

Less is more: zero-shot learning from online textual documents with noise suppression*

Ruizhi Qiao, Lingqiao Liu, Chunhua Shen, Anton van den Hengel
School of Computer Science, The University of Adelaide, Australia

{ruizhi.qiao, lingqiao.liu, chunhua.shen, anton.vandenhengel}@adelaide.edu.au

Abstract

Classifying a visual concept merely from its associated online textual source, such as a Wikipedia article, is an attractive research topic in zero-shot learning because it alleviates the burden of manually collecting semantic attributes. Recent work has pursued this approach by exploring various ways of connecting the visual and text domains. In this paper, we revisit this idea by going further to consider one important factor: the textual representation is usually too noisy for the zero-shot learning application. This observation motivates us to design a simple yet effective zero-shot learning method that is capable of suppressing noise in the text.

Specifically, we propose an $l_{2,1}$ -norm based objective function which can simultaneously suppress the noisy signal in the text and learn a function to match the text document and visual features. We also develop an optimization algorithm to efficiently solve the resulting problem. By conducting experiments on two large datasets, we demonstrate that the proposed method significantly outperforms those competing methods which rely on online information sources but with no explicit noise suppression. Furthermore, we make an in-depth analysis of the proposed method and provide insight as to what kind of information in documents is useful for zero-shot learning.

1. Introduction

Unlike traditional object classification tasks in which the training and test categories are identical, zero-shot learning aims to recognize objects from classes not seen at the training stage. It is recognized as an effective way for large scale visual classification since it alleviates the burden of collecting sufficient training data for every possible class. The key component ensuring the success of zero-shot learning is to find an intermediate semantic representation to bridge the

gap between seen and unseen classes. In a nutshell, with this semantic representation we can first learn its connection with image features and then transfer this connection to unseen classes. So once the semantic representation of an unseen class is given, one can easily classify the image through the learned connection.

Attributes, which essentially represent the discriminative properties shared among both seen and unseen categories, have become the most popular semantic representation in zero-shot learning [7, 8, 24, 12, 26]. Although the recent use of attributes has led to exciting advances in zero-shot learning [10, 2, 27], the creation of attributes still relies on much human labour. This is inevitably discouraging since the motivation for zero-shot learning is to free large-scale recognition tasks from cumbersome annotation requirements.

To remedy this drawback and move towards the goal of fully automatic zero-shot learning, several recent works [23, 9, 16] have explored the possibility of using the easily accessed online information sources to create the intermediate semantic representation. One possible choice is to directly use online textual documents, e.g., those found in Wikipedia, to build such a representation [6, 3]. This is promising because online text documents can be easily obtained and contain rich information about the object. To conduct zero-shot learning with textual documents, existing works [2, 10] develop various ways to measure the similarity between text and visual features. Our work is also based on this idea. We take a step further, however, to consider one additional important factor: the document representation is much more noisy than the human specified semantic representation and negligence of this fact would inevitably lead to inferior performance. For example, when the bag-of-words model is adopted as the document representation, the occurrence of every word in a document will trigger a signal in one dimension of the document representation. However, it is clear that most words in a document are not directly relevant for identifying the object category. Thus it is necessary to design a noise suppression mechanism to down weight the importance of those less relevant words for zero-shot learning. This mechanism is closely related to

*The first two authors contributed to this work equally. Correspondence should be addressed to C. Shen.

feature selection. However, it is not exactly the same. As will be discussed in the following sections, the solution of our method does not discard the less relevant dimensions of the document representation but only suppress their impact for zero-shot learning.

To this end, we propose a zero-shot learning method which particularly caters for the need for noise suppression. More specifically, we proposed a simple yet effective $l_{2,1}$ -norm based objective function which simultaneously suppresses the noisy signal within text descriptions and learns a function to match the visual and text domains. Furthermore, we develop an efficient optimization algorithm to solve this problem. By conducting experiments on two large scale zero-shot learning evaluation benchmarks, we demonstrate the benefit of the proposed noise suppression mechanism as well as its superior performance over other zero-shot learning methods which also rely on online information sources. In addition, we also conduct an in-depth analysis of the proposed method which provides an insight as to what kinds of information within a document are useful for zero-shot learning.

2. Related work

Most zero-shot learning approaches rely on human specified attributes. As one of the earliest attempts in zero-shot learning, Lampert *et al.* [12] adopted a set of attributes obtained from a psychology study. By learning probabilistic predictors of those attributes, they developed a framework to estimate the posterior of the test class. Later, a number of works has been proposed to improve the way of learning the connection between attributes and object categories. For example, the work in [11] addresses unreliability of attributes by exploring the idea of random forest. The work in [1] converted the zero-shot learning into a cross-domain matching problem and they proposed to learn a matching function to compare the attribute and the image feature. Built upon this idea, Romera-Paredes and Torr [20] proposed a simpler but more effective objective function to learn the matching function. Zhang and Saligrama [27] advocated the benefits of using attribute-attribute relationships, termed semantic similarity, as the intermediate semantic representation and they learn a function to match the image features with the semantic similarity.

To go beyond the human specified attributes, recent works also explore the use of other forms of semantic representations which can be easily obtained [14, 2, 9, 10]. For example, the co-occurrence statistics of words has been explored in [14, 2] to capture the semantic relevance of two concepts. The distributed word representation, e.g., word2vec, has been utilized as a substitution of attributes [9] and more recently the word2vec representation has been shown to be complementary to the human specified attributes [10].

The other information source for creating the semantic representation is the online textual document, such as Wikipedia articles. In an earlier work, Berg *et al.* [4] attempt to discover attribute representation from a noisy web source by ranking the visual-ness scores of attribute candidates. Rohrbach *et al.* [18, 19] mine semantic relatedness for attribute-class association from different internet sources. More recent works [6, 3] directly learn a function to measure the compatibility between documents and visual features. However, compared with the state-of-the-art zero-shot learning methods, their performance seems to be disappointing even though some advanced technologies, such as deep learning, has been applied [3].

3. Our approach

3.1. Overview

The overview of our method is depicted in Figure 1. It starts with a raw document representation which is simply a binarized histogram of words. This representation is fed into our zero-shot learning algorithm to generate a classifier to detect relevant images. In the process of generating this classifier, the noise suppression regularizer in our method will automatically suppress the impact of less relevant words (illustrated as the red words in Figure 1).

3.2. Text representation

We extract our text representation based on a simple bag-of-words model. We start by a preprocessing step of tokenizing the words and removing stop words and punctuations. Then a histogram of the remaining word occurrences is calculated and is subsequently binarized as the text representation. In other words, once a word appears in a document, its corresponding dimension within the text representation is set to “1”. One more commonly used choice for the text representation is based on TF-IDF as in [6, 3]. However, we find that it produces worse performance¹ than directly using the binarized representation. This is probably because the weighting calculated of TF-IDF is not suitable for our zero-shot learning although it is considered to be less noisy for applications like document classification. In the binarized histogram we essentially treat each word in a document equally and this inevitably introduces a lot of noisy signals. However, thanks to our noise suppressing zero-shot learning algorithm, we can substantially down-weight the less relevant words and achieve good performance even with a noisy document representation.

3.3. Learning to match text and visual features

We first formally define our problem and introduce the notation used in the following sections. At the training

¹Using TF-IDF is about 7% and 5% inferior to binarized representations on AwA and CUB, respectively.

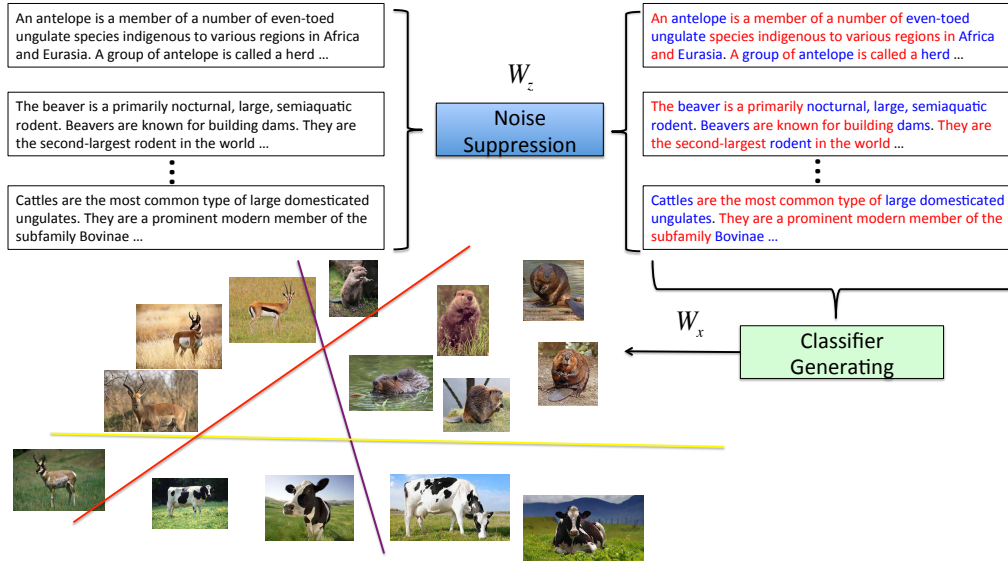


Figure 1. Overview of our zero-shot learning approach. The text representations are processed by the noise suppression mechanism to generate a classifier to detect relevant images and the noisy components of text representations are suppressed to gain better performance.

stage, both image features and document descriptions for C seen categories are available. Let $\mathbf{X} \in \mathbb{R}^{d \times N}$ denote the image features of N training examples and $\mathbf{Z} \in \{0, 1\}^{\hat{d} \times C}$ the aforementioned document representations for C seen classes, where \hat{d} and d are the dimensionality of the document representation and the image features respectively. We also define $\mathbf{Y} \in \{0, 1\}^{N \times C}$ as the indicator matrix for the C seen classes. Each row of \mathbf{Y} has a unique “1” indicating its corresponding class label. At the test stage, the document representations of the \hat{C} unseen classes are given and our task is to assign \hat{C} unseen class labels to the test images.

3.4. Formulation

Our method is inspired by a recently proposed zero-shot learning approach [20] which has demonstrated impressive performance despite a very simple learning process. More specifically, it learns a matrix \mathbf{V} which optimizes the following objective function.

$$\min_{\mathbf{V}} \|\mathbf{X}^T \mathbf{V} \mathbf{S} - \mathbf{Y}\|_F^2 + \lambda \|\mathbf{V} \mathbf{S}\|_F^2 + \gamma \|\mathbf{X}^T \mathbf{V}\|_F^2 + \lambda \gamma \|\mathbf{V}\|_F^2 \quad (1)$$

where \mathbf{S} denotes the semantic attribute matrix and it can be either a binary matrix or a real value matrix. The scalars γ and λ are weights controlling the prominence of the various terms. The underlying idea of this algorithm can be understood as follows. If the task is to classify \mathbf{X} into C categories, we can simply learn a linear classifier by fitting to \mathbf{Y} , that is, $\min_{\mathbf{W}} \|\mathbf{X}^T \mathbf{W} - \mathbf{Y}\|_F^2$. However, in this case \mathbf{W} cannot be transferred to unseen classes. Thus we fur-

ther impose that $\mathbf{W} = \mathbf{V} \mathbf{S}$. In other words, the classifier of a class is generated from its attributes. With this requirement, the classifier of an unseen class can be easily obtained and utilized to predict the category of a test image. Similarly, we can also treat $\mathbf{X}^T \mathbf{V}$ as the classifier operated on the attributes \mathbf{S} . The above understanding naturally gives rise to the regularization terms $\lambda \|\mathbf{V} \mathbf{S}\|_F^2$ and $\gamma \|\mathbf{X}^T \mathbf{V}\|_F^2$, which play the same role of the Frobenius norm regularizer as commonly introduced in multi-class classification or regression.

Since our document representation can also be seen as an attribute vector, the method in [20] can be readily applied to our problem by simply setting $\mathbf{S} = \mathbf{Z}$. However, this naive solution ignores an important fact that the document representation is much more noisy than the human specified attribute vectors. To handle this issue, we introduce a noise suppression mechanism into Eq. (1). More specifically, we first decompose \mathbf{V} into two terms:

$$\mathbf{V} = \mathbf{W}_x^T \mathbf{W}_z, \quad (2)$$

where $\mathbf{W}_x \in \mathbb{R}^{m \times d}$ and $\mathbf{W}_z \in \mathbb{R}^{m \times \hat{d}}$. These two matrices will play different roles in our method. \mathbf{W}_z is used to suppress the noisy components of \mathbf{Z} and transform \mathbf{Z} into a $m \times C$ intermediate representation. \mathbf{W}_x is used to generate the image classifier from the noise-suppressed intermediate representation. Thus, two different regularization terms are imposed to suit these two different roles. The first term is the $l_{2,1}$ -norm of \mathbf{W}_z^T which achieves the noise suppression effect. The second term is the Frobenius norm of $\mathbf{W}_x^T \mathbf{W}_z \mathbf{Z}$ which is similar to the $\lambda \|\mathbf{V} \mathbf{S}\|_F^2$ term in Eq. (1).

The formulation of our method is expressed as follows:

$$\min_{\mathbf{W}_x, \mathbf{W}_z} L(\mathbf{W}_x, \mathbf{W}_z) + \lambda_1 \|\mathbf{W}_x^\top \mathbf{W}_z \mathbf{Z}\|_F^2 + \lambda_2 \|\mathbf{W}_z^\top\|_{2,1}, \quad (3)$$

$$L(\mathbf{W}_x, \mathbf{W}_z) = \|\mathbf{X}^\top \mathbf{W}_x^\top \mathbf{W}_z \mathbf{Z} - \mathbf{Y}\|_F^2.$$

The $l_{2,1}$ -norm is defined as $\|\mathbf{W}_z^\top\|_{2,1} = \sum_{i=1}^{\hat{d}} \|\mathbf{w}_z^i\|_2$, where \mathbf{w}_z^i denotes the i -th column of \mathbf{W}_z . It is known that the $l_{2,1}$ -norm will encourage the column vectors of \mathbf{W}_z to have few large values, which means that the impact of noisy dimensions of \mathbf{Z} will be substantially suppressed or even completely eliminated. In fact, if λ_2 becomes sufficient large, it achieves the effect of feature selection on the document representation. However, by cross-validating λ_1 and λ_2 , our method does not lead to an exactly sparse solution as it seems that the algorithm prefers to keep the majority of the dimensions in \mathbf{Z} for zero-shot learning. This is probably due to the joint regularization effect of $\|\mathbf{W}_x^\top \mathbf{W}_z \mathbf{Z}\|_F^2$ or the fact that dimensions corresponding to lower values of $\|\mathbf{w}_z^i\|_2$ are still useful for zero-shot learning. Therefore we consider the use of the $l_{2,1}$ -norm here as a noise suppression mechanism rather than a feature selection mechanism. We drop out the other regularization terms in Eq. (1) since we find them have little impact on performance.

Similar to [20], once \mathbf{V} , in our case $\mathbf{V} = \mathbf{W}_x^\top \mathbf{W}_z$, is learned, we can infer the class label of a test image \mathbf{x} using the following rule:

$$c^* = \max_c \mathbf{x}^\top \mathbf{W}_x^\top \mathbf{W}_z \mathbf{z}_c, \quad (4)$$

where \mathbf{z}_c is the document representation for the c -th candidate test class.

3.5. Optimization

Eq. (3) is convex for \mathbf{W}_x and \mathbf{W}_z individually but not convex for both of them. Therefore we can solve it using an alternating method, that is, we first fix \mathbf{W}_x and solve for \mathbf{W}_z ; then fix \mathbf{W}_z and solve for \mathbf{W}_x .

(1) Fix \mathbf{W}_x and solve for \mathbf{W}_z :

This sub-problem is a regression problem with $l_{2,1}$ -norm regularization. Nie *et al.* [15] proposes an iterative framework to efficiently solve it. It has been shown that the original problem is equivalent to sequentially solving the following problem until convergence

$$\min_{\mathbf{W}_z, \mathbf{D}} L(\mathbf{W}_x, \mathbf{W}_z) + \lambda_1 \|\mathbf{W}_x^\top \mathbf{W}_z \mathbf{Z}\|_F^2 + \lambda_2 \text{Tr}(\mathbf{W}_z \mathbf{D}^t \mathbf{W}_z^\top), \quad (5)$$

where \mathbf{D}^t is a diagonal matrix whose i -th diagonal element is $1/(2\|(\mathbf{w}_z^i)^{(t-1)}\|_2)^2$ at the t -th iteration, where

²In practice, we relax $1/(2\|(\mathbf{w}_z^i)^{(t-1)}\|_2)$ to $1/(2\sqrt{\mathbf{w}_z^{i\top} \mathbf{w}_z^i + \sigma})$, $\sigma \rightarrow 0$,

Algorithm 1 Fix \mathbf{W}_x and solve \mathbf{W}_z

Input: \mathbf{W}_x ; \mathbf{X} of seen classes; \mathbf{Z} of seen classes; λ_1 and λ_2 ; maximum number of iterations τ .

Initialize \mathbf{D}^0 as identity matrix $\mathbf{I} \in \mathbb{R}^{\hat{d} \times \hat{d}}$.

for $t = 1 \dots \tau$ **do**

- Solve Sylvester equation (6) for \mathbf{W}_z^t with \mathbf{D}^{t-1} .
- Update the diagonal matrix \mathbf{D}^t with its i -diagonal element as $1/(2\|(\mathbf{w}_z^i)^{(t)}\|_2)$, where $(\mathbf{w}_z^i)^{(t)}$ is the i -th column of \mathbf{W}_z^t .

if Converge then

- | · Break.

end

end

Output: \mathbf{W}_z .

$(\mathbf{w}_z^i)^{(t-1)}$ is the i -th column of the optimal \mathbf{W}_z solved at the $(t-1)$ -th iteration. The problem in Eq. (5) further reduces to a Sylvester equation of \mathbf{W}_z

$$\mathbf{A} \mathbf{W}_z + \mathbf{W}_z \mathbf{B} = \mathbf{C}, \quad (6)$$

$$\mathbf{A} = \lambda_2 (\mathbf{W}_x \mathbf{X} \mathbf{X}^\top \mathbf{W}_x^\top + \lambda_1 \mathbf{W}_x \mathbf{W}_x^\top)^{-1},$$

$$\mathbf{B} = \mathbf{Z} \mathbf{Z}^\top (\mathbf{D})^{-1},$$

$$\mathbf{C} = \frac{1}{\lambda_2} \mathbf{A} \mathbf{W}_x \mathbf{X} \mathbf{Y} \mathbf{Z}^\top (\mathbf{D})^{-1}.$$

The Sylvester equation has a unique solution if and only if \mathbf{A} and $-\mathbf{B}$ do not share any eigenvalues. Many state-of-the-art toolboxes are able to solve it efficiently. In our setting, since both \mathbf{A} and \mathbf{B} are positive definite, \mathbf{A} has only positive eigenvalues and $-\mathbf{B}$ has only negative eigenvalues. Therefore Eq. (6) has a unique solution. In summary, the sub-problem of fixing \mathbf{W}_x to solve \mathbf{W}_z can be solved via the algorithm listed in Algorithm 1.

(2) Fix \mathbf{W}_z and solve for \mathbf{W}_x :

This sub-problem is a conventional least squares minimization problem which has the following closed-form solution

$$\mathbf{W}_x^\top = (\mathbf{X} \mathbf{X}^\top + \lambda_1 \mathbf{I})^{-1} \mathbf{X} \mathbf{Y} \mathbf{Z}^\top \mathbf{W}_z^\top (\mathbf{W}_z \mathbf{Z} \mathbf{Z}^\top \mathbf{W}_z^\top)^{-1}. \quad (7)$$

By alternating between the above two matrices, the overall alternating optimization algorithm for Eq. (3) is listed in Algorithm 2.

as the i -th diagonal element to avoid the case of zero columns, and the $l_{2,1}$ norm is therefore approximated by $\sum_{i=1}^{\hat{d}} \sqrt{\mathbf{w}_z^{i\top} \mathbf{w}_z^i + \sigma}$. It has been proved in [15] that this approximation guarantees the convergence and the result approaches to that of $l_{2,1}$ -norm as $\sigma \rightarrow 0$.

Algorithm 2 Alternating algorithm for solving Eq. (3)

Input: \mathbf{X} of seen classes; \mathbf{Z} of seen classes; λ_1 and λ_2 ;
maximum number of iterations τ .
Initialize \mathbf{W}_x^0 with Gaussian distribution.
for $t = 1 \dots \tau$ **do**
 · Solve (5) iteratively for \mathbf{W}_z^t with \mathbf{W}_x^{t-1} according
 to Algorithm 1.
 · Solve (7) for \mathbf{W}_x^t with \mathbf{W}_z^t .
 if *Converge* **then**
 | · Break.
 end
end
Output: $\mathbf{W}_x, \mathbf{W}_z$.

4. Experiments

We divide our experiments into two parts. In the first part we evaluate the proposed method and compare it against both of the methods utilizing online textual sources and human-specified semantic attributes. In the second part we analyse in-depth the noise suppression effect of the proposed method and provide insight into what kind of information in a document is useful for zero-shot learning.

4.1. Experimental setting

Datasets: We test our approach on two widely used benchmarks for attribute learning and zero-shot learning: Animals with Attributes [12] (AwA) and Caltech-UCSD birds-200-2011 [25] (CUB-200-2011). AwA consists of 30,475 images of 50 mammals classes with 85 attributes including color, skin texture, body size, body part, affordance, food source, habitat, and behaviour. CUB-200-2011 contains 11,788 images of 200 categories of bird subspecies with 312 fine-grained attributes such as color/shape/texture of body parts. We follow the train/test split according to [12] and [25], where 10 and 50 testing classes are treated as unseen for AwA and CUB-200-2011, respectively.

Textual document sources: We extract the text representation according to scheme introduced in Section 3.2. The raw textual sources are collected from Wikipedia articles describing each of the categories. When constructing the vocabulary, we use the articles of seen classes only. The dimensionality of the text representation is 3506 for AwA and 6815 for CUB-200-2011, respectively.

Image features: To make fair a comparison, two types of image features, the low-level features in [19] and the fully connected layer activations from the “imagenet-vgg-verydeep-19” [22] CNN are used in our experiments.

Implementation details: The Sylvester equation in Eq. (6) is solved by a MATLAB built-in function, which takes only around 5 seconds on an Intel Core i7 CPU at 3.40GHz. The number of rows of matrices \mathbf{W}_x and \mathbf{W}_z is equal to the number of seen classes. We choose the

Method	Top-1 Acc	Top-5 Acc
Ba <i>et al.</i> [3] (BCE)	1	17.6
Ba <i>et al.</i> [3] (Hinge)	0.6	18.2
Ba <i>et al.</i> [3] (Euclidean)	12	42.8
ESZSL [20]	23.80	59.90
Ours	29.00 ± 0.28	61.76 ± 0.22

Table 1. Zero-shot learning classification results on CUB-200-2011, measured by top 1 and top 5 accuracy. 3 different loss functions are used in [3] for their CNN structure: binary cross entropy (BCE), hinge loss (Hinge), and Euclidean distance (Euclidean). All methods in this table use the same text sources from Wikipedia.

Method	Mean Accuracy
Rohrbach <i>et al.</i> [19] (Wikipedia)	19.7
Rohrbach <i>et al.</i> [19] (WordNet)	17.8
Rohrbach <i>et al.</i> [19] (Yahoo Web)	19.5
Rohrbach <i>et al.</i> [19] (Yahoo Img)	23.6
Rohrbach <i>et al.</i> [19] (Flickr Img)	22.9
ESZSL [20] (Wikipedia)	24.82
Ours (Wikipedia)	29.12 ± 0.07

Table 2. Zero-shot learning classification results of AwA, measured by mean accuracy. In [19], the approach mines attributes names from WordNet and additionally mines class-attribute from online sources of Wikipedia, WordNet, Yahoo, and Flickr. All methods in this table use the same low-level features in [19].

hyper-parameters with a five-fold cross-validation on the seen classes, where 20% (5 for AwA and 30 for CUB-200-2011) of the seen classes are held out for validation and the remaining seen classes are used for training. The hyper-parameters are tuned within the range of all cases of 10^b , where $b = \{-2, -1, \dots, 5, 6\}$. Once the hyper-parameters are selected, we use all seen classes to train the final model. All of our reported results are averaged over 10 trials.

4.2. Performance evaluation

We first compare our method against [3] and [19]. The former is most relevant to our work in the sense that it learns a mapping to match images and textual documents. The work in [19] is a comprehensive comparison study of various information sources for zero-shot learning. Besides these two method, we also treat $\mathbf{S} = \mathbf{Z}$ in Eq. (1), and apply the ESZSL method in [20] to our zero-shot learning problem. To make a fair comparison, we use the same low-level features in [19] when comparing with it and then use the “imagenet-vgg-verydeep-19” to compare with [3]. The comparison results are given in Table 1 and Table 2. As can be seen in Table 1, the proposed method significantly outperforms the methods in [3], although they have used a more complicated deep learning framework. Also, we find

that our baseline ESZSL achieves good performance. However, it is still 5% inferior to our approach, which clearly demonstrates the advantage of the noise suppression mechanism introduced in this paper. The results in Table 2 further show that our method is superior over other approaches which rely on automatically mined information from the web. Again, our method achieves a significant improvement (more than 4%) over ESZSL.

We now compare our work with a few other state-of-the-art approaches on zero-shot learning, even though some of them are not based on online information sources. The results are summarized in Table 3. Results [18, 11, 14, 1] listed in the upper part of the table utilize hand-crafted features and not surprisingly their performance is much inferior to that of the proposed method. The lower part of Table 3 are methods with visual features extracted from a pre-trained CNN and thus are more comparable to our method. In this setting, we find that our method is comparable to most of the state-of-the-art results on AwA and results better than ours are all obtained from the methods using cleaner human defined attributes. The work in [2] evaluates various semantic representations such as Word2Vec embedding, GloVe word co-occurrence from Wikipedia sources, taxonomy embedding inferred from WordNet Hierarchy, and pre-defined binary and real-valued attributes. Our approach outperforms all methods that use online text sources. This shows that although online text sources provide transferable semantic representations, their discriminative ability is affected by the inherent noise and our method is better at handling the noisy information source for zero-shot learning.

Similar results are observed on the CUB-200-2011 dataset. Our approach again outperforms the methods using online sources and those methods that beat ours are all based on human specified fine-grained attributes. Note that many of the bird categories in CUB-200-2011 have very subtle differences which may not be well captured in Wikipedia articles. However, better performance may be expected by using a higher quality text corpus, such as bird watching articles.

4.3. In-depth analysis of the proposed method

In this section we provide an in-depth analysis of the proposed method by examining its noise suppression mechanism and the words that are most discriminative in the view of our method.

4.3.1 Effectiveness of the noise suppression method

In our method, the $l_{2,1}$ -norm is expected to allow only a few dimensions of the document representation to have large values. The importance of each individual dimension of the document representation can therefore be measured by the l_2 -norm of each column of learned \mathbf{W}_z (we call it the im-

portance weight in the following). We visualize this measurement for each dimension of the document representation in the top two subfigures in Figure 2. As can be seen, most of the importance weights are not exactly zero as one might expect given that the $l_{2,1}$ -norm is applied. In fact, there are only 702 zero columns (out of 3506) for AwA and 949 (out of 6815) for CUB-200-2011. As also mentioned in Section 3, this is probably because of the joint regularization effect of $\|\mathbf{W}_x^T \mathbf{W}_z \mathbf{Z}\|_F^2$ in Eq. (3) and/or because by cross-validation most dimensions are still identified as being useful although their weighting should be very low. The second postulate might be supported by the observation that poorer performance will be obtained if we manually remove the dimensions which have low importance weights.

Although our formulation does not achieve the feature selection effect, it does only assign large importance weights to a small number of dimensions. To visually compare its effect, we replace the $l_{2,1}$ -norm and with the Frobenius norm and carry out our learning algorithm again. The resulting importance weights are shown in the two subfigures at the bottom of Figure 2. As can be seen, large importance weights appear in more dimensions in this case. This observation verifies the noise suppression effect of the regularizer introduced in Eq. (3) and explains the superior performance of our method over other text-based zero-shot learning approaches.

4.3.2 Understanding the important dimensions of the document representation

Since each individual dimension of the textual document representation corresponds to an unique word, we can visualize the dimensions/words with large importance weights for better understanding our zero-shot learning algorithm. Table 4 lists at most 15 top scored words for 15 out of 40 seen classes in AwA and we could make several observations from it: (1) even though the document representations are extremely noisy, most of the top-ranked words are semantically meaningful to describe discriminative properties of a category (an animal in this case), such as body parts, habitat, behaviour, affordance, taxonomy, and environment. In fact, we find many top weighted words are consistent with some of the human specified attributes in AwA. (2) Many top-ranked words are not explicitly “visualizable” but they imply visual information of a category. For example, the abstract concept “ruminant” implicitly tells that the creature with this property is “deer-like” or “cattle-like” and builds a visual connection between antelope and deer in Table 4. This observation has also been made in the literature [12, 13, 17, 4, 21]. (3) Interestingly, we also notice that although some concepts are not commonly considered as attributes, they exhibit large importance weight as inferred by our algorithm. By taking a close examina-

Method/Dataset	AwA	CUB
Rohrbach <i>et al.</i> [18]	42.7	
Jayaraman & Grauman [11]	43.01	
Mensink <i>et al.</i> [14]		14.4
Akata <i>et al.</i> [1]	43.5	18.0
Lampert <i>et al.</i> [13] (attr. real)	57.5	
Deng <i>et al.</i> [5] (hierarchy)	44.2	
ESZSL [20] (attr. bin)	62.85	
Akata <i>et al.</i> [2] (Word2Vec)	51.2	28.4
Akata <i>et al.</i> [2] (GloVe)	58.8	24.2
Akata <i>et al.</i> [2] (WordNet)	51.2	20.6
Akata <i>et al.</i> [2] (attr. bin)	52.0	37.8
Akata <i>et al.</i> [2] (attr. real)	66.7	50.1
Fu <i>et al.</i> [10] (attr. & words)	66.0	
Zhang & Saligrama [27] (attr. real)	76.33	30.41
ESZSL [20] (Wikipedia)	58.53	23.80
Ours (Wikipedia)	66.46 ± 0.42	29.00 ± 0.28

Table 3. Zero-shot learning classification results on AwA and CUB-200-2011. Blank spaces indicate these methods are not tested on the corresponding datasets. Contents in braces indicate the semantic sources which these methods use for zero-shot learning. Methods in the upper part of the table use low-level features and the remaining methods in the lower part use deep CNN features.

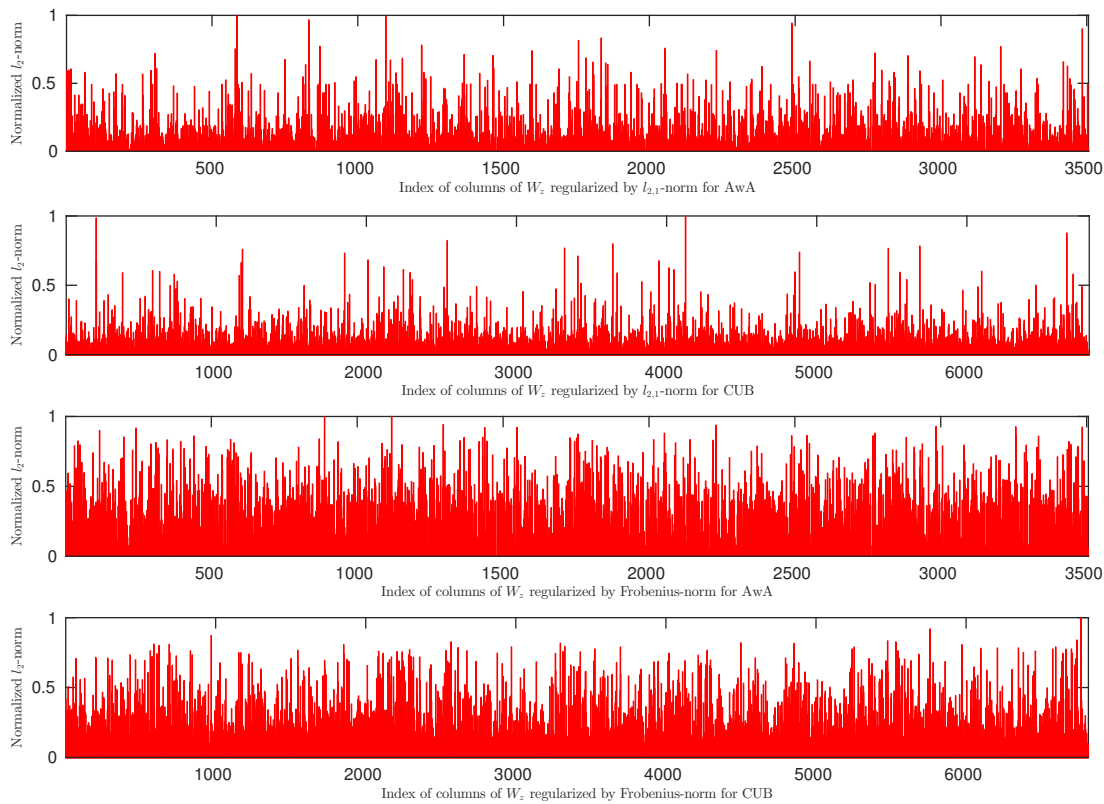


Figure 2. The two subfigures at the top show column-wise l_2 -norms of \mathbf{W}_z learned with $l_{2,1}$ -norm regularization. The two subfigures at the bottom show column-wise l_2 -norms of \mathbf{W}_z learned with Frobenius-norm regularization.

tion, we categorize these words into two types. The first (labelled green in Table 4) are some concepts that are more likely to co-occur with meaningful attributes. For example,

the word “stomach” is only shared by antelope and deer in Table 4, despite its existence in all mammals. This is probably because “stomach” is more likely to be co-occurred with

Class	Top Ranked Words/Dimensions
Antelope	antler, woodland, fight, stomach, spike, antelope, escape, mate, night, variety, ruminant, ridge, broad, scent, herd
Beaver	river, protect, semiaquatic, web, branch, eurasian, american, land, insular, hunt, fur, extant, adult, stream, pond
Blue Whale	ton, whale, flipper, kilometre, marine, ocean, belong, mph, shape, dive, earth, worldwide, indian, travel, pacific
Buffalo	climate, extant, herd, indian, cattle, dairy, animate, bc, trade, behaviour, human, milk, northern, southeast, field
Cow	draft, milk, cattle, widespread, product, meat, domestic, strong, cart, plow, oxen, bullock, cow, animate, india
Deer	antler, fight, mate, elk, palmate, moose, wolf, season, bear, woodland, herd, ruminant, deer, stomach, spike
Moose	herd, elk, palmate, moose, wolf, fight, deer, compete, alces, temperate, climate, aggressive, sedentary, season
Mouse	rodent, house, eat, avoid, burrow, general, genetic, popular, breed, wild, small, tail, vermin, nocturnal, prey
Dolphin	flipper, whale, ton, kilometre, indian, dive, mph, earth, shape, blubber, belong, marine, ocean, capture, prevent
Horse	draft, strong, milk, meat, ungulate, equip, widespread, loose, past, history, compete, endure, technique, style, flee
Hamster	mix, underground, fragile, house, bear, seed, worm, silky, rapid, classify, general, tail, flexible, dwarf, pouch
Killer Whale	ton, whale, dolphin, click, dive, killer, pollution, belong, capture, vocal, calf, tail, threat, fish, fin
Otter	semiaquatic, branch, eurasian, lake, engage, bed, play, trap, river, deplete, giant, cetacean, mink, weasel, web
Rabbit	fragile, house, classify, general, introduce, underground, pad, vegetarian, companionship, defensive, shelf, detect
S. Monkey	agile, arm, walk, tropic, rainforest, primate, source, primary, bark, passage, balance, thumb, moist, threaten

Table 4. Category-wisely top ranked words, sorted by average importance weights within each class. The blue words are generally considered as meaningful attributes of this class. The green words are concepts somewhat related to this class, but are less informative to define it. The red words are concepts that are not semantically related to the corresponding class.

“ruminant”, a discriminative property of ruminant animals. Another type of words (labelled red in Table 4) are not sufficiently meaningful for human interpreter. For example, “belong” and “general” are assigned with high importance weight for all cetaceans (blue whale, dolphin, killer whale etc.) and rodents (mouse, rabbit, hamster etc.), respectively. We suspect the reason is due to the dataset bias of documents. For example, documents of similar categories may be edited by authors from the same background who prefer a certain word choice. In sum, we find most of the top ranked words carry weak information by their own, but it seems that using them collaboratively produces impressive discriminative power for zero-shot learning.

5. Conclusion

In this paper, we have introduced a noise suppression mechanism to text-based zero-shot learning. The proposed $l_{2,1}$ -norm based objective function generates classifiers that are robust against textual noise and achieve state-of-the-art zero-shot learning performance. We have made several findings in the experiments. (1) The inherent noise within text sources has a significant impact on zero-shot learning performance. As all the text-methods without noise suppression are inferior to our approach, we speculate that noise in a component of the mid-level representation decreases its discriminative power. (2) Most noisy components are suppressed rather than completely eliminated by our mechanism. Some words, although unimportant individually, can produce meaningful discriminative power when put together. (3) We find three kinds of words in the de-noised representation that can provide useful information for zero-shot learning. The first kind are the attribute-like words that explicitly describe the category. The second

are words that are weakly related to the category. They usually occur with definitive words. The last kind of words is non-informative to humans, but shows certain distribution patterns among related categories.

Overall, this paper points out an important factor in text-based zero-shot learning that has been previously ignored. By dealing directly with the inevitable variations in human expression, and suppressing words that contain little or no value, the performance of text-based automatic zero-shot learning can be significantly improved.

References

- [1] Z. Akata, F. Perronnin, Z. Harchaoui, and C. Schmid. Label-embedding for attribute-based classification. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, June 2013.
- [2] Z. Akata, S. Reed, D. Walter, H. Lee, and B. Schiele. Evaluation of output embeddings for fine-grained image classification. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.* IEEE, IEEE, June 2015.
- [3] J. L. Ba, K. Swersky, S. Fidler, and R. Salakhutdinov. Predicting deep zero-shot convolutional neural networks using textual descriptions. In *Proc. IEEE Int. Conf. Comp. Vis.* IEEE, 2015.
- [4] T. L. Berg, A. C. Berg, and J. Shih. Automatic attribute discovery and characterization from noisy web data. In *Proc. Eur. Conf. Comp. Vis.*, pages 663–676, Berlin, Heidelberg, 2010. Springer-Verlag.
- [5] J. Deng, N. Ding, Y. Jia, A. Frome, K. Murphy, S. Bengio, Y. Li, H. Neven, and H. Adam. Large-scale object classification using label relation graphs. In *Proc. Eur. Conf. Comp. Vis.*, pages 48–64. Springer International Publishing, 2014.
- [6] M. Elhoseiny, B. Saleh, and A. Elgammal. Write a classifier: Zero-shot learning using purely textual descriptions. In *Proc. IEEE Int. Conf. Comp. Vis.*, pages 2584–2591, Washington, DC, USA, 2013. IEEE.

- [7] A. Farhadi, I. Endres, D. Hoiem, and D. Forsyth. Describing objects by their attributes. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2009.
- [8] V. Ferrari and A. Zisserman. Learning visual attributes. In *Proc. Advances in Neural Inf. Process. Syst.*, 2008.
- [9] A. Frome, G. Corrado, J. Shlens, S. Bengio, J. Dean, M. Ranzato, and T. Mikolov. Devise: A deep visual-semantic embedding model. In *Proc. Advances in Neural Inf. Process. Syst.*, 2013.
- [10] Z. Fu, T. Xiang, E. Kodirov, and S. Gong. Zero-shot object recognition by semantic manifold distance. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, June 2015.
- [11] D. Jayaraman and K. Grauman. Zero-shot recognition with unreliable attributes. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger, editors, *Proc. Advances in Neural Inf. Process. Syst.*, pages 3464–3472. Curran Associates, Inc., 2014.
- [12] C. H. Lampert, H. Nickisch, and S. Harmeling. Learning to detect unseen object classes by betweenclass attribute transfer. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2009.
- [13] C. H. Lampert, H. Nickisch, and S. Harmeling. Attribute-based classification for zero-shot visual object categorization. *IEEE Trans. Pattern Anal. Mach. Intell.*, 36(3):453–465, 2014.
- [14] T. Mensink, E. Gavves, and C. G. Snoek. Costa: Co-occurrence statistics for zero-shot classification. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, June 2014.
- [15] F. Nie, H. Huang, X. Cai, and C. H. Q. Ding. Efficient and robust feature selection via joint $l_{2,1}$ -norms minimization. In *Proc. Advances in Neural Inf. Process. Syst.*, pages 1813–1821, 2010.
- [16] M. Norouzi, T. Mikolov, S. Bengio, Y. Singer, J. Shlens, A. Frome, G. Corrado, and J. Dean. Zero-shot learning by convex combination of semantic embeddings. In *Proc. Int. Conf. Learning Representations*, 2014.
- [17] D. N. Osherson, J. Stern, O. Wilkie, M. Stob, and E. E. Smith. Default probability. *Cognitive Science*, 15(2):251–269, 1991.
- [18] M. Rohrbach, S. Ebert, and B. Schiele. Transfer learning in a transductive setting. In *Proc. Advances in Neural Inf. Process. Syst.*, December 2013.
- [19] M. Rohrbach, M. Stark, G. Szarvas, I. Gurevych, and B. Schiele. What helps where - and why? semantic relatedness for knowledge transfer. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, pages 910–917, 2010.
- [20] B. Romera-Paredes and P. H. Torr. An embarrassingly simple approach to zero-shot learning. *Proc. Int. Conf. Mach. Learn.*, 2015.
- [21] J. Shao, K. Kang, C. C. Loy, and X. Wang. Deeply learned attributes for crowded scene understanding. In *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2015.
- [22] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *Proc. Int. Conf. Learning Representations*, 2014.
- [23] R. Socher, M. Ganjoo, C. D. Manning, and A. Ng. Zero-shot learning through cross-modal transfer. In C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Weinberger, editors, *Proc. Advances in Neural Inf. Process. Syst.*, pages 935–943. Curran Associates, Inc., 2013.
- [24] L. Torresani, M. Szummer, and A. Fitzgibbon. Efficient object category recognition using classemes. In *Proc. Eur. Conf. Comp. Vis.*, pages 776–789, Sept. 2010.
- [25] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011.
- [26] B. Yao, X. Jiang, A. Khosla, A. L. Lin, L. J. Guibas, and L. Fei-Fei. Action recognition by learning bases of action attributes and parts. In *Proc. IEEE Int. Conf. Comp. Vis.*, Barcelona, Spain, November 2011.
- [27] Z. Zhang and V. Saligrama. Zero-shot learning via semantic similarity embedding. In *Proc. IEEE Int. Conf. Comp. Vis. IEEE*, 2015.