

A New Database and Protocol for Image Reuse Detection

Furkan Isikdogan^{1,2}, Ilhan Adiyaman², Alkim Almila Akdag Salah³, and Albert Ali Salah²

¹Department of Electrical and Computer Engineering,
The University of Texas at Austin, USA

²Department of Computer Engineering, Bogazici University, Istanbul, Turkey

³College of Communication, Istanbul Sehir University, Istanbul, Turkey
`isikdogan@utexas.edu {ilhan.adiyaman,salah}@boun.edu.tr`
`alnilasalah@sehir.edu.tr`

Abstract. The use of visual elements of an existing image while creating new ones is a commonly observed phenomenon in digital artworks. The practice, which is referred to as image reuse, is not an easy one to detect even with the human eye, less so using computational methods. In this paper, we study the automatic image reuse detection in digital artworks as an image retrieval problem. First, we introduce a new digital art database (BODAIR) that consists of a set of digital artworks that re-use stock images. Then, we evaluate a set of existing image descriptors for image reuse detection, providing a baseline for the detection of image reuse in digital artworks. Finally, we propose an image retrieval method tailored for reuse detection, by combining saliency maps with the image descriptors.

Keywords: Image database, Digital Art, Image Retrieval, Feature Extraction, DeviantArt, Image Reuse, BODAIR

1 Introduction

One of the main focus of art historical research is the detection of stylistic similarities between works of art. As early as 1915, Heinrich Wölfflin, who is deemed by many as the “father” of art history as a discipline, introduced the notion of comparing artworks to define the style of a period [1]. Art historians, connoisseurs, and art critics are trained to detect whether certain features of an artwork are apparent in another one, and whether two artworks belong to the same artist or not. The experts not only use their visual understanding for such detection, but also rely heavily on historical records and archival information, which are

not always sufficiently clear or available. Hence, for decades, art historical research has applied scientific methods such as infrared and x-ray photographic techniques (among others) to help in different instances where the trained eye faltered. Using computational approaches in detecting stylistic traditions of artworks is a relatively new addition to the field [2]. In this paper, we introduce a new digital image database that consists of original artworks that are re-used to create new artworks. We use this database to examine approaches for image reuse detection. In the long run, trying to detect which image is reused with computational methods will help in detecting stylistic similarities between artworks in general [3].

In Western tradition, artists learned their trade by joining ateliers of masters as apprentices. With the introduction of printing press and the wider availability of paper, and especially due to the replacement of etchings on woodblocks with engravings on metal proliferated the art education in ateliers. Metal engravings started to be widely used to teach apprentices drawing, by copying known forms and designs. Novices used these models as the basis of their new artworks, and in that sense, these designs might be the first ones that were massively re-used in visual art. Today, a similar tendency is to use the so called “stock-images” for the same purpose: to help facilitate the design of a new artwork. These images are made freely available online, and can be found in repositories and dedicated websites. With the help of multimedia technologies and digital drawing tools, as well as the availability of free stock images, it has become a common approach in digital image creation to reuse existing images. The digital re-use scenarios are on the one hand quite different than their forefathers from centuries ago: they heavily rely on photo manipulation tools to generate a desired effect or design. On the other hand, certain photo manipulation tools offer the same (basic) design changes that were commonly used centuries ago. Unlike early archives with erroneous and missing data, today, we may have access to precise information about who has reused which image for which artwork. Social networks and online communities for digital artworks, such as DeviantArt¹, and 500px², help us to follow the interaction between artists to minute detail, and build a reliable database of artworks which have re-used other images.

Image reuse detection in digital art is a high-level semantic task, which can be challenging even for humans. Despite the advances in image retrieval and image copy detection techniques, automatic detection of image reuse remains a challenge due to the lack of annotated data and specific tools designed for the analysis of reuse in digital artworks. Image reuse detection differs from general-purpose image retrieval in its scale and amount of the reused pictorial elements. A small object in an artwork can constitute a major part in another composition. An image can be featured in another image in a variety of forms. Developing a global method that addresses all types of image reuse is challenging, as the types of image reuse and modifications vary greatly among different artists and genres of digital art. Another challenge in reuse detection is that the images

¹ <http://www.deviantart.com>

² <http://www.500px.com>



Fig. 1: Example images from the BODAIR database in the animal, food, nature, place, plant, and premade background categories.

can have similar content without actually reusing parts from each other. For example, a famous architectural structure can be depicted by several artists. An ideal image reuse detection system should be able to detect even a small amount of reuse without retrieving false positive images. To develop a robust framework for image reuse detection, it is essential to develop tools and datasets that are designed for the task.

The prolific expansion in the reuse of pictorial elements introduces problems related to the detection and analysis of image reuse. Automatic detection of image reuse would be useful in numerous tasks, including source image location, similar image retrieval [4, 5], popularity and influence analysis [6], image manipulation and forgery detection [7–9], and copyright violation detection [10, 11, 8]. Information about the sources of elements in an image could be used in image search as a semantic variable in addition to low-level image features. Furthermore, such information would be useful for image influence analysis, for discovering the relationships between different genres of (digital) artworks, measuring the popularity of a specific piece of art, and detecting possible copyright violations.

In this paper, we first introduce a novel database called BODAIR (Bogazici-DeviantArt Image Reuse Database³). The BODAIR database is open for research use under a license agreement. To annotate BODAIR, we introduce a taxonomy in image re-use types and techniques. Next, we evaluate a set of baseline image retrieval methods on this database, discussing their strengths and weaknesses. Finally, we propose a saliency-based image retrieval approach to detect reuse on images.

The rest of the paper is organized as follows: Section 2 introduces the BODAIR database. Section 3 describes the methods that we employ in reuse detection and Section 4 presents the experimental results. Finally, Section 5 summarizes our contributions and conclusions.



2 The Bogazici-DeviantArt Image Reuse Database

DeviantArt is a social network for artists and art enthusiasts with more than 38 million registered users. DeviantArt members post over 160,000 images every

³ Available from <http://www.cmpe.boun.edu.tr/~salah/bodair.html>




(a) Partial reuse



(b) Remake/inspiration



(c) Direct reuse



(d) Use as a background

Fig. 2: Examples to the types of reuse: Source images (left), destination images (right).

day. Images posted under the *stock image* category are usually published under an open license and are free to use by others. Using these images, we built an image reuse database. Being artistic creations, the images in our database pose a real challenge for reuse detection.

The etiquette of DeviantArt requires members to leave each other comments if they reuse an image. This tradition helped us to track down which stock images are used in which new works, by performing link and text analysis on the stock image comments. We used regular expressions in the comments to detect any reference to another artwork and crawled more than 16,000 images in the following six subcategories of stock images: animals, food, nature, places, plants, and premade backgrounds (e.g., Fig. 1). Our image crawler uses a depth limited recursive search to download the reused images (children images) and relates them to their source images (parent images). In addition to the automatically extracted parent-children relationships between images, we manually annotated a total of 1,200 images for four reuse types and nine manipulation types:

Reuse Types

Partial reuse: superimposition of a selected area in an image on another one.

Direct reuse: use of an image as a whole, such as insertion or removal of objects, addition of frames or captions, color and texture filters, and background manipulations.

Remake: remake or inspirational use of an image, such as paintings, sketches, and comics based on another artwork.

Use as a background: use of an artwork as a part of the background in another image.

Manipulation Types

Color manipulations: brightness and contrast change, color replacement, hue and saturation shift, tint and shades, and color balance change.

Translation: moving the visual elements in an image.

Texture manipulations: altering the texture of the image, such as excessive blurring/sharpening, overlaying a texture, or tiling a pattern.

Text overlay: image captions, motivational posters, and flyer designs.

Rotation: rotation of elements in an image.

Aspect ratio change: non-proportional scaling of images.

Alpha blending: partially transparent overlay of visual elements.

Mirroring: horizontal or vertical flipping of images.

Duplicative use: using a visual element more than one time.

Each image in the database has an ID and a reference to the ID of the original work if the image is a reused one. The manually annotated images also include the information about the aforementioned types of reuse and manipulation. The manually annotated images include 200 original images, selected among the most popular posts, and their derivatives in each of the six subcategories. The distribution of the partial reuse, direct reuse, use as a background, and remake/inspiration among the manually annotated images are 27%, 47%, 44%, and 6%, respectively. The direct reuse and background categories have a considerable overlap, since the background images are generally used as a whole without excessive cropping. In this classification, only the direct and partial reuse categories are considered to be mutually exclusive.

Examples of the types of reuse are shown in Fig. 2. Fig. 3 shows the overlaps between different categories of image reuse in the database. The matrix is symmetric, and the diagonals show the total number of images annotated with a given reuse or manipulation category. Only *remake* is not included in this figure. There are 75 exemplars of *remake*, and since they are very different than the original stock images, they are not annotated with any manipulations. Subsequently, the *remake* category has no overlaps with other categories. Manipulation examples are shown in Fig. 4.

3 Methods

In this section, we describe the methods we apply for image reuse detection. We first summarize several image description methods that are used in matching-based tasks in computer vision, such as content-based image retrieval, image copy detection, and object recognition. Then, we discuss how saliency maps could be combined with image descriptors to improve matching accuracy and reduce computation time in image reuse detection.

Representing an image by its most discriminant properties is an important factor to achieve higher accuracies. Different feature descriptors extract different

	Partial reuse	Direct reuse	Background	Color Manipulation	Translation	Texture Manipulation	Text Overlay	Rotation	Aspect Ratio Chg.	Alpha Blending	Mirroring	Duplication
Partial reuse	323	0	37	235	308	74	28	62	29	43	26	34
Direct reuse	0	568	400	409	246	165	142	7	33	4	13	11
Background	37	400	526	395	283	157	116	11	37	15	21	4
Color Manipulation	235	409	395	717	471	229	105	51	51	38	33	32
Translation	308	246	283	471	628	163	88	71	63	48	46	46
Texture Manipulation	74	165	157	229	163	264	38	21	10	14	6	8
Text Overlay	28	142	116	105	88	38	180	5	5	6	8	3
Rotation	62	7	11	51	71	21	5	74	10	7	4	13
Aspect ratio chg.	29	33	37	51	63	10	5	10	70	2	2	3
Alpha Blending	43	4	15	38	48	14	6	7	2	50	3	8
Mirroring	26	13	21	33	46	6	8	4	2	3	47	4
Duplication	34	11	4	32	46	8	3	13	3	8	4	46

Fig. 3: Overlaps between reuse types in the BODAIR database.

features from the images to achieve invariance to certain conditions such as color, illumination or viewpoint changes. Traditional image recognition methods usually involve sampling keypoints, computing descriptors from the keypoints, and matching the descriptors [12]. The image descriptors can also be computed over the entire image without sampling keypoints. However, these global features usually perform poorly in detecting partial correspondences between images where a small portion of an image constitutes a major part of another image. Local descriptors, such as SIFT [13] and its color variants CSIFT [14], Opponent-SIFT [14], on the other hand, are more robust in the detection of partial matches. Although the local descriptors usually perform better than global approaches, computation of local descriptors can be computationally expensive as they usually produce a high-dimensional representation of the image which may create a bottleneck in large-scale image retrieval tasks. As suggested by Bosch et al. [12], Bag-of-visual-words (BoW) methods reduce the high-dimensional representation to a fixed-size feature vector, sacrificing some accuracy [14].



More recent approaches make use of convolutional neural networks (CNNs) to learn powerful models from the data itself [15], [16], [17], [18]. Training these models usually requires a large dataset, such as the ImageNet [19] which consists of over 15 million images in more than 22,000 categories. However, it has been shown that the models trained on a set of natural images can be generalized to other datasets [20], and the features learned by a model can be transferred to another model with another task [21].

In this work, we evaluate five image descriptors that are commonly used in image matching and content-based image retrieval problems for image reuse detection: color histograms, Histogram of Oriented Gradients (HOG) [22], Scale Invariant Feature Transform (SIFT) [13], and the SIFT-variants OpponentSIFT and C-SIFT, which are shown to have a better overall performance than the original SIFT and many other color descriptors [14]. In addition, we also use a CNN model [15] pretrained on the ImageNet [19] as a feature extractor, using the fully connected layer outputs (FC6 and FC7) as feature vectors.



(a) Aspect Ratio Change, Color Manipulation and Translation


(b) Rotation, Color and Texture Manipulations



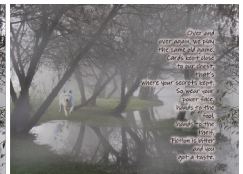

(c) Mirroring and Translation



(d) Alpha Blending and Translation



(e) Translation



(f) Texture Overlay and Color Manipulation

Fig. 4: Examples of manipulations.

Different strategies in image description can lead to a fixed or variable size description of an image. A fixed-size vector representation of images allows the use of vector distance metrics, such as the Euclidean distance, to measure image similarity. As the color histograms and HOG features produce fixed-size image descriptions, candidate matches for a query image can be ranked in order of ascending standardized Euclidean distance. On the other hand, the local descriptors, SIFT and its variants, can extract features from a different number of keypoints on each image, resulting in a variable-size representation. Variable-size representations of images usually require computation-intensive pairwise match-

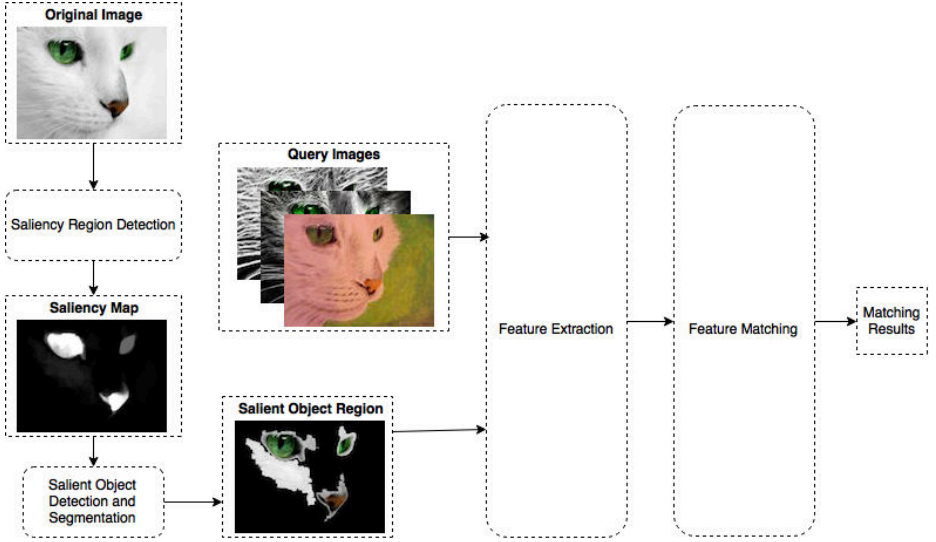


Fig. 5: Proposed framework for image reuse detection.

ing processes. Such a matching process can be improved using an inlier selection algorithm, such as RANSAC, which selects random feature pairs and keep the largest set of inliers to find corresponding image matches [23].

Image saliency can help narrow down the areas of interest in image reuse detection. In our earlier work [24], we showed the effectiveness of using saliency maps in image description for image reuse detection. The purpose of the saliency map is to represent the conspicuity or *saliency* at every spatial location in the visual field by a scalar quantity and to guide the selection of attended locations [25]. Many stock images feature a foreground object that is more likely to be used in other artworks. Therefore, features can be extracted only from the salient regions, which will reduce the processing time and improve the matching accuracy. We use the saliency maps only in the stock images, assuming that each stock image provides such a region of interest to the composition images. We extract features from the query images as a whole, as the use of saliency maps could exclude some references completely.

The overall proposed framework (see Fig. 5) consists of four modules: salient region detection, salient object segmentation, feature extraction, and feature matching. For saliency map estimation, we use a recently proposed Boolean Map based Saliency (BMS) model [26], which is an efficient and simple-to-implement estimator of saliency. Despite its simplicity, BMS achieved state-of-the-art performance on five eye tracking datasets. To segment salient objects, we threshold the saliency maps at their mean intensity to create a binary segmentation mask.




Fig. 6: Cumulative matching accuracies for BoW model with different parameters.

4 Experimental Results

To assess the feature descriptors for detecting image reuse, we designed several experiments. In the experiments, we divided the BODAIR database into a gallery containing a set of stock images and a query set with images that reuse stock images in the gallery. In each experiment, we evaluated the usefulness of the descriptors with a retrieval paradigm. Given a query image I , we ranked the images in the gallery in descending order of probability that the stock image is used in the query image.

4.1 Tuning the Model Parameters

We chose the keypoint sampling strategy and the number of visual words experimentally. For 144 stock and 1,056 query images in the database, we ran the SIFT descriptor with two sampling strategies: sparse salient keypoint detection, and dense sampling. For sparse sampling, we used the default keypoint detector of SIFT, and for dense sampling, we sampled every 8th pixel. Then, we generated a BoW codebook with different vocabulary sizes. We selected the number of clusters for the BoW model as 160, 320, 640, 1,280, and 2,560. The first 20 rank retrieval accuracies for the above-mentioned parameters are shown in Fig. 6.

When the BoW framework is used, dense sampling worked better, as also shown in Nowak et al.’s work on the evaluation of sampling strategies [27]. Thus, we selected uniform dense sampling as our default sampling strategy for the BoW methods. However, in the experiments where we use SIFT with RANSAC without the BoW framework, we selected sparse sampling as our default sampling strategy after some preliminary experiments. Dense sampling increases the outliers in the matching results, which furthermore increases the complexity of finding inliers using RANSAC. Furthermore, sparse sampling results in a smaller set of features, reducing the computational cost.




Fig. 7: Top-1 retrieval accuracies on the BODAIR database for the four types of reuse.




Fig. 8: Top-5 retrieval accuracies on the BODAIR database for the four types of reuse.

The accuracy increased parallel to an increase in the number of clusters, i.e. visual words. Increasing the number of clusters did not improve the performance significantly after a point of saturation (Fig. 6). Therefore, we selected the number of visual words as 1,280 in the rest of the experiments.

4.2 Evaluation of the Methods

We ran experiments on the BODAIR database to evaluate the image description methods for image reuse detection. We compared the methods for all four types of reuse and nine types of manipulations. We calculated and compared Top-1 and Top-5 retrieval accuracies for all of the methods.

Figs. 7 and 8 show the top-1 and top-5 retrieval accuracies, respectively, for the four types of reuse. As the figures show, methods using sparse sampling outperformed dense sampling methods for SIFT with RANSAC on all types of reuse, except *remake*, in the BODAIR database. In the *direct reuse* category, SIFT-based methods produced the best retrieval results. This is in line with Mikolajczyk and Schmid’s earlier results on the use of SIFT for object recognition [28]. The methods that rely on the BoW framework failed to outperform the RANSAC-based methods. Color-based variants of SIFT descriptors gave better results than standard SIFT descriptor. In the *partial reuse* category, the local




Fig. 9: Examples of matching with RANSAC.

descriptors produced the most accurate results. Using saliency to reduce the matched area in the gallery image also marginally improved the performance in the SIFT approach with RANSAC. Fig. 9 shows an example RANSAC matching. In this example, the query image partially reuses the stock image with *color manipulation* and *translation*.

The *remake* category is relatively less restricted, therefore more challenging, than the other types of reuse. Images in this category can be similar to their source images in color, texture, edge distribution, or another aspect. None of the compared methods provided a holistic approach that could recognize all types of artistic remake. Therefore, all of the methods performed poorly on remake images.

We also evaluated these methods and how they perform when it comes to the nine classified image manipulation types: color manipulation, translation, texture manipulation, text overlay, rotation, aspect ratio change, alpha blending, mirroring, and duplication. Overall, the use of saliency maps improved the Top-1 accuracies, although it caused a small decrease in the Top-5 accuracies. HOG features showed poor performance on cropped and translated images, since HOG is not robust to translations when computed globally. All descriptors seem




Fig. 10: Top-1 retrieval accuracies on the BODAIR database for nine different types of manipulations.




Fig. 11: Top-5 retrieval accuracies on the BODAIR database for nine different types of manipulations.

to have a poor performance on images involving rotations, alpha blending, mirroring, and duplication. However, these types of manipulations are frequently observed in tandem with other manipulations in our database. Therefore, the performance of the descriptors is likely to be affected by more than a single type of manipulation. Experimental results for each of the nine types of manipulations are shown in Figs. 10 and 11.

Overall, SIFT and its color-based variants resulted in a higher accuracy without using the BoW framework. Saliency-based approaches provided a better Top-1 retrieval accuracy almost in all types of reuse and manipulations, when they are applied to the original images only. Even though CNN-based approaches failed to outperform SIFT and its color-based variants, the results are promising. With its overall high performance, we recommend using the Opponent SIFT descriptors with RANSAC as a baseline model for the future use of the BODAIR database.

To investigate the poor performance on the *rotation* manipulations, we took 12 rotated versions of each query image, extracted SIFT descriptors, and used them in comparisons with the gallery. We took the best matching rotation for each gallery image. The Top-1 accuracy showed some slight improvement (from 0.09 to 0.11), but the Top-5 accuracy did not change. The reason, we figure, is that *rotation* is often used together with other manipulation and reuse types. The database contains 74 images with *rotation*, of which 71 contain a *translation*, 62 contain *partial reuse*, and 51 contain *color manipulation*.

5 Conclusions

In this work, we focused on how to detect image reuse in digitally created artworks. To that end, we first collected stock images from the social network site DeviantArt, where these are shared freely for reuse purposes, and built the BO-DAIR database. Using automatic link and text analysis in the images' comment sections, as well as manual labeling, we made available a database that has two sets of images: stock images, and images that reuse these stock images. We furthermore made the distinction between "type of reuse", and "type of manipulation", i.e. we highlighted the difference between the contextual approach, and technical approach in reuse. We have detected four type of reuse scenarios, and nine ways of manipulations. We evaluated methods for image reuse detection that are widely used in related tasks, such as image retrieval and object recognition. Lastly, we improved the performance of these methods by using saliency maps. The methods we evaluated provide a baseline for the future research on image reuse detection.

References

1. Wölfflin, H.: *Kunstgeschichtliche Grundbegriffe: das Problem der Stilentwicklung in der neueren Kunst*. München, H. Bruckmann (1915)
2. Stork, D.G.: Computer vision and computer graphics analysis of paintings and drawings: An introduction to the literature. In: *International Conference on Computer Analysis of Images and Patterns*, Springer (2009) 9–24
3. Akdag Salah, A.A., Salah, A.A.: Flow of innovation in deviantart: following artists on an online social network site. *Mind & Society* **12**(1) (2013) 137–149
4. Smeulders, A., Worring, M., Santini, S., Gupta, A., Jain, R.: Content-based image retrieval at the end of the early years. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **22**(12) (2000) 1349–1380
5. Liu, Y., Zhang, D., Lu, G., Ma, W.Y.: A survey of content-based image retrieval with high-level semantics. *Pattern Recognition* **40**(1) (2007) 262 – 282
6. Buter, B., Dijkshoorn, N., Modolo, D., Nguyen, Q., van Noort, S., van der Poel, B., Akdag Salah, A.A., Salah, A.A.: Explorative visualization and analysis of a social network for arts: the case of deviantART. *Journal of Convergence* **2**(2) (2011) 87–94
7. Bayram, S., Avcibas, I., Sankur, B., Memon, N.: Image manipulation detection. *Journal of Electronic Imaging* **15**(4) (2006) 041102–041102
8. Ke, Y., Sukthankar, R., Huston, L., Ke, Y., Sukthankar, R.: Efficient near-duplicate detection and sub-image retrieval. In: *ACM Multimedia*. (2004) 869–876
9. Fridrich, A.J., Soukal, B.D., Lukáš, A.J.: Detection of copy-move forgery in digital images. In: *Proceedings of Digital Forensic Research Workshop*, Citeseer (2003)
10. Kim, C.: Content-based image copy detection. *Signal Processing: Image Communication* **18**(3) (2003) 169–184
11. Zhao, W.L., Ngo, C.W.: Scale-rotation invariant pattern entropy for keypoint-based near-duplicate detection. *Image Processing, IEEE Transactions on* **18**(2) (2009) 412–423
12. Bosch, A., Muoz, X., Mart, R.: Which is the best way to organize/classify images by content? *Image and Vision Computing* **25**(6) (2007) 778 – 791

13. Lowe, D.: Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision* **60**(2) (2004) 91–110
14. Van De Sande, K., Gevers, T., Snoek, C.: Evaluating color descriptors for object and scene recognition. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **32**(9) (2010) 1582–1596
15. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: *Advances in Neural Information Processing Systems*. 2012
16. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proc. IEEE CVPR*. (2016) 770–778
17. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014)
18. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2015) 1–9
19. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)* **115**(3) (2015) 211–252
20. Zeiler, M.D., Fergus, R.: Visualizing and understanding convolutional networks. In: *Computer vision–ECCV 2014*. Springer (2014) 818–833
21. Yosinski, J., Clune, J., Bengio, Y., Lipson, H.: How transferable are features in deep neural networks? In: *Advances in neural information processing systems*. (2014) 3320–3328
22. Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*. Volume 1., IEEE (2005) 886–893
23. Fischler, M.A., Bolles, R.C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM* **24**(6) (1981) 381–395
24. Isikdogan, F., Salah, A.: Affine invariant salient patch descriptors for image retrieval. In: *International Workshop on Image and Audio Analysis for Multimedia Interactive Services, Paris, France* (2013)
25. Itti, L., Koch, C., Niebur, E.: A model of saliency-based visual attention for rapid scene analysis. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **20**(11) (Nov 1998) 1254–1259
26. Zhang, J., Sclaroff, S.: Saliency detection: A boolean map approach. In: *Proceedings of the IEEE International Conference on Computer Vision*. (2013) 153–160
27. Nowak, E., Jurie, F., Triggs, B.: Sampling strategies for bag-of-features image classification. In: *ECCV*. Springer (2006) 490–503
28. Mikolajczyk, K., Schmid, C.: A performance evaluation of local descriptors. *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **27**(10) (2005) 1615–1630