

Visualizing Bag-of-Features Image Categorization Using Anchored Maps

Gao Yi

The University of Tokyo
Chiba 277-8651, Japan
gaoyi@visual.k.u-tokyo.ac.jp

Hsiang-Yun Wu

The University of Tokyo
Tokyo 113-8565, Japan
hsiang.yun.wu@computer.org

Kazuo Misue

University of Tsukuba
Tsukuba 305-8577, Japan
misue@cs.tsukuba.ac.jp

Kazuyo Mizuno

The University of Tokyo
Tokyo 113-8565, Japan
mizunok@visual.k.u-tokyo.ac.jp

Shigeo Takahashi

The University of Tokyo
Tokyo 113-8565, Japan
takahashis@computer.org

ABSTRACT

The bag-of-features models is one of the most popular and promising approaches for extracting the underlying semantics from image databases. However, the associated image categorization based on machine learning techniques may not convince us of its validity since we cannot visually verify how the images have been classified in the high-dimensional image feature space. This paper aims at visually rearrange the images in the projected feature space by taking advantage of a set of representative features called visual words obtained using the bag-of-features model. Our main idea is to associate each image with a specific number of visual words to compose a bipartite graph, and then lay out the overall set of images using anchored map representation in which the ordering of anchor nodes is optimized through a genetic algorithm. For handling relatively large image datasets, we adaptively merge a pair of most similar images one by one to conduct the hierarchical clustering through the similarity measure based on the weighted Jaccard coefficient. Voronoi partitioning has been also incorporated into our approach so that we can visually identify the image categorization based on support vector machine. Experimental results are finally presented to demonstrate that our visualization framework can effectively elucidate the underlying relationships between images and visual words through the anchored map representation.

ACM Classification Keywords

I.3.8 [Computer Graphics]: Applications

General Terms

Algorithms

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

VINCI 2014, August 5–8, 2014, Sydney, Australia.
Copyright 2014 ACM 978-1-4503-2765-7...\$10.00

Author Keywords

bag-of-features, bipartite graphs, anchored maps, hierarchical clustering, weighted Jaccard similarity coefficients

1. INTRODUCTION

Sophisticating methods for image categorization becomes more crucial in content-based retrieval of image databases due to the rapid increase in their data sizes. While the associated techniques have been significantly improved until recently, it is still labor intensive to sufficiently infer the underlying semantics from images. This problem primarily arises from the fact that we cannot precisely identify specific objects embedded in the images regardless of possible variations in their view, lighting, and occlusion conditions. The *bag-of-features* (BoF) model [21, 3] successfully alleviates this problem for effective image retrieval. A main idea behind the BoF model is to seek an analogy of methods for inferring text categorization based on the bag-of-words model, where each document is represented as a sparse vector of representative words by referring to their occurrence without worrying about their associated orders. In practice, the BoF model allows us to associate an individual image with a small weighted set of *visual words*, each of which stands for a group of local features in the high-dimensional feature space and thus corresponds to some specific image content in the image.

Nonetheless, the correctness of the image categorization based on such BoF models is not always convincing even with the help of classification methods based on machine learning algorithms, since the actual mechanism for the associated image categorization has not been fully visualized due to the high-dimensionality of the image feature space. In this study, we solve this problem by encoding the relationship between images and visual words as a bipartite graph first, and then employing *anchored map* representation [14] to rearrange the image set on the 2D screen space, as shown in Figure 1(a). Here, we have coin and eyeglass images as the input set of images and try to group them into the

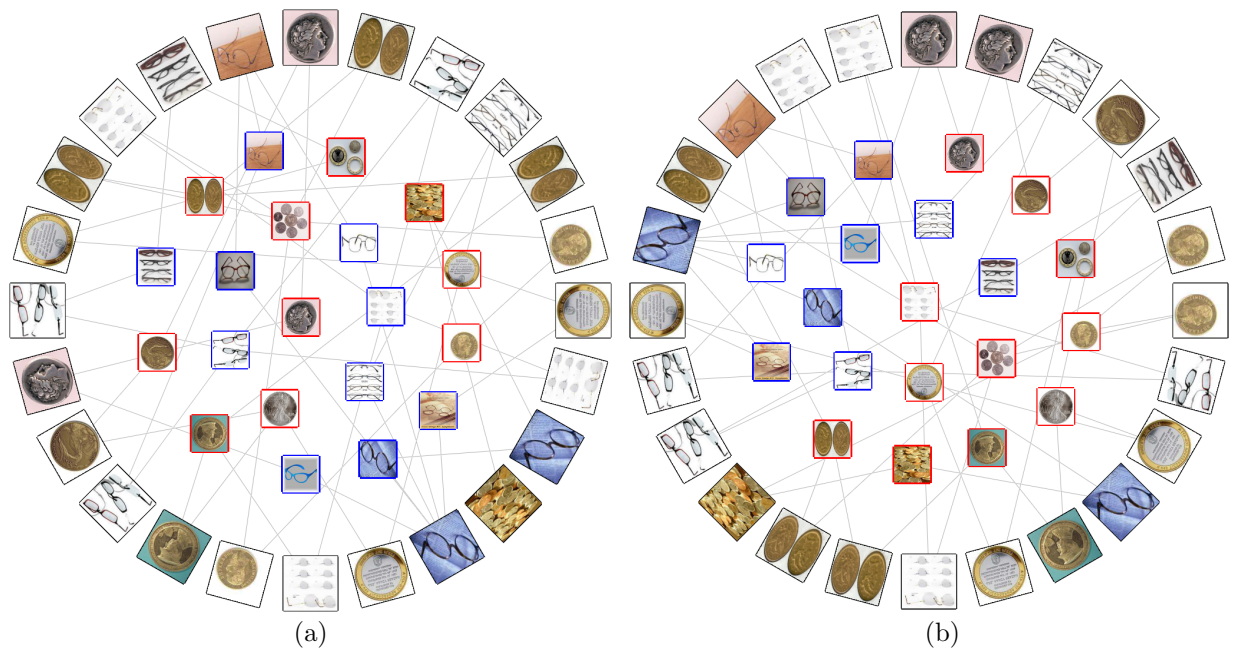


Figure 1. Using anchored maps to visualize the bag-of-features image categorization for coin and eyeglass images. (a) Original layout. (b) Enhanced layout with an optimized circular ordering of visual words annotated with representative images. Images in same category are brought closer to each other. ($\#\{\text{input images}\} = 20$, $\#\{\text{visual words}\} = 24$.)

two categories according to the image contents. Genetic algorithms have also been employed to optimize the circular ordering of visual words around the image feature space, so that we can reduce the number of edge crossings and thus visually elucidate the underlying relationship between images and visual words, as shown in Figure 1(b). Furthermore, we introduced hierarchical representation of the images by adaptively merging images according to their similarity values for effectively handling a large set of images. This hierarchical representation also facilitates users to conduct the image categorization according to their preference, by interactively selecting a training set of images for machine learning techniques.

The remainder of this paper is organized as follows: Section 2 first provides a brief survey on conventional techniques for image categorization together as well as visualization of high-dimensional space. Section 3 then explains how we can extract image features and construct the dictionary of visual words by extracting low-level image features. This is followed by our main contribution of this paper, where we present a novel approach to transforming the high-dimensional image feature space to an anchored map representation by referring to the bipartite relationships between images and visual words in Section 4. After having presenting several experimental results to demonstrate the feasibility of our prototype system in Section 5, we conclude this paper and refer to future work in Section 6.

2. RELATED WORK

Content-based image retrieval has been a hot topic in the research on image processing, computer graphics, and multimedia. For effective search for specific contents, it is important to classify images into several categories by inferring semantics of visual features embedded in them. The *bag-of-features* (BoF) model is a well-known approach for such image representation and helps us categorize images by computing the number of occurrence of particular visual features contained in each image [21, 3]. This idea originates from the concept of *bag-of-words* that naturally allows us to classify documents by counting the number of particular words defined in the dictionary [8]. Indeed, this concept has been extended to the image databases where a set of local features called *visual words* is employed as the dictionary for the analysis of image contents.

In the early stage of approaches of this type, several studies focused on detecting global image features for encoding the image as a whole. Nonetheless, these features appeared to be inappropriate for the purpose of categorizing images because they are too sensitive to image transformations including translation, scaling, and rotation together with lighting conditions and occlusions. Lowe presented a feature detection technique called *scale-invariant feature transform* (SIFT) [12], which allows us to extract local image features in a way that they are robust enough to the prescribed conditions. In practice, the visual words were obtained by collecting the SIFT features from a set of training images and employing the conventional *k*-means clustering to identify the corresponding cluster centers as the

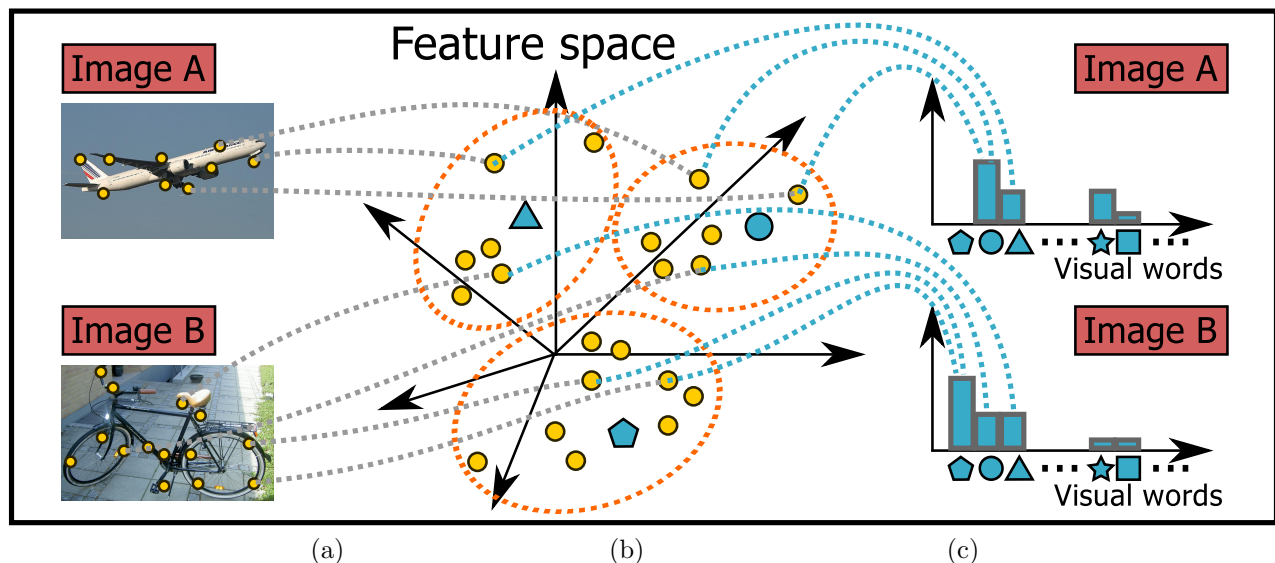


Figure 2. The BoF model. (a) SIFT feature vectors extracted from images are plotted in the 128 dimensional feature space. (b) The k -means clustering algorithm is employed to identify a visual word as the center of each cluster. (c) Each image is encoded as normalized histogram coordinates in terms of the visual words.

visual words.

As for practical approaches for image categorization, we first encode each image in the database as a weighted sum of the relevant visual words. Indeed, the BoF model facilitates us to assign a sparse vector representation of visual words to each image by quantizing it in terms of its associated visual words [21, 3]. *Support vector machine* (SVM) has been often employed as a standard classifier since it produces high accuracy in image categorization [3, 4]. As an extension, Bosch et al. [1] revisited the recognition scheme and apply it to the video by employing probabilistic latent semantic analysis (pLSA) followed by k -nearest neighbor (k -NN) classification. Over the years, a wide range of methods have been developed to improve the quality of the image categorization. A state-of-the-art technique is spatial pyramid matching (SPM) proposed by Lazebnik et al. [11], where they incorporated spatial gradient information of images at multiple scales into the BoF model. More studies also focused on improving the discriminative power of the visual words dictionary. For example, Winn et al. [24] introduced a statistical measure for the optimization framework to make the dictionary of the visual words more compact, while Perronnin [18] combined local and global feature detection frameworks to exhibit higher performance. However, the space of image features extracted by these approaches is always high-dimensional and too abstract to understand the meaningful structures hidden behind that space.

Visualizing high-dimensional feature space often successfully elucidates the image classification obtained through machine learning techniques. A dimensionality reduction technique called *multidimensional scaling* (MDS) [23, 10] is one of the common techniques to

project the high-dimensional space onto a 2D screen space for better readability. The MDS has been also adapted as a visualization tool for exploring feature subspaces from high-dimensional data [22, 25]. Sedlmair et al. [20] conducted an empirical data study to seek the reasonable guidance on the choice of dimensionality reduction techniques and scatterplot representations. Recently, Paulovich et al. [17], Joia et al. [9], and Mamani et al. [13] developed dimensionality reduction frameworks that allows us to interactively edit the underlying structures of the high-dimensional space through screen-space manipulations. Furthermore, Mizuno et al. [16] presented a framework for interactively exploring feature space that is specific to the BoF models, by referring to the relationships between images and visual words. In our approach, we also focus on the such relationships specific to the BoF models and encode them as anchored map representations [14, 19] for visualization purposes. Technical details of the present approach will be detailed in what follows.

3. BOF MODELS FOR IMAGE CATEGORIZATION

This section first provides a brief overview of the BoF model for encoding images as feature vectors, and describes how images are categorized using machine learning techniques.

3.1 Image Representation Based on BoF Models

In general, the BoF model consists of the three steps: *feature extraction*, *visual words dictionary formation* and *image-histogram representation*. The first step of the BoF construction is the *feature extraction*, where we extract SIFT features from the respective images. Here, the SIFT features are described as 128-dimensional feature vectors and plotted within the corresponding feature space as shown in Figure 2(a). For conducting the

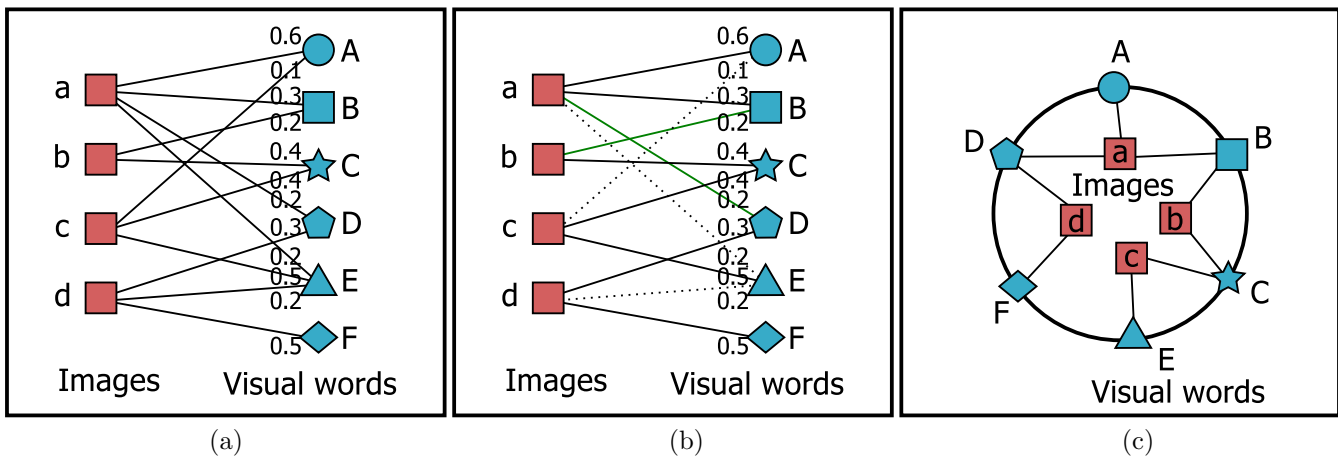


Figure 3. Bipartite relationships between images and visual words in the BoF model. (a) An original bipartite graph. (b) A sparse bipartite graph after edge pruning. (c) The corresponding anchored map representation.

second step for the *visual words dictionary formation*, all the SIFT features are grouped into a specific number of clusters. The simplest technique for this purpose is the conventional k -means clustering algorithm, where the number of clusters k is predefined. Now we are ready to identify the center of each cluster as a representative feature called a *visual word*, and compose the list of k visual words as the dictionary as exhibited in Figure 2(b). Our last step is *image histogram representation*, where we encode each image as a histogram coordinates in terms of the visual words. This is accomplished by quantizing each SIFT feature vector contained in the image to the visual word of the corresponding cluster in the 128-dimensional feature space first, and then counting the occurrence of each visual word to construct the histogram. This process usually allows us to represent each image as a sparse vector of visual words, by normalizing the bins of the histogram to compose the normalized histogram coordinates as shown in Figure 2(c).

3.2 Image Categorization Using Support Vector Machine

In the BoF model, the support vector machine (SVM) is employed as the simplest learning models for classifying images by partitioning the high-dimensional space spanned by the extracted visual words [3]. In practice, the classifier finds the maximum marginal hypersurface that separates positive and negative samples in the training dataset, and further classifies each of the unknown samples by referring to the separating hypersurface. In this paper, we introduce the SVM-based image categorization process proposed by Csurka et al. [3] and visualize how the bounding hypersurface encloses the images of specific type according to the input training samples provided by users. In our approach, we employed radial basis functions (RBFs) kernels for representing such separating hyperplanes to better classify the complicated configuration of images in the high-dimensional space, and visualize the associated image classification in the screen space for more convincing

representation.

4. HIERARCHICAL BIPARTITE GRAPH VISUALIZATION

In this section, we describe how to visualize image categorization via an anchored map representation by referring to the bipartite relationships between images and visual words. We also introduce the weighted Jaccard similarity index for adaptively clustering images so that we can hierarchically represent large scale image sets within the framework of the anchored maps.

4.1 Bipartite Network Composition

The most common way of visualizing the high-dimensional image feature space is to employ dimensionality reduction techniques. Nonetheless, it is often the case that we still cannot fully discriminate each image category from others if the images are simply projected onto the low-dimensional space as described previously. Our novel idea for alleviating this problem is to extract bipartite relationships between images and visual words from the BoF model first, and then transform them into a network structure so that we can take advantage of existing graph drawing techniques for better visualization.

For this purpose, we first establish edge connections between each image and its relevant visual words if they correspond to non-zero histogram coordinates of that image. Note that here we represent images and visual words as nodes of the bipartite graph, while we associate each normalized histogram coordinate value with the corresponding edge as its weight value as shown in Figure 3(a). Furthermore, we would like to make the bipartite graph as sparse as possible for better readability of the resulting graph visualization. Thus, we sort the edges in an ascending order according to the weight values, and prune the edge having the minimum weight one by one until we cannot remove edges any more without decomposing the graph into multiple connected components, as shown in Figure 3(b). In this

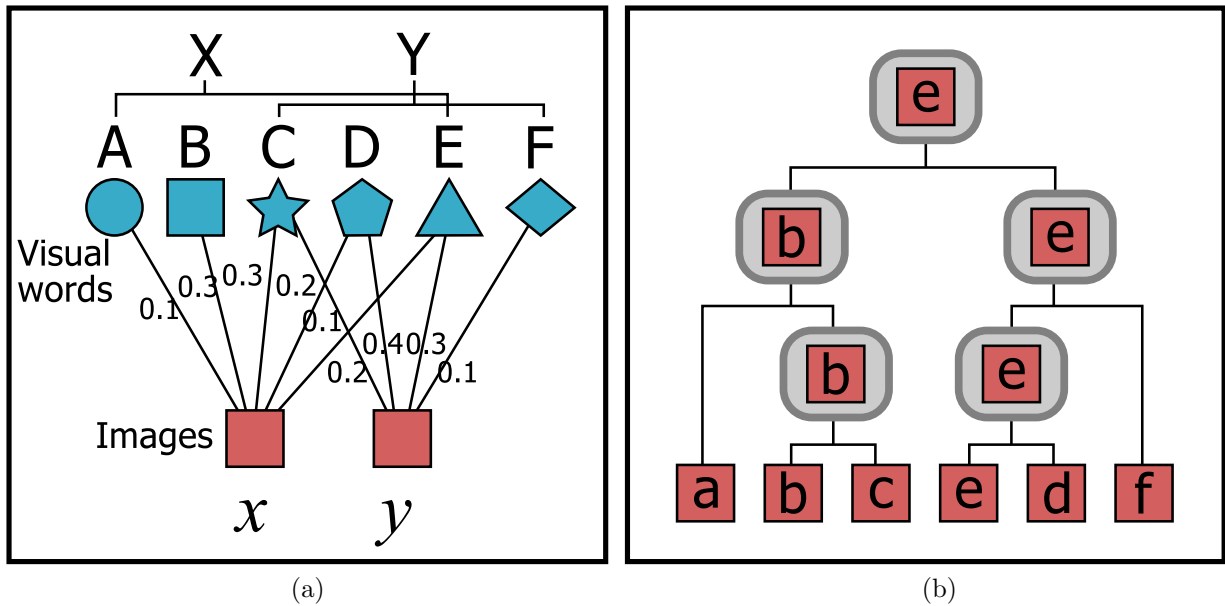


Figure 4. Hierarchical structure of bipartite graph visualization. (a) An example bipartite graph between images and visual words. (b) Dendrogram-based representation of clustered images.

way, we construct a sparse representation of the bipartite graph over the image and visual word nodes.

4.2 Anchored Map Representation

As for the visualization of the bipartite relationships, we employ anchored map representations formulated by Misue [14, 15]. In the anchored map representation, nodes in one of the two disjoint sets of the bipartite graph are equally spaced along the boundary of a disk region, while nodes of the other set are free to move within the disk, as shown in Figure 3(c). In the figure, we release the image nodes within the central disk region of the anchored map and fixed the visual word nodes along its circular boundary. The conventional spring embedder algorithm is also applied to the free nodes to avoid unnecessary overlap among images in the central region, where we also incorporate edge weights into our formulation so that each image will be brought closer to its relevant visual words according to their corresponding normalized histogram coordinates.

In our sparse representation of the bipartite graph, each image usually depends on a small number of visual words. This means that our scheme is more likely to bring images of the same category close to each other in the anchored map representation since they usually share almost the same set of visual words in their histogram representation. Furthermore, this visual readability of the image categorization can be enhanced if we carefully reorder the visual word nodes along the circular boundary of the disk to make each image node have its neighbor visual word nodes within its vicinity. This is accomplished by devising genetic-based algorithms for optimizing the circular ordering of visual words, where we define a chromosome as a value-encoding se-

quence of visual word IDs. For fully discriminating between image categories, we optimize the chromosome sequence by defining the cost function so that, for each image node, every pair of its adjacent visual word nodes become closer to each other. This amounts to calculating the circular distance between adjacent visual word nodes for each image node, and summing up the squared distances except for the largest one [14]. This genetic-based optimization provides us with better anchored maps in the sense that images in the same category will be closer to each other in the central disk region as shown in Figure 1(b). Note that, in our implementation, we replace each visual word with an image having the highest occurrence of that word, so that we can enhance the visual readability of the resulting image categorization as shown in Figure 1. Furthermore, the genetic-based optimization also potentially reduces the number of edge crossings in order to alleviate the distracting visual clutter.

4.3 Hierarchical Clustering of Images

As the number of input images increases, the central disk region of the anchored map will be more crowded with the images. For improving the scalability of the anchored map representation, we also introduced hierarchical representation of the anchored map by adaptively clustering images according to their image similarities. More specifically, we compose a dendrogram tree structure of images by merging a pair of the most similar images one by one iteratively [19]. For evaluating the similarity among images, we employ the conventional Jaccard similarity index, which is the most popular similarity measure between a pair of sets [2]. Let us consider two sets X and Y for example. The conventional Jaccard index is defined as $J(X, Y) = |X \cap Y| / |X \cup Y|$,

where $|Z|$ represents the number of elements contained in the set Z . However, in our case, the weighted Jaccard similarity index [7, 2] is more appropriate in the sense that we can incorporate the importance of each relevant visual word when evaluating the image similarities, rather than simply counting the number of relevant visual words in the union and intersection of the two sets.

As described earlier, our bipartite graph is composed by connecting an image with its relevant visual words, and the weight of each edge is equivalent to the normalized histogram coordinate value of the corresponding visual word with respect to that image. Thus we can easily compute the weighted Jaccard similarity index between a pair of images by referring to their corresponding sets of visual words X and Y , together with their corresponding edge weights, as follows:

$$\text{WJ}(X, Y) = \frac{\sum_{i=1}^n \min(X_i, Y_i)}{\sum_{i=1}^n \max(X_i, Y_i)} \quad (1)$$

where n denotes the total number of visual words contained in the union of X and Y . Note that the numerator is obtained by summing up the minimum values between two weights of the edges emanating from visual words in X and Y , while the denominator is the sum of the maximum values. Figure 4(a) shows an example, where the X_i and Y_i are defined as normalized histogram coordinates for the image nodes x and y , and thus we can set $(X_i) = (0.1, 0.3, 0.3, 0.2, 0.1, 0.0)$ and $(Y_i) = (0.0, 0.0, 0.2, 0.4, 0.3, 0.1)$. This means that we can compute the weighted Jaccard similarity index between the image nodes x and y as

$$\text{WJ}(X, Y) = \frac{0.0 + 0.0 + 0.2 + 0.2 + 0.1 + 0.0}{0.1 + 0.3 + 0.3 + 0.4 + 0.3 + 0.1} = \frac{1}{3}.$$

Using the weighted Jaccard measure, we can iteratively merge a pair of the most similar images into a group one by one, and encode the clustering process as a dendrogram tree representation as shown in Figure 4(b). As illustrated in this figure, we incorporate an image node having a smaller number of child nodes into the other image node representing more child nodes in our implementation.

4.4 Visualizing SVM-based Image Classification

We also equip our prototype system with an interface for classifying images using support vector marching (SVM). In practice, users are allowed to interactively specify a subset of images as a training set for SVM-based classifier together with the tags that represent whether the corresponding images are classified into a specific category or not. Nonetheless, conventional BoF models just present the classification results only and do not provide us with any information about how the images are classified in the high-dimensional image feature space. When projecting the high-dimensional image categorization onto the central disk region within the anchored map, we introduced the Voronoi tessellation technique in order to clarify how the region is parti-

tioned according to the image categorization. Here, we employ the position of each image node as a seed point for the Voronoi cell, and assign a specific color to that cell according to its image category obtained through the SVM classification. This is reasonable because the image nodes are uniformly distributed at the final stage with the help of the spring-embedder algorithm and thus the associated Voronoi tessellation usually produces visually plausible partitioning of the central region of the anchored map. The overall classification process makes us convinced with the image categorization provided by the SVM-based classifier by visualizing the associated image categorization within the anchored map representation. Note that, in our implementation, we incorporated a hardware-assisted algorithm for computing Voronoi diagrams [6] and restrict the drawing area to the central disk region of the anchored map using the stencil buffer.

5. RESULTS

Our prototype system has been implemented on a laptop PC with an Intel Core i7 CPU (2GHz, 4MB cache) and 8GB RAM, and the source code has been written in C++ using the OpenGL library for drawing graph layouts, OpenCV library for SIFT feature extraction and SVM learning models, Boost Graph Library for constructing the graph data structures, and GALib library for the implementation of the genetic-based algorithm. The images datasets used in this paper were collected from Caltech256 [5].

Figure 1 exemplifies how the underlying image categorization can be better visualized by taking advantage of the optimal ordering of visual words around the circular boundary of the anchored map representation. Here, Figure 1(a) shows the initial ordering of visual words and layout of images in the dataset where images of coins and eyeglasses are intricately mixed. On the other hand, images of two categories are sufficiently discriminated in Figure 1(b) when we rearrange the ordering of the visual words using genetic-based optimization.

The image set exhibited in Figure 5 contains images of three different objects, i.e., cars, tomatoes, and grapes from which we try to discriminate car images specifically from the others. For effectively handling a large number of images, we first compute a small number of image groups through hierarchical clustering of images, and distinguish car images from the others as our target using the SVM-based image categorization. Note that here the images outlined in red are labeled as example images within the specific category (i.e. car images), while those in blue are images that are out of our target. We then gradually decompose each image cluster into smaller clusters, and adjust the image categorization by interactively labeling a small number of images as the training set according to their categories. This successfully allows us to enclose car images within yellow background region from the coarsest level to the finest (i.e. original) level as shown in Figure 5.




Figure 5. Discriminating car images using the support vector machine at multiple hierarchical levels. Images of the training set are labeled as red (car images) and blue (others). The inferred region of the car images is rendered in yellow through the Voronoi tessellation. (a) 10%, (b) 30%, (c) 40%, and (d) 100% of images. ($\#\{\text{input images}\} = 240$, $\#\{\text{visual words}\} = 100$.)

Figure 6 demonstrates how we can categorize images of a specific category even when we train our image classifier indirectly with similar looking images. In this case, we represent each image in terms of visual words obtained from training images containing tomatoes, coins, and cars and try to collect images of round shapes. However, we also take as input images of additional categories such as CDs and glasses in this example,

while we still can categorize images of round objects into our target category using the SVM-based classifier, and clearly visualize the associated image categorization both at coarse and fine levels through the anchored map representation as shown in the figure.

A user study has been also conducted for evaluating the feasibility of the proposed visualization approach in the



(a)



(b)

Figure 6. Categorizing images of round objects from images of five different categories (tomatoes, coins, cars, CDs, and glasses). (a) Coarse level. (b) Fine level. ($\#\{\text{input images}\} = 420$. $\#\{\text{visual words}\} = 100$.)

Table 1. Results of a user study. Each column represents the number of participants who returned the corresponding response.

	Figure 1		Figure 5	
	Yes	No	Yes	No
Q1. Does the anchored map provide a better layout than MDS?	10	0	10	0
Q2. Do the optimal circular orders improve the layout?	10	0	10	0
Q3. Is the hierarchical image clustering useful?	NA	NA	9	1
Q4. Are you satisfied with our scheme for interactively categorizing images?	10	0	10	0

cases of Figures 1 and 5. We recruited 10 participants (4 females and 6 males) ranging in age from 22 to 30, where all of them worked on research on visualization and image synthesis. In our user study, we asked them the following yes/no questions:

1. Does the anchored map provide a better layout than MDS?

We prepared the low-dimensional layout obtained using MDS and asked them to compare it with the anchored map representations generated by our approach.

2. Do the optimal circular orders improve the layout?

We generated an initial anchored map layout and its improved version by optimizing the circular ordering of anchored nodes for each case, and asked them to compare between these two layouts.

3. Is the hierarchical image clustering useful?

Especially in the second case of Figure 5 where we take more images as input, we demonstrated that how we can change the level of details in the anchored map representation and asked them to evaluate its effectiveness.

4. Are you satisfied with our scheme for interactively categorizing images?

We also asked users to interactively annotate images according to the prescribed image categorization, and then evaluate whether they were satisfied with our proposed scheme.

Table 1 summarizes the results we obtained through the user study. Overall, the participants gave us positive responses to all the questions, which basically supports our visualization scheme for image categorization based on the BoF model.

6. CONCLUSION

In this paper, we have presented an approach to visualizing image categorization within the high-dimensional feature space by taking advantage of the characteristics

of the BoF model. The idea behind our approach is to extract the bipartite relationships between the input images and visual words first and then visualize them as a network using the anchored map representation. This new type of dimensionality reduction framework successfully convinces us of the plausibility of resulting image categorization based on the BoF model. The readability of the anchored map representation has been further enhanced by seeking the optimal circular ordering of visual words and dendrogram-based hierarchical representation of images. Voronoi-based partitioning has been also incorporated into the central disk region of the anchored map to visualize the border of some specific image category.

Fully classifying images of multiple categories according to users' preference remains to be tackled. The readability of the anchored map representations also depends on the quality of the sparse vector representations of the images in terms of the extracted visual words. Applying hierarchical clustering to the set of visual words together with an input set of images will further alleviate the scalability issue of the present approach. It is interesting to seek more visually plausible layout of image nodes by incorporating the underlying Voronoi partitioning of the central disk region. Enhancing the interactivity of the present image retrieval system is also left as a future research theme.

Acknowledgments

This work has been partially supported by MEXT KAKENHI under Grants-in-Aid for Scientific Research on Innovative Areas No. 25120014 and JSPS KAKENHI under Grants-in-Aid for Scientific Research (B) No. 24300033.

REFERENCES

1. BOSCH, A., ZISSERMAN, A., AND MUÑOZ, X. Scene classification via pLSA. In *Proc. 9th European Conference on Computer Vision (ECCV 2006)* (2006), vol. 3954 of *Springer Lecture Notes in Computer Science*, pp. 517–530.
2. CHIERICHETTI, F., KUMAR, R., PANDEY, S., AND VASSILVITSKII, S. Finding the jaccard median. In *Proc. 21st Annual ACM-SIAM Symposium on Discrete Algorithms* (2010), pp. 293–311.
3. CSURKA, G., DANCE, C. R., FAN, L., WILLAMOWSKI, J., AND BRAY, C. Visual categorization with bags of keypoints. In *ECCV'04 Workshop on Statistical Learning in Computer Vision* (2004), pp. 1–22.
4. EICHORN, J., AND CHAPPELLE, O. Object categorization with SVM: Kernels for local features. In *Advances in Neural Information Processing Systems (NIPS)* (2004).
5. GRIFFIN, G., HOLUB, A. D., AND PERONA, P. The caltech 256. Tech. rep., California Institute of

Technology, 2006.

6. HOFF III, K. E., CULVER, T., KEYSER, J., LIN, M., AND MANOCHA, D. Fast computation of generalized voronoi diagrams using graphics hardware. In *Proc. SIGGRAPH '99* (1999), pp. 277–286.
7. IOFFE, S. Improved consistent sampling, weighted minhash and l1 sketching. In *Proc. 10th IEEE International Conference on Data Mining 2010* (2010), pp. 246–255.
8. JOACHIMS, T. Text categorization with support vector machines: Learning with many relevant features. In *Proc. 10th European Conference on Machine Learning*, vol. 1398 of *Springer Lecture Notes in Computer Science*. Springer, 1998, pp. 137–142.
9. JOIA, P., PAULOVICH, F. V., COIMBRA, D., CUMINATO, J. A., AND NONATO, L. G. Local affine multidimensional projection. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (2011), 2563–2571.
10. KRUSKAL, J. B. Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. *Psychometrika* 29, 1 (1964), 1–27.
11. LAZEBNIK, S., SCHMID, C., AND PONCE, J. Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (2006), vol. 2, pp. 2169–2178.
12. LOWE, D. G. Object recognition from local scale-invariant features. In *Proc. 7th IEEE International Conference on Computer Vision* (1999), vol. 2, pp. 1150–1157.
13. MAMANI, G. M. H., FATORE, F. M., NONATO, L. G., AND PAULOVICH, F. V. User-driven feature space transformation. *Computer Graphics Forum* 32, 3 (2013), 291–299.
14. MISUE, K. Drawing bipartite graphs as anchored maps. In *Proc. Asia-Pacific Symposium on Information Visualisation 2006 (APVis '06)* (2006), pp. 169–177.
15. MISUE, K. Anchored map: Graph drawing technique to support network mining. *IEICE Transactions 91-D*, 11 (2008), 2599–2606.
16. MIZUNO, K., WU, H.-Y., AND TAKAHASHI, S. Manipulating bilevel feature space for category-aware image exploration. In *Proc. of the 7th IEEE Pacific Visualization Symposium (PacificVis 2014)* (2014), pp. 217–224.
17. PAULOVICH, F. V., ELER, D. M., POCO, J., BOTHA, C. P., MINGHIM, R., AND NONATO, L. G. Piecewise laplacian-based projection for interactive data exploration and organization. *Computer Graphics Forum* 30, 3 (2011), 1091–1100.
18. PERRONNIN, F. Universal and adapted vocabularies for generic visual categorization. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 30, 7 (2008), 1243–1256.
19. SATO, S., MISUE, K., AND TANAKA, J. Readable representations for large-scale bipartite graphs. In *Proc. Knowledge-Based Intelligent Information and Engineering Systems* (2008), vol. 5178 of *Springer Lecture Notes in Computer Science*, pp. 831–838.
20. SEDLMAIR, M., MUNZNER, T., AND TORY, M. Empirical guidance on scatterplot and dimension reduction technique choices. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2634–2643.
21. SIVIC, J., AND ZISSERMAN, A. Video google: a text retrieval approach to object matching in videos. In *Proc. 9th IEEE International Conference on Computer Vision* (2003), pp. 1470–1477.
22. TATU, A., MAASS, F., FARBER, I., BERTINI, E., SCHRECK, T., SEIDL, T., AND KEIM, D. Subspace search and visualization to make sense of alternative clusterings in high-dimensional data. In *Proc. IEEE Conference on Visual Analytics Science and Technology (VAST)* (2012), pp. 63–72.
23. TORGERSON, W. S. Multidimensional scaling: I. theory and method. *Psychometrika* 17, 4 (1952), 401–419.
24. WINN, J., CRIMINISI, A., AND MINKA, T. Object categorization by learned universal visual dictionary. In *Proc. 10th IEEE International Conference on Computer Vision* (2005), vol. 2, pp. 1800–1807.
25. YUAN, X., REN, D., WANG, Z., AND GUO, C. Dimension projection matrix/tree: Interactive subspace visual exploration and analysis of high dimensional data. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2625–2633.