



HAL
open science

Integrating Visual and Semantic Similarity Using Hierarchies for Image Retrieval

Aishwarya Venkataramanan, Martin Laviale, Cédric Pradalier

► **To cite this version:**

Aishwarya Venkataramanan, Martin Laviale, Cédric Pradalier. Integrating Visual and Semantic Similarity Using Hierarchies for Image Retrieval. Computer Vision Systems. ICVS 2023, Sep 2023, Vienne (AUT), Austria. pp.422-431, 10.1007/978-3-031-44137-0_35 . hal-04423201

HAL Id: hal-04423201


<https://hal.univ-lorraine.fr/hal-04423201>

Submitted on 22 Feb 2024

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

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

Integrating Visual and Semantic Similarity Using Hierarchies for Image Retrieval

Aishwarya Venkataramanan^{1,2,3}[0000-0002-6100-0034] , Martin Laviale^{1,3}[0000-0002-9719-7158], and Cédric Pradalier^{2,3}[0000-0002-1746-2733]

¹ Laboratoire Interdisciplinaire des Environnements Continentaux, Université de Lorraine, Metz, France aishwarya.venkataramanan@univ-lorraine.fr

² Georgia Tech Europe- International Research Lab Georgia Tech - CNRS IRL 2958, Metz, France

³ LTER- “Zone Atelier Moselle”, France

Abstract. Most of the research in content-based image retrieval (CBIR) focus on developing robust feature representations that can effectively retrieve instances from a database of images that are visually similar to a query. However, the retrieved images sometimes contain results that are not semantically related to the query. To address this, we propose a method for CBIR that captures both visual and semantic similarity using a visual hierarchy. The hierarchy is constructed by merging classes with overlapping features in the latent space of a deep neural network trained for classification, assuming that overlapping classes share high visual and semantic similarities. Finally, the constructed hierarchy is integrated into the distance calculation metric for similarity search. Experiments on standard datasets: CUB-200-2011 and CIFAR100, and a real-life use case using diatom microscopy images show that our method achieves superior performance compared to the existing methods on image retrieval.

Keywords: Similarity Search · Visual Hierarchy · Content-based image retrieval

1 Introduction

In recent years, the exponential growth of image data has made effective management of such data increasingly important. One way to manage this data is through Content-based Image Retrieval (CBIR), where a query image is used to retrieve images from the database that are visually similar to it [5]. The comparison is performed on the features extracted from the images, represented by a set of numerical descriptors or vectors. The search is performed on these descriptors using a similarity measure such as the Euclidean or cosine distance. The challenge is to extract robust feature representations, such that objects belonging to the same category are close by, despite the changes in view-point, background and illumination. Recent methods employing Convolutional Neural Network (CNN) for feature extraction have proven to be robust and effective over the classical hand-crafted feature extractors [16]. A commonly adopted approach is to extract feature vectors from the latent space of a trained CNN

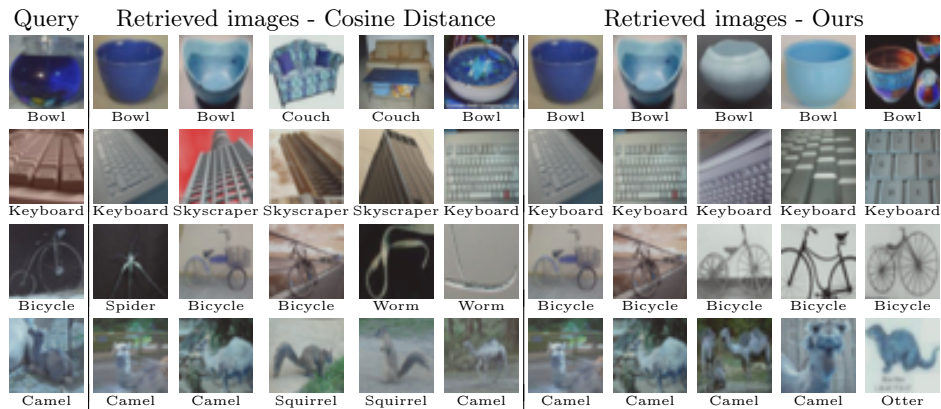


Fig. 1. Examples of top-5 images retrieved using cosine distance and our method on CIFAR100. Our method incorporates hierarchy while ranking image similarity, with retrieves images that are more visually and semantically relevant.

classifier. Deep classifiers use cross-entropy loss that learns feature representations to group together objects belonging to the same category, while separating them from the rest. Thus, given a query image, the retrieved data from the local neighbourhood in the latent space should ideally consist of images belonging to the same classification category.

Despite the ability of the CNNs to learn discriminative representations, it has been observed that sometimes, images unrelated to the query are retrieved [1,10]. These images are typically visually similar to the query, but semantically very different. Fig. 1 illustrates this for the images retrieved using cosine distance on CIFAR100 [8]. While the retrieved images have certain visual similarities with the query, some of them are semantically very different. One way to address this is by incorporating domain knowledge while training, so that the network learns feature representations that are both visually and semantically relevant [1]. The semantic data is obtained from expert domain knowledge or lexical databases such as WordNet [9]. However, incorporating external domain knowledge during training can be challenging since it would require significant changes in the training procedure, and can be detrimental when there are inconsistencies between visual and semantic similarities [3]. Also, it assumes that the specialized domain knowledge is available, which can be hard to obtain or irrelevant in certain situations, such as, querying the images of people or vehicles. In this paper, we demonstrate that the visual hierarchies obtained from the representations learned by the classification network are rich in visual and semantic information, achieving superior retrieval performance. Contrary to other approaches, our method does not require specialized domain knowledge or modifications to the training procedure, making it easy to integrate into an off-the-shelf classification network.

Our method relies on two key aspects for the retrieval: construction of a hierarchy that is visually and semantically meaningful, and integration of hierarchy into the distance calculation metric for similarity search. Our method for hierarchy construction involves identifying the classes in the latent space whose feature representations overlap with one another. This is based on the assumption that the overlapping classes are those that are very similar to each other, both visually and semantically. The clusters of overlapping classes are then grouped together into the same class at a higher hierarchy level. This process is repeated at different levels of the hierarchy until there are no overlaps in the feature representations, resulting in a hierarchical structure that captures both visual and semantic similarities between images. Once the hierarchy is obtained, we combine the cosine distance and a hierarchy-based distance to obtain a metric for ranking the images for retrieval. Experimental results show that our method improves image retrieval performance on two standard benchmarks: CIFAR100 [8] and CUB-200-2011 [15], and a real life diatom image retrieval use case. Additionally, we demonstrate the robustness of our method by evaluating on images subject to perturbations that mimic practical difficulties encountered during acquisition.

Our contributions are as follows: (i) we develop a method to construct a hierarchy that is visually and semantically relevant and show that incorporating this hierarchy in the distance metric enhances the image retrieval performance, (ii) we design the method such that it does not require any additional domain knowledge and introduces no modifications to the model architecture or training procedure. The code is available in <https://github.com/vaishwarya96/Hierarchy-image-retrieval.git>

2 Related Works

The role of visual attributes and semantics is crucial in describing images for identification or retrieval. Hierarchies are an intuitive way to express this, and several methods have leveraged it for image categorization and retrieval [11]. This is done by integrating hierarchy obtained from the biological taxonomy or from lexical databases to learn feature representations that capture semantic similarity between the categories. [4] uses WordNet to construct hierarchies and incorporates the domain knowledge obtained from it into a CNN classifier to improve the network’s performance. [2] extends the traditional cross-entropy loss used in flat classification to work on hierarchies. To do so, they compute the conditional probability of the nodes at each layer of the hierarchy and calculate the cross-entropy loss at each layer. DeVISE [6] transfers semantic knowledge from text to visual object recognition. It involves pre-training a language model to learn word embeddings and re-training a deep visual model using these embeddings for image label prediction. [1] uses an embedding algorithm that maps images onto a hypersphere. The distances on this hypersphere represent similarities that are derived from the lowest common ancestor height in a given hierarchy tree. One limitation of these methods is their reliance on an exist-

ing domain knowledge, which can be difficult to obtain or irrelevant in some contexts.

A line of work relies on the visual features learnt by the network. [18] uses metric loss while training the network for image retrieval. Metric loss tries to pull closer the representations of images belonging to the same class, while pushing apart the inter-class representations. While this reinforces the condition that the intra-class representations are clustered tightly, it does not impact the semantic relationship between the instances. [19] constructs a visual hierarchy from the distance values between the image features and uses it to improve classification. Our method is similar to [10], which uses Ward clustering for constructing the hierarchy and uses it for re-ranking the images retrieved using a distance metric. A drawback of using Ward clustering is that they are sensitive to outliers and can sometimes group classes with low semantic similarities. In contrast, our approach groups classes only when their features overlap, which results in a more meaningful and semantically relevant hierarchy.

3 Method

This section explains the two main steps of our method: the construction of the hierarchy from the learned feature representations, and the calculation of the distance metric for ranking the images.

Notations: Consider a CNN classifier trained on K classes, each class represented as $\{C_1, \dots, C_K\}$. The feature vectors are extracted from the penultimate layer of the CNN. We apply dimensionality reduction on these feature vectors using PCA to reduce the redundant features and speed up the computation [17]. We automatically determine the number of principal components required to explain 95% of the original data variance by examining the eigenvalues in decreasing order. Let the feature vectors corresponding to the training images be x_{train} . Assume that the class features follow multivariate Gaussian distributions with mean μ_c and covariance Σ_c , where $c = \{1, \dots, K\}$. Note that this assumption can be broken when the intra-class variance is large. In such cases, one can constrain the latent space to be Gaussian using the methodology in [13]. Given a query image q , its feature vector is denoted as x_q . Suppose the database consists of N images, their corresponding feature vectors are denoted as $\{x_1, \dots, x_N\}$.

3.1 Hierarchy Construction

The hierarchy construction involves identifying the classes with overlapping distributions of feature vectors in x_{train} . Inter-class overlapping typically occurs when there are visual and semantic similarities in the class images. To construct the hierarchy, a bottom-up approach is used, where the overlapping classes in the latent space are merged in the higher hierarchy level. The overlapping classes are identified using the Bhattacharya Coefficient (BC), which is a statistical metric to measure the amount of overlap between two statistical populations.

The Bhattacharya distance between two classes C_i and C_j is given by:

$$D_B(C_i, C_j) = \frac{1}{8}(\mu_i - \mu_j)^T \Sigma^{-1}(\mu_i - \mu_j) + \frac{1}{2} \ln \left(\frac{\det \Sigma}{\sqrt{\det \Sigma_i \det \Sigma_j}} \right) \quad (1)$$

where $\Sigma = \frac{\Sigma_i + \Sigma_j}{2}$

The BC is calculated as:

$$BC(C_i, C_j) = e^{-D_B(C_i, C_j)} \quad (2)$$

The value of BC ranges from 0 to 1 inclusive, where a value of 0 signifies no overlap and 1 means full overlap. If the value of the BC between two classes is greater than or equal to a specified threshold t , then they are considered to be overlapping. Every cluster of classes that overlap with each other are identified. All classes within each cluster are merged into a unified class at a higher level of the hierarchy. On the other hand, if the BC is less than t , then these classes are considered to be distinct and are kept separate in the higher hierarchy level. The t value is a hyperparameter that depends on the dataset. The process of BC estimation is repeated at each level on the new set of classes, and the overlapping classes are merged to create a hierarchy structure. This is performed until there are no overlaps between the classes, or all the classes have been merged. At the higher hierarchy levels, some classes include feature vectors from multiple small classes that were merged. At every level, the μ and Σ values are recalculated from the feature vectors of the merged classes.

3.2 Distance Calculation

After constructing the hierarchy, we use it to perform the similarity search for image retrieval. The distances are calculated between the query feature x_q and the image features x_i in the database. Our final distance metric for ranking the images is a weighted sum of two distance values: cosine distance and a hierarchy based distance.

Cosine distance: The cosine distance is a similarity measure between two feature vectors and is popularly used in image retrieval tasks. The cosine distance between x_q and x_i is given by:

$$D_C(x_q, x_i) = 1 - \frac{x_q \cdot x_i}{\|x_q\| \|x_i\|} \quad (3)$$

where \cdot is the dot product and $\| \cdot \|$ is the L2 norm.

Hierarchical distance: The hierarchy is a tree structure $\mathcal{H} = (V, E)$ where V represents the nodes and E , the edges. \mathcal{H} consists of K leaf nodes. We want to identify the position of any database image in the hierarchy, which means finding the leaf closest to each image. To identify the leaf nodes, we calculate

the Mahalanobis Distance (MD) between the feature vectors x_i and the feature vectors x_{train} corresponding to each leaf node. The MD between x_i and node c , is given as:

$$MD_c(x_i) = \sqrt{(x_i - \mu_c)^T \Sigma_c^{-1} (x_i - \mu_c)} \quad (4)$$

x_i is assigned to the leaf node n_i corresponding to the smallest MD value.

Similarly, the leaf node n_q corresponding to x_q is identified. The distance calculation uses the height of the lowest common ancestor (LCA), where the LCA of two nodes is the parent node that is an ancestor of both nodes and does not have any child that is also an ancestor of both nodes. The hierarchical distance [1] between x_q and x_i is calculated as:

$$D_H(x_q, x_i) = \frac{\text{height}(\text{LCA}(n_i, n_q))}{\text{height}(\mathcal{H})} \quad (5)$$

The final distance is given as:

$$D = D_C(x_q, x_i) + \alpha D_H(x_q, x_i) \quad (6)$$

where α is a hyperparameter.

The images are sorted in ascending order of D to rank the images based on similarity and retrieve them.

3.3 Robustness Analysis

Robustness analysis demonstrates the ability of the image retrieval methods to maintain accurate and reliable performance even when the images are distorted or corrupted. The distortions introduced during the robustness analysis mimic real-world scenarios where images can be affected by factors like blurring, occlusions, and changes in lighting conditions.

We consider the following scenarios: (i) Blurring – The images are subject to Gaussian blurring, simulating the effect of an out-of-focus camera or a low-quality image (ii) Saturation change – The saturation values of the original images are modified to simulate variations in colour intensity. (iii) Occlusions – Random occlusions are introduced in different parts of the images, simulating objects or regions being partially or completely blocked.

4 Experiments

4.1 Baselines

We compare our method against the following approaches: (i) Image retrieval using Euclidean Distance, (ii) Image retrieval using Cosine Distance, (iii) Our method using hierarchical distance only, (iv) Semantic Embeddings [1], (v) Our method applied on the features extracted from semantic embeddings in (iv), (vi) Distance calculation using the method in Sec. 3.2 on hierarchy constructed using Ward clustering [10], (vii) Distance calculation on semantic hierarchy obtained using domain knowledge.

4.2 Experimental Setup

We perform our experiments on three datasets with fine-grained visual features:

CUB-200-2011 [15]. This is a widely used dataset for image retrieval, consisting of 200 classes of birds and a total of 11,788 images. The train data consists of 5,994 images and the test data consists of 5,794 images. The images were resized to 224×224 , and random horizontal flipping and rotation ($\pm 30^\circ$) were used for data augmentation. Following [1], we used ResNet-50 [7] for training with an SGD optimizer. The network was trained for 100 epochs with a batch size of 12. The values of $t = 0.30$ and $\alpha = 3$ are obtained using hyperparameter search.

CIFAR100 [8]. It consists of 60,000 images, of size 32×32 belonging to 100 classes of wide range of objects. The training dataset consists of 50,000 images and the remaining are test data. Following [1], we used ResNet-110 [7] for training with an SGD optimizer. The network was trained for 400 epochs with a batch size of 24. The values of $t = 0.20$ and $\alpha = 5$ are obtained using hyperparameter search.

Diatom Dataset. The diatom dataset consists of 9895 individual diatom images belonging to 166 classes [14]. The network was trained using EfficientNet [12]. An SGD optimizer with a learning rate of 0.0002 was used. The images were resized to 256×256 . This dataset contains high intra-class variance with distinct feature representations per class, and hence the latent space is not Gaussian, making it incompatible for BC calculation. The network was trained using the method in [13], which ensures a Gaussian latent space. The training was performed for 100 epochs with a batch size of 12. The values of $t = 0.25$ and $\alpha = 3$ are obtained using hyperparameter search.

For the CIFAR100 and CUB-200-2011, 80% of the train data was used for training and 20% for validation. For the diatoms, 70% of the dataset was used for training, 10% for validation and 20% for testing. The networks were trained on GeForce RTX 3090 with 24 Gb of RAM. The semantic hierarchy of CUB-200-2011 and CIFAR100 for the baseline experiments were obtained from [1]. The diatom hierarchy was obtained from the taxonomy of the species, consisting of two levels: 'genus' and 'species'. The same model architectures were used for comparing all the baselines.

Setup for Robustness Analysis. The analysis is performed on the CUB-200-2011 dataset with the following settings. (i) A Gaussian Blur with kernel size (11,11) was applied to all the images, (ii) The saturation of images were adjusted with randomly chosen values in the range [0,1,0.5,1,2,5] and (iii) A random crop of dimension (70,70) was replaced by black patch in every image. The networks were retrained on the modified images, with the same training parameters.

4.3 Evaluation Metrics

We use MAP@k values for evaluating the image retrieval performance. MAP@k measures the average precision (AP) of the top k retrieved images, and then takes the mean across all queries. A higher value indicates better performance.




Fig. 2. Retrieval performance on CUB-200-2011, CIFAR100 and Diatom dataset. Our method achieves state-of-the-art performance. The results suggest that incorporating hierarchy into content-based image retrieval is effective in improving the retrieval performance.

5 Results

5.1 Retrieval Results

Fig. 2 shows the plots of $mAP@k$ values for different values of k for the three datasets. The results suggest that combining the visual and semantic cues improves the retrieval. Overall, our method achieves the best performance over the other baselines.

Effect of distances: The cosine distance retrieves images that are visually similar to the query, whereas, the hierarchical distance takes into account the semantic relationships between classes. Both these aspects are important when retrieving the images, and removing one of them deteriorates the performance. There is a significant drop in the MAP values in CIFAR100 when using only the hierarchical distance. One explanation for this is that CIFAR100 contains classes that are visually diverse, and hence, visual cues are more important over the semantic ones here. This could also explain the marginal improvement in the metrics when combining both the visual and semantic attributes than when using only the visual information for CIFAR100 over the other datasets.

Effect of hierarchy: In the hierarchy-based methods, our approach outperforms the baselines. This is closely followed by our approach of distance calculation using semantic hierarchy. While semantic embeddings [1] also relies on visual and semantic features for the retrieval, its drop in performance could be attributed to the errors introduced due to disagreements between the visual and semantic relationships while training [3]. Our approach of hierarchy construction merges the classes only when there is a correlation between visual and semantic attributes, and thus it is robust against these errors. The Ward hierarchy approach may sometimes group classes that are semantically unrelated, resulting in less effective performance.

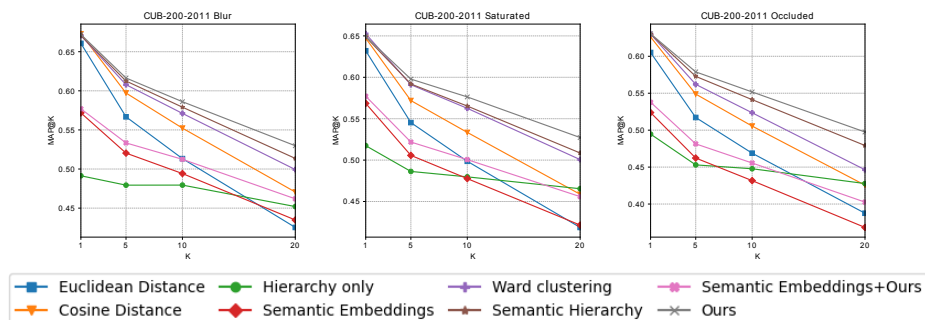


Fig. 3. Robustness analysis on the CUB-200-2011 dataset. Blurring, saturation change and occlusion were applied to test the impact on the retrieval. Results show that our method still achieves the best performance.

Qualitative Analysis: Fig. 1 shows some images retrieved using cosine distance and our method. While cosine distance retrieves images that are visually similar, some of them are semantically very different. Whereas, our method retrieves images that are visually and semantically meaningful.

Robustness Analysis: Fig. 3 shows the MAP@K plots for the robustness analysis on CUB-200-2011. Compared to the plot in Fig. 2, one can observe a drop in the metrics. However, our method still achieves the best performance over the other baselines, indicating its robustness even in challenging conditions.

6 Discussion and Conclusion

In this paper, we presented a method to tackle the problem of semantically dissimilar images retrieved in CBIR. By leveraging the learned feature representations from a CNN classifier, we construct a meaningful hierarchy that captures both visual and semantic information. This hierarchy is then integrated into the distance metric for similarity search, resulting in superior image retrieval performance, as demonstrated on fine-grained visual datasets. One limitation of our method arises when classes exhibit high variance in their features, resulting in potential overlaps with several or all other classes. Consequently, the network may incorrectly group them all into a single class when constructing the hierarchy. To mitigate this, we only group classes when the degree of overlap is above a threshold. One could also train the network using a metric loss to alleviate this, which we leave for future work.

References

1. Barz, B., Denzler, J.: Hierarchy-based image embeddings for semantic image retrieval. In: 2019 IEEE WACV. pp. 638–647. IEEE (2019)

2. Bertinetto, L., Mueller, R., Tertikas, K., Samangoeei, S., Lord, N.A.: Making better mistakes: Leveraging class hierarchies with deep networks. In: Proceedings of the IEEE/CVF CVPR. pp. 12506–12515 (2020)
3. Brust, C.A., Denzler, J.: Not just a matter of semantics: The relationship between visual and semantic similarity. In: 41st DAGM GCPR 2019, Dortmund, Germany, September 10–13, 2019, Proceedings 41. pp. 414–427. Springer (2019)
4. Brust, C.A., Denzler, J.: Integrating domain knowledge: using hierarchies to improve deep classifiers. In: ACPR 2019, Auckland, New Zealand. pp. 3–16. Springer (2020)
5. Chen, W., Liu, Y., Wang, W., Bakker, E.M., Georgiou, T., Fieguth, P., Liu, L., Lew, M.S.: Deep learning for instance retrieval: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2022)
6. Frome, A., Corrado, G.S., Shlens, J., Bengio, S., Dean, J., Ranzato, M., Mikolov, T.: Devise: A deep visual-semantic embedding model. *Advances in neural information processing systems* **26** (2013)
7. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE CVPR. pp. 770–778 (2016)
8. Krizhevsky, A., Hinton, G., et al.: Learning multiple layers of features from tiny images (2009)
9. Miller, G.A.: Wordnet: a lexical database for english. *Communications of the ACM* **38**(11), 39–41 (1995)
10. Park, G., Baek, Y., Lee, H.K.: A ranking algorithm using dynamic clustering for content-based image retrieval. In: CIVR 2002 London, UK, July 18–19, 2002 Proceedings. pp. 328–337. Springer (2002)
11. Silla, C.N., Freitas, A.A.: A survey of hierarchical classification across different application domains. *Data Mining and Knowledge Discovery* **22**, 31–72 (2011)
12. Tan, M., Le, Q.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: International conference on machine learning. pp. 6105–6114. PMLR (2019)
13. Venkataramanan, A., Benbihi, A., Laviale, M., Pradalier, C.: Self-supervised gaussian regularization of deep classifiers for mahalanobis-distance-based uncertainty estimation. *arXiv preprint arXiv:2305.13849* (2023)
14. Venkataramanan, A., Faure-Giovagnoli, P., Regan, C., Heudre, D., Figus, C., Usseglio-Polatera, P., Pradalier, C., Laviale, M.: Usefulness of synthetic datasets for diatom automatic detection using a deep-learning approach. *Engineering Applications of Artificial Intelligence* **117**, 105594 (2023)
15. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The caltech-ucsd birds-200-2011 dataset (2011)
16. Wan, J., Wang, D., Hoi, S.C.H., Wu, P., Zhu, J., Zhang, Y., Li, J.: Deep learning for content-based image retrieval: A comprehensive study. In: Proceedings of the 22nd ACM international conference on Multimedia. pp. 157–166 (2014)
17. Wu, P., Manjunath, B., Shin, H.: Dimensionality reduction for image retrieval. In: Proceedings of ICIP. vol. 3, pp. 726–729. IEEE (2000)
18. Yang, J., She, D., Lai, Y.K., Yang, M.H.: Retrieving and classifying affective images via deep metric learning. In: Proceedings of the AAAI Conference. vol. 32 (2018)
19. Zhang, C., Cheng, J., Tian, Q.: Image-level classification by hierarchical structure learning with visual and semantic similarities. *Information Sciences* **422**, 271–281 (2018)