niu, and Songcan Chen. Active Feature Acquisition with Supervised Matrix Completion. *arXiv*, cs.LG:1802.05380, 2018.
- [18] Sergey Ioffe and Christian Szegedy. Batch Normalization - Accelerating Deep Network Training by Reducing Internal Covariate Shift. In *ICML*, 2015.
- [19] Rodolphe Jenatton, Julien Mairal, Guillaume Obozinski, and Francis R Bach. Proximal Methods for Sparse Hierarchical Dictionary Learning. In *ICML*, 2010.
- [20] Julie Josse, Nicolas Prost, Erwan Scornet, and Gaël Varoquaux. On the consistency of supervised learning with missing values. *arXiv.org*, page arXiv:1902.06931, Feb. 2019.
- [21] A Krizhevsky. *Learning multiple layers of features from tiny images*. PhD thesis, Toronto University, Toronto, 2009.
- [22] Yann LeCun, Bernhard E Boser, John S Denker, Donnie Henderson, Richard E Howard, Wayne E Hubbard, and Lawrence D Jackel. Handwritten Digit Recognition with a Back-Propagation Network. In *NIPS*, 1989.
- [23] Yann LeCun, Patrick Haffner, Léon Bottou, and Yoshua Bengio. Object Recognition with Gradient-Based Learning. In *Shape, Contour and Grouping in Computer Vision*, 1999.
- [24] Honglak Lee, Alexis Battle, Rajat Raina, and Andrew Y Ng. Efficient sparse coding algorithms. In *NIPS*, 2006.
- [25] Xuejun Liao, Hui Li, and Lawrence Carin. Quadratically gated mixture of experts for incomplete data classification. In *ICML*, 2007.
- [26] Tsung I Lin, Jack C Lee, and Hsiu J Ho. On fast supervised learning for normal mixture models with missing information. *Pattern Recognition*, 39(6):1177–1187, 2006.
- [27] Roderick J A Little and Donald B Rubin. *Statistical Analysis with Missing Data*. John Wiley & Sons, Aug. 2014.
- [28] Julien Mairal, Francis R Bach, Jean Ponce, and Guillermo Sapiro. Online dictionary learning for sparse coding. In *ICML*, 2009.
- [29] J Mairal, J Ponce, G Sapiro, A Zisserman Advances in neural, and 2009. Supervised dictionary learning. In *NIPS*, 2008.
- [30] SG Mallat. *A Wavelet Tour of Signal Processing*. The Sparse Way. Academic Press, 2009.
- [31] BA Olshausen and DJ Field. Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vision research*, 37(23):3311–3325, 1997.
- [32] B A Olshausen and D J Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607–609, 1996.
- [33] Ignacio Ramirez, Pablo Sprechmann, and Guillermo Sapiro. Classification and clustering via dictionary learning with structured incoherence and shared features. In *CVPR*, 2010.
- [34] Shai Shalev-Shwartz and Ambuj Tewari. Stochastic Methods for ℓ_1 -regularized Loss Minimization. *Journal of Machine Learning Research (JMLR)*, 2011.
- [35] Marek Smieja, Lukasz Struski, Jacek Tabor, Bartosz Zielinski, and Przemyslaw Spurek. Processing of missing data by neural networks. In *NeurIPS*, 2018.
- [36] Pablo Sprechmann and Guillermo Sapiro. Dictionary learning and sparse coding for unsupervised clustering. In *ICASSP*, 2010.
- [37] Ivana Tosic and Pascal Frossard. Dictionary Learning. *IEEE Signal Processing Magazine*, 28(2):27–38, 2011.
- [38] Paul Tseng and Sangwoon Yun. A coordinate gradient descent method for nonsmooth separable minimization. *Mathematical Programming*, 117(1-2):387–423, 2007.
- [39] Jonathan Weed. Approximately Certifying the Restricted Isometry Property is Hard. *IEEE Transactions on Information Theory*, 64(8):5488–5497, 2017.

- [40] David Williams, Xuejun Liao, Ya Xue, and Lawrence Carin. Incomplete-data classification using logistic regression. *ICML*, 2005.
- [41] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-MNIST: a Novel Image Dataset for Benchmarking Machine Learning Algorithms. *arXiv*, cs.LG:arXiv:1708.07747, 2017.