

# An analysis on the use of autoencoders for representation learning: Fundamentals, learning task case studies, explainability and challenges

David Charte<sup>a,\*</sup>, Francisco Charte<sup>b</sup>, María J. del Jesus<sup>b</sup>, Francisco Herrera<sup>a</sup>

<sup>a</sup> Department of Computer Science and A.I., University of Granada, 18071 Granada, Spain

<sup>b</sup> Department of Computer Science, University of Jaén, Jaén 23071, Spain

## ARTICLE INFO

### Article history:

Received 31 January 2020

Revised 30 March 2020

Accepted 8 April 2020

Available online 19 May 2020

Communicated by Dr. Nianyin Zeng

### MSC:

68T05

68T10

### Keywords:

Representation learning

Autoencoders

Deep learning

Feature extraction

## ABSTRACT

In many machine learning tasks, learning a good representation of the data can be the key to building a well-performant solution. This is because most learning algorithms operate with the features in order to find models for the data. For instance, classification performance can improve if the data is mapped to a space where classes are easily separated, and regression can be facilitated by finding a manifold of data in the feature space. As a general rule, features are transformed by means of statistical methods such as principal component analysis, or manifold learning techniques such as Isomap or locally linear embedding. From a plethora of representation learning methods, one of the most versatile tools is the autoencoder. In this paper we aim to demonstrate how to influence its learned representations to achieve the desired learning behavior. To this end, we present a series of learning tasks: data embedding for visualization, image denoising, semantic hashing, detection of abnormal behaviors and instance generation. We model them from the representation learning perspective, following the state of the art methodologies in each field. A solution is proposed for each task employing autoencoders as the only learning method. The theoretical developments are put into practice using a selection of datasets for the different problems and implementing each solution, followed by a discussion of the results in each case study and a brief explanation of other six learning applications. We also explore the current challenges and approaches to explainability in the context of autoencoders. All of this helps conclude that, thanks to alterations in their structure as well as their objective function, autoencoders may be the core of a possible solution to many problems which can be modeled as a transformation of the feature space.

© 2020 Elsevier B.V. All rights reserved.

## 1. Introduction

Creating new representations of data is a fundamental task in most machine learning tasks. First off, certain types of problems that require a classifier or a regressor will certainly benefit from transformations of the features which facilitate their work [1]. In addition to this, there exists a variety of problems whose solution relies strongly on finding an appropriate representation of the data. Although the use of representation learning techniques is mainly used as a complement to other learners in the former case, in the latter one these methods become the focus. This work highlights some of these situations, with specific applications that can be modeled as representation learning problems.

The features that are used as input conform one of the most important factors when building machine learning models. When

the training set contains intact data from its collection or measurements, it may not be ready for treatment yet. Instead, it is common for data to be expressed with redundant or uninformative variables and for it to include some level of noise. These and other obstacles presented by the data [2] are the reason why most of the manual work of building machine learning models is spent in the preprocessing stage [3].

The success of a classifier, a regressor or other models will greatly depend on the quality of the features it can learn from. For instance, decision trees, regardless of whether the task is classification or regression, attempt to find the most informative variables to branch at each step [4]; support vector machines calculate the hyperplane that best separates classes in a feature space originating from specific transformations of the original one [4], and k-means clustering computes distances among pairs of instances and thus depends strongly on the input domain [5]. As a result, it is of vital importance that the features provided to these learners are useful and as independent as possible.

\* Corresponding author.

E-mail addresses: [fdavidcl@ugr.es](mailto:fdavidcl@ugr.es) (D. Charte), [fcharte@ujaen.es](mailto:fcharte@ujaen.es) (F. Charte), [mjjesus@ujaen.es](mailto:mjjesus@ujaen.es) (M.J. del Jesus), [herrera@decsai.ugr.es](mailto:herrera@decsai.ugr.es) (F. Herrera).

However, finding alternative representations for data is not only a medium to build classification and regression models, but it may be an end in itself in many applications. For example, finding compact binary codes that represent text documents [6], compressing signals to a lower resolution without losing information [7], transforming the problem domain to a different one [8], or producing filtered versions of images with less distortions [9].

Learning representations usually consists in feature engineering [1] or feature extraction [10], depending on whether new features are computed manually by human intervention (either by selection [11] or simple arithmetic operations) or they are generated, evaluated and selected by the machine. Feature engineering leverages expert knowledge and human creativity in order to select features and operate with them in a way that results in a new feature set which seems appropriate for predictors to work with. Nowadays there exist many automatic approaches to feature learning, which relieve users from the tedious task of engineering new features [12]. These methods range from probabilistic to topological and from shallow to deep: principal component analysis [13], Isomap [14], locally linear embedding [15] and Laplacian eigenmaps [16], among others.

With the introduction of deep neural networks, the representation learning stage became integrated within the predictors themselves [17]. These techniques iteratively optimize the classification performance by modifying the weights in several layers of individual neurons which compute a hierarchy of abstractions over the original data. For this purpose, the backpropagation algorithm [18] allows to efficiently accumulate gradients along the network, so that an optimizer such as Stochastic Gradient Descent [19] or one of its derivatives [20–23] may compute each weight update. Since neural networks can be structured as needed for each kind of problem, they are able to function as standalone feature learners as well. This is the case of autoencoders (AEs) [24], neural architectures whose objective is to find the best representation for the data according to the criterium defined by their loss function.

The objective of this paper is to analyze how AEs can serve as the main basis for solving a wide variety of learning tasks and demonstrate this with concrete applications and experimental results. Throughout the paper, we examine several case studies that expose the adaptability of AEs to these problems.

- First, an example of data embedding onto a very low dimensional space for visualization and exploratory analysis.
- Then, a case where noisy signals are to be repaired by the model.
- Later, a different example where very high dimensional sparse data, such as text documents, is to be compressed onto compact binary codes in a semantic way.
- Additionally, we study anomaly detection, the situation where abnormal patterns are to be detected in sequences but no anomalies are available to learn from.
- As a last case study, we propose the generation of new instances which do not belong to the training set.

Other applications are also briefly discussed: image superresolution, image compression, transfer learning, human pose recovery and recommender systems.

As a starting point, we provide the reader with the necessary background knowledge about the field of representation learning, as well as a summary of the main features of AEs that make them a good candidate model to solve the different problems later approached. The solutions to these tasks using AEs as the only automatic learner highlight their potential and flexibility as feature extraction techniques.

Following the current increase in search for developing explainable models [25], the main approaches for obtaining interpretable

predictions are summarized, finding that quality features can be the key to explainable solutions. AE models which can build helpful features are also highlighted.

The rest of this paper is structured as follows. Section 2 describes the background of the problems and techniques above introduced. Section 3 details the inner workings of AEs. Section 4 further develops on several case studies where AEs resolve feature learning tasks and outlines other existing learning applications, and Section 5 describes the current state of the art in explainable AI and how AEs are involved. Lastly, Section 6 concludes the text.

## 2. Background: feature learning and deep representation learning

This section explains some well-known methods that can extract features from data. Afterwards, it introduces deep learning techniques.

### 2.1. Classical feature learning methods

Traditionally, feature extraction methods have been developed with linear as well as nonlinear transformations of the variables [10]. They can be considered nonconvex or convex, according to whether the objective function presents local optima or not, respectively [26]. Many of these techniques perform unsupervised learning, but others are supervised [27–29] or even semi-supervised [30]. Next, a summary of typical feature learning methods is provided.

### 2.2. Linear methods

The most common linear feature extraction methods are the following. Principal component analysis (PCA) consists in extracting successive variables or *principal components* with maximum variance while being uncorrelated with the previous components. It is a statistical technique developed geometrically by Pearson [31] and algebraically by Hotelling [32], but an analytical derivation can be found in [13]. Factor analysis [33] is a similar procedure to PCA which considers a set of latent variables or *factors* that are not observed but are linearly combined to produce the final variables. Linear discriminant analysis [27] is a supervised statistical technique which attempts to find linear combinations of features to project samples onto new coordinates that best discriminate classes, albeit making some assumptions about the distribution of the data.

### 2.3. Nonlinear methods

Some well known nonlinear approaches to feature extraction are kernel PCA, restricted Boltzmann machines and manifold learning methods. Kernel PCA [34] extends PCA to nonlinear combinations of features by projecting samples onto higher-dimensional spaces and using the kernel trick [35]. Restricted Boltzmann machines are undirected graphical probabilistic models, also known as harmoniums [36], with one visible layer and one hidden layer that acts as the set of extracted features. They can be trained using the contrastive divergence algorithm [37]. Many nonlinear feature learning methods attempt to find coordinates for a lower dimensional structure embedded in the original features, namely, a manifold. Multidimensional scaling (MDS) is one of the first techniques that can be considered manifold learning, as its objective is projecting samples in a low-dimensional space while translating as much information of pairwise distances as possible. There are several variants of MDS, one of them is Sammon mapping [38], which improves on MDS by using a different cost function which

stresses large distances similarly to small ones. Isomap [14] is a more recent extension of MDS which looks for the coordinates that describe the actual degrees of freedom of the data while preserving distances among neighbors and geodesic distances (the length of the shortest path that connects two points in the manifold). Locally Linear Embedding [15] also seeks a manifold which preserves neighbors but, in order to maintain the local structure, it linearly reconstructs each point from its neighbors. Laplacian eigenmaps [16] is a procedure that builds a graph based on the neighborhood structure of the data, and from it a weight matrix whose eigenvectors can be used to compute new coordinates for each point.

#### 2.4. Deep representation learning

Deep learning architectures are hierarchies of abstractions of the input feature space and, as such, they compute several transformations of the features before reaching a response. In some cases, these can be seen as learned representations, since they must be able to capture the relevant information from each instance in order to output an accurate result. This effect can be observed especially in convolutional neural network classifiers, which are usually split into a feature extraction component formed by convolutional layers and a decision module composed by fully connected layers [39]. Apart from neural networks with other objectives such as supervised classification or regression, there have been different approaches to shallow as well as deep neural structures for unsupervised feature learning [40], such as self-organizing Kohonen maps [41,42], predictability minimization [43], restricted Boltzmann machines [44], deep belief networks [45–47] and AEs [48,49]. There have been many instances of these unsupervised techniques being used to either pre-train or provide feature transformations for supervised models [50].

AEs are probably the most versatile unsupervised neural network models. They essentially combine some kind of bottleneck or restriction in the learned data representations with the objective of reconstructing and repairing the original input from that representation [24]. There are several ways to impose restrictions that produce interesting representations, and the reconstruction objective will cause the network to retain all invariant feature information along its weights, so that the representation or encoding holds mainly instance-specific traits. For example, undercomplete AEs project inputs into lower-dimensional encodings, sparse AEs obtain representations with very few activated neurons, and denoising AEs attempt to repair partially corrupted data.


Their versatility is demonstrated by the amount of applications AEs have and their diversity. Across the rest of this work, we focus on certain applications of representation learning that are solved with AEs and we analyze how each model is built and trained.

### 3. Autoencoder fundamentals

AEs are neural network structures designed with the purpose of learning new features. Throughout the following subsections, their main characteristics and differentiating aspects are outlined, and some ways to influence the encoded variables are discussed.

#### 3.1. Origin and essentials of autoencoders

AEs were originally conceived as a way of initializing neural networks [51] and continued fulfilling that purpose for some time, serving as a starting point for training of deep networks as well [52]. Over the last years, other applications for AEs have been emerging and at the same time other approaches to neural network training and regularization have succeeded over AEs [53,54]. As a consequence, the common uses for AEs have shifted from



**Fig. 1.** Illustration of the general structure of a basic AE: an encoder and a decoder connected by the encoding layer.

**Table 1**  
Interpretation of symbols used in the formulae.

Symbol	Interpretation
$\theta$	Full set of parameters of the AE (weights and biases)
$\mathcal{X}$	Set of input instances
$\mathcal{Z}$	Set of instances in encoding space
$n$	Dimension of input space
$k$	Dimension of encoding space
$f$	Encoder mapping
$g$	Decoder mapping
$d$	Distance function in input space
$r$	Regularization function

helping train other neural networks to other applications of their own.

In general, the training process required to learn an AE can be unsupervised, that is, it does not need labels or class information in order to generate a model for the data. Instead, it extracts useful information from each instance by feeding its feature vector through some transformations which impose a bottleneck or restriction on the possible representations it can compute. Then, the representation is mapped to the original feature space through a similar set of transformations, and the AE is evaluated according to the fidelity of the reconstruction. This feedback allows to modify the parameters iteratively until convergence is reached.

AEs take the form of a neural network with at least one hidden layer and two components, an encoder and a decoder, which are connected by the coding layer [24]. These components are usually symmetric in layer shapes to each other, especially if they are implemented as fully connected neural networks. In certain occasions, even the weights of each layer in the decoder are tied to the corresponding layer in the encoder. In general terms, however, it suffices with the input layer of the encoding and the output layer having the same shape. Fig. 1 shows how the architecture of an AE may look like.

In summary, an AE can be seen as the composition of an encoding map  $f$  which projects inputs onto a different feature space, and a decoding map  $g$  which operates inversely (see Table 1 for the meaning of all symbols used below). The main objective of the AE is to recover as much information as possible of the original input, so it will attempt to minimize a distance between the inputs and the outputs:

$$\min_{\theta} \sum_{x \in \mathcal{X}} d(x, g_{\theta}(f_{\theta}(x))) \quad (1)$$

The distance function  $d$  used in the loss function is usually either the mean squared error, see Eq. (2), or the cross entropy, shown in Eq. (3). In the first case, data may not be normalized and the output units should use an unbounded activation function.

For a cross entropy loss, each input and output variable is modeled as following a Bernoulli distribution, so data should be scaled to the  $[0,1]$  interval and output units could make use of a sigmoid activation.

The mean squared error for an input  $x$  and output  $x'$  of length  $n$  is defined as:

$$d(x, x') = \frac{1}{n} \sum_{i=1}^n (x_i - x'_i)^2 \quad (2)$$

Similarly, the binary cross entropy for the same input and output is computed as:

$$d(x, x') = -(x \bullet \log(x') + (1 - x) \bullet \log(1 - x')), \quad (3)$$

where  $\bullet$  denotes element-wise product and all other operations are also performed element-wise.

### 3.2. Modeling the coding layer

The main objective of the AE (Eq. (1)) only promotes faithful reconstructions without explicitly considering any aspect about the codes used. This can be enough in many cases where the codes are low dimensional and they can capture only the relevant information of the instances just by training to reconstruct accurately. Notwithstanding, there are situations that require considering a more general case of the objective, which allows penalizing certain behaviors of the encoding found by the network, or even the values of the parameters themselves (Eq. (4)).

$$\min_{\theta} \sum_{x \in \mathcal{X}} d(x, g_{\theta}(f_{\theta}(x))) + r_1(f_{\theta}(\mathcal{X})) + r_2(\theta) \quad (4)$$

A straightforward example of this kind of restrictions is the sparse AE [55,56], which adds a penalty for high activation rates in the neurons of the code layer (Eqs. (5) and (6)):

$$r(\mathcal{Z}) = \sum_{j=1}^k (\rho - \rho_j)^2, \text{ or} \quad (5)$$

$$r(\mathcal{Z}) = \sum_{j=1}^k \rho \log \frac{\rho}{\rho_j} + (1 - \rho) \log \frac{1 - \rho}{1 - \rho_j}, \quad (6)$$

where  $\rho_j = \frac{1}{|\mathcal{Z}|} \sum_{z \in \mathcal{Z}} z_j$  is the average activation vector,  $k$  is the length of the code and  $\rho$  is the desired activation rate.

Other, more sophisticated variations on the AE with different penalties are the contractive AE [57,58], which promotes finding and preserving any local structure from the original feature space, and the variational AE [59], which uses a penalty to impose a distribution to the codes computed by the encoder.

Penalties on the codes are not, however, the only way of incentivizing a behavior on the encoder mapping. Denoising AEs [60,61] establish a slightly different criterion to evaluate the performance of the reconstruction: the network must be able to repair any noise or corruption from the input. Robust AEs [62] use another objective function, correntropy [63], which has a similar effect in repairing several kinds of noise from the input data.

### 3.3. Evaluation metrics

The quality of learned features can be evaluated by the model's ability to project instances back to the original feature space. For this purpose, regression metrics can be used. Some common metrics which serve to assess the usefulness of the learned features are the following, where  $x$  is the original feature vector and  $x'$  is the reconstruction, mapped from the encoding space back onto the input space:

- Mean squared error (Eq. (2)) and root mean squared error:

$$\text{RMSE}(x, x') = \sqrt{\frac{1}{n} \sum_{i=1}^n (x_i - x'_i)^2}$$

- Mean absolute error:

$$\text{MAE}(x, x') = \frac{1}{n} \sum_{i=1}^n |x_i - x'_i|$$

- Mean absolute percentage error

$$\text{MAPE}(x, x') = \frac{1}{n} \sum_{i=1}^n \left| \frac{x_i - x'_i}{x_i} \right|$$

In certain cases, the encoded features can also be evaluated independently from the original features, by assessing their quality with respect to their complexity, class separability and overlap [2]. This usually requires that data belongs to a classification problem so that a class is defined for each instance.

### 3.4. Beyond unsupervised autoencoders

Although the objective of an AE usually does not involve direct prediction of labels, it can sometimes learn from classified examples. The most straightforward way to introduce class information into the AE is to modify the loss function so it propagates different errors according to the class of each instance. For example, we could weight each class differently. Assuming the classes are binary, dividing the dataset into  $\mathcal{X}^+$  for positive instances and  $\mathcal{X}^-$  for negative ones, and  $\alpha$  is a parameter in  $[0,1]$ , the objective in Eq. (7)

$$\min_{\theta} (1 - \alpha) \sum_{x \in \mathcal{X}^-} d(x, g_{\theta}(f_{\theta}(x))) + \alpha \sum_{x \in \mathcal{X}^+} d(x, g_{\theta}(f_{\theta}(x))) \quad (7)$$

would give more importance to reconstructing one of the classes, which may help if the aim is to find a manifold for that class and the other one is less relevant.

Several uses of label information can be found in the proposal of the adversarial AE [29]. This AE has a similar behavior to the variational AE in that it also forces the codes to follow a given distribution. Instead of using just a loss penalty, it adds a generator which samples the distribution, and a discriminator which attempts to distinguish distribution samples from codes belonging to actual instances, analogous to a generative adversarial network [64]. The label information can be used then to locate each label in a region of the distribution, by feeding labels as well as codes to the discriminator. Alternatively, labels can be feeded to the decoder, which causes the codes to discard label information and instead model style in the data.

Another step forward in introducing label information in AEs would be for them to be able to predict labels as well. Some work has been already done along these lines, by training an encoder and decoder simultaneously to reconstruct and to produce codes as similar as possible to the labels in a one-hot format [65].

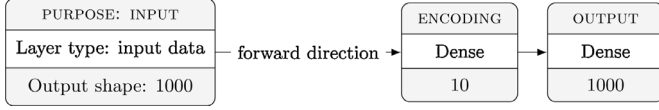
## 4. Learning task case studies

The following subsections detail several real examples of application of AEs: embedding data onto a very low-dimensional space for visualization purposes, reducing the noise in images, computing semantic hashes for large text documents, finding anomalous behaviors in sequences and generating new instances outside the training set. For each application, a relevant dataset has been selected and a model has been specifically designed to solve the problem. The basic traits of all chosen datasets can be found in Table 2.



**Table 2**  
Main traits of datasets used for the experiments.

Dataset	Application	Input features	Training examples	Test examples
CPU Activity	Visualization	21	6553	1639
Satellite image	Visualization	36	5142	1288
STL10 [66]	Noise reduction	$96 \times 96 \times 3$	5000	8000
Bibtex[67]	Semantic hashing	1836	5916	1479
UNSW-NB15[68]	Anomaly detection	187	37000	175341
AT&T faces	Instance generation	$64 \times 64$	400	–



**Fig. 2.** Example AE architecture. Each block represents a layer and is splitted into three parts: the meaning or purpose of the layer, the type of operation performed and its output shape (size of each dimension).

The models described below are each associated to a diagram describing the layer structure of the corresponding AE and the purpose of each layer. Please refer to Fig. 2 for an example of how each model is detailed.

All examples have been implemented and executed employing the following setup: Tensorflow [69] 1.14.0 and Keras [70] 2.2.4 on top of Python 3.7 and R 3.6, running on an Intel Core i5-8400 CPU and a NVIDIA GeForce RTX 2060 GPU. The associated software can be found at the following GitHub repository: <https://github.com/ari-dasci/autoencoder-case-studies/>.

#### 4.1. Data visualization


Most of the data collected nowadays, either from industries or from the web, is high-dimensional. Visualization techniques can help its interpretability, but the data generally needs to be summarized for this purpose. Traditionally, an alternative representation would be a subset of its features or its principal components [71]. An AE, however, is able to automatically compute a representation that fits each dataset. This representation can be 2 or 3-dimensional if the AE is configured conveniently [72], or if another embedding technique (such as t-SNE [73]) is used after a higher-dimensional encoding.

In particular, if our dataset consists of instances  $(x, y)$  where  $x$  is a feature vector and  $y$  is its associated label, we can use a training subset to learn an autoencoder model with an encoding  $f: \mathbb{R}^n \rightarrow \mathbb{R}^2$  resulting of the composition of the hidden layers up to the code layer. Then, encoded examples can be colored in a scatter plot according to their class.

Although a simple AE could fulfill the embedding task, it can be convenient to restrict or modify its behavior so as to influence the projection to the embedding space, in a way that improves how the populated regions in the original space are modeled. Along these lines, there are several approaches: denoising criteria [60], contractive regularizations [57] and embedding regularizations [72].

##### 4.1.1. Denoising criterion

A denoising AE [60] trains, as briefly explained in Section 3.2, by reconstructing partially corrupted inputs. In order to do this, a corruption or noise function introduces alterations on the input data: for example, a Gaussian noise  $\xi \sim N(0, \sigma)$  would be used to produce the input  $v(x) = x + \xi$ . The reconstruction error is now computed as  $\sum_{x \in \mathcal{X}} d(x, g(f(v(x))))$ . During the training process, the AE is forced to distinguish useful information from mere perturbations of the data. If the instances lie on a manifold in the original feature space, this can effectively train the AE to “push back” instances to the manifold by discarding small displacements



**Fig. 3.** AE architectures for visualization

from it. This can remove noise in the inputs as well as reconstruct some missing values if inputs are just an estimation [9,74]. As a result, the encoding can serve as a set of coordinates for the manifold.

##### 4.1.2. Contractive regularization

The contractive AE [57] uses an additional penalty in the training objective which promotes local invariance to displacements in many directions around the training samples, i.e., it is less sensitive to small perturbations especially in directions that lead outside the manifold. The penalty consists in the squared Frobenius norm of the Jacobian matrix of the encoder, that is, the sum of the squares of all first-order partial derivatives applied to all inputs:  $\sum_x \|J_f(x)\|^2$ . This can be seen as a generalization of L2 weight decay to the case where the encoder is nonlinear. This regularization favors encodings where all dimensions are contracted, but the reconstruction error prevents the AE from contracting dimensions along the manifold.

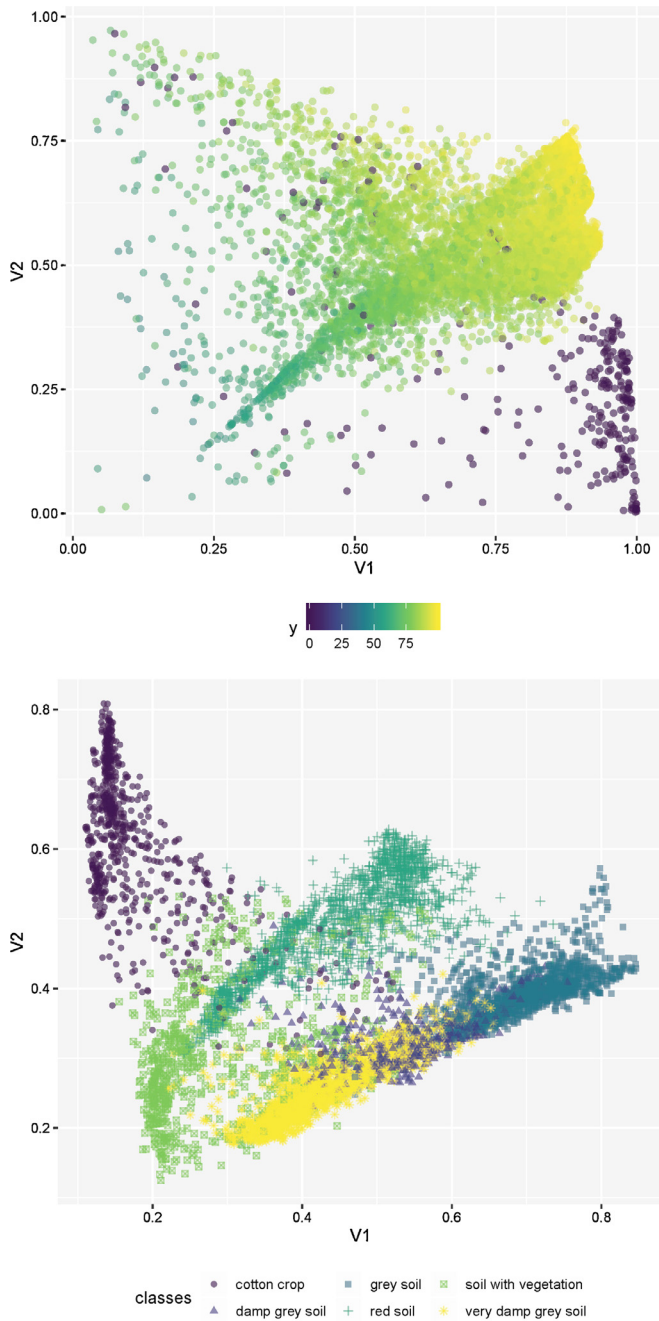
##### 4.1.3. Embedding regularization

An alternative objective function for AEs can be the same loss function from other embedding techniques. This is the idea behind embeddings with AE regularization [72], which combines the reconstruction error with one of several possible embedding loss functions coming from Laplacian eigenmaps [16], multidimensional scaling [75] and margin-based embedding [76]. These loss functions evaluate the embedding by taking pairs of instances, and the AE is adapted the same way, by computing the embedding loss across all pairs of instances and the reconstruction loss across all instances.

For the purposes of demonstrating the capacity of AEs to find manifolds and appropriate embeddings, we have selected a regression dataset, CPU activity<sup>1</sup>, and a classification dataset, Satellite image<sup>2</sup>. The AE used to find embeddings is the contractive AE. AEs for both datasets have been designed using the same criteria: three hidden layers, the encoding layer having 2 variables and the rest having as much variables as needed so that the compression ratio from the input to the first hidden layer is the same than from the hidden layer to the encoding layer. The resulting architectures are detailed in Fig. 3. The AEs have found the projections shown in Fig. 4, where the label of each instance is used to color each

<sup>1</sup> CPU activity dataset is available at <https://www.openml.org/d/573>.

<sup>2</sup> Satellite image dataset can be found at <https://www.openml.org/d/294>.



**Fig. 4.** Embeddings learned by an unsupervised contractive AE. The top image shows the projection of the CPU Activity dataset where each point has been shaded according to the level of user activity. The bottom image displays the projected samples of the Satellite Image dataset, each one colored according to its class.

point. Notice that the AEs have trained without the respective target variables, but there appears to be some degree of separability of classes and different values of the regression variable in each graph.

In order to verify to a certain degree that these embeddings, in addition to producing meaningful visualizations, contain the necessary information about the data, the mean squared error between each instance and its reconstruction through the AE can be computed. As a reference for comparison purposes, the same reconstruction error can be computed from the two first principal components of the data and from the encoding found by a basic AE. Table 3 holds these results, which are very favourable to the contractive AE, since the error is lower in every case. The difference

**Table 3**

Mean squared error comparison between the reconstructions of a contractive AE with a 2-variable encoding and the projections to the original feature space from the two principal components of the data. Lower values are better.

Method	Mean squared error			
	CPU Activity		Satellite	
	Train	Test	Train	Test
PCA	0.5577	0.5097	0.1475	0.1483
Basic AE	0.5238	0.4729	0.1136	0.1160
Contractive AE	<b>0.5053</b>	<b>0.4546</b>	<b>0.1132</b>	<b>0.1157</b>

among both AEs is small, but it serves to deduce that the contractive penalty in the AE does not hinder the reconstruction objective, instead it helps obtain useful low-dimensional embeddings.

#### 4.2. Noise reduction

Similar to searching for interesting representations of data in the encodings of an AE, we can look for a reconstruction that adds value to the input data. One way an AE can help with this is to remove noise from its inputs. This is especially useful when dealing with images [9], sound [77] and other kinds of signals [78], since capture methods usually may introduce some noise and it would be desirable to have a clearer and sharper output.

In general, an AE can be trained to be resilient to input perturbations with a mere random additive noise at the input. Throughout the optimization stage, the AE only takes as input partially corrupted versions of the training examples and attempts to reconstruct the original ones. Once trained, this AE does not necessarily expect more noisy data, but instead it will have learned to be robust against small changes in its inputs. This type of AE is usually called a denoising AE, and performs well in many scenarios that do not necessarily involve treatment of noisy data [61].

In this case, nonetheless, the goal is to eliminate potential perturbations in the inputs. Unlike a generic noise reduction filter, which will perform similar operations no matter what data it receives, a denoising AE can be fitted to a specific training set and may thus be more reliable with different kinds of data. More formally, we consider a noise function  $\nu$ , which generates the corrupted data that the autoencoder trains with to minimize

$$\sum_{x \in \mathcal{X}} d(x, g(f(\nu(x))))$$

The following are some possible noise functions that may be applied:

- $\nu(x) = x + \xi$  where  $\xi$  is sampled from a Gaussian distribution with small variance
- $\nu(x) = x + \xi'$  where  $\xi'$  is sampled from a Cauchy distribution with small scale
- $\nu(x) = \begin{cases} 0 & \text{with low probability } x \\ \text{otherwise} \end{cases}$
- $\nu(x) = \begin{cases} 0 & \text{with low probability } 1 \\ \text{with low probability } x \\ \text{otherwise} \end{cases}$

Notice that the Gaussian and Cauchy distributions will usually induce small changes to most inputs, while the zero and zero-one noises will leave most values intact but the change in the corrupted ones will be more drastic. Thus, for a given application, a specific type of corruption function can be selected so that it fits best to the types of noise the samples could have.

When using denoising AEs, it is also convenient to adapt the type of layers used to the kind of data. For instance, a convolutional AE would be best for noisy images, and an LSTM AE for corrupted signals or sequences. Fig. 5 details a possible encoder-decoder structure for a denoising AE which uses convolutional layers in the encoding phase as well as deconvolution operations during decoding.




Fig. 5. Denoising AE architecture for noise reduction.



Fig. 6. Random selection of test examples (first and third rows) and their reconstructions (second and fourth rows) via forward passes through the denoising AE.

Table 4

Summary of results for noise reduction (average values and standard deviations are provided). Original images without noise are the reference for measuring the mean squared error, and the noise reduction is computed for each image as the percentage decrease in this error. Images are represented by their RGB values from 0 to 255.

Images	Mean squared error	Noise reduction
Reference	0	100%
Noisy	1656.08 $\pm$ 696.31	0%
Basic AE	576.68 $\pm$ 156.53	62.14% $\pm$ 9.54
Denoising AE	<b>159.74 <math>\pm</math> 74.55</b>	<b>88.94% <math>\pm</math> 6.38</b>

When this AE is trained with data from the STL10 dataset [66], a subset of the ImageNet dataset, the objective function will force it to configure its weights so that input noise is reduced along the network. The noise used in this case has been zeros with a probability of 0.1. The test images measure  $96 \times 96$  pixels and have also been corrupted with around 10% of noisy values, which can affect any color channel, so each pixel has a 30% likelihood of having any of its 3 values altered. The AE was trained during 10 epochs with the training data using optimizer Adam.

The results can be analyzed in Table 4, which shows the designed AE achieves a reduction in the mean squared error of about 89%. For comparison purposes, a basic AE has also been trained with Fig. 6 displays some of the test inputs together with their reconstruction by the network. The resulting reconstructions remove most of the noise and appear slightly softer than the originals.

#### 4.3. Semantic hashing

Hashing usually refers to the process of summarizing large batches of data in smaller or simpler codes. Hashes are employed in data structures for fast search times, they can be used to find duplicates and to protect data against corruption and manipulation.

This task in particular, semantic hashing [6], involves finding binary codes which form buckets of similar data, i.e. when two data points are similar to each other, there is high probability that they will be assigned the same hash. Furthermore, if two similar data

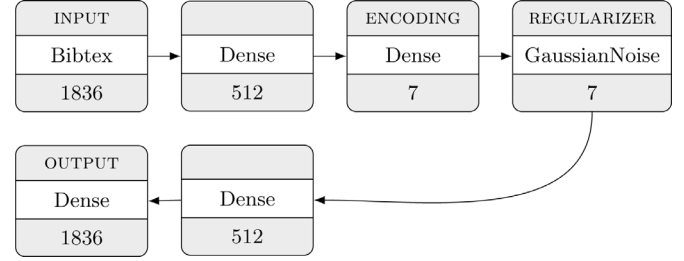


Fig. 7. AE architecture for semantic hashing.

points are not hashed identically, their hashes will likely differ in only a few digits. In consequence, a way of finding instances similar to a query instance is to hash it and look for those whose hashes are the same or almost identical. This is the opposite of cryptographic hashing [79], where the likelihood of two similar entries obtaining the same hash is almost zero and there is no way of retrieving a document from its hash.

The idea of finding semantic relations between data points is especially useful in document searches: if a query document is provided, then the search method should find those documents in the dataset which match as closely as possible. It is also of application in an image domain, where finding matching binary sequences is much more efficient than comparing two pictures [80].

The approach described in [6] uses a very simple AE architecture, with an added noise generator after the encoding which forces the encoder to polarize its outputs.

In this case, the Bibtex dataset [67] was selected to illustrate the application. Fig. 7 shows the AE architecture that was defined for this purpose. The input data provides 1836 binary features which are then projected onto a smaller feature space and lastly onto a 7-dimensional encoding, which is in turn slightly corrupted before decoding. The noise introduced in the encoding during training requires it to take extreme values, for the noise not to affect the reconstruction.


In order to assess whether the trained model serves the purpose of semantic hashing, we can group all possible pairs of hashes according to their Hamming distance (e.g. 0001000 and 001001 are 1 digit away from each other, while 1010101 and 0101010 are separated by a Hamming distance of 7). Then, we measure the inter-cluster distance between those pairs of hashes, computed as the mean cosine distance from each instance in the first cluster to each one in the second. Assuming the clusters group similar instances, the intercluster distance should increase along with the Hamming distance. The distances for this example are illustrated in Fig. 8, which indeed shows simultaneous growth of both.

In addition to quantitatively evaluating the quality of the model, it can be qualitatively analyzed in order to verify whether semantic hashing indeed groups topics in similar hashes. One way of doing this is computing the term frequency-inverse document frequency (tf-idf) [81] of the words for each cluster. This way, words that are frequent within a cluster but uncommon along the rest of the test set are considered the most relevant words. Table 5 shows a truncated list of hashes used by the AE to cluster documents, along with their most relevant words ranked by tf-idf.

#### 4.4. Anomaly detection

Sometimes the objective of a machine learning task is to find unusual behaviors or abnormalities in data, for example, detecting a possible security attack by analyzing server logs, or identifying rare patterns in medical checks. This is known as anomaly detection because the cases of interest are few in contrast to the amount of normal instances, and even in some cases there are no anoma-





**Fig. 8.** Intercluster cosine distance boxplot according to the hamming distance between hashes. Blue diamonds indicate the mean cosine distance among all pairs of clusters that differ in  $k$  digits where  $k$  is a Hamming distance. Gray dots indicate outlier cosine distances.


**Table 5**

The first 15 hashes used as semantic codes for clusters found by the AE, ordered in Gray code. The most relevant words are selected according to tf-idf computed for each cluster. They show some common topics between hashes 0001110 and 0001010, and between 0000001, 0001001 and 0001000.


Hash	Relevant words
0000001	thermodynamic, transitions, induced, generalized, completely, interacting
0000011	relaxation, barrier, mainly, contribute, surfaces, rights
0000010	lipoproteins, capacity, oxidation, apo, receptor, recognized
0000110	identifying, amino, united, capable, matrix, region
0000111	carbon, storage, enzymes, assessed, notes, roles
0000101	infrastructure, configuration, challenge, location, qualitative, improvement
0000100	innovation, construction, ontologies, communities, 1999, located
0001100	mining, advances, bioinformatics, er, solved, intelligence
0001101	reuse, object, perspectives, intelligent, notes, logic
0001111	trans, reading, behavioral, cultural, 1997, gap
0001110	ss, siamese, betta, splendens, male, fighting
0001010	siamese, ss, fighting, male, display, fish
0001011	treated, barrier, combines, electrostatic, solvent, molecule
0001001	thermal, boltzmann, origin, bulk, fluctuations, disorder
0001000	numerically, temperatures, exact, magnetic, glass, zero

lies to train with. In this situation, a traditional classifier cannot solve the problem since it will not be able to assign a class it has not seen before.

An approach to anomaly detection without previously observed anomalous cases is to model those considered typical, and mark as anomalies those instances which do not fit the model. An AE can be used for this purpose, since it can be trained to accurately encode and reconstruct instances following a certain distribution. When the AE is feeded new instances, it is assumed that reconstruction of anomalous data will not be as accurate, since it should follow a different distribution [82–84]. More formally, the hypothesis of this methodology is that, when trained with normal data,  $d(x, g(f(x)))$  will be very small when  $x$  is normal and very high when  $x$  is anomalous.



**Fig. 9.** Denoising AE architecture for anomaly detection.



**Fig. 10.** Precision-recall curve for the detection of individual anomalies in the UNSW dataset.

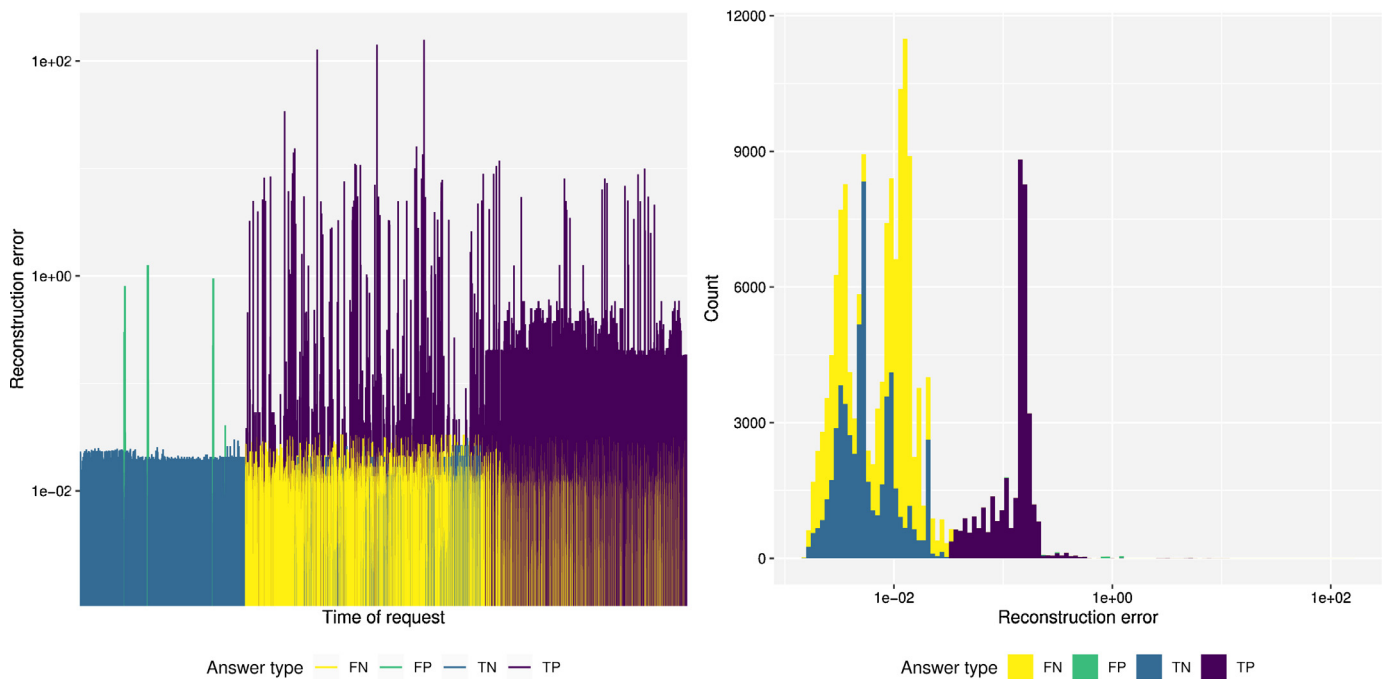
An useful application of anomaly detection where real world data will generally lack anomalies is network intrusion [85,86], that is, the detection of potential security attacks and malicious accesses to a server. The straightforward approach is to continuously log server accesses, and extract data from a period of time where usage has been normal. By means of these data, an AE can be trained to recognize typical usage parameters. Then, new log accesses are constantly feeded to the AE in order to predict their reconstruction error. In the case that several successive errors are much higher than the mean, an attack may be underway.

The AE used for this purpose will work as follows: the encoding layer will perform a drastic dimensionality reduction in order for it to model the most essential information from the training data, which does not include any anomaly. This should help have low error rates on normal data, similar to training instances, but very high ones on anomalous data. In general, this may not work well for uncommon, isolated anomalies, but it is useful when anomalies are several in sequence, so this strategy is especially designed for time series data.

The dataset treated in this example is UNSW-NB15 [68], which has 3 nominal variables and 42 numerical descriptors. Since AEs cannot work directly with nominal variables, these have been converted into dummy binary variables. In addition, any anomalous data from the training subset has been removed. In total, 37000 instances with 187 features are being introduced as the training input of the AE, whose architecture is shown in 9. The extraction of two features is sufficient to model an approximation of most of the normal data, but cannot preserve enough information for the reconstruction of most anomalies.

The results of training this model are summarized in Figs. 10 and 11. The first is a precision-recall curve which gives details about the fraction of detections which are actually anomalies and the ratio of detected anomalies among all of them. We find that it is possible to detect more than half the anomalies without obtaining too many false alarms. Since the test dataset contains many more anomalies than normal instances and the objective is to detect abnormal sections more than to find every individual anomaly,





**Fig. 11.** Reconstruction error of the AE during test. The graph on the left shows the reconstruction error of each request in sequence, where the detection threshold is set to the mean training error plus 6 times its standard deviation. The histogram on the right shows the amount of hits and misses according to the reconstruction error.

a recall of around 50% could be enough as long as the precision is high so that few false alarms are raised.

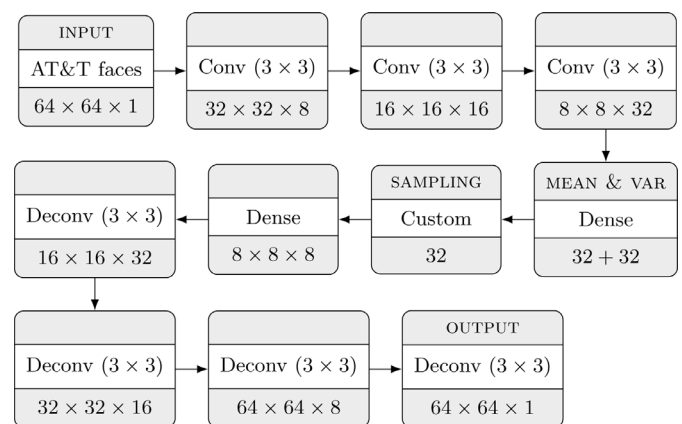
Indeed, Fig. 11 graphs the reconstruction error for each test instance and shows that when an adequate threshold is chosen, anomalous sections can be easily detected with very few isolated false alarms that can be discarded. In this case, the chosen threshold is the mean reconstruction error plus 6 times its standard deviation, but it could be tuned high or low in order to adjust the sensitivity of the detection.

#### 4.5. Instance generation

The representation learned by an AE may be useful to encode or reconstruct individual instances from a training set, but in certain cases it will be very convenient to ensure that this representation is actually attempting to perform some kind of manifold learning, mapping the feature space onto a smaller space in a way that makes sense to work with the whole encoding space. This encoding space would allow to predict a reconstruction for encodings that do not come from an instance in the original feature space, and still produce a coherent result. For instance, an useful application would be to generate new images of faces similar to those in a training set but not identical to any of them. This is usually harder to achieve with simple operations such as interpolation, because they would compute many images that do not represent faces.

There are several variants of AEs that can fulfill this purpose, namely variational [59], adversarial and contractive AEs. Variational as well as adversarial AEs force a prior distribution in the encodings in different ways, which allows to sample new instances by taking points from this space and projecting them onto the original feature space via reconstruction (g). The contractive AE, on the contrary, only imposes a regularization which promotes instances to be mapped to encodings near their neighbors. This helps the autoencoder perform transformations that find manifolds in the data, since local structure is preserved. The manifold can then be traversed in order for the decoder to generate new instances.

Variational AEs are stochastic in the sense that they do not map each instance to a single point in the embedding space, but a dis-




**Fig. 12.** Variational AE architecture for instance generation. The sampling layer draws a sample from the vector of normal distributions with means and variances given by the previous layer.

tribution instead. This is usually a normal distribution, defined by its mean and standard deviation. Then, a reconstruction is produced by sampling that distribution and propagating the results through the decoder network. The objective function in this AE combines the clustering behavior of the reconstruction loss function with a regularization loss which forces the distribution to be as similar as possible to, generally, a multivariate unit Gaussian. This helps the AE extract a very compact representation which only preserves the necessary information to provide a reconstruction of the input.

In this example, a variational AE following the structure in Fig. 12 is trained to generate human faces that do not belong to any person, since they will not be present in the training dataset. The input data used during training belong to the AT&T faces dataset<sup>3</sup>, also known as Olivetti faces dataset. The resulting model can be sampled by feeding arbitrary values to the generator

<sup>3</sup> AT&T faces dataset is available at <https://www.openml.org/d/41083>.



**Fig. 13.** Faces sampled from the encoding space of a variational AE, using interpolations between the projections of images in the original dataset.

component, which then outputs previously unseen images. Fig. 13 shows some representative examples of the generated faces using this AE.

#### 4.6. Other applications

Apart from the previous selection of applications approached with representation learning techniques based on AEs, there are many other situations where AEs can be applied to extract features from data. The following are learning applications present in the literature that fell out of the scope of this article.

##### 4.6.1. Image superresolution

This problem consists in building a high resolution image from a low resolution sample, such as a thumbnail. By using an AE trained with low resolution images and another with the high resolution ones, a map can be trained from the first encoding to the second [87]. This way, the encoder from the first AE can be connected to the decoder from the second AE and the resulting network can be fine-tuned. During prediction it suffices with feeding a low resolution image through the new network, which will encode it and decode it through the high resolution decoder, producing a higher quality image.

##### 4.6.2. Image compression

Images are usually compressed with algorithms designed for this specific purpose, e.g. the JPEG standard [88]. Since a compression mechanism must include a component which compresses the image and another which performs decompression, AEs can be trained in different ways to treat this problem as well [7,89,90], even surpassing the capacity of JPEG2000 especially at low bit rates.

##### 4.6.3. Transfer learning

In a transfer learning task, the learner must make use of the knowledge extracted from data in a given domain to apply it to a different domain. This may consist in using pre-trained networks

with a large dataset to use them with a small dataset by a fine-tuning process. However, when labels for the large dataset are not available, the first stage will necessarily be unsupervised [50], in which case an AE can be trained and its extracted features can initialize a network for a supervised problem with a dataset from other domain.

##### 4.6.4. Human pose and facial features

Human pose recovery is an application specific to image and video data where people appear and the aim is to recognize the pose of each person from the visual information, i.e., to generate a skeleton describing the position and orientation of the legs, arms and the rest of the body. One of the challenges is to model this skeleton as a 3D object while images are only 2D. AEs have been used as the core of a human pose recovery model [91] for extracting an inner pose representation which then maps onto a representation of the 3D pose and is decoded as a 3D pose. This process is, in fact, achieved with two AEs, one for each inner representation required, which are then connected through the representation mapping. In a similar way, facial expression recognition aims to identify the human emotional state from facial images. An approach based on deep sparse autoencoders [92], which are used to extract robust and discriminative features, has been developed to tackle this task.

##### 4.6.5. 3D shape learning

Extracting features from three-dimensional shapes usually has a high computational cost but it is fundamental for tasks such as 3D object retrieval and matching. There are several AE-based models for automatic feature extraction that can help model this type of data [93–95]. These range from simple stacked AEs to combinations of convolutional AEs and extreme learning machines. In general, retrieving similar objects to a given input consists in encoding the input and comparing the result to the codes of known objects in order to find the nearest or most similar ones.

##### 4.6.6. Recommender systems and tagging systems

Recommender systems are filters that seek to predict user preferences for products, taking into account previous choices or ratings. Collaborative filters for recommendation combine the information of different users to build predictions. In [96], a collaborative variational AE for recommendation is developed. It models the implicit relationships among items and users by making use of a shared latent representation and the variational regularization. A task similar to recommendation is tagging, since tags can be ranked for an item according to its similarity to other items. AEs have been also used as the core of tagging systems [97] using denoising AEs and relational denoising AEs.

## 5. Challenges for autoencoder progress and prospects on explainability

Along this section, several difficulties and consequences of using AEs in machine learning are explored. Some brief comments are provided beforehand on the current state of explainability and transparency in artificial intelligence (AI), in order to understand how they could affect the way AEs are designed and used.

First, we introduce the most popular approaches to finding transparent and explainable machine learning models. Later, we develop on the ways AEs can help build interpretable solutions to different problems, by learning disentangled and fair features.

### 5.1. State and prospects on explainability

Explainable AI [25] encompasses many concepts around the idea that people should be able to understand how trained machine learning models work and why they make their decisions.

The recent surge in interest in explainable models derives from the bias found in existing models as well as the search for AI safety [98]. The first issue involves models that make decisions potentially affecting human beings and those decisions can discriminate against certain population groups, e.g. people of color or women. For instance, a prediction model for criminal recidivism was found to be heavily biased against African-American people [99]. The second concept relates to the presence of relatively autonomous agents, such as robots, which execute the actions computed by a machine learning model. These models sometimes find unexpected ways to optimize their reward function (reward hacking) [100], even without completing the objective or having other potential consequences (side effects).

### 5.1.1. Model transparency

The issues above reflect the fact that we should not completely trust trained models unless we can comprehend the ways they are making decisions and predictions. This has attracted the interest of researchers, domain experts and users to more explainable models and strategies to explain *black-box* models, a category which includes most deep learning techniques.

The variety of algorithms to fit machine learning models to data presents a tradeoff between performance and explainability: usually, a simpler, more explainable model is less performant than an opaque model. As a consequence most simple models, such as decision trees, rule-based learners and k-nearest neighbors, are considered transparent, since they provide an interpretable behavior out of the box.

When a model is not transparent enough, there are two main ways to approach explainability: one can use different post-hoc explainability approaches, or modify the model to facilitate our understanding of its decision process. Some new models derived from deep feature learners are designed to improve transparency and be more self-explanatory.

One way deep learning models can increase transparency is by highlighting which input features are causing their predictions. For example, attention-based models [101] have an embedded scoring technique which highlights the zones in the input that are being taken into account to make predictions. This works for image classification and object detection [102] as well as for document processing [103].

A different proposal for transparent image classification is a convolutional neural network-based classifier which identifies prototypes in similar images [104], that is, it provides examples on images of the same class that justify the prediction.

### 5.1.2. Explainability techniques

When an opaque model is used, there are still ways to improve our understanding of its inner workings or its predictions. In many cases, a post-hoc explainability technique may be applied. The different methods that can render a model more interpretable are usually categorized into two groups. They can be either model-agnostic, if they work independently of the model used, or model-specific, otherwise.

Some examples of model-agnostic approaches and tools are the following:

- Local approximations. LIME [105] this is a method which linearly performs a local approximation of a classifier or regressor, in a way which is interpretable. An AE-based variant of LIME has been developed to improve its stability [106].
- FairML [107]. The FairML toolbox can find strong dependencies between model outputs and the input features.
- Sensitivity analysis [108]. It is a computation based on the derivative of the conditional probability of not predicting a class given the input features. This defines a vector field where each

vector indicates the direction an instance needs to be moved to, so as to be classified differently.

- Auditing. Trained models can be repeatedly tested against different inputs in order to analyze how the outputs are affected. These inputs, however, need to be provided according to some criteria. As a way to compute direct and indirect influence of each feature in the output of a model, there is a procedure which obscures the effect of a variable in the data [109]. It works without retraining the model, and can assess the degree in which a feature is relevant to a classifier. There are several other approaches to analyzing direct influence of a feature in the output of a model [110,111].
- Counterfactuals [112–114]. This is an approach with a similar objective to auditing but from a different perspective. Finding a counterfactual consists in detecting the smallest possible change in feature values that causes an alteration to the prediction of the model. These serve as an explanation for the “closest possible world” where the prediction would have been different, without providing further insight into the decision process.

There are several specific techniques for explaining the outputs of deep learning models:

- Layer-wise relevance propagation [115]. This is a methodology for visualization of pixel-wise contributions to predictions, where classifiers are decomposed into several layers of computation, so the relevance of each pixel is found by propagating relevance backwards through the network.
- Saliency map generation [116]. Saliency maps are heatmaps where the most relevant features from the input are highlighted. These are usually applied to convolutional neural networks in order to obtain the image regions that cause the output for each instance.
- DeepLIFT [117]. This is a technique for computing relevance for each input feature to a neural network, by assigning contribution scores to each neuron according to its activation given a specific input.
- SHAP [118]. This tool provides several model-specific techniques which find local explanations for different models based on Shapley values from game theory. In particular, it includes DeepExplainer and GradientExplainer, which apply to deep learning models.
- Traceability [119]. This is a more theoretical concept from the field of software development that could be applied to deep neural models. It seeks to describe how each component of a final inference model is related back to its training model, the dataset, hyperparameters and all the way up to some high level requirements on what task the model should carry out. Being able to trace every item in the development of a deep neural networks to a higher level cause could serve to ensure that all choices such as hyperparameters and architecture are well justified.

## 5.2. Current challenges and influence in future work

As discussed in the previous section, most of the well-known explainability techniques involve analysis of features in one way or another. The contribution that AEs can provide in this field is, therefore, substantial. This is due to the fact that AEs can transform a set of highly dependent, correlated features in a different set of independent, interpretable ones, by using adequate regularizations. In this section, we comment on different ways to learn features that are *meaningful* and *fair*, and on recent developments for also improving the explainability of the feature extraction process itself.

### 5.2.1. Improving features: disentanglement and fairness

One way extracted features can improve their quality is by holding an understandable meaning by themselves, e.g. a model could train with face pictures and extract a feature for hair color, another one for nose size, etc. These new features would be much more useful than the original ones which represent individual pixels. This task is usually known as *feature disentanglement*.

Some recent AE models whose objective is to disentangle features are Total Correlation VAE [120], Wasserstein AE [121] and InfoGAN [122]. All of these are generative models, so, as a result, extracted features not only provide interpretable meaning to instances, but can also be sampled in order to generate unseen examples in a way that resembles the manipulation of existing instances: for example, a model could generate a realistic face similar to an existing image but changing blonde hair to black.

Another step forward in improving learned representations is forcing these to become fair [123], which means that the extracted features obfuscate information about membership to potentially discriminated groups, e.g. gender or ethnicity. Fairness usually applies only in contexts where model predictions affect human lives, e.g. job applications, legal proceedings, etc. A statistic can be defined to measure the discrimination of a classifier with respect to a binary variable. The objective is then to optimize a tradeoff between classification accuracy and discrimination [124].

There already exist AE-based models for learning fair representations. In [125], an adversarial AE-based classifier is proposed where the adversary attempts to predict the sensitive (potentially discriminatory) attributes from the encoding, but its prediction ability is minimized by the AE and classifier. The objective function can be adjusted according to the desired type of fairness. Another model in [126] consists in a variational AE which disentangles sensitive information from the non-sensitive latent features and is flexible in the sense that potentially sensitive information can be retained or removed from the encoding during inference.

### 5.2.2. Explainable feature learning

The described approaches provide the possibility of explaining the end predictions of other models, as well as rendering them fairer. However, as has been extensively discussed in this work, the extracted features can be the actual core of a solution to many problems. As a consequence, it would be necessary as well to develop strategies which facilitate the explainability of the transformations an AE can perform in order to learn features. This is an area only explored very recently, but there are already some developments.

Variational AEs can be used to detect anomalies, similarly to the denoising AE explained in Section 4.4. In addition, they enable another, more explainable way of detecting anomalies: computing the gradients of the reconstruction error with respect to the inputs [127]. This allows to notice which input features are contributing to the error, and to cluster anomalies according to this same criterion.

A different approach to improving the explainability of the embedding consists in restricting the operations each neuron performs to just logical AND/OR operators [128], which limits the origin of each extracted feature to a relatively simple logical combination of the input features, thus facilitating its interpretability.

### 5.2.3. Influence in future works

There is currently much to be researched in the area of explainable AEs as well as AEs which help explain other models by extracting better features. The current trends focus especially on generative models such as variational AEs for these purposes, and will probably continue to do so, even if some diversification is achieved as new works appear.

The adaptability of AEs to many different problems, illustrated in previous sections, together with the possibility of producing in-

terpretable and fair features, may lead to an increase in usage of these models throughout all kinds of machine learning applications.

In our future work, we intend to approach explainable feature learning in the context of AEs, that is, find AE-based models that extract features and at the same time provide an understandable meaning to the mapping from the original features to the encoded ones. Ideally, an explainable feature learner should not be restricted to one end application, but could be used for many purposes, as common AEs already can.

## 6. Conclusions

Throughout this text, we have summarized the traditional alternatives for learning representations, the origins and essential characteristics of AEs, including how to introduce certain behaviors into the coding layer.

Later, we have thoroughly examined several case studies of AE applications in unstructured data as well as images and sequences: data visualization, image denoising, semantic hashing, anomaly detection and instance generation. Other applications have also been briefly discussed: image superresolution, image compression, transfer learning, human pose recovery and recommender systems.

An introduction to the state of the art in explainable AI and its application to the field of AEs has been provided as well. AEs have notoriously contributed to the areas of feature disentanglement and fair representations, and there have been some recent developments on explainable feature learning as well.

We can conclude that AEs are a versatile framework for solving a wide variety of problems where a central task is to learn representations of the data. They can adapt to a given problem in structure as well as in the objective they optimize. This way, if the solution to a problem can be modeled with a transformation of the feature space onto another space, there will be many instances where the parameters of the transformation can be learned by an AE.

## Declaration of Competing Interest

None.

## Acknowledgments

D. Charte is supported by the Spanish Ministry of Science, Innovation and Universities under the FPU National Program (Ref. FPU17/04069). This paper is partially supported by the Spanish National Research Projects TIN2015-68854-R and TIN2017-89517-P and the project DeepSCOP Ayudas Fundación BBVA a Equipos de Investigación Científica en Big Data 2018.

## References

- [1] P. Domingos, A few useful things to know about machine learning, *Commun. ACM* 55 (10) (2012) 78–87, doi:[10.1145/2347736.2347755](https://doi.org/10.1145/2347736.2347755).
- [2] A.C. Lorena, L.P. Garcia, J. Lehmann, M.C. Souto, T.K. Ho, How complex is your classification problem?: A survey on measuring classification complexity, *ACM Comput. Surv. (CSUR)* 52 (5) (2019) 107, doi:[10.1145/3347711](https://doi.org/10.1145/3347711).
- [3] S. García, J. Luengo, F. Herrera, *Data Preprocessing in Data Mining*, Springer, 2015, doi:[10.1007/978-3-319-10247-4](https://doi.org/10.1007/978-3-319-10247-4).
- [4] S.B. Kotsiantis, I. Zaharakis, P. Pintelas, *Supervised machine learning: a review of classification techniques*, *Emerg. Artif. Intell. Appl. Comput. Eng.* 160 (2007) 3–24.
- [5] A.K. Jain, M.N. Murty, P.J. Flynn, Data clustering: a review, *ACM Comput. Surv.* 31 (1999) 264–323, doi:[10.1145/331499.331504](https://doi.org/10.1145/331499.331504).
- [6] R. Salakhutdinov, G. Hinton, Semantic hashing, *Int. J. Approx. Reason.* 50 (7) (2009) 969–978, doi:[10.1016/j.ijar.2008.11.006](https://doi.org/10.1016/j.ijar.2008.11.006).
- [7] L. Theis, W. Shi, A. Cunningham, F. Huszár, Lossy image compression with compressive autoencoders, in: *Proceedings of the Fifth International Conference on Learning Representations*, 2017, <https://openreview.net/forum?id=rJiNwv9gg>.



- [8] J. Deng, Z. Zhang, E. Marchi, B. Schuller, Sparse autoencoder-based feature transfer learning for speech emotion recognition, Humaine Association Conference on Affective Computing and Intelligent Interaction, IEEE (2013) 511–516, doi:10.1109/ACII.2013.90, 2013.
- [9] J. Xie, L. Xu, E. Chen, Image denoising and inpainting with deep neural networks, Adv. Neural Inf. Process. Syst. (2012) 341–349.
- [10] I. Guyon, S. Gunn, M. Nikravesh, L.A. Zadeh, Feature extraction, Found. Appl. (2008) 207, doi:10.1007/978-3-540-35488-8, Springer.
- [11] M. Dash, H. Liu, F.s. f. classification, Intelligent data analysis, 1997, 1, 1–4, 131–156, 10.1016/S1088-467X(97)00008-5.
- [12] Y. Bengio, A. Courville, P. Vincent, Representation learning: a review and new perspectives, IEEE Trans. Pattern Anal. Mach. Intell. 35 (8) (2013) 1798–1828, doi:10.1109/TPAMI.2013.50.
- [13] I.T. Jolliffe, Introduction, in: Principal Component Analysis, Springer, 1986, pp. 1–7, doi:10.1007/978-1-4757-1904-8.
- [14] J.B. Tenenbaum, V. De Silva, J.C. Langford, A global geometric framework for nonlinear dimensionality reduction, Science 290 (5500) (2000) 2319–2323, doi:10.1126/science.290.5500.2319.
- [15] S.T. Roweis, L.K. Saul, Nonlinear dimensionality reduction by locally linear embedding, Science 290 (5500) (2000) 2323–2326, doi:10.1126/science.290.5500.2323.
- [16] M. Belkin, P. Niyogi, Laplacian eigenmaps for dimensionality reduction and data representation, Neural Comput. 15 (2003) 1373–1396, doi:10.1162/08997660321780317.
- [17] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, Nature 521 (2015) 7553, doi:10.1038/nature14539436.
- [18] D.E. Rumelhart, G.E. Hinton, R.J. Williams, Learning representations by back-propagating errors, Cogn. Model. 5 (3) (1988), doi:10.1038/323533a0, 1.
- [19] H. Robbins, S. Monro, A stochastic approximation method, Ann. Math. Stat. (1951) 400–407, doi:10.1007/978-1-4612-5110-1\_9.
- [20] J. Duchi, E. Hazan, Y. Singer, Adaptive subgradient methods for online learning and stochastic optimization, J. Mach. Learn. Res. 12 (2011) 2121–2159.
- [21] M.D. Zeiler, Adadelta: an adaptive learning rate method, ArXiv preprint arXiv:1212.5701.
- [22] D. Kingma, J. Ba, Adam: A method for stochastic optimization, in: Proceedings of the Third International Conference on Learning Representations, 2015.
- [23] T. Tieleman, G. Hinton, Lecture 6.5-RMSProp, COURSE: Neural Netw. Machine Learn. 4 (2) (2012) 26–31.
- [24] D. Charte, F. Charte, S. García, M.J. del Jesus, F. Herrera, A practical tutorial on autoencoders for nonlinear feature fusion: Taxonomy, models, software and guidelines, Inf. Fusion 44 (2018) 78–96, doi:10.1016/j.inffus.2017.12.007.
- [25] A.B. Arrieta, N. Díaz-Rodríguez, J.D. Ser, A. Bennetot, S. Tabik, A. Barbado, S. García, S. Gil-López, D. Molina, R. Benjamins, Explainable artificial intelligence (xai): concepts, taxonomies, opportunities and challenges toward responsible AI, Inf. Fusion 58 (2020) 82–115, doi:10.1016/j.inffus.2019.12.012.
- [26] L.V.D. Maaten, E. Postma, J.V.D. Herik, Dimensionality reduction: a comparative review, Tech. Rep. (2009).
- [27] R.A. Fisher, The use of multiple measurements in taxonomic problems, Ann. Eugen. 7 (2) (1936) 179–188, doi:10.1111/j.1469-1809.1936.tb02137.x.
- [28] H. Zhao, S. Sun, Z. Jing, J. Yang, Local structure based supervised feature extraction, Pattern Recogn. 39 (8) (2006) 1546–1550, doi:10.1016/j.patcog.2006.02.023.
- [29] A. Makhzani, J. Shlens, N. Jaitly, I. Goodfellow, Adversarial autoencoders, in: Proceedings of the Fourth International Conference on Learning Representations, 2016.
- [30] J. Zhao, K. Lu, X. He, Locality sensitive semi-supervised feature selection, Neurocomputing 71 (10–12) (2008) 1842–1849, doi:10.1016/j.neucom.2007.06.014.
- [31] K. Pearson, LIII. On lines and planes of closest fit to systems of points in space, Philos. Mag. Ser. 6 (2) (1901) 559–572, doi:10.1080/14786440109462720, 11.
- [32] H. Hotelling, Analysis of a complex of statistical variables into principal components, J. Educ. Psychol. 24 (6) (1933), doi:10.1037/h0071325, 417.
- [33] I.T. Jolliffe, Principal component analysis and factor analysis, in: Principal Component Analysis, Springer, 1986, pp. 115–128, doi:10.1007/978-1-4757-1904-8.
- [34] B.K. Schölkopf, A. Smola, R. Müller, Nonlinear component analysis as a kernel eigenvalue problem, Neural Comput. 10 (5) (1998) 1299–1319, doi:10.1162/089976698300017467.
- [35] B. Schölkopf, in: Statistical Learning and Kernel Methods, Springer, Vienna, 2001, pp. 3–24, doi:10.1007/978-3-7091-2580-9\_1.
- [36] P. Smolensky, Information processing in dynamical systems: foundations of harmony theory, Technical Report, Colorado University at Boulder, Department of Computer Science (1986).
- [37] G.E. Hinton, Training products of experts by minimizing contrastive divergence, Neural Comput. 14 (8) (2002) 1771–1800, doi:10.1162/089976602760128018.
- [38] J.W. Sammon, A nonlinear mapping for data structure analysis, IEEE Trans. Comput. 100 (5) (1969) 401–409, doi:10.1109/T-C.1969.222678.
- [39] A. Krizhevsky, I. Sutskever, G.E. Hinton, Imagenet classification with deep convolutional neural networks, in: Advances in Neural Information Processing Systems, 25, 2012, pp. 1097–1105.
- [40] J. Schmidhuber, Deep learning in neural networks: an overview, Neural Netw. 61 (2015) 85–117, doi:10.1016/j.neunet.2014.09.003.
- [41] T. Kohonen, The self-organizing map, Proc. IEEE 78 (9) (1990) 1464–1480, doi:10.1109/5.58325.
- [42] P. Koikkalainen, E. Oja, Self-organizing hierarchical feature maps, in: Proceedings of the IJCNN International Joint Conference on Neural Networks, IEEE, 1990, pp. 279–284, doi:10.1109/IJCNN.1990.137727, 1990.
- [43] J. Schmidhuber, M. Eldracher, B. Foltin, Semilinear predictability minimization produces well-known feature detectors, Neural Comput. 8 (4) (1996) 773–786, doi:10.1162/neco.1996.8.4.773.
- [44] I. Goodfellow, Y. Bengio, A. Courville, Deep learning, in: Deep Generative Models, MIT Press, 2016, pp. 651–716. <http://www.deeplearningbook.org>.
- [45] G.E. Hinton, R.R. Salakhutdinov, Reducing the dimensionality of data with neural networks, Science 313 (5786) (2006) 504–507, doi:10.1126/science.1127647.
- [46] N. Zeng, Z. Wang, H. Zhang, W. Liu, F.E. Alsaadi, Deep belief networks for quantitative analysis of a gold immunochromatographic strip, Cogn. Comput. 8 (4) (2016) 684–692.
- [47] N. Zeng, Z. Wang, H. Zhang, K.-E. Kim, Y. Li, X. Liu, An improved particle filter with a novel hybrid proposal distribution for quantitative analysis of gold immunochromatographic strips, IEEE Trans. Nanotechnol. 18 (2019) 819–829.
- [48] M.A. Kramer, Nonlinear principal component analysis using autoassociative neural networks, AIChE J. 37 (2) (1991) 233–243, doi:10.1002/aic.690370209.
- [49] E. Oja, Data compression, feature extraction, and autoassociation in feedforward neural networks, Artif. Neural Netw. 1 (1991) 737–745.
- [50] Y. Bengio, Deep learning of representations for unsupervised and transfer learning, in: Proceedings of ICLR Workshop on Unsupervised and Transfer Learning, 2012, pp. 17–36.
- [51] D.H. Ballard, Modular learning in neural networks, AAAI (1987) 279–284.
- [52] G.E. Hinton, S. Osindero, Y.W. Teh, A fast learning algorithm for deep belief nets, Neural Comput. 18 (7) (2006) 1527–1554, doi:10.1162/neco.2006.18.7.1527.
- [53] X. Glorot, A. Bordes, Y. Bengio, Deep sparse rectifier neural networks, in: Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics, 2011, pp. 315–323.
- [54] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, R. Salakhutdinov, Dropout: a simple way to prevent neural networks from overfitting, J. Mach. Learn. Res. 15 (1) (2014) 1929–1958.
- [55] H. Lee, C. Ekanadham, A.Y. Ng, Sparse deep belief net model for visual area v2, in: Advances in Neural Information Processing Systems, 20, 2008, pp. 873–880.
- [56] A. Ng, Sparse autoencoder, CS294A Lecture notes 72 (2011) (2011) 1–19.
- [57] S. Rifai, P. Vincent, X. Muller, X. Glorot, Y. Bengio, Contractive auto-encoders: explicit invariance during feature extraction, in: Proceedings of the Twenty-eighth International Conference on Machine Learning (ICML-11), 2011, pp. 833–840.
- [58] S. Rifai, G. Mesnil, P. Vincent, X. Muller, Y. Bengio, Y. Dauphin, X. Glorot, Higher order contractive auto-encoder, in: Proceedings of the Joint European Conference on Machine Learning and Knowledge Discovery in Databases, Springer, 2011, pp. 645–660, doi:10.1007/978-3-642-23783-6\_41.
- [59] D.P. Kingma, M. Welling, Auto-encoding variational bayes, arXiv preprint ArXiv:1312.6114.
- [60] P. Vincent, H. Larochelle, Y. Bengio, P.A. Manzagol, Extracting and composing robust features with denoising autoencoders, in: Proceedings of the Twenty-fifth International Conference on Machine Learning, ACM, 2008, pp. 1096–1103, doi:10.1145/1390156.1390294.
- [61] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, P.A. Manzagol, Stacked denoising autoencoders: learning useful representations in a deep network with a local denoising criterion, J. Mach. Learn. Res. 11 (2010) 3371–3408.
- [62] Y. Qi, Y. Wang, X. Zheng, Z. Wu, Robust feature learning by stacked autoencoder with maximum correntropy criterion, in: Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE, 2014, pp. 6716–6720, doi:10.1109/ICASSP.2014.6854900.
- [63] W. Liu, P.P. Pokharel, J.C. Principe, Correntropy: a localized similarity measure, in: Proceedings of the IEEE International Joint Conference on Neural Networks, 2006, pp. 4919–4924, doi:10.1109/IJCNN.2006.247192, IJCNN, IEEE.
- [64] I. Goodfellow, J. Pouget, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair-Abadie, A. Courville, Y. Bengio, Generative adversarial nets, in: Advances in Neural Information Processing Systems, 27, 2014, pp. 2672–2680.
- [65] F. Zhuang, X. Cheng, P. Luo, S.J. Pan, Q. He, Supervised representation learning: Transfer learning with deep autoencoders, in: Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, 2015, pp. 4119–4125.
- [66] A. Coates, A. Ng, H. Lee, An analysis of single-layer networks in unsupervised feature learning, in: Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics, 2011, pp. 215–223.
- [67] I. Katakis, G. Tsoumakas, I. Vlahavas, Multilabel text classification for automated tag suggestion, in: Proceedings of the ECML/PKDD, 2008, pp. 75–83.
- [68] N. Moustafa, J. Slay, Unsw-nb15: a comprehensive data set for network intrusion detection systems (unsw-nb15 network data set), in: Proceedings of the Communications and Information Systems Conference (MilCIS), 2015, pp. 1–6, doi:10.1109/MilCIS.2015.7348942, Military.
- [69] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, Tensorflow: a system for large-scale machine learning, in: Proceedings of the Twelfth USENIX Symposium on Operating Systems Design and Implementation (OSDI 16), 2016, pp. 265–283.
- [70] F. Chollet, Keras, 2015, <https://keras.io>.
- [71] I. Jolliffe, Principal Component Analysis, Springer, Berlin, Heidelberg, 2011, pp. 1094–1096, doi:10.1007/978-3-642-04898-2\_455.
- [72] W. Yu, G. Zeng, P. Luo, F. Zhuang, Q. He, Z. Shi, Embedding with autoencoder

- regularization, in: Proceedings of the Joint European Conference on Machine Learning and Knowledge Discovery in Databases, Springer, 2013, pp. 208–223, doi:10.1007/978-3-642-40994-3\_14.
- [73] L.v.d. Maaten, G. Hinton, Visualizing data using t-SNE, *J. Mach. Learn. Res.* 9 (2008) 2579–2605.
- [74] J. Li, Z. Struzik, L. Zhang, A. Cichocki, Feature learning from incomplete eeg with denoising autoencoder, *Neurocomputing* 165 (2015) 23–31, doi:10.1016/j.neucom.2014.08.092.
- [75] W.S. Torgerson, Multidimensional scaling: I. Theory and method, *Psychometrika* 17 (4) (1952) 401–419, doi:10.1007/BF02288916.
- [76] R. Hadsell, S. Chopra, Y. LeCun, Dimensionality reduction by learning an invariant mapping, in: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06), 2, 2006, pp. 1735–1742, doi:10.1109/CVPR.2006.100. IEEE
- [77] X. Lu, Y. Tsao, S. Matsuda, C. Hori, Speech enhancement based on deep denoising autoencoder, *Interspeech* (2013) 436–440.
- [78] P. Xiong, H. Wang, M. Liu, S. Zhou, Z. Hou, X. Liu, Ecg signal enhancement based on improved denoising auto-encoder, *Eng. Appl. Artif. Intell.* 52 (2016) 194–202, doi:10.1016/j.engappai.2016.02.015.
- [79] J. Katz, Y. Lindell, *Introduction to Modern Cryptography*, 2014.
- [80] M.A. Carreira-Perpinán, R. Raziperchikolaei, Hashing with binary autoencoders, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2015, pp. 557–566, doi:10.1109/CVPR.2015.7298654.
- [81] S. Robertson, Understanding inverse document frequency: on theoretical arguments for IDF, *J. Document.* 60 (5) (2004) 503–520, doi:10.1108/00220410410560582.
- [82] T. Petsche, A. Marcantonio, C. Darken, S.J. Hanson, G.M. Kuhn, N.I. Santos, A neural network autoassociator for induction motor failure prediction, in: *Advances in Neural Information Processing Systems*, 9, 1996, pp. 924–930.
- [83] M. Sakurada, T. Yairi, Anomaly detection using autoencoders with nonlinear dimensionality reduction, in: Proceedings of the MLSDA(2014) Second Workshop on Machine Learning for Sensory Data Analysis, ACM, 2014, pp. 4–11, doi:10.1145/2689746.2689747.
- [84] S. Park, M. Kim, S. Lee, Anomaly detection for http using convolutional autoencoders, *IEEE Access* 6 (2018) 70884–70901, doi:10.1109/ACCESS.2018.2881003.
- [85] Y. Mirsky, T. Doitsman, Y. Elovici, A. Shabtai, Kitsune: an ensemble of autoencoders for online network intrusion detection, in: Proceedings of the Network and Distributed Systems Security (NDSS) Symposium, 2018, pp. 1–15, doi:10.14722/ndss.2018.23204.
- [86] N. Shone, T.N. Ngoc, V.D. Phai, Q. Shi, A deep learning approach to network intrusion detection, *IEEE Trans. Emerg. Top. Comput. Intell.* 2 (1) (2018) 41–50, doi:10.1109/TETCI.2017.2772792.
- [87] K. Zeng, J. Yu, R. Wang, C. Li, D. Tao, Coupled deep autoencoder for single image super-resolution, *IEEE Trans. Cybern.* 47 (1) (2015) 27–37, doi:10.1109/TCYB.2015.2501373.
- [88] G.K. Wallace, The JPEG still picture compression standard, *IEEE Trans. Consum. Electron.* 38 (1) (1992), doi:10.1145/103085.103089. xviii–xxxiv
- [89] J. Ballé, V. Laparra, E.P. Simoncelli, End-to-end optimized image compression, in: *Proceedings of the Fifth International Conference on Learning Representations*, 2017.
- [90] Z. Cheng, H. Sun, M. Takeuchi, J. Katto, Deep convolutional autoencoder-based lossy image compression, in: Proceedings of the Picture Coding Symposium (PCS), IEEE, 2018, pp. 253–257, doi:10.1109/PCS.2018.8456308.
- [91] C. Hong, J. Yu, J. Wan, D. Tao, M. Wang, Multimodal deep autoencoder for human pose recovery, *IEEE Trans. Image Process.* 24 (12) (2015) 5659–5670, doi:10.1109/TIP.2015.2487860.
- [92] N. Zeng, H. Zhang, B. Song, W. Liu, Y. Li, A.M. Dobaie, Facial expression recognition via learning deep sparse autoencoders, *Neurocomputing* 273 (2018) 643–649.
- [93] Z. Zhu, X. Wang, S. Bai, C. Yao, X. Bai, Deep learning representation using autoencoder for 3d shape retrieval, *Neurocomputing* 204 (2016) 41–50, doi:10.1016/j.neucom.2015.08.127.
- [94] Y. Wang, Z. Xie, K. Xu, Y. Dou, Y. Lei, An efficient and effective convolutional auto-encoder extreme learning machine network for 3d feature learning, *Neurocomputing* 174 (2016) 988–998, doi:10.1016/j.neucom.2015.10.035.
- [95] B. Leng, S. Guo, X. Zhang, Z. Xiong, 3d object retrieval with stacked local convolutional autoencoder, *Signal Process.* 112 (2015) 119–128, doi:10.1016/j.sigpro.2014.09.005.
- [96] X. Li, J. She, Collaborative variational autoencoder for recommender systems, in: Proceedings of the Twenty-third ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, 2017, pp. 305–314, doi:10.1145/3097983.3098077.
- [97] H. Wang, X. Shi, D.Y. Yeung, Relational stacked denoising autoencoder for tag recommendation, in: Proceedings of the Twenty-ninth AAAI conference on artificial intelligence, 2015.
- [98] D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, D. Man, Concrete problems in ai safety arXiv preprint, ArXiv:1606.06565. <https://arxiv.org/pdf/1606.06565.pdf>.
- [99] J. Angwin, J. Larson, S. Mattu, L. Kirchner, Machine Bias, *ProPublica* (23)
- [100] J. Bird, P. Layzell, The evolved radio and its implications for modelling the evolution of novel sensors, in: Proceedings of the 2002 Congress on Evolutionary Computation. CEC'02 (Cat. No. 02TH8600), 2, 2002, pp. 1836–1841, doi:10.1109/CEC.2002.1004522. IEEE
- [101] C. Cao, X. Liu, Y. Yang, Y. Yu, J. Wang, Z. Wang, Y. Huang, L. Wang, C. Huang, W. Xu, Look and think twice: capturing top-down visual attention with feed-back convolutional neural networks, in: Proceedings of the IEEE International Conference on Computer Vision, 2015, pp. 2956–2964, doi:10.1109/ICCV.2015.338.
- [102] F. Wang, M. Jiang, C. Qian, S. Yang, C. Li, H. Zhang, X. Wang, X. Tang, Residual attention network for image classification, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2017, pp. 3156–3164, doi:10.1109/CVPR.2017.683.
- [103] Z. Yang, D. Yang, C. Dyer, X. He, A. Smola, E. Hovy, Hierarchical attention networks for document classification, in: Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, 2016, pp. 1480–1489.
- [104] C. Chen, O. Li, D. Tao, A. Barnett, C. Rudin, J.K. Su, This looks like that: deep learning for interpretable image recognition, in: *Advances in Neural Information Processing Systems* 32, Curran Associates, Inc., 2019, pp. 8928–8939.
- [105] M.T. Ribeiro, S. Singh, C. Guestrin, “Why should i trust you?”: Explaining the predictions of any classifier, in: Proceedings of the Twenty-second ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16, Association for Computing Machinery, New York, NY, USA, 2016, pp. 1135–1144, doi:10.1145/2939672.2939778.
- [106] S.M. Shankaranarayana, D. Runje, Alime: autoencoder based approach for local interpretability, in: Proceedings of the International Conference on Intelligent Data Engineering and Automated Learning, Springer, 2019, pp. 454–463, doi:10.1007/978-3-030-33607-3\_49.
- [107] J.A. Adebayo, FairML: Toolbox for Diagnosing Bias in Predictive Modeling, (Ph.D. thesis), Massachusetts Institute of Technology, 2016.
- [108] D. Baehrens, T. Schroeter, S. Harmeling, M. Kawanabe, K. Hansen, K.R. Müller, How to explain individual classification decisions, *J. Mach. Learn. Res.* 11 (2010) 1803–1831.
- [109] P. Adler, C. Falk, S.A. Friedler, T. Nix, G. Rybeck, C. Scheidegger, B. Smith, S. Venkatasubramanian, Auditing black-box models for indirect influence, *Knowl. Inf. Syst.* 54 (1) (2018) 95–122, doi:10.1007/s10115-017-1116-3.
- [110] A. Henelius, K. Puolamäki, H. Boström, L. Asker, P. Papadimitrou, A peek into the black box: exploring classifiers by randomization, *Data Min. Knowl. Discov.* 28 (5–6) (2014) 1503–1529, doi:10.1007/s10618-014-0368-8.
- [111] A. Datta, S. Sen, Y. Zick, Algorithmic transparency via quantitative input influence: theory and experiments with learning systems, in: Proceedings of the IEEE Symposium on Security and Privacy (SP), IEEE, 2016, pp. 598–617, doi:10.1109/SP.2016.42.
- [112] C. Molnar, Interpretable machine learning, Ch. 6.1. Counterfactual Explanations, 2019. <https://christophm.github.io/interpretable-ml-book/>.
- [113] S. Wachter, B. Mittelstadt, C. Russell, Counterfactual explanations without opening the black box: automated decisions and the gdpr, *Harv. J. L. & Tech.* 31 (2017) 841.
- [114] A. Barredo-Arrieta, J. Del Ser, Plausible counterfactuals: auditing deep learning classifiers with realistic adversarial examples, (2020) arXiv e-prints, ArXiv:2003.11323. <http://arxiv.org/abs/2003.11323>
- [115] S. Bach, A. Binder, G. Montavon, F. Klauschen, K.-R. Müller, W. Samek, On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation, *PLoS One* 10 (7) 10.1371/journal.pone.0130140.
- [116] K. Simonyan, A. Vedaldi, A. Zisserman, Deep inside convolutional networks: visualising image classification models and saliency maps, arXiv preprint, ArXiv:1312.6034.
- [117] A. Shrikumar, P. Greenside, A. Kundaje, Learning important features through propagating activation differences, in: Proceedings of the Thirty-fourth International Conference on Machine Learning, 70, 2017, pp. 3145–3153.
- [118] S.M. Lundberg, S.-I. Lee, I. Guyon, U.V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. Garnett, A unified approach to interpreting model predictions, in: *Advances in Neural Information Processing Systems* 30, Curran Associates, Inc., 2017, pp. 4765–4774. <http://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions.pdf>.
- [119] V. Aravantis, F. Diehl, Traceability of deep neural networks, arXiv preprint, ArXiv:1812.06744.
- [120] T.Q. Chen, X. Li-Bianchi, R.B. Grosse, D.K. Duvenaud, S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa, R. Garnett, Isolating sources of disentanglement in variational autoencoders, in: *Advances in Neural Information Processing Systems* 31, Curran Associates, Inc., 2018, pp. 2610–2620.
- [121] P.K. Rubenstein, B. Schölkopf, I. Tolstikhin, Learning disentangled representations with wasserstein auto-encoders, in: Proceedings of the Sixth International Conference on Learning Representations, 2018. <https://openreview.net/forum?id=Hy79-UJPM>
- [122] X. Chen, Y. Duan, R. Houthoofd, J. Schulman, I. Sutskever, P. Abbeel, Infogan: interpretable representation learning by information maximizing generative adversarial nets, in: *Advances in Neural Information Processing Systems* 29, Curran Associates, Inc., 2016, pp. 2172–2180.
- [123] R. Zemel, Y. Wu, K. Swersky, T. Pitassi, C. Dwork, Learning fair representations, in: Proceedings of the International Conference on Machine Learning, 2013, pp. 325–333.
- [124] H. Edwards, A. Storkey, Censoring representations with an adversary, arXiv, Preprint arXiv:1511.05897.
- [125] D. Madras, E. Creager, T. Pitassi, R. Zemel, Learning adversarially fair and transferable representations, arXiv preprint, ArXiv:1802.06309.
- [126] E. Creager, D. Madras, J.-H. Jacobsen, M. Weis, K. Swersky, T. Pitassi, R. Zemel, Flexibly fair representation learning by disentanglement, in: Proceedings of the International Conference on Machine Learning, 2019, pp. 1436–1445.
- [127] Q.P. Nguyen, K.W. Lim, D.M. Divakaran, K.H. Low, M.C. Chan, Gee: a gradient-based explainable variational autoencoder for network anomaly detection, in:

Proceedings of the IEEE Conference on Communications and Network Security (CNS), IEEE, 2019, pp. 91–99, doi:[10.1109/CNS.2019.8802833](https://doi.org/10.1109/CNS.2019.8802833).

- [128] R. Al-Hmouz, W. Pedrycz, A. Balamash, A. Morfeq, Logic-driven autoencoders, Knowl.-Based Syst. 183 (2019) 104874, doi:[10.1016/j.knosys.2019.104874](https://doi.org/10.1016/j.knosys.2019.104874).



**David Charte** received his B.Sc. degrees in Mathematics and Computer Science (2017), and M.Sc. degree in Data Science and Computer Engineering (2018) from the University of Granada (Spain). He is currently a PhD student with the Department of Computer Science and Artificial Intelligence at this same university. He is coauthor of 5 JCR papers and several pieces of software for Data Science. His research interests include deep feature learning techniques and their applications, model interpretability and fairness.



**Francisco Charte** Tech. Eng. Computer Science (2008) and B. Eng. Computer Science (2010) from the Universidad de Jaén, with Extraordinary Award in both degrees and 1st National Award for Excellence in Academic Performance from the MEC (2010). M.Sc. Soft Computing and Computational Intelligence (2011) and PhD from the Universidad de Granada (2015). He is currently an Assistant Professor with the Computer Science Dept., Universidad de Jaén, Spain. He is the author of more than 120 books, including the title “Multilabel Classification. Problem analysis, metrics and techniques” published by Springer, as well as author of more than 20 JCR research papers and 25 contributions to international conferences. His main research

interests include multilabel learning, imbalanced and high-dimensionality problems and representation learning through deep learning techniques.



**Maria J. Del Jesus** received the M.Sc. and Ph.D. degrees in Computer Science from the University of Granada, Granada, Spain, in 1994 and 1999, respectively. She is a Professor with the Department of Computer Science at University of Jaen (Spain). She has been the supervisor of 8 Ph.D. students. She has published more than 65 journal papers and she is coauthor of the book “Multilabel Classification. Problem analysis, metrics and techniques” (Springer, 2016). Her current main research interests include fuzzy rule-based systems, subgroup discovery, emerging pattern mining, multilabel classification, data preparation, radial basis neural networks, knowledge extraction based on evolutionary algorithms, and data sci-

ence and big data.



**Francisco Herrera** (SM'15) received his M.Sc. in Mathematics in 1988 and Ph.D. in Mathematics in 1991, both from the University of Granada, Spain. He is a Professor in the Department of Computer Science and Artificial Intelligence at the University of Granada and Director of the Andalusian Research Institute in Data Science and Computational Intelligence (DaSCI). He is an academician in the Royal Academy of Engineering (Spain). He has been the supervisor of 48 Ph.D. students. He has published more than 500 journal papers, receiving more than 81,000 citations (Scholar Google, H-index 139). He has been selected as a Highly Cited Researcher <http://highlycited.com/> (in the fields of Computer Science

and Engineering, respectively, 2014 to present, Clarivate Analytics). He currently acts as Editor in Chief of the international journal “Information Fusion” (Elsevier). He acts as editorial member of a dozen of journals. He received the several honors and awards, among others: ECCAI Fellow 2009, IFSA Fellow 2013, 2010 Spanish National Award on Computer Science ARITMEL to the “Spanish Engineer on Computer Science”, International Cajastur “Mamdani” Prize for Soft Computing (Fourth Edition, 2010), IEEE Transactions on Fuzzy System Outstanding 2008 and 2012 Paper, 2011 Lotfi A. Zadeh Prize Best paper Award (IFSA Association), 2013 AEPIA Award to a scientific career in Artificial Intelligence, 2014 XV Andaluca Research Prize Maimonides, 2017 Andalucía Medal (by the regional government of Andalucía), 2018 “Granada: Science and Innovation City” award. His current research interests include among others, Computational Intelligence (including fuzzy modeling, computing with words, evolutionary algorithms and deep learning), information fusion and decision making, and data science (including data preprocessing, prediction, singular problems, and big data).