

Hierarchy-based Image Embeddings for Semantic Image Retrieval

Björn Barz

Joachim Denzler

Friedrich Schiller University Jena
Computer Vision Group

`bjoern.barz@uni-jena.de`

Abstract

Deep neural networks trained for classification have been found to learn powerful image representations, which are also often used for other tasks such as comparing images w.r.t. their visual similarity. However, visual similarity does not imply semantic similarity. In order to learn semantically discriminative features, we propose to map images onto class embeddings whose pair-wise dot products correspond to a measure of semantic similarity between classes. Such an embedding does not only improve image retrieval results, but could also facilitate integrating semantics for other tasks, e.g., novelty detection or few-shot learning. We introduce a deterministic algorithm for computing the class centroids directly based on prior world-knowledge encoded in a hierarchy of classes such as WordNet. Experiments on CIFAR-100, NABirds, and ImageNet show that our learned semantic image embeddings improve the semantic consistency of image retrieval results by a large margin.

Source code: <https://github.com/cvjena/semantic-embeddings>

1. Introduction

During the past few years, deep convolutional neural networks (CNNs) have continuously advanced the state-of-the-art in image classification [19, 30, 13, 15] and many other tasks. The intermediate image representations learned by such CNNs trained for classification have also proven to be powerful image descriptors for retrieving images from a database that are visually or semantically similar to one or more query images given by the user [2, 29, 1]. This task is called content-based image retrieval (CBIR) [31].

Usually, the categorical cross-entropy loss after a softmax activation is used for training CNN classifiers, which results in well-separable features. However, these features are not necessarily discriminative, *i.e.*, the inter-class variance of the learned representations may be small compared to the intra-class variance [35]. The average distance be-

tween the classes “cat” and “dog” in feature space may well be equally large as the distance between “cat” and “forest”. For image retrieval, that situation is far from optimal, since the nearest neighbors of a certain image in feature space may belong to completely different classes. This can, for instance, be observed in the upper row of each example in Fig. 1, where we used features extracted from the global average pooling layer of ResNet-110 [13] for image retrieval on the CIFAR-100 dataset [18]. While the first results often belong to the same class as the query, the semantic relatedness of the results at later positions deteriorates significantly. Those seem to be mainly visually similar to the query with respect to shape and color, but not semantically.

Many authors have therefore proposed a variety of metric learning losses, aiming to increase the separation of classes, while minimizing the distances between samples from the same class. Popular instances of this branch of research are the contrastive loss [6], the triplet loss [27], and the quadruplet loss [5]. However, these methods require sampling of hard pairs, triplets, or even quadruplets of images, making training cumbersome and expensive. On the other hand, they still do not impose any constraints on inter-class relationships: Though images of the same class should be tightly clustered together in feature space, neighboring clusters may still be completely unrelated.

This harms the semantic consistency of image retrieval results. Consider, for example, a query image showing a poodle. It is unclear, whether the user is searching for images of other poodles only, for images of any other dog breed, or even for images of animals in general. Current CBIR methods would—ideally—first retrieve all images of the same class as the query but then continue with mostly unrelated, only *visually* similar images from other classes.

In this work, we propose a method for learning *semantically* meaningful image representations, so that the Euclidean distance in the feature space directly corresponds to the semantic dissimilarity of the classes. Since “semantic similarity” is, in principle, arbitrary, we rely on prior knowledge about the physical world encoded in a hierarchy



Figure 1: Comparison of image retrieval results on CIFAR-100 [18] for 3 exemplary queries using features extracted from a variant of ResNet-110 [13] trained for classification and semantic embeddings learned by our method. The border colors of the retrieved images correspond to the semantic similarity between their class and the class of the query image (with dark green being most similar and dark red being most dissimilar). It can be seen that hierarchy-based semantic image embeddings lead to much more semantically consistent retrieval results.

of classes. Such hierarchies are readily available in many domains, since various disciplines have been striving towards organizing the entities of the world in ontologies for years. WordNet [10], for example, is well-known for its good coverage of the world with over 80,000 semantic concepts and the Wikispecies project¹ provides a taxonomy of living things comprising more than half a million nodes.

We make use of the knowledge explicitly encoded in such hierarchies to derive a measure of semantic similarity between classes and introduce an algorithm for explicitly computing target locations for all classes on a unit hypersphere, so that the pair-wise dot products of their embeddings equal their semantic similarities. We then learn a transformation from the space of color images to this semantically meaningful feature space, so that the correlation between image features and their class-specific target embedding is maximized. This can easily be done using the negated dot product as a loss function, without having to

consider pairs or triplets of images or hard-negative mining.

Exemplary retrieval results of our system are shown in the bottom rows of each example in Fig. 1. It can be seen that our semantic class embeddings lead to image features that are much more invariant against superficial visual differences: An orange is semantically more similar to a green apple than to orange tulips and a palm tree is more similar to an oak than to a spider. The results follow the desired scheme described above: For a query image showing an orange, all oranges are retrieved first, then all apples, pears, and other fruits. As a last example, incorporating semantic information about classes successfully helps avoiding the not so uncommon mistake of confusing humans with apes.

We discuss related work on learning semantic image representations in Section 2. Our approach for computing class embeddings based on a class hierarchy is presented in Section 3, and Section 4 explains how we learn to map images onto those embeddings. Experimental results on three popular datasets are presented in Section 5, and Section 6 concludes this work.

¹<https://species.wikimedia.org/>

2. Related Work

Since the release of the ImageNet dataset [8] in 2009, many authors have proposed to leverage the WordNet ontology, which the classes of ImageNet have been derived from, for improving classification and image retrieval: The creators of ImageNet, Deng et al. [7], derived a deterministic bilinear similarity measure from the taxonomy of classes for comparing image feature vectors composed of class-wise probability predictions. This way, images assigned to different classes can still be considered as similar if the two classes are similar to each other.

Regarding classification, Zhao et al. [37] modify multinomial logistic regression to take the class structure into account, using the dissimilarity of classes as misclassification cost. Verma et al. [33], on the other hand, learn a specialized Mahalanobis distance metric for each node in a class hierarchy and combine them along the paths from the root to the leafs. This results in different metrics for all classes and only allows for nearest-neighbor-based classification methods. In contrast, Chang et al. [4] learn a global Mahalanobis metric on the space of class-wise probability predictions, where they enforce margins between classes proportional to their semantic dissimilarities. HD-CNN [36] follows an end-to-end learning approach by dividing a CNN into a coarse and several fine classifiers based on two levels of a class hierarchy, fusing predictions at the end.

All these approaches exploit the class structure for classification or retrieval, but only at the classifier level instead of the features themselves. Our approach, in contrast, embeds images into a semantically meaningful space where the dot product corresponds to the similarity of classes. This does not only make semantic image retrieval straightforward, but also enables the application of a wide variety of existing methods that rely on metric feature spaces, *e.g.*, clustering or integration of relevance feedback into the retrieval [9].

A very similar approach has been taken by Weinberger et al. [34], who propose “taxonomy embeddings” (“taxem”) for categorization of documents. However, they do not specify how they obtain semantic class similarities from the taxonomy. Moreover, they learn a linear transformation from hand-crafted document features onto the class embedding space using ridge regression, whereas we perform end-to-end learning using neural networks.

More recently, several authors proposed to jointly learn embeddings of classes and images based solely on visual information, *e.g.*, using the “center loss” [35] or a “label embedding network” [32]. However, semantics are often too complex for being derived from visual information only. For example, the label embedding network [32] learned that pears are similar to bottles, because their shape and color is often similar and the image information alone is not sufficient for learning that fruits and man-made containers are fundamentally different concepts.

To avoid such issues, Frome et al. (“DeViSE” [11]) and Li et al. [20] propose to incorporate prior world-knowledge by mapping images onto word embeddings of class labels learned from text corpora [22, 24]. To this end, Frome et al. [11] need to pre-train their image embeddings for classification initially, and Li et al. [20] first perform region mining and then use three sophisticated loss functions, requiring the mining of either hard pairs or triplets.

In contrast to this expensive training procedure relying on the additional input of huge text corpora, we show how to explicitly construct class embeddings based on prior knowledge encoded in an easily obtainable hierarchy of classes, without the need to learn such embeddings approximately. These embeddings also allow for straightforward learning of image representations by simply maximizing the dot product of image and class embeddings.

A broader overview of research aiming for incorporating prior knowledge into deep learning of image representations can, for instance, be found in [28].

3. Hierarchy-based Class Embeddings

In the following, we first describe how we measure semantic similarity between classes based on a hierarchy and then introduce our method for explicitly computing class embeddings based on those pair-wise class similarities.

3.1. Measuring Semantic Similarity

Let $\mathcal{G} = (V, E)$ be a directed acyclic graph with nodes V and edges $E \subseteq V \times V$, specifying the hyponymy relation between semantic concepts. In other words, an edge $(u, v) \in E$ means that v is a sub-class of u . The actual classes of interest $\mathcal{C} = \{c_1, \dots, c_n\} \subseteq V$ are a subset of the semantic concepts. An example for such a graph, with the special property of being a tree, is given in Fig. 2a.

A commonly used measure for the dissimilarity $d_{\mathcal{G}} : \mathcal{C} \times \mathcal{C} \rightarrow \mathbb{R}$ of classes organized in this way is the height of the sub-tree rooted at the *lowest common subsumer (LCS)* of two classes, divided by the height of the hierarchy [7, 33]:

$$d_{\mathcal{G}}(u, v) = \frac{\text{height}(\text{lcs}(u, v))}{\max_{w \in V} \text{height}(w)}, \quad (1)$$

where the height of a node is defined as the length of the longest path from that node to a leaf. The LCS of two nodes is the ancestor of both nodes that does not have any child being an ancestor of both nodes as well.

Since $d_{\mathcal{G}}$ is bounded between 0 and 1, we can easily derive a measure for semantic similarity between semantic concepts as well:

$$s_{\mathcal{G}}(u, v) = 1 - d_{\mathcal{G}}(u, v). \quad (2)$$

For example, the toy hierarchy in Fig. 2a has a total height of 3, the LCS of the classes “dog” and “cat”

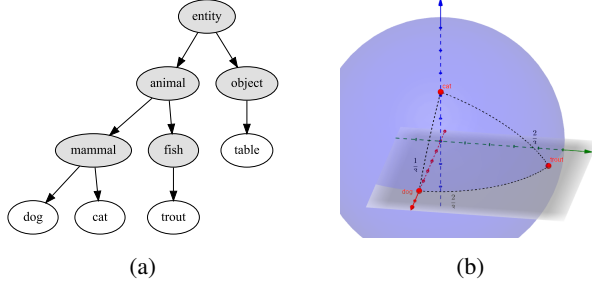


Figure 2: (a) A toy hierarchy and (b) an embedding of 3 classes from this hierarchy with their pair-wise d_G .

is “mammal” and the LCS of “dog” and “trout” is “animal”. It follows that $d_G(\text{“dog”}, \text{“cat”}) = \frac{1}{3}$ and $d_G(\text{“dog”}, \text{“trout”}) = \frac{2}{3}$.

Note that though d_G is symmetric and non-negative, it is only guaranteed to be a proper metric if \mathcal{G} is a tree with all classes of interest being leaf nodes (the proof can be found in Appendix A). Some well-known ontologies such as WordNet [10] do not have this property and violate the triangle inequality. For instance, the WordNet synset “golf-cart” is a hypernym of both “vehicle” and “golf equipment”. It is hence similar to cars and golf balls, while both are not similar to each other at all.

Our goal is to embed classes onto a unit hypersphere so that their dot product corresponds to s_G . Thus, the Euclidean distance between such embeddings of the classes equals $\sqrt{2} \cdot d_G$, which hence has to be a metric. Therefore, we assume the hierarchy \mathcal{G} to be given as tree in the following. In case that this assumption does not hold, approaches from the literature for deriving tree-shaped hierarchies from ontologies such as WordNet could be employed. For instance, YOLO-9000 [25] starts with a tree consisting of the root-paths of all concepts which have only one such path and then successively adds paths to other concepts that result in the least number of nodes added to the existing tree.

3.2. Class Embedding Algorithm

Consider n classes $\mathcal{C} = \{c_1, \dots, c_n\}$ embedded in a hierarchy \mathcal{G} as above. Our aim is to compute embeddings $\varphi(c_i) \in \mathbb{R}^n$ of all classes $c_i, i = 1, \dots, n$, so that

$$\begin{aligned} \forall_{1 \leq i, j \leq n} : \quad & \varphi(c_i)^\top \varphi(c_j) = s_G(c_i, c_j), \quad (3) \\ \forall_{1 \leq i \leq n} : \quad & \|\varphi(c_i)\| = 1. \quad (4) \end{aligned}$$

In other words, the correlation of class embeddings should equal the semantic similarity of the classes and all embeddings should be L2-normalized. Eq. (4) actually is a direct consequence of (3) in combination with the fact that $s_G(u, u) = 1$ for any class $u \in \mathcal{C}$, but we formulate it as an explicit constraint here for clarity, emphasizing that all class embeddings lie on a unit hypersphere. This does

Algorithm 1 Hierarchy-based Class Embeddings

Input: hierarchy $\mathcal{G} = (V, E), E \subseteq (V \times V)$, set $\mathcal{C} = \{c_1, \dots, c_n\} \subseteq V$ of n classes

Output: embeddings $\varphi(c_i) \in \mathbb{R}^n$ for $i = 1, \dots, n$

- 1: $\varphi(c_1) \leftarrow [1, 0, \dots, 0]^\top$
 - 2: **for** $i = 2$ **to** n **do**
 - 3: $\hat{\varphi}(c_i) \leftarrow$ solution of (5)
 - 4: $x \leftarrow$ maximum of the solutions of (6)
 - 5: $\varphi(c_i) \leftarrow [\hat{\varphi}(c_i)^\top, x, 0, \dots, 0]^\top$
 - 6: **end for**
-

not only allow us to use the negated dot product as a substitute for the Euclidean distance, but also accounts for the fact that L2 normalization has proven beneficial for CBIR in general, because the direction of high-dimensional feature vectors often carries more information than their magnitude [17, 16, 14].

We follow a step-wise approach for computing the embeddings $\varphi(c_i), i = 1, \dots, n$, as outlined in Algorithm 1: We can always choose an arbitrary point on the unit sphere in an n -dimensional space as embedding for the first class c_1 , because the constraint (3) is invariant against arbitrary rotations and translations. Here, we choose $\varphi(c_1) = [1, 0, \dots, 0]^\top$, which will simplify further calculations.

The remaining classes $c_i, i = 2, \dots, n$, are then placed successively so that each new embedding has the correct dot product with any other already embedded class:

$$\forall_{1 \leq j < i} : \varphi(c_j)^\top \varphi(c_i) = s_G(c_j, c_i). \quad (5)$$

This is a system of $i - 1$ linear equations, where $\varphi(c_i)$ is the vector of unknown variables. As per our construction, only the first j dimensions of all $\varphi(c_j), j = 1, \dots, i - 1$, are non-zero, so that the effective number of free variables is $i - 1$. The system is hence well-determined and in lower-triangular form, so that it has a unique solution that can be computed efficiently with $\mathcal{O}(i^2)$ floating-point operations using forward substitution.

This solution $\hat{\varphi}(c_i) = [\varphi(c_i)_1, \dots, \varphi(c_i)_{i-1}]^\top$ for the first $i - 1$ coordinates of $\varphi(c_i)$ already fulfills (3), but not (4), i.e., it is not L2-normalized. Thus, we use an additional dimension to achieve normalization:

$$\varphi(c_i)_i = \sqrt{1 - \|\hat{\varphi}(c_i)\|^2}. \quad (6)$$

Without loss of generality, we always choose the non-negative solution of this equation, so that all class embeddings lie entirely in the positive orthant of the feature space.

Due to this construction, exactly n feature dimensions are required for computing a hierarchy-based embedding of n classes. An example of such an embedding for 3 classes is given in Fig. 2b. The overall complexity of the algorithm outlined in Algorithm 1 is $\mathcal{O}(n^3)$.

3.3. Low-dimensional Approximation

In settings with several thousands of classes, the number of features required by our algorithm from Section 3.2 might become infeasible. However, it is possible to obtain class embeddings of arbitrary dimensionality whose pair-wise dot products best *approximate* the class similarities.

Let $\mathcal{S} \in \mathbb{R}^{n \times n}$ be the matrix of pair-wise class similarities, i.e., $S_{ij} = s_{\mathcal{G}}(c_i, c_j)$, $1 \leq i, j \leq n$, and $\Phi \in \mathbb{R}^{n \times n}$ be a matrix whose i -th row is the embedding $\varphi(c_i)$ for class c_i . Then, we can reformulate (3) as

$$\Phi \cdot \Phi^\top = \mathcal{S} = Q\Lambda Q^\top, \quad (7)$$

where $Q \in \mathbb{R}^{n \times n}$ is a matrix whose rows contain the eigenvectors of \mathcal{S} and $\Lambda \in \mathbb{R}^{n \times n}$ is a diagonal matrix containing the corresponding eigenvalues.

Thus, we could also use the eigendecomposition of \mathcal{S} to obtain the class embeddings as $\Phi = Q\Lambda^{1/2}$. However, we have found our algorithm presented in Section 3.2 to provide better numerical accuracy, resulting in a maximum error of pairwise distances of 1.7×10^{-15} for the 1000 classes of ILSVRC [26], while the eigendecomposition only obtains 1.7×10^{-13} . Moreover, eigendecomposition does not guarantee all class embeddings to lie in the positive orthant of the feature space. However, we have found this to be a beneficial regularization in practice, resulting in slightly better performance.

On the other hand, when dealing with a large number of classes, the eigendecomposition can be useful to obtain a low-dimensional embedding that does not reconstruct \mathcal{S} exactly but approximates it as best as possible by keeping only the eigenvectors corresponding to the largest eigenvalues. However, the resulting class embeddings will not be L2-normalized any more. Experiments in Appendix D show that our method can still provide superior retrieval performance even with very low-dimensional embeddings, which is also advantageous regarding memory consumption when dealing with large datasets.

4. Mapping Images onto Class Centroids

Knowing the target embeddings for all classes, we need to find a transformation $\psi : \mathcal{X} \rightarrow \mathbb{R}^n$ from the space \mathcal{X} of images into the hierarchy-based semantic embedding space, so that image features are close to the centroid of their class.

For modeling ψ , we employ convolutional neural networks (CNNs) whose last layer has n output channels and no activation function. Since we embed all classes onto the unit hypersphere, the last layer is followed by L2 normalization of the feature vectors. The network is trained on batches $B \in (\mathcal{X} \times \mathcal{Y})^m$ of m images $I_b \in \mathcal{X}$ with class labels $y_b \in \mathcal{Y} = \{1, \dots, n\}$, $b = 1, \dots, m$, under the supervision of a simple loss function $\mathcal{L}_{\text{CORR}} : (\mathcal{X} \times \mathcal{Y})^m \rightarrow \mathbb{R}$

that enforces similarity between the learned image representations and the semantic embedding of their class:

$$\mathcal{L}_{\text{CORR}}(B) = \frac{1}{m} \sum_{b=1}^m (1 - \psi(I_b)^\top \varphi(c_{y_b})) . \quad (8)$$

Note that the class centroids $\varphi(c_1), \dots, \varphi(c_n)$ are computed beforehand using the algorithm described in Section 3 and fixed during the training of the network. This allows our loss function to be used stand-alone, as opposed to, for example, the center loss [35], which requires additional supervision. Moreover, the loss caused by each sample is independent from all other samples, so that expensive mining of hard pairs [6] or triplets [27] of samples is not necessary.

Features learned this way may not only be used for retrieval but also for classification by assigning a sample to the class whose embedding is closest in the feature space. However, one might as well add a fully-connected layer with softmax activation on top of the embedding layer, producing output $\rho : \mathcal{X} \rightarrow [0, 1]^n$. The network could then simultaneously be trained for computing semantic image embeddings and classifying images using a combined loss

$$\mathcal{L}_{\text{CORR+CLS}} = \mathcal{L}_{\text{CORR}} + \lambda \cdot \mathcal{L}_{\text{CLS}}, \quad (9)$$

where $\mathcal{L}_{\text{CLS}} : (\mathcal{X} \times \mathcal{Y})^m \rightarrow \mathbb{R}$ denotes the categorical cross-entropy loss function

$$\mathcal{L}_{\text{CLS}}(B) = \frac{1}{m} \sum_{b=1}^m \log(\rho(I_b)_{y_b}) . \quad (10)$$

Since we would like the embedding loss $\mathcal{L}_{\text{CORR}}$ to dominate the learning of image representations, we set λ to a small value of $\lambda = 0.1$ in our experiments.

5. Experiments

In the following, we present results on three different datasets and compare our method for learning semantic image representations with features learned by other methods that try to take relationships among classes into account.

5.1. Datasets and Setups

5.1.1 CIFAR-100

The extensively benchmarked CIFAR-100 dataset [18] consists of 100 classes with 500 training and 100 test images each. In order to make our approach applicable to this dataset, we created a taxonomy for the set of classes, mainly based on the WordNet ontology [10] but slightly simplified and with a strict tree structure. A visualization of that taxonomy can be found in Appendix E.

We test our approach with 3 different architectures designed for this dataset [18], while, in general, any neural network architecture can be used for mapping images onto the semantic space of class embeddings:

- Plain-11 [3], a recently proposed strictly sequential, shallow, wide, VGG-like network consisting of only 11 trainable layers, which has been found to achieve classification performance competitive to ResNet-110 when trained using cyclical learning rates.
- A variant of ResNet-110 [13], a deep residual network with 110 trainable layers. We use 32, 64, and 128 instead of 16, 32, and 64 channels per block, so that the number of features before the final fully-connected layer is greater than the number of classes. To make this difference obvious, we refer to this architecture as “ResNet-110w” in the following (“w” for “wide”).
- PyramidNet-272-200 [12], a deep residual network whose number of channels increases with every layer and not just after pooling.

Following [3], we train these networks using SGD with warm restarts (SGDR [21]), starting with a base learning rate of 0.1 and smoothly decreasing it over 12 epochs to a minimum of 10^{-6} using cosine annealing. The next cycle then begins with the base learning rate and the length of the cycles is doubled at the end of each one. All architectures are trained over a total number of 5 cycles (= 372 epochs) using a batch-size of 100. To prevent divergence caused by the initially high learning rate, we use gradient clipping [23] and restrict the norm of gradients to a maximum of 10.0.

5.1.2 North American Birds

The North American Birds (NAB) dataset² comprises 23,929 training and 24,633 test images showing birds from 555 different species. A 5 levels deep hierarchy of those classes is provided with the dataset.

Due to the comparatively small amount of training data, we do not only report results for training from scratch on the NAB dataset only, but also for fine-tuning models pre-trained on ILSVRC 2012 (see below).

During training, the images are randomly resized so that their smaller side is between 256 and 480 pixels wide and a random crop of size 224×224 pixels is extracted. Moreover, we apply random horizontal flipping and random erasing [38] for data augmentation.

We use the ResNet-50 architecture and train it for 4 cycles of SGDR (180 epochs) as described above, using a base learning rate of 0.5 and a batch size of 128 images. Only in case of our method with the $\mathcal{L}_{\text{CORR}}$ loss (*i.e.*, without classification objective), one additional SGDR cycle was required to achieve convergence (*i.e.*, 372 epochs). For fine-tuning, we did not change the learning rate schedule for our semantic embeddings, but chose the number of cycles and the learning rate individually for all competitors to achieve good performance without overfitting.

²<http://dl.allaboutbirds.org/nabirds>

5.1.3 ILSVRC 2012

We also conduct experiments on data from the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) 2012, which comprises over 1.1 million training and 50,000 test images from 1,000 classes. The classes originate from the WordNet ontology. Since this taxonomy is not a tree, we used the method of Redmon & Farhadi [25] described in Section 3.1 for deriving a tree-shaped hierarchy from WordNet. For the evaluation, however, we used the full WordNet taxonomy for computing class similarities.

We employ the ResNet-50 architecture [13] with a batch size of 128 images but only use a single cycle of SGDR, *i.e.*, just the cosine annealing without warm restarts, starting with a learning rate of 0.1 and annealing it down to 10^{-6} over 80 epochs. We used the same data augmentation techniques as He *et al.* [13] (flipping, scale augmentation, random cropping), except color augmentation, and use the weights provided by them as classification-based baseline to compare our method with.

5.2. Performance Metrics

Image retrieval tasks are often evaluated using the precision of the top k results ($P@k$) and mean average precision (mAP). However, these metrics do not take similarities among classes and varying misclassification costs into account. Thus, we introduce variants of them that are aware of the semantic relationships among classes.

Let $x_q \in \mathcal{X}$ be a query image belonging to class $y_q \in \mathcal{Y} = \{1, \dots, n\}$ and $R = ((x_1, y_1), \dots, (x_m, y_m))$ denote an ordered list of retrieved images $x_i \in \mathcal{X}$ and their associated classes $y_i \in \mathcal{Y}$. Following Deng *et al.* [7], we define the *hierarchical precision at k* ($HP@k$) for $1 \leq k \leq m$ as

$$HP@k(R) = \frac{\sum_{i=1}^k s_G(y_q, y_i)}{\max_{\pi} \sum_{i=1}^k s_G(y_q, y_{\pi_i})}, \quad (11)$$

where π denotes any permutation of the indices from 1 to m . Thus, the denominator in (11) normalizes the hierarchical precision through division by the sum of the precisions of the best possible ranking.

For several values of k , $HP@k$ can be plotted as a curve for gaining a more detailed insight into the behavior of a retrieval system regarding the top few results (short-term performance) and the overall ranking (long-term performance). We denote the area under that curve ranging from $k = 1$ to $k = K$ as *average hierarchical precision at K* ($AHP@K$). It can be used to compare the semantic performance of image retrieval systems by means of a single number. In the following, we always report the $AHP@250$, because we do not expect the typical user to inspect more than 250 results.




Figure 3: Hierarchical precision on CIFAR-100 at several cutoff points k for various network architectures.

5.3. Competitors

We evaluate two variants of our hierarchy-based semantic image embeddings: trained with $\mathcal{L}_{\text{CORR}}$ only (eq. (8)) or trained with the combination of correlation and categorical cross-entropy loss, $\mathcal{L}_{\text{CORR}+\text{CLS}}$, according to (9).

As a baseline, we compare them with features extracted from the last layer right before the final classification layer of the same network architecture, but trained solely for classification. We also evaluate the performance of L2-normalized versions of these classification-based features, which are usually used for CBIR (*cf.* Section 1).

Moreover, we compare with DeViSE [11], the center loss [35], and label embedding networks [32]. All methods have been applied to identical network architectures and trained with exactly the same optimization strategy, except DeViSE, which requires special training.

Since DeViSE [11] learns to map images onto word embeddings learned from text corpora, we were only able to apply it on CIFAR-100, since about one third of the 1,000 ILSVRC classes could not be matched to word embeddings automatically and the scientific bird names of the NAB dataset are not part of the vocabulary either. For CIFAR-100, however, we used 100-dimensional pre-computed GloVe word embeddings [24] learned from Wikipedia.

5.4. Semantic Image Retrieval Performance

For all datasets, we used each of the test images as individual query, aiming to retrieve semantically similar images from the remaining ones. Retrieval is performed by ranking the images in the database decreasingly according to their dot product with the query image in the feature space.

The mean values of HP@ k on CIFAR-100 over all queries at the first 250 values of k are reported in Fig. 3. It can clearly be seen that our hierarchy-based semantic embeddings achieve a much higher hierarchical precision than all competing methods. While the differences are rather

small when considering only very few top-scoring results, our method maintains a much higher precision when taking more retrieved images into account.

There is an interesting turning point in the precision curves at $k = 100$, where only our method suddenly starts increasing hierarchical precision. This is because there are exactly 100 test images per class, so that retrieving only images from exactly the same category as the query is not sufficient any more after this point. Instead, a semantically-aware retrieval method should continue retrieving images from the most similar classes at the later positions in the ranking, of which only our method is capable. It is possible for the HP@ k to improve after this point because previously missed good matches can be retrieved that result in a larger increase of the nominator of (11) than the denominator.

Additionally taking a classification objective into account during training improves the semantic retrieval performance on CIFAR-100 only slightly compared to using our simple $\mathcal{L}_{\text{CORR}}$ loss function alone. However, it is clearly beneficial on NAB and ILSVRC. This can be seen from the mAHP@250 reported for all datasets in Table 1. Our approach outperforms the second-best method for every dataset and network architecture by between 5% and 23% relative improvement of mAHP@250 on CIFAR-100, 36%-74% on NAB, and 16% on ILSVRC.

We also found that our approach outperforms all competitors in terms of classical mAP, which does not take class similarities into account. The detailed results as well as qualitative examples for ILSVRC are provided in Appendices B and C. Meanwhile, Fig. 1 shows some qualitative results on CIFAR-100.

5.5. Classification Performance

While our method achieves superior performance in the scenario of content-based image retrieval, we also wanted to make sure that it does not sacrifice classification perfor-

Method	CIFAR-100			NAB		ILSVRC
	Plain-11	ResNet-110w	PyramidNet	from scratch	fine-tuned	
Classification-based	0.5287	0.7261	0.6775	0.2538	0.5988	0.6831
Classification-based + L2 Norm	0.5821	0.7468	0.7334	0.2696	0.6259	0.7132
DeViSE [11]	0.6117	0.7348	0.7116	—	—	—
Center Loss [35]	0.6587	0.6815	0.6227	0.4185	0.4916	0.4094
Label Embedding [32]	0.6678	0.7950	0.7888	0.4251	0.5735	0.4769
Semantic Embeddings ($\mathcal{L}_{\text{CORR}}$) [ours]	0.8207	0.8290	0.8653	0.7157	0.7849	0.7902
Semantic Embeddings ($\mathcal{L}_{\text{CORR}+\text{CLS}}$) [ours]	0.8205	0.8329	0.8638	0.7399	0.8146	0.8242

Table 1: Retrieval performance of different image features in mAHP@250. The best value per column is set in bold font.

Method	CIFAR-100			NAB		ILSVRC
	Plain-11	ResNet-110w	PyramidNet	from scratch	fine-tuned	
Classification-based	72.18%	76.95%	81.44%	33.35%	70.17%	73.87%
DeViSE [11]	69.24%	74.66%	77.32%	—	—	—
Center Loss [35]	73.27%	75.18%	76.83%	56.07%	61.72%	70.05%
Label Embedding [32]	74.38%	76.96%	79.35%	48.05%	61.87%	70.94%
Semantic Embeddings ($\mathcal{L}_{\text{CORR}}$) [ours]	73.99%	75.03%	79.87%	57.52%	63.05%	48.97%
Semantic Embeddings ($\mathcal{L}_{\text{CORR}+\text{CLS}}$) [ours]	74.10%	76.60%	80.49%	59.46%	69.49%	69.18%

Table 2: Balanced classification accuracy of various methods. The best value in every column is set in bold font.

mance. Though classification was not the objective of our work, we have hence also compared the balanced classification accuracy obtained by all tested methods on all datasets and show the results in Table 2. It can be seen that learning to map images onto the semantic space of class embeddings does not impair classification performance unreasonably.

For DeVISE, which does not produce any class predictions, we have trained a linear SVM on top of the extracted features. The same issue arises for hierarchy-based semantic embeddings trained with $\mathcal{L}_{\text{CORR}}$, where we performed classification by assigning images to the class with the nearest embedding in the semantic feature space. Though doing so gives fair results, the accuracy becomes more competitive when training with $\mathcal{L}_{\text{CORR}+\text{CLS}}$ instead, *i.e.*, a combination of the embedding and the classification objective.

In the case of the NAB dataset, our method even obtains the best accuracy when training from scratch. Incorporating prior knowledge about class similarities obviously facilitates learning on a limited amount of training data.

6. Conclusions and Future Work

We proposed a novel method for integrating basic prior knowledge about the semantic relationships between classes, given as class taxonomy, into deep learning. Our hierarchy-based semantic embeddings preserve the semantic similarity of the classes in the joint space of image

and class embeddings and thus allow for retrieving images from a database that are not only *visually*, but also *semantically* similar to a given query image. This avoids unrelated matches and improves the quality of content-based image retrieval results significantly compared with other recent representation learning methods.

In contrast to other often used class representations such as text-based word embeddings, the hierarchy-based embedding space constructed by our method allows for straightforward training of neural networks by learning a regression of the class centroids using a simple loss function involving only a dot product.

The learned image features have also proven to be suitable for image classification, providing performance similar to that of networks trained explicitly for that task only.

Since the semantic target feature space is, by design, very specific to the classes present in the training set, generalization w.r.t. novel classes is still an issue. It thus seems promising to investigate the use of activations at earlier layers in the network, which we expect to be more general.

Acknowledgements

This work was supported by the German Research Foundation as part of the priority programme “Volunteered Geographic Information: Interpretation, Visualisation and Social Computing” (SPP 1894, contract number DE 735/11-1). We also gratefully acknowledge the support of NVIDIA Corporation with the donation of Titan Xp GPUs used for this research.

References

- [1] A. Babenko and V. Lempitsky. Aggregating local deep features for image retrieval. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1269–1277, 2015.
- [2] A. Babenko, A. Slesarev, A. Chigorin, and V. Lempitsky. Neural codes for image retrieval. In *European Conference on Computer Vision (ECCV)*, pages 584–599. Springer, 2014.
- [3] B. Barz and J. Denzler. Deep learning is not a matter of depth but of good training. In *International Conference on Pattern Recognition and Artificial Intelligence (ICPRAI)*, pages 683–687. CENPARMI, Concordia University, Montreal, 2018.
- [4] J. Y. Chang and K. M. Lee. Large margin learning of hierarchical semantic similarity for image classification. *Computer Vision and Image Understanding*, 132:3–11, 2015.
- [5] W. Chen, X. Chen, J. Zhang, and K. Huang. Beyond triplet loss: a deep quadruplet network for person re-identification. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 403–412, 2017.
- [6] S. Chopra, R. Hadsell, and Y. LeCun. Learning a similarity metric discriminatively, with application to face verification. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 539–546. IEEE, 2005.
- [7] J. Deng, A. C. Berg, and L. Fei-Fei. Hierarchical semantic indexing for large scale image retrieval. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 785–792. IEEE, 2011.
- [8] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255. IEEE, 2009.
- [9] T. Deselaers, R. Paredes, E. Vidal, and H. Ney. Learning weighted distances for relevance feedback in image retrieval. In *International Conference on Pattern Recognition (ICPR)*, pages 1–4. IEEE, 2008.
- [10] C. Fellbaum. *WordNet*. Wiley Online Library, 1998.
- [11] A. Frome, G. S. Corrado, J. Shlens, S. Bengio, J. Dean, T. Mikolov, et al. Devise: A deep visual-semantic embedding model. In *Advances in Neural Information Processing Systems (NIPS)*, pages 2121–2129, 2013.
- [12] D. Han, J. Kim, and J. Kim. Deep pyramidal residual networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5927–5935, 2017.
- [13] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- [14] S. Horiguchi, D. Ikami, and K. Aizawa. Significance of softmax-based features in comparison to distance metric learning-based features. *arXiv preprint arXiv:1712.10151*, 2017.
- [15] G. Huang, Z. Liu, L. van der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. 2017.
- [16] S. S. Husain and M. Bober. Improving large-scale image retrieval through robust aggregation of local descriptors. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 39(9):1783–1796, 2017.
- [17] H. Jégou and A. Zisserman. Triangulation embedding and democratic aggregation for image search. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3310–3317. IEEE, 2014.
- [18] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- [19] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1097–1105, 2012.
- [20] D. Li, H.-Y. Lee, J.-B. Huang, S. Wang, and M.-H. Yang. Learning structured semantic embeddings for visual recognition. *arXiv preprint arXiv:1706.01237*, 2017.
- [21] I. Loshchilov and F. Hutter. Sgdr: Stochastic gradient descent with warm restarts. In *International Conference on Learning Representations (ICLR)*, 2017.
- [22] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems (NIPS)*, pages 3111–3119, 2013.
- [23] R. Pascanu, T. Mikolov, and Y. Bengio. On the difficulty of training recurrent neural networks. In *International Conference on Machine Learning (ICML)*, pages 1310–1318, 2013.
- [24] J. Pennington, R. Socher, and C. D. Manning. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, 2014.
- [25] J. Redmon and A. Farhadi. Yolo9000: better, faster, stronger. *arXiv preprint arXiv:1612.08242*, 2016.
- [26] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015.
- [27] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 815–823, 2015.
- [28] F. Setti. To know and to learn – about the integration of knowledge representation and deep learning for fine-grained visual categorization. In *VISIGRAPP (5: VISAPP)*, pages 387–392, 2018.
- [29] A. Sharif Razavian, H. Azizpour, J. Sullivan, and S. Carlsson. Cnn features off-the-shelf: an astounding baseline for recognition. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPR-WS)*, pages 806–813, 2014.
- [30] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [31] A. W. Smeulders, M. Worring, S. Santini, A. Gupta, and R. Jain. Content-based image retrieval at the end of the early years. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 22(12):1349–1380, 2000.
- [32] X. Sun, B. Wei, X. Ren, and S. Ma. Label embedding network: Learning label representation for soft training of deep networks. *arXiv preprint arXiv:1710.10393*, 2017.

- [33] N. Verma, D. Mahajan, S. Sellamanickam, and V. Nair. Learning hierarchical similarity metrics. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2280–2287. IEEE, 2012.
- [34] K. Q. Weinberger and O. Chapelle. Large margin taxonomy embedding for document categorization. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1737–1744, 2009.
- [35] Y. Wen, K. Zhang, Z. Li, and Y. Qiao. A discriminative feature learning approach for deep face recognition. In *European Conference on Computer Vision (ECCV)*, pages 499–515. Springer, 2016.
- [36] Z. Yan, H. Zhang, R. Piramuthu, V. Jagadeesh, D. DeCoste, W. Di, and Y. Yu. Hd-cnn: hierarchical deep convolutional neural networks for large scale visual recognition. In *IEEE International Conference on Computer Vision (ICCV)*, pages 2740–2748, 2015.
- [37] B. Zhao, F. Li, and E. P. Xing. Large-scale category structure aware image categorization. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1251–1259, 2011.
- [38] Z. Zhong, L. Zheng, G. Kang, S. Li, and Y. Yang. Random erasing data augmentation. *arXiv preprint arXiv:1708.04896*, 2017.

Appendices

Appendix A. d_G applied to trees is a metric

Theorem 1. Let $\mathcal{G} = (V, E)$ be a directed acyclic graph whose edges $E \subseteq V \times V$ define a hyponymy relation between the semantic concepts in V . Furthermore, let \mathcal{G} have exactly one unique root node $\text{root}(\mathcal{G})$ with indegree $\deg^-(\text{root}(\mathcal{G})) = 0$. The lowest common subsumer $\text{lcs}(u, v)$ of two concepts $u, v \in V$ hence always exists. Moreover, let $\text{height}(u)$ denote the maximum length of a path from $u \in V$ to a leaf node, and $\mathcal{C} \subseteq V$ be a set of classes of interest.

Then, the semantic dissimilarity $d_G : \mathcal{C} \times \mathcal{C} \rightarrow \mathbb{R}$ between classes given by

$$d_G(u, v) = \frac{\text{height}(\text{lcs}(u, v))}{\max_{w \in V} \text{height}(w)} \quad (12)$$

is a proper metric if

- (a) \mathcal{G} is a tree, i.e., all nodes $u \in V \setminus \{\text{root}(\mathcal{G})\}$ have indegree $\deg^-(u) = 1$, and
- (b) all classes of interest are leaf nodes of the hierarchy, i.e., all $u \in \mathcal{C}$ have outdegree $\deg^+(u) = 0$.

Proof. For being a proper metric, d_G must possess the following properties:

- (i) Non-negativity: $d_G(u, v) \geq 0$.
- (ii) Symmetry: $d_G(u, v) = d_G(v, u)$.
- (iii) Identity of indiscernibles: $d_G(u, v) = 0 \Leftrightarrow u = v$.
- (iv) Triangle inequality: $d_G(u, w) \leq d_G(u, v) + d_G(v, w)$.

The conditions (i) and (ii) are always satisfied since $\text{height} : V \rightarrow \mathbb{R}$ is defined as the length of a path, which cannot be negative, and the lowest common subsumer (LCS) of two nodes is independent of the order of arguments.

The proof with respect to the remaining properties (iii) and (iv) can be conducted as follows:


- (b)→(iii)** Let $u, v \in \mathcal{C}$ be two classes with $d_G(u, v) = 0$. This means that their LCS has height 0 and hence must be a leaf node. Because leaf nodes have, by definition, no further children, $u = \text{lcs}(u, v) = v$. On the other hand, for any class $w \in \mathcal{C}$, $d_G(w, w) = 0$ because $\text{lcs}(w, w) = w$ and w is a leaf node according to (b).
- (a)→(iv)** Let $u, v, w \in \mathcal{C}$ be three classes. Due to (a), there exists exactly one unique path from the root of the hierarchy to any node. Hence, $\text{lcs}(u, v)$ and $\text{lcs}(v, w)$ both lie on the path from $\text{root}(\mathcal{G})$ to v and they are, thus, either identical or one is an ancestor of the other. Without loss of generality, we assume that $\text{lcs}(u, v)$ is an ancestor of $\text{lcs}(v, w)$ and thus lies on the root-paths to u, v , and w . In particular, $\text{lcs}(u, v)$ is a subsumer of u and w and, therefore, $\text{height}(\text{lcs}(u, w)) \leq \text{height}(\text{lcs}(u, v))$. In general, it follows that $d_G(u, w) \leq \max\{d_G(u, v), d_G(v, w)\} \leq d_G(u, v) + d_G(v, w)$, given the non-negativity of d_G .

□


Remark regarding the inversion

If d_G is a metric, all classes $u \in \mathcal{C}$ of interest must necessarily be leaf nodes, since $d_G(u, u) = 0 \Rightarrow \text{height}(\text{lcs}(u, u)) = \text{height}(u) = 0$.

However, (iv)→(a) does not hold in general, since d_G may even be a metric for graphs \mathcal{G} that are not trees. An example is given in Fig. 4a. Nevertheless, most such graphs violate the triangle inequality, like the example shown in Fig. 4b.



(a) A non-tree hierarchy where d_G is a metric.



(b) A non-tree hierarchy where d_G violates the triangle inequality.

Figure 4: Examples for non-tree hierarchies.

Appendix B. Further Quantitative Results




Figure 5: Hierarchical precision on NAB and ILSVRC 2012.

Method	CIFAR-100			NAB		ILSVRC
	Plain-11	ResNet-110w	PyramidNet	from scratch	fine-tuned	
Classification-based	0.2078	0.4870	0.3643	0.0283	0.2771	0.2184
Classification-based + L2 Norm	0.2666	0.5305	0.4621	0.0363	0.3194	0.2900
DeViSE	0.2879	0.5016	0.4131	—	—	—
Center Loss	0.4180	0.4153	0.3029	0.1591	0.1802	0.1285
Label Embedding	0.2747	0.6202	0.5920	0.1271	0.2417	0.2683
Semantic Embeddings ($\mathcal{L}_{\text{CORR}}$) [ours]	0.5660	0.5900	0.6642	0.4249	0.5246	0.3037
Semantic Embeddings ($\mathcal{L}_{\text{CORR}+\text{CLS}}$) [ours]	0.5886	0.6107	0.6808	0.4316	0.5768	0.4508

Table 3: Classical mean average precision (mAP) on all datasets. The best value per column is set in bold font. Obviously, optimizing for a classification criterion only leads to sub-optimal features for image retrieval.

Appendix C. Qualitative Results on ILSVRC 2012

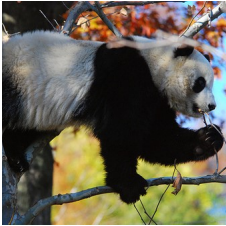








































































































Query	#1	#11	#21	#31	#41	#51	#61	#71	#81	#91	
 giant panda											Classification Features
											Semantic Embeddings
 great grey owl											Classification Features
											Semantic Embeddings
 monarch											Classification Features
											Semantic Embeddings
 ice bear											Classification Features
											Semantic Embeddings
 ice cream											Classification Features
											Semantic Embeddings

Figure 6: Comparison of a subset of the top 100 retrieval results using L2-normalized classification-based and our hierarchy-based semantic features for 3 exemplary queries on ILSVRC 2012. Image captions specify the ground-truth classes of the images and the border color encodes the semantic similarity of that class to the class of the query image, with dark green being most similar and dark red being most dissimilar.

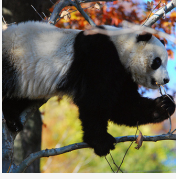

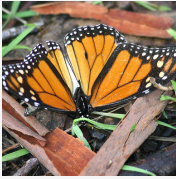


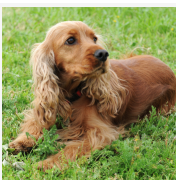
Image	Classification-based		Semantic Embeddings (ours)	
	1. <u>Giant Panda</u>	(1.00)	1. <u>Giant Panda</u>	(1.00)
	2. American Black Bear	(0.63)	2. Lesser Panda	(0.89)
	3. Ice Bear	(0.63)	3. Colobus	(0.58)
	4. Gorilla	(0.58)	4. American Black Bear	(0.63)
	5. Sloth Bear	(0.63)	5. Guenon	(0.58)
	1. <u>Great Grey Owl</u>	(1.00)	1. <u>Great Grey Owl</u>	(1.00)
	2. Sweatshirt	(0.16)	2. Kite	(0.79)
	3. Bonnet	(0.16)	3. Bald Eagle	(0.79)
	4. Guenon	(0.42)	4. Vulture	(0.79)
	5. African Grey	(0.63)	5. Ruffed Grouse	(0.63)
	1. <u>Monarch</u>	(1.00)	1. <u>Monarch</u>	(1.00)
	2. Earthstar	(0.26)	2. Cabbage Butterfly	(0.84)
	3. Coral Fungus	(0.26)	3. Admiral	(0.84)
	4. Stinkhorn	(0.26)	4. Sulphur Butterfly	(0.84)
	5. Admiral	(0.84)	5. Lycaenid	(0.84)
	1. <u>Ice Bear</u>	(1.00)	1. <u>Ice Bear</u>	(1.00)
	2. Arctic Fox	(0.63)	2. Brown Bear	(0.95)
	3. White Wolf	(0.63)	3. Sloth Bear	(0.95)
	4. Samoyed	(0.63)	4. Arctic Fox	(0.63)
	5. Great Pyrenees	(0.63)	5. American Black Bear	(0.95)
	1. <u>Ice Cream</u>	(1.00)	1. <u>Ice Cream</u>	(1.00)
	2. Meat Loaf	(0.63)	2. Ice Lolly	(0.84)
	3. Bakery	(0.05)	3. Trifle	(0.89)
	4. Strawberry	(0.32)	4. Chocolate Sauce	(0.58)
	5. Fig	(0.32)	5. Plate	(0.79)
	1. <u>Cocker Spaniel</u>	(1.00)	1. <u>Cocker Spaniel</u>	(1.00)
	2. Irish Setter	(0.84)	2. Sussex Spaniel	(0.89)
	3. Sussex Spaniel	(0.89)	3. Irish Setter	(0.84)
	4. Australien Terrier	(0.79)	4. Welsh Springer Spaniel	(0.89)
	5. Clumber	(0.89)	5. Golden Retriever	(0.84)

Figure 7: Top 5 classes predicted for several example images by a ResNet-50 trained purely for classification and by our network trained with $\mathcal{L}_{\text{CORR}+\text{CLS}}$ incorporating semantic information. The correct label for each image is underlined and the numbers in parentheses specify the semantic similarity of the predicted class and the correct class. It can be seen that class predictions made based on our hierarchy-based semantic embeddings are much more relevant and consistent.

Appendix D. Low-dimensional Semantic Embeddings




Figure 8: Hierarchical precision of our method for learning image representations based on class embeddings with varying dimensionality, compared with the usual baselines.

As can be seen from the description of our algorithm for computing class embeddings in Section 3.2, an embedding space with n dimensions is required in general to find an embedding for n classes that reproduces their semantic similarities exactly. This can become problematic in settings with a high number of classes.

For such scenarios, we have proposed a method for computing low-dimensionality embeddings of arbitrary dimensionality approximating the actual relationships among classes in Section 3.3. paper. We experimented with this possibility on the NAB dataset, learned from scratch, to see how reducing the number of features affects our algorithm for learning image representations and the semantic retrieval performance.

The results in Fig. 8 show that obtaining low-dimensional class embeddings through eigendecomposition is a viable option for settings with a high number of classes. Though the performance is worse than with the full amount of required features, our method still performs better than the competitors with as few as 16 features. Our approach hence also allows obtaining very compact image descriptors, which is important when dealing with huge datasets.

Interestingly, the 256-dimensional approximation even gives slightly better results than the full embedding after the first 50 retrieved images. We attribute this to the fact that fewer features leave less room for overfitting, so that slightly lower-dimensional embeddings can generalize better in this scenario.

Appendix E. Taxonomy used for CIFAR-100

