

MODELLING SEMANTIC CONTEXT FOR NOVELTY DETECTION IN WILDLIFE SCENES

Sui-Peng Yong, Jeremiah D. Deng, and Martin K. Purvis

Department of Information Science, University of Otago

PO Box 56, Dunedin, New Zealand

{spyong, ddeng, mpurvis}@infoscience.otago.ac.nz

ABSTRACT

Novelty detection is an important functionality that has found many applications in information retrieval and processing. In this paper we propose a novel framework that deals with novelty detection for multiple-scene image sets. Working with wildlife image data, the framework starts with image segmentation, followed by feature extraction and classification of the image blocks extracted from image segments. The labelled image blocks are then scanned through to generate a co-occurrence matrix of object labels, representing the semantic context within the scene. The semantic co-occurrence matrices then undergo binarization and principal component analysis for dimension reduction, forming the basis for constructing one-class models for each scene category. An algorithm for outlier detection that employs multiple one-class models is proposed. An advantage of our approach is that it can be used for scene classification and novelty detection at the same time. Our experiments show that the proposed approach algorithm gives favourable performance for the task of detecting novel wildlife scenes, and binarization of the label co-occurrence matrices helps to significantly increase the robustness in dealing with the variation of scene statistics.

Keywords— context, co-occurrence matrix, semantics, novel image, multi-class

1. INTRODUCTION

Something out of context in our perception is considered as either abnormal or interesting to us. According to psychological studies, novelty seems intuitively tied to interest [1], but short-term novelty can also be referred to the contrasts with recent experience [2]. The modeling of this novelty detection mechanism, together with image analysis techniques, can be useful in a number of applications. For instance, on a social network such as Flickr or Facebook, images can be automatically analyzed, and abnormal or inappropriate contents can be detected and then either removed or hidden from certain users or groups. Similarly, when browsing a large photo collection, it is desirable to use novelty detection to highlight or select images with novelty or ‘interestingness’.

For the detection of novel scenes, a straightforward idea would be employing traditional statistical methods and machine learning methods [3], and applying them to low-level visual features extracted from images. This is however rather questionable, since different objects may produce quite similar visual

features and the ambiguity, long revealed by content-based image retrieval studies, cannot be resolved unless semantic analysis is conducted and objects within a scene are recognized. Recent psychophysical studies have revealed the importance of high-level semantics in object recognition and scene interpretation. It has been shown that top-down facilitation in recognition is triggered by early information about an object, and also by contextual associations between the object and others with which it typically appears [4, 5]. In light of these findings, it is reasonable to model the novelty of image scenes based on semantics within the scene.

In this paper, we consider the problem of novelty detection for wildlife scenes. A computer system for this purpose would not only require the ability of image analysis and object classification, but also the modelling of image semantics. For wildlife images, normal expectation is that wildlife animals reside to their habitats. So dolphins are usually in water, and zebras have land in the background. When a zebra appears with a dolphin in water, a novelty or anomaly can be detected. On the other hand, an image with similar objects may be classified into the same category, however, the different spatial context of the objects may create novelty. For example the two images in Fig.1 have similar occurrence of object classes ('elephant' and 'sky'), but the co-occurrence contexts are different, so Fig.1(b) is treated as novel because elephants are normally under the sky but not above it.




Fig. 1. Images with similar objects but different context.

We therefore propose a computational framework that models the semantic context, where novelty detection can be conducted based on comparing the co-occurrence of semantic labels. In the remainder of the paper, we first briefly review some relevant work in Section 2. Our computational framework and the novelty detection algorithm are introduced in Section 3, and the experiment results are presented in Section 4, followed by conclusion and discussion on future work.

2. RELATED WORK

Novelty detection has received most interest in textual information retrieval tasks [6, 7]. Novelty detection in video sequence has also been studied with the growing demand in surveillance applications [8, 9]. In the area of wildlife surveillance [10], image analysis techniques have not found much utilization yet. Our approach differs from these in that we do not consider motion and trajectory information acquirable from videos, but rather concentrate on conducting semantic and statistical analysis of static scenes.

Semantic representations of images have been explored for the tasks of content-based image retrieval and scene classification [11, 12, 13]. The general approach is to identify semantics as the set of objects that appear in the image (e.g. sky, grass and mountain), and the scene is described by statistical modelling on the semantic objects [14, 15, 16, 17]. The relevant models are mainly applied in scene classification or understanding, but do not include novelty detection.

3. COMPUTATIONAL FRAMEWORK

To achieve the goal of detecting novel or interesting images from a wildlife image collection, we propose a framework as shown in Fig.2. The input images are first segmented into homogeneous segments, after which colour and texture features are extracted from image blocks within the image segments. These blocks can then be classified using classifiers trained on the visual features. To achieve the goal of novelty detection, we construct co-occurrence matrices of block labels and train a number of one-class classifiers, each modelled for a scene category. If a scene is rejected by all these one-class classifier, then it is considered as novel or interesting. Hence, by going through the same process, both scene classification and novelty detection tasks can be accomplished.

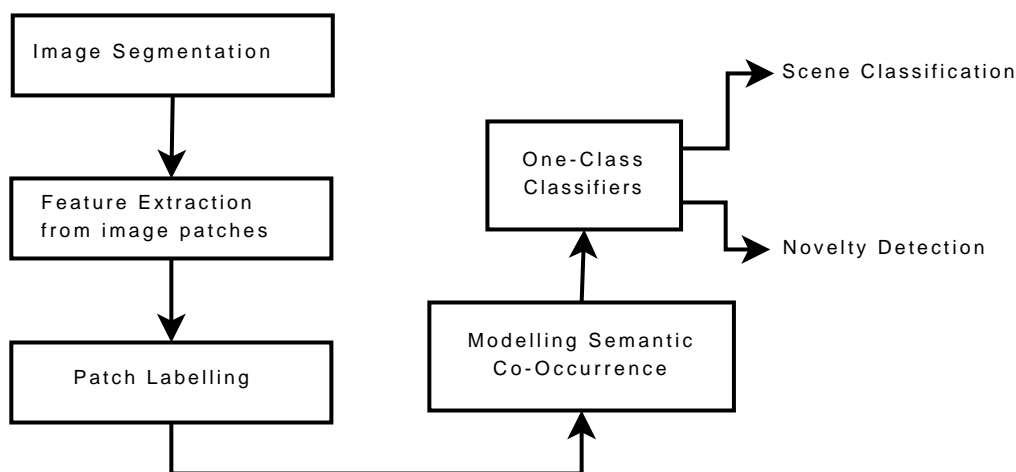


Fig. 2. The computational framework for image novelty detection and scene classification.

3.1. Image segmentation

A scene usually contains multiple objects of different visual characteristics. By segmenting an image into homogeneous regions, it facilitates detection or classification of these objects. We use the JSEG algorithm [18] for this purpose. In JSEG, colours in the image are first quantized to several representing classes that are used to separate the regions in the image. Image pixel colours are then replaced by their corresponding colour class labels to form a class-map of the image. In order to get good segmentation, the high and low values correspond to possible boundaries and interiors of colour-texture regions are adjusted until the object and background can be differentiated. Next, those segments that exceed a threshold will be taken out and stored as individual segment images respectively. Those segments will then be classified manually as the ground truth for training classifiers. Each segment image will be tiled into blocks of size 32×32 . The image blocks that fall out of the segment edges will be ignored. For semantic analysis of new images, we have found it is more effective to train classifiers on segment blocks instead of on segment images directly.

3.2. Feature extraction

Visual features are then extracted from image segment blocks. First we employ the LUV colour histogram to encode the colour information of image blocks. Colour histograms are found to be robust to resolution

and rotation changes. The LUV colour space is adopted because it models human’s perception of colour similarity very well, and it is also machine independent [19]. Each channel is quantized with the same interval, thus L channel has 20-bins, U 70-bins and V 52-bins respectively. The standard deviation of the LUV histogram values is also calculated. The LUV histogram feature therefore has takes 143 dimensions.

Texture features extracted from the image blocks are also included as local features. We compared Edge Histogram Descriptor and Gabor filtering features [20] with the Haralick features [21], the latter performed the best in our test and are therefore adopted in further experiments. The Haralick texture features consist of a few statistical properties based on the gray scale co-occurrence matrix. The image block is first converted to gray scale. The co-occurrence matrix is a two-dimensional histogram composed from the pair-wise statistics of gray scale co-occurrence among adjacent pixels. Four orientations are considered, each giving a co-occurrence matrix for an image block. A total of 13 statistical measures can be calculated for each co-occurrence matrix. The mean and deviation values of each of these 13 measures over the four orientations, form a feature vector of 26 dimensions.

Finally, the colour and texture features are concatenated together, giving a feature vector of 169 dimensions to represent an image block. Through manual labelling of image segments, semantic ground truth is assigned to the training images. The image blocks inherit semantic labels from their corresponding segments. Their feature codes along with the relevant class labels are used to train object classifiers.

3.3. Modelling the semantic context

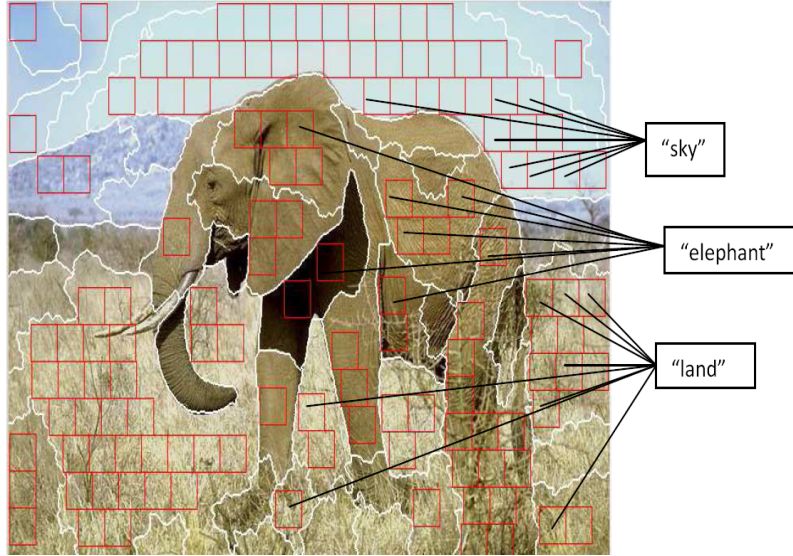
To model the semantic context within a scene, we further generate a *block label co-occurrence matrix* (BLCM). First the labels of all image blocks should be obtained - either by manually annotation for training images, or automatic classification when testing. Then the image blocks are scanned from right to left and top to bottom, the co-occurrence of labels for blocks within a distance d is collected. The co-occurrence statistics is gathered across the entire image and normalized by the total number of image blocks. The variation on the object sizes does affect the matrix values of BLCM, so to reduce this effect the BLCM is binarized. Fig.3 shows an ‘elephant’ image as an example, with blocks labels and the relevant binary BLCM shown below. There are 8 object labels: ‘coast’, ‘land’, ‘sky’, ‘water’, ‘dolphin’, ‘elephant’, ‘penguin’, and ‘zebra’, giving an 8×8 BLCM.

Since the scanning of image blocks is directional, the BLCM is consequently asymmetric. This is however a desirable feature as we intend to keep the spatial location information in the semantic context representation. For example, a scene context of ‘sky’ \rightarrow ‘land’ is normal (scanning top-down), while a ‘land’ \rightarrow ‘sky’ BLCM entry may suggest some novelty. On the other hand, images with similar characteristics, i.e, with similar types of foreground and background settings will produce quite similar BLCM features, which is helpful for scene classification.

In images with salient foreground objects, normally the BLCM is quite sparse. We concatenate the BLCM matrix rows into a 1-D vector, and principle component analysis (PCA) can then be used for dimension reduction. The dimension-reduced BLCM vectors are then ready to be further processed by classifiers.

3.4. Building classifiers

The classification task is two-fold. Given an image, if it is of an ordinary scene, the classifiers need to classify it into the right scene category; otherwise, novelty should be reported. This is however not a typical multi-class classification scenario, since there is no restriction of the appearance of novel images,



class	coast	land	sky	water	dolphin	elephant	penguin	zebra
coast	0	0	0	0	0	0	0	0
land	0	1	0	0	0	1	0	0
sky	0	1	1	0	0	1	0	0
water	0	0	0	0	0	0	0	0
dolphin	0	0	0	0	0	0	0	0
elephant	0	1	1	0	0	1	0	0
penguin	0	0	0	0	0	0	0	0
zebra	0	0	0	0	0	0	0	0

Fig. 3. Image blocks in an ‘elephant’ image and its corresponding co-occurrence matrix.

and it is usually hard to find sufficient training data especially for the ‘novelty’ class. Hence we resort to one-class classification. By building a one-class classifier for each of the scene types, the classifiers need only normally labelled images for training. Testing images are assessed by calculating their feature vectors’ distance to the trained one-class centers. If the distance is greater than a defined threshold for that class, it is rejected by that one-class classifier. A data instance rejected by all one-class classifiers is then reported as an outlier. Hence, such an approach carries out scene classification and novelty detection at the same time.

Our algorithm, called multiple one-class classification with distance thresholding (MOC-DT), is summarized as follows:

1. During training, for images of each image scene type, extract their BLCM code and calculate the mean, denoted as μ_i , of the image group. Here $i = 1, 2, \dots, N$, N is the number of scene types.
2. Calculate the Manhattan distance towards μ_i for all instances of the same scene and obtain the distance mean m_i , and the standard deviation, σ_i ;
3. Set the outlier distance threshold as $T_i = m_i + 0.5\sigma_i$ for the scene type;

4. Given a query image, calculate its BLCM code and the Manhattan distance towards each training model, d_i ;
5. If $d_i < \min(T_i)$, assign the scene label k to the image, $k = \arg_i \min(T_i)$; otherwise (i.e., $d_i > T_i, \forall i$), label the image as 'novel'.

4. EXPERIMENTS AND RESULTS

4.1. Experiment settings

To validate the proposed approach, we conduct a few experiments using wildlife images and have the performance on scene classification and novelty detection measured.

For the training phase, 172 images are taken from ImageNet [22] and from Google Image Search. These wildlife images usually feature one animal in scene and contain relatively simple semantic context, each belonging to one of the following 4 scene types (number of instances for each type bracked as follows): 'dolphin' (39), 'elephant' (43), 'penguin' (45) and 'zebra' (45). Image sizes vary from 200×154 to 1024×768 . Some sample images are shown in Fig.4. Four background classes are included: 'coast', 'land', 'sky' and 'water'. For block labelling we have therefore 8 classes, including both the background and animals in foreground. After segmentation, there are over 13,000 image blocks extracted and used to train the classifier for the eight object classes, using the feature schemes presented in Section 3.

For testing, we have 12 'normal' images plus 27 'novel' images. Novel images basically carry some different semantic context as compared with normal ones. For instance, scenes with zebra and dolphin together, penguin and land, elephant and coast etc., are considered 'novel'. Sample testing images are shown as in Fig.5.

The classification of image blocks achieves an accuracy rate of 91% using the k-nearest neighbour classifier. This we believe gives a sound basis for constructing semantic context from labelled blocks.

4.2. Results

The semantic context of image scenes are calculated, undergoing binarization and PCA. We select the first 16 principal components to feed into classification. To visualize the generated BLCM codes for different scene types, a 2-D projection is produced using the first two principal components of the data, as shown in Fig.6. Note that the projected data display a clustered structure with good separability. The 'elephant' and 'zebra' data points are quite close to each other, probably due to the fact that the relevant images have very similar background (i.e., mainly 'land').

First, we assess the performance of BLCM for scene classification using 10-fold cross validation on the training dataset. Two classifiers are employed: k-Nearest Neighbours (k-NN) and Random Forest (RF), both validated on two feature schemes: global scene descriptor (GSD) and BLCM. The GSD uses the same low-level features as used to construct BLCM: LUV colour histograms and Haralick texture features, but these are extracted globally (rather than locally in BLCM). The results, presented in Table 1, show that the binary BLCM with or without PCA achieves much higher accuracy in scene classification compared with using visual features directly.

For novelty detection, the testing images are processed by the proposed MOC-DT algorithm. The results of their respective distance towards each training model were shown in Fig.7. It can be seen that the error instances as instance number 4, 5, 16 and 35. Some of these are shown in Fig.8. The strong context contributed by abnormal background may have caused these errors.




Fig. 4. Samples of training images from four scene classes: ‘dolphin’, ‘elephant’, ‘penguin’ and ‘zebra’.

Table 1. Accuracy Results for Scene Classification

Feature schemes	Classifiers	
	k-NN	RF
GSD	66.86%	60.47%
Binary BLCM	94.77%	95.93%
Binary BLCM + PCA	96.51%	94.77%

The proposed MOC-DT algorithm is compared with the probability density based one-class classifier (PDOC) [23]. Each class model is treated as target class while the rest which are not classified into any particular class are considered as outliers. In other words, PDOC and MOC-DT are compared using the same treatment on the BLCM data, but their difference is only on the one-class classification algorithm. We also compare the algorithms’ performance using normalized and binarized BLCM with or without PCA. The results are shown in Table 2. We can see that with BLCM works better through binarization, and the best result is obtained using MOC-DT with binary BLCM plus PCA. PDOC performs worse than MOC-DT in all situations.

For the normal images (3 for each class) in the testing set, the MOC-DT algorithm manages classify all but one into the correct scene class, therefore giving an accuracy of 91.7%.



(a) 'Novel' images



(b) Normal images

Fig. 5. Sample images used for testing.

5. CONCLUSION

Novelty detection of image scenes can be better achieved by working with semantic context modelling. In this paper we have proposed a simple but effective computational framework that conducts semantic context modelling and employs multiple one-class classifiers to carry out both scene classification and novelty detection. Our experiments on a set of images of four scene classes have given some promising results, especially the proposed MOC-DT algorithm performs favourably compared with other algorithms. In the future we will take more object/scene categories into consideration and conduct more experiments in larger scales.

On the other hand, the novelty or 'interestingness' of a scene is tightly coupled with visual experience and therefore can be dynamic and even subjective to the viewers. It is our intention to investigate using incremental models and algorithms to model the semantic context and its classification, and also incorporating user feedback to fine-tune these incremental models.

6. REFERENCES

- [1] Paul J. Silvia, *Exploring the Psychology of Interest*, vol. 56, Oxford University Press, 2006.
- [2] D.E. Berlyne, *Stimulus Selection and Conflict*, McGraw-Hill Book Company, 1960.




Fig. 6. The 2-D projection of BLCM data for normal and novel images. Normal images are marked by their classes and novel images as ‘outliers’.

- [3] V. Chandola, A. Banerjee, and V. Kumar, “Anomaly detection: A survey,” *ACM Computing Surveys*, vol. 41, pp. 1–58, 2009.
- [4] Mark J. Fenske, Elissa Aminoff, Nurit Gronau, and Moshe Bar, “Top-down facilitation of visual object recognition: object-based and context-based contributions,” *Progress in Brain Research*, vol. 155, pp. 3–21, 2006.
- [5] J. A. Stirck and G. Underwood, “Low-level visual saliency does not predict change detection in natural scences.,” *Journal of Vision*, vol. 7(10):3, pp. 1–10, 2007.
- [6] Evgeniy Gabrilovich, Susan Dumais, and Eric Horvitz, “Newsjunkie: providing personalized news-feeds via analysis of information novelty,” in *WWW’04: Proceedings of the 13th international conference on World Wide Web*, New York, NY, USA, 2004, pp. 482–490, ACM.

Table 2. Novelty detection results.

Feature schemes	Classifiers	
	MOC-DT	PDOC
Normalized BLCM	79.49%	46.15%
Normalized BLCM + PCA	48.72%	30.77%
Binary BLCM	84.62%	61.54%
Binary BLCM + PCA	89.74%	30.77%




Fig. 7. Test instances with their respective distances towards each training class model. Instances 1-22 belong to ‘novel’ classes, 23-39 are ‘normal’ images.



Fig. 8. Novel Images misclassified as normal.

[7] Xiaoyan Li and W. Bruce Croft, “An information-pattern-based approach to novelty detection,” *Information Processing and Management*, vol. 44, no. 3, pp. 1159 – 1188, 2008.

[8] S. Singh and M. Markou, “An approach to novelty detection applied to the classification of image regions,” *IEEE Transactions on Knowledge and Data Engineering*, vol. 16, no. 4, pp. 396–407, 2004.

[9] Shehzad Khalid, “Motion-based behaviour learning, profiling and classification in the presence of anomalies,” *Pattern Recognition*, vol. 43, no. 1, pp. 173 – 186, 2010.

[10] Christine Connolly, “Wildlife-spotting robots,” *Sensor Reviews*, vol. 27, pp. 282–287, 2007.

[11] J. Wang, J. Li, and G. Wiederhold, “Simplicity: semantics-sensitive integrated matching for picture libraries,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 23(9), pp. 947–963, 2001.

[12] L. Fei-Fei and P. Perona, “A bayesian hierarchical model for learning natural science categories,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2005, pp. 524–531.

[13] F. Jing, M. Li, L. Zhang, H. Zhang, and B. Zhang, “Learning in region-based image retrieval,” in *International Conference on Image and Video Retrieval, Urbana-Champaign, Illinois*, 2003.

[14] David M. Blei, Andrew Y. Ng, and Michael I. Jordan, “Latent dirichlet allocation,” *Journal of Machine Learning Research*, vol. 3, pp. 993–1022, 2003.

[15] A. Bosch, A. Zisserman, and X. Munoz, “Scene classification via pLSA,” in *Proceedings of the European Conference on Computer Vision*, 2006, pp. 517–530.

- [16] R. Fergus, L. Fei-Fei, P. Perona, and A. Zisserman, "Learning object categories from google's image search," in *International Conference on Computer Vision*, 2005, vol. 2, pp. 1816–1823.
- [17] L.-J. Li, R. Socher, and L. Fei-Fei, "Towards total scene understanding: classification, annotation and segmentation in an automatic framework," in *Proc. IEEE Computer Vision and Pattern Recognition (CVPR)*, 2009, pp. 2036–2043.
- [18] Y. Deng, , and B.S. Manjunath, "Unsupervised segmentation of color-texture regions in images and video," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 23, no. 8, pp. 800–810, Aug 2001.
- [19] Yu-Fei Ma and Hong-Jiang Zhang, "Contrast-based image attention analysis by using fuzzy growing," in *Multimedia'03: Proceedings of the 11th ACM International Conference on Multimedia*, New York, NY, USA, 2003, pp. 374–381, ACM.
- [20] B. S. Manjunath, J. R. Ohm, V. V. Vasudevan, and A. Yamada, "Color and texture descriptors," *Circuits and Systems for Video Technology, IEEE Transactions on*, vol. 11, no. 6, pp. 703–715, 2001.
- [21] R. M. Haralick, K. Shanmugam, and I. Dinstein, "Textural features for image classification," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 3, pp. 610–621, 1973.
- [22] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *CVPR'09*, 2009, pp. 248–255.
- [23] K. Hempstalk, E. Frank, and I.H. Witten, "One-class classification by combining density and class probability estimation," in *Proc. ECML/PKDD'08*, Berlin, September 2008, vol. 5211 of *Lecture Notes in Computer Science*, pp. 505–519, Springer.