

Improving Semantic Embedding Consistency by Metric Learning for Zero-Shot Classification

Maxime Bucher, Stéphane Herbin, Frédéric Jurie

► **To cite this version:**

Maxime Bucher, Stéphane Herbin, Frédéric Jurie. Improving Semantic Embedding Consistency by Metric Learning for Zero-Shot Classification. ECCV 2016, Oct 2016, amsterdam, Netherlands. 2016, <<http://www.eccv2016.org/>>. <hal-01348827>

HAL Id: hal-01348827

<https://hal.archives-ouvertes.fr/hal-01348827>

Submitted on 25 Jul 2016

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

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

Improving Semantic Embedding Consistency by Metric Learning for Zero-Shot Classification

Maxime Bucher^{1,2}, Stéphane Herbin¹, Frédéric Jurie²

¹ONERA - The French Aerospace Lab, Palaiseau, France

²Normandie Univ, UNICAEN, ENSICAEN, CNRS

Abstract. This paper addresses the task of zero-shot image classification. The key contribution of the proposed approach is to control the semantic embedding of images – one of the main ingredients of zero-shot learning – by formulating it as a metric learning problem. The optimized empirical criterion associates two types of sub-task constraints: metric discriminating capacity and accurate attribute prediction. This results in a novel expression of zero-shot learning not requiring the notion of class in the training phase: only pairs of image/attributes, augmented with a consistency indicator, are given as ground truth. At test time, the learned model can predict the consistency of a test image with a given set of attributes, allowing flexible ways to produce recognition inferences. Despite its simplicity, the proposed approach gives state-of-the-art results on four challenging datasets used for zero-shot recognition evaluation.

Keywords: zero-shot learning, attributes, semantic embedding

1 Introduction

This paper addresses the question of zero-shot learning (ZSL) image classification, *i.e.*, the classification of images belonging to classes not represented by the training examples. This problem has attracted much interest in the last decade because of its clear practical impact: in many applications, having access to annotated data for the categories considered is often difficult, and requires new ways to increase the interpretation capacity of automated recognition systems. The efficiency of ZSL relies on the existence of an intermediate representation level, effortlessly understandable by human designers and sufficiently formal to be the support of algorithmic inferences. Most of the studies have so far considered this representation in the form of *semantic attributes* mainly because it provides an easy way to describe compact yet discriminative descriptions of new classes.

It has also been observed [1, 2] that attribute representations as provided by humans may not be the ideal embedding space because it can lack the informational quality necessary to conduct reliable inferences: the structure of the attribute manifold for a given data distribution may be rather complex, redundant, noisy and unevenly organized. Attribute descriptions, although semantically meaningful and useful to introduce a new category, are not necessarily isomorphic to image data or to image processing outputs.

To compensate for the shortcomings induced by attribute representations, the recent trends in ZSL studies are aiming at better controlling the classification inference as well as the attribute prediction in the learning criteria. Indeed, if attribute classifiers are learned independently of the final classification task, as in the Direct Attribute Prediction model [3], they might be optimal at predicting attributes but not necessarily at predicting novel classes.

In the work proposed in this paper, we instead suggest that better controlling the structure of the embedding attribute space is at least as important than constraining the classification inference step. The fundamental idea is to empirically disentangle the attribute distribution by learning a metric able to both select and transform the original data distribution according to informational criteria. This metric is obtained by optimizing an objective function based on pairs of attributes/images without assuming that the training images are assigned to categories; only the semantic annotations are used during training. More specifically, we empirically validated the idea that optimizing jointly the attribute embedding and the classification metric, in a multi-objective framework, is what makes the performance better, even with a simple linear embedding and distance to mean attribute classification.

The approach is experimentally validated on 4 recent datasets for zero-shot recognition, *i.e.* the ‘aPascal&aYahoo’, ‘Animals with Attributes’, ‘CUB-200-2011 and ‘SUN attribute’ datasets for which excellent results are obtained, despite the simplicity of the approach.

The rest of the paper is organized as follows: Section 2 presents the related works, Section 3 describes the proposed approach while the experimental validation is given in Section 4.

2 Related work

2.1 Visual features and semantic attributes

Image representation – *i.e.* the set of mechanisms allowing to transform image raw pixel intensities into representations suitable for recognition tasks – plays an important role in image classification. State-of-the art image representations were, a couple of years ago, mainly based on the pooling of hard/soft quantized local descriptors (*e.g.* SIFT [4]) through the bag-of-words [5] of Fisher vectors [6] models. However, the work of Krizhevsky *et al.* [7] has opened a new area and most of the state-of-the-art image descriptors nowadays rely on Deep Convolutional Neural Networks (CNN). We follow this trend in our experiments and use the so-called ‘VGG-VeryDeep-19’ (4096-dim) descriptors of [8].

Two recent papers have exhibited existing links between CNN features and semantic attributes. Ozeki *et al.* [9] showed that some CNN units can predict some semantic attributes of the ‘Animals with Attributes’ dataset fairly accurately. One interesting conclusion of their paper is that the visual semantic attributes can be predicted much more accurately than the non-visual ones by the nodes of the CNN. More recently, [10] showed the existence of Attribute

Centric Nodes (ACNs) within CNNs trained to recognize objects, collectively encoding information pertinent to visual attributes, unevenly and sparsely distributed across all the layers of the network.

Despite these recent findings could certainly make the performance of our method better, we don't use them in our experiments and stick to the use of standard CNN features, with the intention of making our results directly comparable to the recent zero-shot learning papers (e.g. [11]).

2.2 Describing images by semantic and non-semantic attributes

Zero-shot learning methods rely on the use of intermediate representations, usually given as *attributes*. This term can, however, encompass different concepts. For Lampert *et al.* [12] it denotes the presence/absence of a given object property, assuming that attributes are *nameable* properties (color or presence or absence of a certain part, *etc.*). The advantage of so-defined attributes is that they can be used easily to define new classes expressed by a shared semantic vocabulary.

However, finding a discriminative and meaningful set of attributes can sometimes be difficult. [13, 14] addressed this issue by proposing an interactive approach that discovers local attributes both discriminative and semantically meaningful, employing a recommender system that selects attributes through human interactions. An alternative for identifying attribute vocabulary without human labeling is to mine existing textual description of images sampled from the Internet, such as proposed by [15]. In the same line of thought, [16] presented a model for classifying unseen categories from their (already existing) textual description. [17] proposed an approach for zero-shot learning where the description of unseen categories comes in the form of typical text such as an encyclopedia entries, without the need to explicitly define attributes.

Another drawback of human generated attributes is that they can be redundant or not adapted to image classification. These issues have been addressed by automatically designing discriminative category-level attributes and using them for tasks of cross-category knowledge transfer, such as in the work of Yu *et al.* [18]. Finally, the attributes can also be structured into hierarchies [19–21] or obtained by text mining or textual descriptions [21, 16, 17].

Beside these papers which all consider attributes as meaningful for humans, some authors denoted by attributes any latent space providing an intermediate representation between image and hidden descriptions that can be used to transfer information to images from unseen classes. This is typically the case of [22] which jointly learn the attribute classifiers and the attribute vectors, with the intention of obtaining a better attribute-level representation, converting undetectable and redundant attributes into discriminative ones while retaining the useful semantic attributes. [23] has also introduced the concept of *discriminative attributes*, taking the form of random comparisons. The space of class names can also constitute an interesting embedding, such as in the works of [24, 25, 11, 26] which represent images as mixtures of known classes distributions.

Finally, it worth pointing out that the aforementioned techniques are restricting attributes to categorical labels and dont allow the representation of

more general semantic relationships. To counter this limitation, [27] proposed to model relative attributes by learning a ranking function.

2.3 Zero shot learning from semantic and attribute embedding

As defined by [28], the problem of zero-shot learning can be seen as the problem of learning a classifier $f : x \rightarrow y$ that can predict novel values of y not available in the training set. Most of existing methods rely on the computation of a similarity or *consistency function* linking image descriptors and the semantic description of the classes. These links are given by learning two embeddings – the first from the image representation to semantic space and the second from the class space to the semantic space – and defining a way to describe the constraints between the class space and the image space, the two being strongly interdependent.

DAP (Direct Attribute Prediction) and Indirect attribute prediction (IAP), first proposed by [3], use the between layer of attributes as variables decoupling the images from the layer of labels. In DAP, independent attribute predictors are used to build the embedding of the image, the similarity between two semantic representations (one predicted from the image representation, one given by the class) being given as the probability of the class attribute knowing the image. In IAP, attributes form a connecting layer between two layers of labels, one for classes that are known at training time and one for classes that are not known. In this case, attributes are predicted from (known) class predictions. Lampert *et al.* [3] concluded that DAP gives much better performance than IAP.

However, as mentioned in the introduction, DAP has two problems: first, it does not model any correlation between attributes, each being predicted independently. Second, the mapping between classes and the attribute space does not weight the relative importance of the attributes nor the correlations between them. Inspired by [29], which learns a linear embedding between image features and annotations, [30] and [2] tried to overcome this limitation. The work of Akata *et al.* [30] introduced a function measuring the consistency between an image and a label embedding, the parameters of this function is learned to ensure that, given an image, the correct classes rank higher than the incorrect ones. This consistency function has the form of a bilinear relation W associating the image embedding $\theta(x)$ and the label representation $\phi(y)$ as $S(x, y; W) = \theta(x)^t W \phi(y)$. Romera *et al.* [2] proposed a simple closed form solution for W , assuming a specific form of regularization is chosen. In comparison to our work none of the two papers ([30], [2]) use a metric learning framework to control the statistical structure of the attribute embedding space.

The coefficients of the consistency constraint W can be also predicted from a semantic textual description of the image. As an example, the goal of [17] is to predict a classifier for a new category based only on the learned classes and a textual description of this category. They solve this problem as a regression function, learnt from the textual feature domain to the visual classifier domain. [16] builds on these ideas, extending them by using a more expressive regression function based on a deep neural network. They take advantage of the architecture of CNNs and learn features at different layers, rather than just learning

an embedding space for both modalities. The proposed model provides means to automatically generate a list of pseudo-attributes for each visual category consisting of words from Wikipedia articles.

In contrast with the aforementioned methods, Hamm *et al.* [31] introduced the idea of ordinal similarity between classes (eg. $d(\text{'cat'}, \text{'dog'}) < d(\text{'cat'}, \text{'automobile'})$), claiming that not only this type of similarity may be sufficient for distinguishing cat and truck, but also that it seems a more natural representation since the ordinal similarity is invariant under scaling and monotonic transformation of numerical values. It is also worth mentioning the work of Jayaraman *et al.* [32] which proposed to leverage the statistics about each attribute error tendencies within a random forest approach, allowing to train zero-shot models that explicitly account for the unreliability of attribute predictions.

Wu *et al.* [33] exploit natural language processing technologies to generate event descriptions. They measure their similarity to images by projecting them into a common high-dimensional space using text expansion. The similarity is expressed as the concatenation of L_2 distances of the different modalities considered. Strictly speaking, there is no metric learning involved but a concatenation of L_2 distances. Finally, Frome *et al.* [34] aim at leveraging semantic knowledge learned in the text domain, and transfer it to a model trained for visual object recognition by learning a metric aligning the two modalities. However, in contrast to our work, [34, 33] do not explicitly control the quality of the embedding.

2.4 Zero-shot learning as transductive and semi-supervised learning

All the previously mentioned approaches consider that the embedding and the consistency function have to be learned from a set of training data of known classes, used in a second time to infer predictions about the images of new classes not available during training. However, a different problem can be addressed when images from the unknown classes are already available at training time and can hence be used to produce a better embedding. In this case, the problem can be cast as a transductive learning problem *i.e.* the inference of the correct labels for the given unlabeled data only, or a semi-supervised learning problem *i.e.* the inference of the best embedding using both labeled and unlabeled data.

Wang and Forsyth [35] proposed MIL framework for jointly learning attributes and object classifiers from weakly annotated data. Wang and Mori [22] treated attributes of an object as latent variables and captured the correlations among attributes using an undirected graphical model, allowing to infer object class labels using the information of both the test image and its (latent) attributes. In [1], the class information is incorporated into the attribute classifier to get an attribute-level representation that generalizes well to unseen examples of known classes as well as those of the unseen classes, assuming unlabeled images are available for learning. [36] considered the introduction of unseen classes as a novelty detection problem in a multi-class classification problem. If the image is of a known category, a standard classifier can be used. Otherwise, images are assigned to a class based on the likelihood of being an unseen category.

Fu *et al.* [37] rectified the projection domain shift between auxiliary and target datasets by introducing a multi-view semantic space alignment process to correlate different semantic views and the low-level feature view, by projecting them onto a latent embedding space learnt using multi-view Canonical Correlation Analysis. More recently Li *et al.* [38] learned the embedding from the input data in a semi-supervised large-margin learning framework, jointly considering multi-class classification over observed and unseen classes. Finally, [39] formulated a regularized sparse coding framework which used the target domain class label projections in the semantic space to regularize the learnt target domain projection, with the aim of overcoming the projection domain shift problem.

2.5 Zero-shot learning as a metric learning problem

Two contributions, [40] and [41], exploit metric learning to Zero shot class description. Mensink *et al.* [40] learn a metric adapted to measure the similarity of images, in the context of k -nearest neighbor image classification, and apply it in fact to One Shot Learning to show it can generalize well to new classes. They don't use any attribute embedding space nor consider ZSL in their work. Kuznetsova *et al.* [41] learn a metric to infer pose and object class from a single image. They use the expression *zero-shot* to actually denote a (new) transfer learning problem when data are unevenly sampled in the joint pose and class space, and not a Zero-Shot Classification problem where new classes are only known from attribute descriptions.

As far as we know, zero-shot learning has never been addressed explicitly as a metric learning problem in the attribute embedding space, which is one of the key contributions of this paper.

3 Method

3.1 Embedding consistency score

Most of the inference problems can be cast into an optimal framework of the form:

$$\mathbf{Y}^* = \arg \min_{\mathbf{Y} \in \mathcal{Y}} S(\mathbf{X}, \mathbf{Y})$$

where $\mathbf{X} \in \mathcal{X}$ is a given sample from some modality, *e.g.* an image or some features extracted from it, \mathbf{Y}^* is the most consistent association from another modality \mathcal{Y} , *e.g.* a vector of attribute indicators or a textual description, and S is a measure able to quantify the joint consistency of two observations from the two modalities. In this formulation, the smaller the score, the more consistent the samples. One can think of this score as a negative likelihood.

When trying to design such a consistency score, one of the difficult aspects is to relate meaningfully the two modalities. One usual approach consists in

embedding them into a common representational space \mathcal{A}^1 where their heterogeneous nature can be compared. This space can be abstract, *i.e.* its structure can be obtained from some optimization process, or semantically interpretable *e.g.* a fixed list of attributes or properties each indexed by a tag referring to some shared knowledge or ontology, leading to a p -dimensional vector space. Let $\hat{\mathbf{A}}_X(\mathbf{X})$ and $\hat{\mathbf{A}}_Y(\mathbf{Y})$ be the two embeddings for each modality X and Y , taking values in \mathcal{X} and \mathcal{Y} and producing outputs in \mathcal{A} .

In this work, it is proposed to define the consistency score as a metric on the common embedding space \mathcal{A} . More precisely, we use the Mahalanobis like description of a metric parametrized by a linear mapping \mathbf{W}_A :

$$d_A(\mathbf{A}_1, \mathbf{A}_2) = \left\| (\mathbf{A}_1 - \mathbf{A}_2)^T \mathbf{W}_A \right\|_2,$$

assuming that the embedding space is a vector space, and define the consistency score as:

$$S(\mathbf{X}, \mathbf{Y}) = d_A(\hat{\mathbf{A}}_X(\mathbf{X}), \hat{\mathbf{A}}_Y(\mathbf{Y})) = \left\| (\hat{\mathbf{A}}_X(\mathbf{X}) - \hat{\mathbf{A}}_Y(\mathbf{Y}))^T \mathbf{W}_A \right\|_2.$$

The Mahalanobis mapping \mathbf{W}_A can be interpreted itself as a linear embedding in an abstract m -dimensional vector space where the natural metric is the Euclidean distance, and acts as a multivariate whitening filter. It is expected that this property will improve empirically the reliability of the consistency score (1) by choosing the appropriate linear mapping.

We are now left with two questions: how to define the embedding? How build the Mahalanobis mapping? We see in the following that these two questions can be solved jointly by optimizing a unique criterion.

3.2 Embedding in the attribute space

The main problem addressed in this work is to be able to discriminate a series of new hypotheses that can only be specified using a single modality, the Y one with our notations. In many Zero-Shot Learning studies, this modality is often expressed as the existence or presence of several attributes or properties from a fixed given set. The simplest embedding space one can think of is precisely this attribute space, implying that the Y modality embedding is the identity: $\hat{\mathbf{A}}_Y(\mathbf{Y}) = \mathbf{Y}$ with $\mathcal{A} = \mathcal{Y}$. In this case, the consistency score simplifies as:

$$S(\mathbf{X}, \mathbf{Y}) = \left\| (\hat{\mathbf{A}}_X(\mathbf{X}) - \mathbf{Y})^T \mathbf{W}_A \right\|_2 \quad (1)$$

The next step is to embed the X modality into Y directly. We suggest using a simple linear embedding with matrix \mathbf{W}_X and bias \mathbf{b}_X , assuming that X is in a d -dimensional vector space. This can be expressed as:

$$\hat{\mathbf{A}}_X(\mathbf{X}) = \max(0, \mathbf{X}^T \mathbf{W}_X + \mathbf{b}_X). \quad (2)$$

¹ We use the letters A and \mathcal{A} in our notations since we will focus on the space of *attribute* descriptions as the embedding space.

We use a reLu-type output normalization to keep the significance of the attribute space as property detectors, negative numbers being difficult to interpret in this context.

In the simple formulation proposed here, we do not question the way new hypotheses are specified in the target modality, nor use any external source of information (*e.g.* word vectors) to map the attributes into a more semantically organized space such as in [36]. We leave the problem of correcting the original attribute description to the construction of the metric in the common embedding space.

3.3 Metric learning

The design problem is now reduced to the estimation of three mathematical objects: the linear embedding to the attribute space \mathbf{W}_X of dimensions $d \times p$, a bias \mathbf{b}_X of dimension p , and the Mahalanobis linear mapping \mathbf{W}_A of dimensions $p \times m$, m being a free parameter to choose.

The proposed approach consists in building empirically those objects from a set of examples by applying metric learning techniques. The training set is supposed to contain pairs of data $(\mathbf{X}_i, \mathbf{Y}_i)$ sampling the joint distribution of the two modalities: \mathbf{X}_i is a vector representing an image or some features extracted from it, while \mathbf{Y}_i denotes an attribute-based description. Notice that we do not introduce any class information in this formulation: the link between class and attribute representations is assumed to be specified by the use case considered.

The rationale behind the use of metric learning is to transform the original representational space so that the resulting metric takes into account the statistical structure of the data using pairwise constraints. One usual way to do so is to express the problem as a binary classification on pairs of samples, where the role of the metric is to separate similar and dissimilar samples by thresholding (see [42] for a survey on M.L.). It is easy to build pairs of similar and dissimilar examples from the annotated examples by sampling randomly (uniformly or according to some law) the two modalities \mathcal{X} and \mathcal{Y} and assigning an indicator $Z \in \{-1, 1\}$ stating whether \mathbf{Y}_i is a good attribute description of \mathbf{X}_i ($Z_i = 1$) or not ($Z_i = -1$). Metric learning approaches try to catch a data-dependent way to encode similarity. In general, the data manifold has a smaller intrinsic dimension than the feature space, and is not isotropically distributed.

We are now given a dataset of triplets $\{(\mathbf{X}_i, \mathbf{Y}_i, Z_i)\}_{i=1}^N$, the Z indicator stating that the two modalities are similar, *i.e.* consistent, or not². The next step is to describe an empirical criterion that will be able to learn \mathbf{W}_X , \mathbf{b}_X and \mathbf{W}_A . The idea is to decompose the problem in three objectives: metric learning, good embedding and regularization.

The metric learning part follows a now standard hinge loss approach [43] taking the following form for each sample:

$$l_H(\mathbf{X}_i, \mathbf{Y}_i, Z_i, \tau) = \max(0, 1 - Z_i(\tau - S(\mathbf{X}_i, \mathbf{Y}_i)^2)). \quad (3)$$

² To make notations simpler, we do not rename or re-index from the original dataset the pairs of data for the similar and dissimilar cases.

The extra parameter τ is free and can also be learned from data. Its role is to define the threshold separating similar from dissimilar examples, and should depend on the data distribution.

The embedding criterion is a simple quadratic loss, but only applied to similar data:

$$l_A(\mathbf{X}_i, \mathbf{Y}_i, Z) = \max(0, Z_i) \cdot \left\| \mathbf{Y}_i - \hat{\mathbf{A}}_X(\mathbf{X}_i) \right\|_2^2. \quad (4)$$

Its role is to ensure that the attribute prediction is of good quality, so that the difference $\mathbf{Y} - \hat{\mathbf{A}}_X(\mathbf{X})$ reflects dissimilarity due to modality inconsistencies rather than bad representational issues.

The size of the learning problem ($d \times p + p + p \times m$) can be large and requires regularization to prevent over fitting. We use a quadratic penalization:

$$R(\mathbf{W}_A, \mathbf{W}_X, \mathbf{b}_X) = \|\mathbf{W}_X\|_F^2 + \|\mathbf{b}_X\|_2^2 + \|\mathbf{W}_A\|_F^2 \quad (5)$$

where $\|\cdot\|_F$ is the Frobenius norm.

The overall optimization criterion can now be written as the sum of the previously defined terms:

$$\begin{aligned} \mathcal{L}(\mathbf{W}_A, \mathbf{W}_X, \mathbf{b}_X, \tau) = & \sum_i l_H(\mathbf{X}_i, \mathbf{Y}_i, Z_i, \tau) + \lambda \sum_i l_A(\mathbf{X}_i, \mathbf{Y}_i, Z_i) \\ & + \mu R(\mathbf{W}_A, \mathbf{W}_X, \mathbf{b}_X) \end{aligned} \quad (6)$$

where λ and μ are hyper-parameters that are chosen using cross-validation. Note that the criterion (6) can also be interpreted as a multi-objective learning approach since it mixes two optimal but dependent issues: attribute embedding and metric on the embedding space.

To solve the optimization problem, we do not follow the approach proposed in [43] since we also learn the attribute embedding part \mathbf{W}_X jointly with the metric embedding \mathbf{W}_A . We use instead a global stochastic gradient descent (see section 4 for details).

3.4 Application to image recognition and retrieval

The consistency score (1) is a versatile tool that can be used for several image interpretation problems. Section 4 will evaluate the potential of our approach on three of them.

Zero-shot learning The problem can be defined as finding the most consistent attribute description given the image to classify, and a set of exclusive attribute class descriptors $\{\mathbf{Y}_k^*\}_{k=1}^C$ where k is the index of a class:

$$k^* = \arg \min_{k \in \{1 \dots C\}} S(\mathbf{X}, \mathbf{Y}_k^*) \quad (7)$$

In this formulation, classifying is made equivalent to identifying between the C classes the best attribute description. A variant of this scheme can exploit a voting process to identify the best attribute among a set of k candidates, inspired from a k -nearest neighbor approach.

Few-shot learning Learning a metric in the embedding space can conveniently be used to specialize the consistency score to new data when they are available. We study a simple fine tuning approach using stochastic gradient descent on criterion (6) applied to novel triplets (X, Y, Z) from unseen classes only, starting with the model learned with seen classes. This makes “few-shot learning” possible. The decision framework is identical to the ZSL one.

Zero-shot retrieval The score (1) can also be used to retrieve the data from a given database that have at least a consistent level λ with a given query defined in the Y (or A) modality:

$$\text{Retrieve}(\mathbf{A}, \lambda) = \{\mathbf{X} \in \mathcal{X} / S(\mathbf{X}, \mathbf{A}) < \lambda\}$$

The performance is usually characterized by precision-recall curves.

4 Experiments

This section presents the experimental validation of the proposed method. The section first introduces the 4 datasets evaluated as well as the details of the experimental settings. The method is empirically evaluated on three different tasks as described in section 3: Zero-Shot-Learning (ZSL), Few-Shot Learning (FSL) and Zero-Shot Retrieval (ZSR). The ZSL experiments aim at evaluating the capability of the proposed model to predict unseen classes. This section also evaluates the contribution of the different components of the model to the performance, and makes comparisons with state-of-the-art results. In the FSL experiments, we show how the ZSL model can serve as good prior to learning a classifier when only a few samples of the unknown classes are available. Finally, we evaluate our model on a ZSR task, illustrating the capability of the algorithm to retrieve images using attribute-based queries.

4.1 Datasets and Experimental Settings

The experimental valuation is done on 4 public datasets widely used in the community, allowing to compare our results with those recently proposed in the literature: the aPascal&aYahoo (aP&Y) [23], Animals with Attributes (AwA) [3], CUB-200-2011 (CUB) [44] and SUN attribute (SUN) [45] datasets (see Table 1 for few statistics on their content). These datasets exhibit a large number of categories (indoor and outdoor scenes, objects, person, animals, *etc.*) and attributes (shapes, materials, color, parts, *etc.*)

These datasets have been introduced for training and evaluating ZSL methods and contain images annotated with semantic attributes. More specifically, each image of the aP&Y, CUB and SUN datasets has its own attribute description, meaning that two images of the same class can have different attributes. This is not the case for AwA where all the images of a given class share the same attributes. As a consequence, in the ZSL experiments on aP&Y, CUB and SUN,

Table 1: Dataset statistics

| Dataset | #Training classes | #Test classes | #Instances | #Attributes |
|-----------------------------|-------------------|---------------|------------|-------------|
| aPascal & aYahoo [23] | 20 | 12 | 15,339 | 64 |
| Animals with Attributes [3] | 40 | 10 | 30,475 | 85 |
| CUB 200-2011 [44] | 150 | 50 | 11,788 | 312 |
| SUN Attributes [45] | 707 | 10 | 14,340 | 102 |

the attribute representation of unknown classes, required for class prediction, is taken as their mean attribute frequencies.

In order to make comparisons with previous works possible, we use the same training/testing splits as [23] (aP&Y), [3] (AwA), [21] CUB and [32] (SUN).

Regarding the representation of images, we used both the VGG-VeryDeep-19 [8] and AlexNet [7] CNN models, both pre-trained on imageNet – without fine tuning to the attribute datasets – and use the penultimate fully connected layer (*e.g.*, FC7 4096-d layer for VGG-VeryDeep-19) for representing the images. Very deep CNN models act as generic feature extractors and have been demonstrated to work well for object recognition. They have been also used in many recent ZSL experiments and we use exactly the same descriptors as [21, 11].

One of the key characteristics of our model is that it requires a set of image/attributes pairs for training. Positive (resp. negative) pairs are obtained by taking the training images associated with their own provided attribute vector (resp. by randomly assigning attributes not present in the image) and are assigned to the class label ‘1 (resp. ‘-1). In order to bound the size of the training set we generate only 2 pairs per training image, one positive and one negative.

Our model has three hyper-parameters: the weight λ , the dimensionality of the space in which the distance is computed (m) and the regularization parameters μ . These hyper-parameters are estimated through a grid search validation procedure by randomly keeping 20% of the training classes for cross-validating the hyper-parameters, and choosing the parameters giving best accuracy for these so-obtained validation classes. The parameter are searched in the following ranges: $m \in [20\%, 120\%]$ of the initial attribute dimension, $\lambda \in [0.05, 1.0]$ and $\mu \in [0.01, 10.0]$. τ is a parameter learned during training.

Once the hyper-parameters are tuned, we take the whole training set to learn the final model and evaluate it on the test set (unseen classes in case of ZSL).

The optimization of \mathbf{W}_A and \mathbf{W}_X is done with stochastic gradient descent, the parameters being initialized randomly with normal distribution. The size of the mini-batch is of 100. As the objective function is non-convex, different initializations can give different parameters. We addressed this issue by doing 5 estimations of the parameters starting from 5 different initializations and selecting the best one on a validation set (we keep a part of the train set for doing this and fine-tune the parameters on the whole train set when the best initialization is known). We use the optimizer provided in the TensorFlow framework [46]. Using the GPU mode with a Nvidia 750 GTX GPU, learning a model (\mathbf{W}_A

Table 2: Zero-shot classification accuracy (mean \pm std). We report results both with VGG-verydeep-19 [8] and AlexNet [7] features for fair comparisons, whenever it’s possible (AwA images are not public anymore preventing the computation of their AlexNet representations).

| Feat. | Method | aP&Y | AwA | CUB | SUN |
|------------------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|----------------------------------|
| AlexNet [7] | Akata <i>et al.</i> [21] | - | 61.9 | 40.3 | - |
| | Ours | 46.14 \pm 0.91 | - | 41.98\pm0.67 | 75.48 \pm 0.43 |
| VGG-VeryDeep [8] | Lampert <i>et al.</i> [12] | 38.16 | 57.23 | - | 72.00 |
| | Romera-Paredes <i>et al.</i> [2] | 24.22 \pm 2.89 | 75.32 \pm 2.28 | - | 82.10 \pm 0.32 |
| | Zhang <i>et al.</i> [11] | 46.23 \pm 0.53 | 76.33 \pm 0.83 | 30.41 \pm 0.20 | 82.50 \pm 1.32 |
| | Zhang <i>et al.</i> [47] | 50.35 \pm 2.97 | 80.46\pm0.53 | 42.11 \pm 0.55 | 83.83 \pm 0.29 |
| | Ours w/o ML | 47.25 \pm 0.48 | 73.81 \pm 0.13 | 33.87 \pm 0.98 | 74.91 \pm 0.12 |
| | Ours w/o constraint | 48.47 \pm 1.24 | 75.69 \pm 0.56 | 38.35 \pm 0.49 | 79.21 \pm 0.87 |
| | Ours | 53.15\pm0.88 | 77.32 \pm 1.03 | 43.29\pm0.38 | 84.41\pm0.71 |

and \mathbf{W}_X) takes 5-10 minutes for a given set of hyper-parameters. Computing image/attribute consistency takes around 4ms per pair.

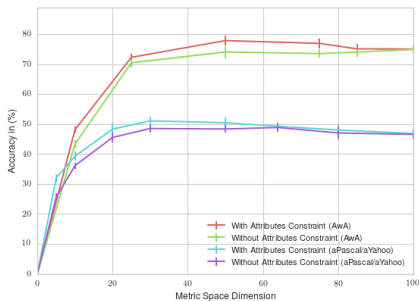
4.2 Zero-Shot Learning Experiments

The experiments follow the standard ZSL protocol: during training, a set of images from known classes is available for learning the model parameters. At test time, images from unseen classes are processed and the goal is to find the class described by an attribute representation most consistent with the images.

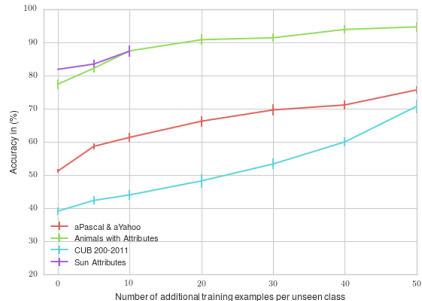
Table 2 gives the performance of our approach on the 4 datasets considered – expressed as multi-class accuracy – and makes comparisons with state-of-the-art approaches. The performances of previous methods are taken from [11, 21, 47]. Performance is reported with 2 different features *i.e.* VGG-VeryDeep-19 [8] and AlexNet [7] for fair comparisons. As images of AwA are not public anymore it is only possible to use the features available for download. On the four datasets our model achieves above state-of-the-art performance (note : [47] was published after our submission), with a noticeable improvement of more than 8% on aP&Y.

As explained in the previous section, our model is based on a multi-objective function trying to maximize metric discriminating capacity as well as attribute prediction. It is interesting to observe how the performance degrades when one of the two terms is missing. In Table 2, the ‘Ours w/o ML’ setting makes use of the Euclidean distance *i.e.* $\mathbf{W}_A = \mathbb{I}$. The ‘Ours w/o constraint’ setting is when the attribute prediction term (Eq. 4) is missing in the criterion. This term gives a 4% improvement, on average.

Figure 1a shows the accuracy as a function of the embedding dimension. This projection maps the original data in a space in which the Euclidean distance is good for the task considered. It can be seen as a way to exploit and select the correlation structure between attributes. We experimented that the best



(a) Classification accuracy



(b) Image retrieval

Fig. 1: (a) ZSL accuracy as a function of the dimensionality of the metric space; best results are obtained when the dimension of the metric embedding is less than 40% of the image space dimension. It also shows the improvement due to the attribute prediction term in the objective function. (b) Few-shot learning: Classification accuracy (%) as a function of the amount of training examples per unseen classes.

performance is generally obtained when the dimension of this space less than 40% smaller than the size of the initial attribute space.

4.3 Few-Shot Learning

Few-shot learning corresponds to the situation where 1 (or more) annotated example(s) from unseen classes are available at test time. In this case our model is first trained using only the seen classes (same as with ZSL), and we introduced the examples from unseen class data one by one before fine-tuning the model parameters by doing a few more learning iterations using these new data only.

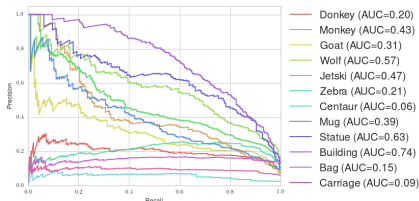
Figure 1b shows the accuracy evolution, given as a function of the number of additional images from the unseen classes. Please note that for the SUN dataset we have used a maximum of 10 additional examples as unseen classes contain only 20 images. We observed that knowing even a very few number of annotated examples significantly improves the performance. It is a very encouraging behavior for large-scale applications where annotations for a large number of categories are hard and expensive to get.

4.4 Zero-Shot Retrieval

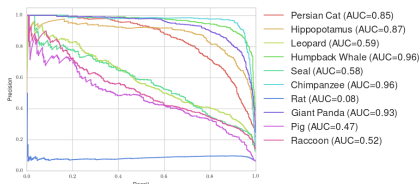
The task of Zero-Shot image Retrieval consists in searching an image database with attribute-based queries. For doing this, we first train our model as for standard ZSL. We then take the attribute descriptions of unseen classes as queries, and rank the images from the unseen classes based on the similarity with the

Table 3: Zero-Shot Retrieval task: Mean Average Precision (%) on the 4 datasets

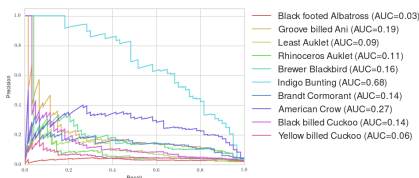
| | aP&Y | AwA | CUB | SUN | Av. |
|--------------------------|--------------|-------------|--------------|--------------|--------------|
| Zhang <i>et al.</i> [11] | 15.43 | 46.25 | 4.69 | 58.94 | 31.33 |
| Ours (VGG features) | 36.92 | 68.1 | 25.33 | 52.68 | 45.76 |



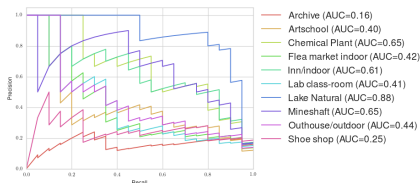
(a) aP&Y



(b) AwA



(c) CUB-200-2011



(d) SUN

Fig. 2: Precision Recall curve for each unseen class by dataset. For CUB dataset we randomly choose 10 classes (best viewed on a computer screen).

query. Table 3 reports the mean average precision on the 4 datasets. Our model outperforms the state-of-the-art SEE method [11] by more than 10% on average.

Figure 2 shows the average precision for each class of the 4 datasets. In the aP&aYa dataset, the ‘donkey’, ‘centaur’ and ‘zebra’ classes have a very low average precision. This can be explained by the strong visual similarity between these classes which only differ by a few attributes.

5 Conclusion

This paper has presented a novel approach for zero-shot classification exploiting multi-objective metric learning techniques. The proposed formulation has the nice property of not requiring any ground truth at the category level for learning a consistency score between the image and the semantic modalities, but only requiring weak consistency information. The resulting score can be used with versatility on various image interpretation tasks, and shows close or above state-of-the-art performance on four standard benchmarks. The formal simplicity of the approach allows several avenues for future improvement. A first one would be to provide a better embedding on the semantic side of the consistency score $\hat{\mathbf{A}}_Y(\mathbf{Y})$. A second one would be to explore more complex functions than the linear mappings tested in this work, and introduce, for instance, deep network architectures.

References

1. Mahajan, D.K., Sellamanickam, S., Nair, V.: A joint learning framework for attribute models and object descriptions. In: IEEE International Conference on Computer Vision (ICCV). (2011)
2. Romera-Paredes, B., Torr, P.H.: An embarrassingly simple approach to zero-shot learning. In: Proceedings of the International Conference on Machine Learning. (2015) 2152–2161
3. Lampert, C.H., Nickisch, H., Harmeling, S.: Learning to detect unseen object classes by between-class attribute transfer. In: IEEE International Conference on Computer Vision and Pattern Recognition (CVPR). (2009)
4. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision (IJCV)* **60**(2) (2004) 91–110
5. Csurka, G., Dance, C., Fan, L., Willamowski, J., Bray, C.: Visual categorization with bags of keypoints. In: Workshop on statistical learning in computer vision, ECCV. (2004) 1–2
6. Sánchez, J., Perronnin, F., Mensink, T., Verbeek, J.: Image classification with the Fisher vector: Theory and practice. *International Journal of Computer Vision (IJCV)* **105**(3) (2013) 222–245
7. Krizhevsky, A., Sutskever, I., Hinton, G.E.: ImageNet Classification with Deep Convolutional Neural Networks. In: Conference on Neural Information Processing Systems (NIPS). (2012) 1106–1114
8. Simonyan, K., Zisserman, A.: Very Deep Convolutional Networks for Large-Scale Image Recognition. In: ICLR. (2014)
9. Ozeki, M., Okatani, T.: Understanding Convolutional Neural Networks in Terms of Category-Level Attributes. In: Asian Conference on Computer Vision (ACCV). (2014)
10. Escorcia, V., Niebles, J.C., Ghanem, B.: On the Relationship between Visual Attributes and Convolutional Networks. In: IEEE International Conference on Computer Vision and Pattern Recognition (CVPR). (2015)
11. Zhang, Z., Saligrama, V.: Zero-Shot Learning via Semantic Similarity Embedding. In: IEEE International Conference on Computer Vision (ICCV). (2015)
12. Lampert, C.H., Nickisch, H., Harmeling, S.: Attribute-Based Classification for Zero-Shot Visual Object Categorization. *IEEE Trans. on Pattern Analysis and Machine Intelligence* **36**(3) (2014) 453–465
13. Parikh, D., Grauman, K.: Interactively building a discriminative vocabulary of nameable attributes. In: IEEE International Conference on Computer Vision and Pattern Recognition (CVPR). (2011)
14. Duan, K., Parikh, D., Crandall, D., Grauman, K.: Discovering localized attributes for fine-grained recognition. In: IEEE International Conference on Computer Vision and Pattern Recognition (CVPR). (2012)
15. Berg, T.L., Berg, A.C., Shih, J.: Automatic attribute discovery and characterization from noisy web data. In: European Conference on Computer Vision (ECCV). (2010)
16. Ba, L.J., Swersky, K., Fidler, S., Salakhutdinov, R.: Predicting deep zero-shot convolutional neural networks using textual descriptions. In: 2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015. (2015) 4247–4255
17. Elhoseiny, M., Saleh, B., Elgammal, A.: Write a Classifier: Zero-Shot Learning Using Purely Textual Descriptions. In: IEEE International Conference on Computer Vision (ICCV). (2013)

18. Yu, F.X., Cao, L., Feris, R.S., Smith, J.R., Chang, S.F.F.: Designing category-level attributes for discriminative visual recognition. In: IEEE International Conference on Computer Vision (ICCV), IEEE (2013)
19. Verma, N., Mahajan, D., Sellamanickam, S., Nair, V.: Learning hierarchical similarity metrics. In: IEEE International Conference on Computer Vision and Pattern Recognition (CVPR). (2012)
20. Rohrbach, M., Stark, M., Schiele, B.: Evaluating knowledge transfer and zero-shot learning in a large-scale setting. In: IEEE International Conference on Computer Vision and Pattern Recognition (CVPR). (2011)
21. Akata, Z., Reed, S., Walter, D., Lee, H., Schiele, B.: Evaluation of Output Embeddings for Fine-Grained Image Classification. In: IEEE International Conference on Computer Vision and Pattern Recognition (CVPR). (2015)
22. Wang, Y., Mori, G.: A Discriminative Latent Model of Object Classes and Attributes. In: European Conference on Computer Vision (ECCV). (2010)
23. Farhadi, A., Endres, I., Hoiem, D., Forsyth, D.: Describing objects by their attributes. In: IEEE International Conference on Computer Vision and Pattern Recognition (CVPR). (2009)
24. Mensink, T., Gavves, E., Snoek, C.G.M.: COSTA: Co-Occurrence Statistics for Zero-Shot Classification. In: IEEE International Conference on Computer Vision and Pattern Recognition (CVPR). (2014)
25. Norouzi, M., Mikolov, T., Bengio, S., Singer, Y., Shlens, J., Frome, A., Corrado, G.S., Dean, J.: Zero-Shot Learning by Convex Combination of Semantic Embeddings. In: International Conference on Learning Representations (ICLR). (December 2013)
26. Fu, Z., Xiang, T.A., Kodirov, E., Gong, S.: Zero-shot object recognition by semantic manifold distance. In: IEEE International Conference on Computer Vision and Pattern Recognition (CVPR). (2015)
27. Parikh, D., Grauman, K.: Relative attributes. In: IEEE International Conference on Computer Vision (ICCV). (2011)
28. Palatucci, M., Pomerleau, D., Hinton, G.E., Mitchell, T.M.: Zero-shot learning with semantic output codes. In: Conference on Neural Information Processing Systems (NIPS). (2009)
29. Weston, J., Bengio, S., Usunier, N.: WSABIE: scaling up to large vocabulary image annotation. In: IJCAI. (2011) 2764–2770
30. Akata, Z., Perronin, F., Harchaoui, Z., Schmid, C.: Label-Embedding for Image Classification. *IEEE Trans. on Pattern Analysis and Machine Intelligence* (2015)
31. Hamm, J., Belkin, M.: Probabilistic Zero-shot Classification with Semantic Rankings. arXiv.org (February 2015)
32. Jayaraman, D., Grauman, K.: Zero-shot recognition with unreliable attributes. In: Conference on Neural Information Processing Systems (NIPS). (2014)
33. Wu, S., Bondugula, S., Luisier, F., Zhuang, X., Natarajan, P.: Zero-Shot Event Detection Using Multi-modal Fusion of Weakly Supervised Concepts. In: IEEE International Conference on Computer Vision and Pattern Recognition (CVPR). (2014)
34. Frome, A., Corrado, G.S., Shlens, J., Bengio, S., Dean, J., Ranzato, M., Mikolov, T.: DeViSE: A Deep Visual-Semantic Embedding Model. In: Conference on Neural Information Processing Systems (NIPS). (2013)
35. Wang, G., Forsyth, D.: Joint learning of visual attributes, object classes and visual saliency. In: IEEE International Conference on Computer Vision (ICCV). (2009)

36. Socher, R., Ganjoo, M., Manning, C.D., Ng, A.: Zero-Shot Learning Through Cross-Modal Transfer. In: Conference on Neural Information Processing Systems (NIPS). (2013)
37. Fu, Y., Hospedales, T.M., Xiang, T., Fu, Z., Gong, S.: Transductive multi-view embedding for zero-shot recognition and annotation. In: European Conference on Computer Vision (ECCV). (2014)
38. Li, X., Guo, Y., Schuurmans, D.: Semi-Supervised Zero-Shot Classification With Label Representation Learning. In: IEEE International Conference on Computer Vision (ICCV). (2015)
39. Kodirov, E., Xiang, T., Fu, Z., Gong, S.: Unsupervised Domain Adaptation for Zero-Shot Learning. In: IEEE International Conference on Computer Vision (ICCV). (2015)
40. Mensink, T., Verbeek, J., Perronnin, F., Csurka, G.: Metric learning for large scale image classification: Generalizing to new classes at near-zero cost. In: Computer Vision—ECCV 2012. Springer (2012) 488–501
41. Kuznetsova, A., Hwang, S.J., Rosenhahn, B., Sigal, L.: Exploiting view-specific appearance similarities across classes for zero-shot pose prediction: A metric learning approach. In: Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA. (2016) 3523–3529
42. Bellet, A., Habrard, A., Sebban, M.: A Survey on Metric Learning for Feature Vectors and Structured Data. Technical Report arXiv:1306.6709v4, University of St Etienne (2013)
43. Shalev-Shwartz, S., Singer, Y., Ng, A.Y.: Online and batch learning of pseudo-metrics. In: Proceedings of the International Conference on Machine learning, ACM (2004) 94
44. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The Caltech-UCSD Birds-200-2011 Dataset. Technical report (July 2011)
45. Patterson, G., Xu, C., Su, H., Hays, J.: The SUN Attribute Database: Beyond Categories for Deeper Scene Understanding. *International Journal of Computer Vision (IJCV)* **108**(1-2) (2014) 59–81
46. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G.S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., Zheng, X.: TensorFlow: Large-scale machine learning on heterogeneous systems (2015) Software available from tensorflow.org.
47. Zhang, Z., Saligrama, V.: Zero-shot learning via joint latent similarity embedding. In: IEEE International Conference on Computer Vision and Pattern Recognition (CVPR). (2016) 6034–6042