

Mid-level Deep Pattern Mining*

Yao Li^{1,2} Lingqiao Liu¹

¹The University of Adelaide

Chunhua Shen^{1,3}

²NICTA

Anton van den Hengel^{1,3}

³Australian Centre for Robotic Vision

Abstract

Mid-level visual element discovery aims to find clusters of image patches that are both representative and discriminative. In this work, we study this problem from the perspective of pattern mining while relying on the recently popularized Convolutional Neural Networks (CNNs). Specifically, we find that for an image patch, activation extracted from the first fully-connected layer of a CNN have two appealing properties which enable its seamless integration with pattern mining. Patterns are then discovered from a large number of CNN activations of image patches through the well-known association rule mining. When we retrieve and visualize image patches with the same pattern (See Fig. 1), surprisingly, they are not only visually similar but also semantically consistent. We apply our approach to scene and object classification tasks, and demonstrate that our approach outperforms all previous works on mid-level visual element discovery by a sizeable margin with far fewer elements being used. Our approach also outperforms or matches recent works using CNN for these tasks. Source code of the complete system is available online.

1. Introduction

Mid-level visual elements, which are clusters of image patches rich in semantic meaning, were proposed by Singh *et al.* [1] with the aim of replacing low-level visual words (play the game in Fig. 1 and then check your answers below¹). In this pioneering work, mid-level visual elements must meet two requirements, that is, representativeness and discriminativeness. Representativeness means mid-level visual elements should frequently occur in the target category, while discriminativeness implies that they should be visually discriminative against the natural world. The discovery of mid-level visual elements has boosted performance in a variety of vision tasks, including image clas-



Figure 1. *Name that Object*: Given the mid-level visual elements discovered by our algorithm from the PASCAL VOC 2007 dataset, can you guess what categories are they from? (answer key below)

sification [1, 2, 3, 4, 5, 6, 7], action recognition [8, 9], discovering stylistic elements [10, 11], geometry estimation [12] and 2D-3D alignment [13, 14].

Originally motivated by market basket analysis, association rule mining is a well-known pattern mining algorithm that aims to discover a collection of if-then rules (*i.e.*, association rules) from a large number of records named transactions. The main advantage of association rule mining lies in its ability to process “Big Data”: association rules can be discovered from millions of transactions efficiently. In the context of mid-level visual element discovery, as noted by Doersch *et al.* [2], finding discriminative patches usually involves searching through tens of thousands of patches, which has become a bottleneck in recent works. In this sense, if appropriately used, association rule mining can be an appealing solution for handling “Big Data” in mid-level visual element discovery.

In this paper, building on the well-known association rule mining, we propose a pattern mining algorithm, *Mid-level Deep Pattern Mining* (MDPM), to study the problem of mid-level visual element discovery. This approach is particularly appealing because the specific properties of activations extracted from the fully-connected layer of a Convolutional Neural Network (CNN) allow them to be seamlessly integrated with association rule mining, which enables the

*The first two authors contributed to this work equally. This work was in part funded by the Data to Decisions Cooperative Research Centre, Australia. Correspondence should be addressed to C. Shen (e-mail: chhshen@gmail.com).

¹Answer key: 1.aeroplane, 2.train, 3.cow, 4.motorbike, 5.bike, 6.sofa.

discovery of category-specific patterns from a large number of image patches. Moreover, we find that two requirements of mid-level visual elements, representativeness and discriminativeness, can be effortlessly fulfilled by association rule mining. When we visualize image patches with the same pattern (mid-level visual element in our scenario), it turns out that they are not only visually similar, but also semantically consistent (see Fig. 1).

To our knowledge, hand-crafted features, typically HOG [15], are used as feature descriptors for image patches in all current methods of mid-level visual element discovery. Vondrick *et al.* [16], however, have illustrated the limitations of HOG, implying that HOG may be too lossy a descriptor to achieve high recognition performance. In this sense, an extra bonus of our formulation lies in that we are now relying on CNN activations, a more appealing alternative than the hand-crafted HOG, as indicated in recent works [17, 18, 19, 20, 21, 22].

One issue must be considered before using any pattern mining algorithms, that is, they have two strict requirements for the transactions that they can process (Sec. 4.1). Thanks to two appealing properties of CNN activations (Sec. 3), these two requirements are effortlessly fulfilled in the proposed MDPM algorithm (Sec. 4).

To show the effectiveness of the proposed MDPM algorithm, we apply it to scene and generic object classification tasks (Sec. 5). Specifically, after retrieving visual elements from the discovered patterns, we train element detectors and generate new feature representations using these detectors. We demonstrate that we achieve classification results which not only outperform all current methods in mid-level visual element discovery by a noticeable margin with far fewer elements used, but also outperform or match the performance of state-of-the-art using CNNs for the same task.

In summary, our contributions are twofold.

1. We formulate mid-level visual element discovery from the perspective of pattern mining, finding that its two requirements, representativeness and discriminativeness, can be easily fulfilled by the well-known association rule mining algorithm.
2. We present two properties of CNN activations that allow seamless integration with association rule mining, avoiding the limitations of pattern mining algorithms.

Preliminaries. To extract CNN activations, we rely on the publicly available `caffe` [23] reference model which was pre-trained on the ImageNet [24]. More specially, given a mean-subtracted patch or image, we resize it to the size of 227×227 and pass it to the `caffe` CNN. We extract its non-negative 4096-dimensional CNN activation from the sixth layer `fc6` (the first fully-connected layer) after the rectified linear unit (ReLU) transformation.

2. Related Work

Mid-level visual element discovery. Mid-level visual element discovery aims to discover clusters of image patches that are both representative and discriminative. Recent studies on this topic have shown that mid-level visual elements can be used for image classification [1, 2, 3, 4, 5, 6, 7]. The process typically proceeds as follows. Firstly, mining visual elements and training element detectors. Secondly, generating new feature representations using the most discriminative detectors. Various methods have been proposed for the first step, such as cross validation training of element detectors [1], ranking and selecting exemplar detectors on the validation set [3] and discriminative mode seeking [2].

Convolutional Neural Networks. Although proposed by LeCun *et al.* [25] for solving the handwritten digit recognition in the late '80s, CNNs have regained popularity mainly because of having shown very promising result in the ILSVRC challenge [26]. In the benchmark CNN architecture of Krizhevsky *et al.* [17], raw pixels first pass through five convolutional layers where responses of filters are max-pooled in sequence, before producing an activation of 4096 dimensions at each of the two fully-connected layers. Recent studies [18, 27] have demonstrated that the 4096-dimensional activation extracted from the fully-connected layer is an excellent representation for general recognition tasks.

Pattern mining in vision. Pattern mining techniques have been studied primarily amongst the data mining community, but a growing number of applications can be found in vision, such as image classification [28, 29, 30, 31], action recognition [32] and recognizing human-object interaction [33]. The main advantage of pattern mining lies its ability to process massive volumes of data, which is particularly appealing in this era of information explosion.

3. Properties of CNN activations of patches

In this section we provide a detailed analysis of the performance of CNN activations on the MIT Indoor dataset [34], from which we are able to conclude two important properties thereof. These two properties are key ingredients in the proposed MDPM algorithm in Sec. 4.

We first sample 128×128 patches with a stride of 32 pixels from each image. Then, for each image patch, we extract its 4096-dimensional CNN activation using `caffe`. To generate final the feature representation for an image, we consider two strategies as follows.

1. **CNN-Sparsified.** For each 4096-dimensional CNN activation of an image patch, we only keep the k largest dimensions (in terms of magnitude) and set the remaining elements to zero. The feature representation for an

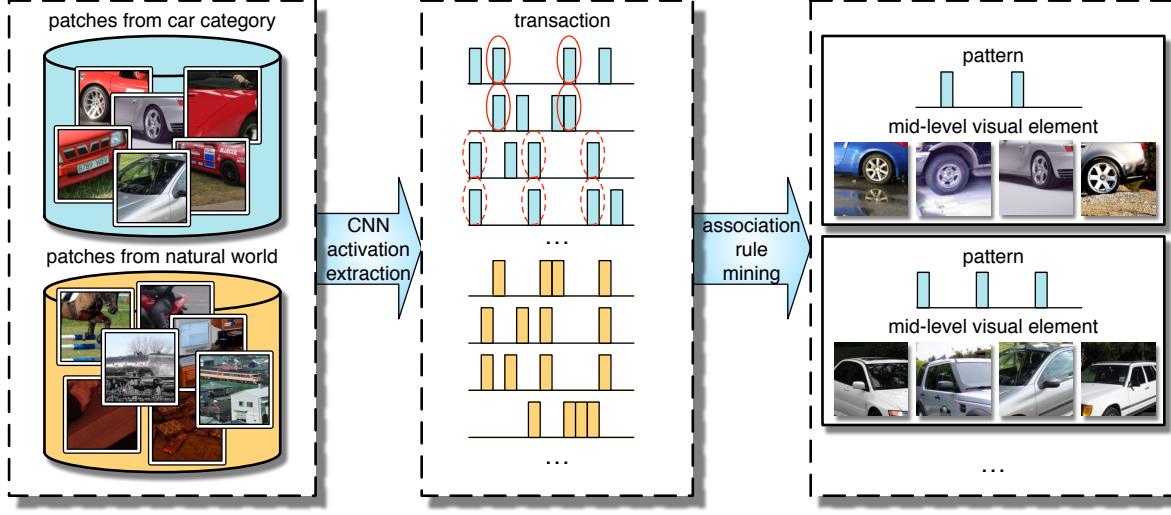


Figure 2. Pipeline of mid-level deep pattern mining. Given image patches sampled from both the target category (e.g. car) and the natural world, we represent each as a transaction after extracting their CNN activation. Patterns are then discovered through association rule mining. Mid-level visual elements are discovered by retrieving image patches with the same patterns.

k	10	20	50	100
CNN-Sparsified	50.10	56.33	60.34	61.68
CNN-Binarized	54.34	59.15	61.35	61.29

Table 1. The classification accuracies achieved by the two proposed strategies for keeping the k largest magnitudes of CNN activations of image patches on the MIT Indoor dataset.

image is the outcome of max pooling on the revised CNN activations.

2. **CNN-Binarized.** For each 4096-dimensional CNN activation of an image patch, we set its k largest dimensions to one and the remaining elements to zero. The feature representation for an image is the outcome of max pooling on these binarized CNN activations.

For each of the above cases, we train a multi-class linear SVM classifier in the one-vs-all fashion. The classification accuracy achieved by each of the above strategies for a range of k values is summarized in Table 1. In comparison, our baseline feature, which is the outcome of max pooling on CNN activations of all patches in an image, gives an accuracy of 65.15%. Analysing the results in Table 1 leads to 2 conclusions:

1. **Sparse.** Comparing the performance of ‘‘CNN-Sparsified’’ with that of the baseline feature, it is clear that accuracy is reasonably high when using sparsified CNN activations with only k non-zero magnitudes out of 4096.
2. **Binary.** Comparing ‘‘CNN-Binarized’’ case with its ‘‘CNN-Sparsified’’ counterpart, it is observed that CNN activations do not suffer from binarization when k is small, accuracy even increases slightly instead.

Conclusion. The above two properties imply that for an image patch, the discriminative information within its CNN activation is mostly embedded in *the dimension indices of the k largest magnitudes*.

4. Mid-level deep pattern mining

In this section, we give details of the proposed MDPM algorithm, an overview of which is provided in Fig. 2. We start by introducing some important concepts and terminology in pattern mining.

4.1. Pattern mining revisited

Frequent itemset. Let $A = \{a_1, a_2, \dots, a_K\}$ denote a set of K items. A transaction T is a subset of A , i.e., $T \subseteq A$. We also define a transaction database $\mathcal{D} = \{T_1, T_2, \dots, T_N\}$ containing N transactions (N is usually very large). Given I , a subset of A , we are interested in the fraction of transactions $T \in \mathcal{D}$ which contain I . The *support* value of I reflects this quantity:

$$\text{supp}(I) = \frac{|\{T | T \in \mathcal{D}, I \subseteq T\}|}{N} \in [0, 1], \quad (1)$$

where $|\cdot|$ measures cardinality. I is called a *frequent itemset* when $\text{supp}(I)$ is larger than a predefined threshold.

Association rule. An *association rule* $I \rightarrow a$ implies a relationship between I and an item a . We are interested in how likely it is that a is present in the transactions which contain I within \mathcal{D} . In a typical pattern mining application this might be taken to imply that customers who bought items in I are also likely to buy item a , for instance. The *confidence* of an association rule $\text{conf}(I \rightarrow a)$ can be taken to reflect

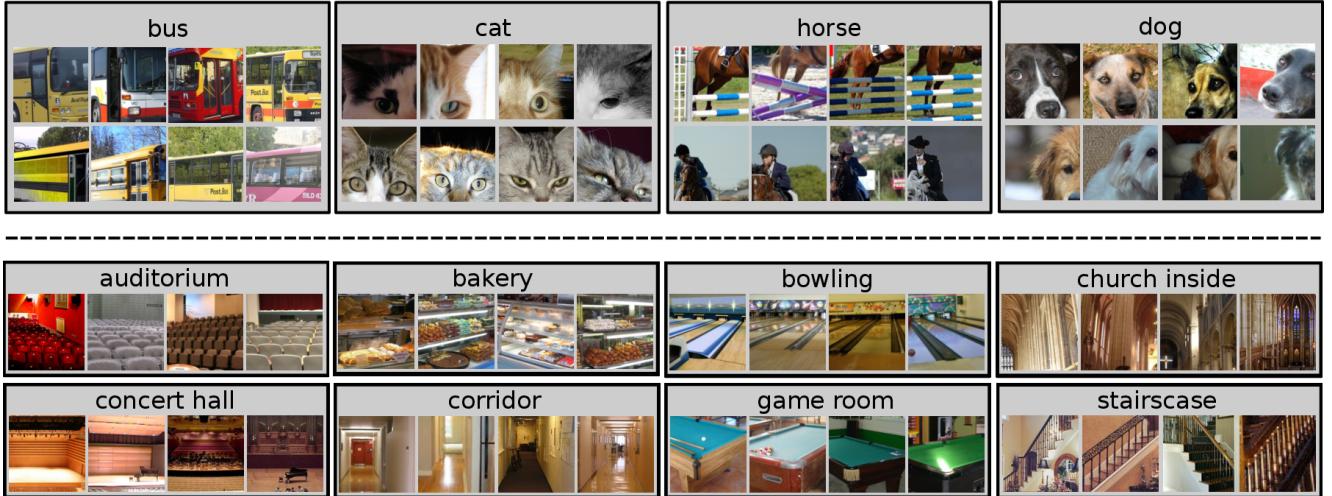


Figure 3. Mid-level visual elements from PASCAL VOC 2007 (top, two per category) and MIT Indoor (bottom, one per category) datasets.

this probability:

$$\begin{aligned} \text{conf}(I \rightarrow a) &= \frac{\text{supp}(I \cup \{a\})}{\text{supp}(I)} \\ &= \frac{|\{T|T \in \mathcal{D}, (I \cup \{a\}) \subseteq T\}|}{|\{T|T \in \mathcal{D}, I \subseteq T\}|} \in [0,1]. \end{aligned} \quad (2)$$

In practice, we are interested in “good” rules, that is, the confidence of these rules should be reasonably high.

Two strict requirements of pattern mining. As noted in [31], there are two strict requirements that must be met if we use pattern mining algorithms.

1. Each transaction can only have a small number of items, as the potential search space grows exponentially with the number of items in each transaction.
2. What is recorded in a transaction is a set of integers, as opposed to the real-valued elements of most vision features (such as SIFT and HOG for example).

4.2. Transaction creation

Transactions must be created before any pattern mining algorithms can process. In our work, as we aim to discover patterns from image patches through pattern mining, *an image patch is utilized to create one transaction*.

The most critical question now is how to transform an image patch into a transaction while maintaining most of its discriminative information. In this work, we rely on its CNN activation which has two appealing properties (Sec. 3). More specifically, we treat *each dimension index of a CNN activation* as an item (4096 items in total). Thanks to two properties in Sec. 3, each transaction is then represented by *the dimension indices of the k largest magnitudes of the corresponding image patch*.

This strategy satisfies both requirements given in Sec. 4.1. Specifically, due to the sparse nature of CNN activations (sparse property in Sec.3), each integer vector transaction calculated as described contains only k items, and k can be set to be small (20 in all of our experiments).

Following the work of [28], at the end of each transaction, we add a *pos* (*neg*) item if the corresponding image patch comes from the target category (natural world). Therefore, each complete transaction has $k + 1$ items, consisting of indices of k largest CNN magnitudes plus one class label. For example, if we set k equals three, given a CNN activation of an image patch from the target category which has 3 largest magnitudes in its 3rd, 100th and 4096th dimensions, the corresponding transaction will be {3,100,4096,*pos*}.

In practice, we first sample a large number of patches from images in both the target category and the natural world. After extracting their CNN activations from `caffe`, a transaction database \mathcal{D} is created containing a large number of transactions created using the proposed technique above. Note that the class labels, *pos* and *neg*, are represented by 4097 and 4098 respectively in the transactions.

4.3. Association rule mining

Given the transaction database \mathcal{D} , we use the Apriori algorithm [35] to discover a set of patterns \mathcal{P} through association rule mining. Each pattern $P \in \mathcal{P}$ must satisfy the following two criteria:

$$\text{supp}(P) > \text{supp}_{\min}, \quad (3)$$

$$\text{conf}(P \rightarrow \text{pos}) > \text{conf}_{\min}, \quad (4)$$

where supp_{\min} and conf_{\min} are thresholds for the support value and confidence.

Representativeness and discriminativeness. We now demonstrate how association rule mining implicitly satisfies the two requirements of mid-level visual element discovery, *i.e.*, representativeness and discriminativeness. Specifically, based on Eq. 3 and Eq. 4, we are able to rewrite Eq. 2 thus

$$\begin{aligned} \text{supp}(P \cup \{pos\}) &= \text{supp}(P) \times \text{conf}(P \rightarrow pos) \\ &> \text{supp}_{\min} \times \text{conf}_{\min}, \end{aligned} \quad (5)$$

where $\text{supp}(P \cup \{pos\})$ measures the fraction of pattern P found in transactions of the target category among all the transactions. Therefore, values of $\text{supp}(P)$ and $\text{conf}(P \rightarrow pos)$ above their thresholds ensure that pattern P is found frequently in the target category, akin to the representativeness requirement. A high value of conf_{\min} (Eq. 4) will also ensure that pattern P is more likely to be found in the target category rather than the natural world, reflecting the discriminativeness requirement.

5. Application to image classification

We now apply our MDPM algorithm to the image classification task. To discover patterns from a particular class, this class is treated as the target category while all other classes in the dataset are treated as the natural world. Note that only training images are used to discover patterns.

5.1. Retrieving mid-level visual elements

Given the set of patterns \mathcal{P} discovered in Sec. 4, finding mid-level visual elements is straightforward. A mid-level visual element V contains the image patches sharing the same pattern P , which can be retrieved efficiently through an inverted index. This process gives rise to a set of mid-level visual elements \mathcal{V} (*i.e.* $V \in \mathcal{V}$).

We provide a visualization of some of the visual elements discovered by the MDPM in Fig. 3. It is clear that image patches in each visual element are visually similar and depict similar semantic concepts. An interesting observation is that visual elements discovered by the MDPM are invariant to horizontal flipping.

5.2. Ensemble merging and training detectors

We note that patches belonging to different elements may overlap or describe the same visual concept. To remove this redundancy, we propose to merge elements in an iterative procedure while training element detectors.

Algorithm 1 summarizes the proposed ensemble merging procedure. At each iteration, we greedily merge overlapping mid-level elements and train the corresponding detector through the `MergingTrain` function in Algorithm 1. In the `MergingTrain` function, we begin by selecting the element covering the maximum number of training images, followed by training a Linear Discriminant

Algorithm 1: Ensemble Merging Pseudocode

```

Input: A set of partially redundant visual elements  $\mathcal{V}$ 
Output: A set of clean mid-level visual elements  $\mathcal{V}'$ 
        and corresponding element detectors  $D$ 
Initialize  $\mathcal{V}' \leftarrow \emptyset$ ,  $D_c \leftarrow \emptyset$ ;
while  $\mathcal{V} \neq \emptyset$  do
     $[\mathcal{V}_t, d] \leftarrow \text{MergingTrain}(\mathcal{V});$ 
     $\mathcal{V} \leftarrow \mathcal{V} \setminus \mathcal{V}_t;$ 
     $\mathcal{V}' \leftarrow \mathcal{V}' \cup \{\bigcup_{V \in \mathcal{V}_t} V\};$ 
     $D \leftarrow D \cup \{d\};$ 
end
return  $\mathcal{V}', D$ ;

Function MergingTrain ( $\mathcal{V}$ )
    Select  $V^* \in \mathcal{V}$  which covers the maximum
    number of training images;
    Initialize  $\mathcal{V}_t \leftarrow \{V^*\}$ ,  $\mathcal{S} \leftarrow \emptyset$ ;
    repeat
         $\mathcal{V}_t \leftarrow \mathcal{V}_t \cup \mathcal{S}$ ;
        Train LDA detector  $d$  using  $\mathcal{V}_t$ ;
         $\mathcal{S} \leftarrow \{V \in \mathcal{V} \setminus \mathcal{V}_t | \text{Score}(V, d) > Th\}$  where
         $\text{Score}(V, d) = \frac{1}{|\mathcal{V}|} \sum_{x \in V} d^T x$  ( $Th$  is a
        pre-defined threshold,  $x$  is a CNN activation of
        an image patch in  $V$ );
    until  $\mathcal{S} = \emptyset$ ;
    return  $\mathcal{V}_t, d$ ;

```

Analysis (LDA) detector [36]. We then incrementally revise this detector. At each step, we run the current detector on the patches of all the remaining visual elements, and retrain it by augmenting the positive training set with positive detections. We repeat this iterative procedure until no more elements can be added into the positive training set. The idea behind this process is using detection score as a similarity metric, much inspired by Exemplar SVM [37, 38].

The final output of the ensemble merging step is a clean set of visual elements and their corresponding detectors.

5.3. Selecting and encoding

We can now use the learned element detectors to encode a new image. There is a computational cost, however, associated with applying each successive learned element, and particular elements may be more informative when applied to particular tasks. We thus now seek to identify those elements of most value to the task at hand.

In practice, we rank all of the elements in a class based on the number of training images that they cover. We then select the detectors corresponding to the elements which cover the maximum number of training images, akin to “maximizing coverage” in [2]. This process is then repeated such that the same number of detectors are selected from

each class and stacked together.

To generate a feature representation for a new image, we evaluate all of the selected detectors at three scales. For each scale, we take the max score per detector per region encoded in a 2-level (1×1 and 2×2) spatial pyramid. The final feature representation of an image is the outcome of max pooling on the features from all three scales.

6. Experiments

We test our algorithm on two image classification tasks, scene classification and generic object classification.

Implementation details. For each image, we resize its smaller dimension to 256 while maintaining its aspect ratio, then we sample 128×128 patches with a stride of 32 pixels, and calculate the CNN activations using `caffe`. Because the number of patches sampled varies in different datasets, two parameters supp_{\min} and conf_{\min} in the association rule mining (Sec. 4.3) are set according to each dataset with the goal that at least 100 patterns are discovered for each category. The merging threshold Th in Algorithm 1 (Sec. 5.2) is set as 150. For training classifiers, we use the Liblinear toolbox [39] with 5-fold cross validation.

6.1. Scene classification

The MIT Indoor dataset [34] contains 67 classes of indoors scenes. Verified by recent works on mid-level visual element discovery, indoor scenes are better characterized by the unique objects that they contain (*e.g.*, computers are more likely to be found in computer room rather than laundry). We follow the standard partition of [34], *i.e.*, approximately 80 training and 20 test images per class. supp_{\min} and conf_{\min} are set as 0.01% and 30% respectively.

Comparison to methods using mid-level visual elements. We first compare our approach against recent works on mid-level visual element discovery (See Table 2). Using only 20 visual elements per class, our approach yields an accuracy of 68.24%. Increasing the number of visual elements to 50 makes our performance increases to 69.69%, outperforming all previous mid-level visual element discovery methods by a sizeable margin. As shown in Table 2, compared with the work of Doersch *et al.* [2] which achieved best performance among previous mid-level visual elements algorithms, our approach uses an order of magnitude fewer elements than [2] (20 vs. 200) while outperforming it by over 4 percent in accuracy. Also, our approach surpasses a very recent work RFDC [7] by over 15% in the same setting (50 elements per class). Thanks to the fact that CNN activations from `caffe` are invariant to horizontal flipping [22], we avoid adding right-left flipped images (*c.f.* [3, 2]).

Ablation study. As we are the first to use CNN activation for mid-level visual elements discovery, a natural question is that what is the performance of previous works if CNN

Method	#elements	Acc(%)
D-patch [1]	210	38.10
BoP [3]	50	46.10
miSVM [5]	20	46.40
MMDL [6]	11	50.15
D-Parts [4]	73	51.40
RFDC [7]	50	54.40
DMS [2]	200	64.03
LDA-Retrained	20	58.78
LDA-Retrained	50	62.30
LDA-KNN	20	59.14
LDA-KNN	20	63.93
Ours	20	68.24
Ours	50	69.69

Table 2. Classification results of mid-level visual element discovery algorithms on the MIT Indoor dataset.

Method	Acc(%)	Comments
CNN-G	57.74	CNN for whole image
CNN-Avg	65.77	average pooling
CNN-Max	65.15	max pooling
CNN-SVM [18]	58.40	OverFeat toolbox
VLAD level2 [22]	65.52	VLAD encoding
VLAD level3 [22]	62.24	VLAD encoding
VLAD level1&2 [22]	66.64	concatenation
MOP-CNN [22]	68.88	concatenation
CNN-jittered [40]	65.50	jittered CNN
CNN-finetune [40]	66.30	fine-tuned CNN
Places-CNN [41]	68.24	Places dataset used
SCFVC [42]	68.20	new Fisher encoding
CL-45C [43]	68.80	cross-layer pooling
Ours	69.69	MDPM (50 element)
Ours+CNN-G	70.46	concatenation

Table 3. Classification results of methods using CNN activations on the MIT Indoor dataset.

activations are adopted. To answer this question, we implement two baselines using CNN activations as image patch representations. “LDA-Retrained” initially trains Exemplar LDA using a sampled patch and then re-trains the detector 10 times by adding top-10 positive detections as positive training samples at each iteration. This is quite similar to the “Expansion” step of [3]. Another baseline “LDA-KNN” retrieves 5-nearest neighbors of an image patch and trains a LDA detector using the retrieved patches (including itself) as positive training data. For both baselines, discriminative detectors are selected based on the Entropy-Rank Curves proposed in [3]. As shown in Table 2, we report much better results than both baselines in the same setting, which verifies that the proposed MDPM is an essential step for achieving good performance when using CNN activations for mid-level visual element discovery.

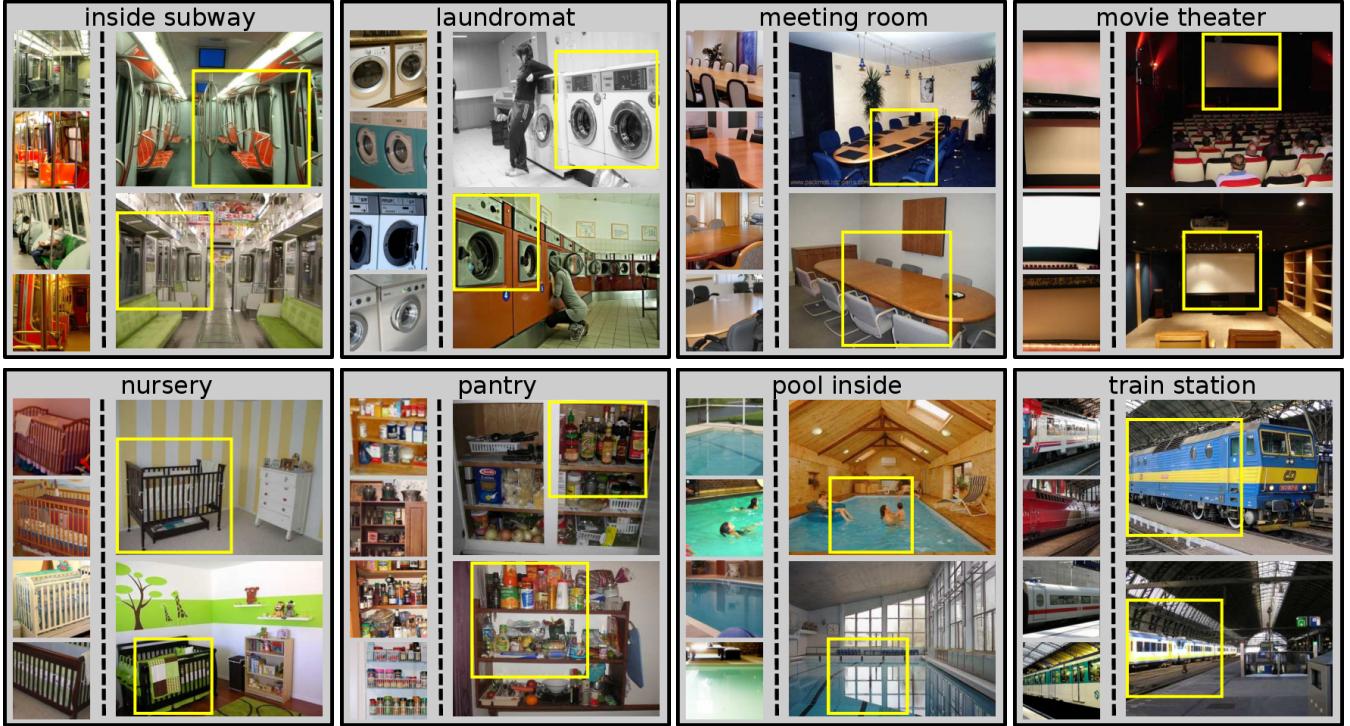


Figure 4. Discovered mid-level visual elements and their corresponding detections on test images on the MIT Indoor dataset.

Comparison to methods using CNN activation. In Table 3, we compare our approach to others in which CNN activation is involved. Our baseline method, using f_{c6} CNN activations extracted from the whole image, gives an accuracy of 57.74%. Our approach (based on the MDPM algorithm) achieves 69.69% accuracy, which is a significant improvement over all the baselines. Our work is most closely related to MOP-CNN [22] and SCFVC [43] in the sense that all these works rely on off-the-shelf CNN activations of image patches. To encode these local CNN activations, MOP-CNN [22] rely on the classical VLAD encoding [44], whereas SCFVC [43] is a new Fisher encoding [45] strategy for encoding high-dimensional local features. Our encoding method, which is based on the discovered visual elements, not only outperforms MOP-CNN [22] on 128×128 and 64×64 patches by a noticeable margin (69.69% vs. 65.52% and 69.69% vs. 62.24%), also slightly bypasses that of SCFVC (69.69% vs. 68.20%).

Fine-tuning has been shown to be beneficial when transferring pre-trained CNN models on the ImageNet to another dataset [19, 20, 21]. Jittered CNN features (*e.g.*, crops, flips) extracted from the fine-tuned network of Azizpour *et al.* [40] offer 66.3% accuracy, which is still below ours.

After concatenating with CNN activations of the whole image (both normalized to unit norm), our performance increases to 70.46%, outperforming all previous works using CNN on this dataset.

Computational complexity. Given pre-computed CNN activation from about 0.2 million patches, the baseline method “LDA-Retrained” takes about 9 hours to find visual elements in a class. However, our approach only takes about 3 minutes (writing transaction file and association rule mining) to discover representative and discriminative rules.

Visualization. We visualize some visual elements discovered and their firings on test images in Fig. 4. It is intuitive that the discovered mid-level visual elements capture the patterns which are often repeated within a scene category. Some of the mid-level visual elements refer to frequently occurring objects configurations, *e.g.*, the configuration between table and chair in the *meeting room* category. Some instead capture a particular object in the scene, such as the baby cot in the *nursery* and screen in the *movie theater*.

6.2. Object classification

The Pascal VOC 2007 dataset [46] contains 9,963 images from 20 object classes. For this dataset, training and validation sets are utilized to discover patterns and training final classifiers. The parameters supp_{\min} and conf_{\min} are set as 0.01% and 60% respectively.

Comparison to methods using CNN activation. Table 4 reports our results along with those of other recent methods based on CNN activation. On this dataset, when using 50 visual elements per class, the proposed method achieves

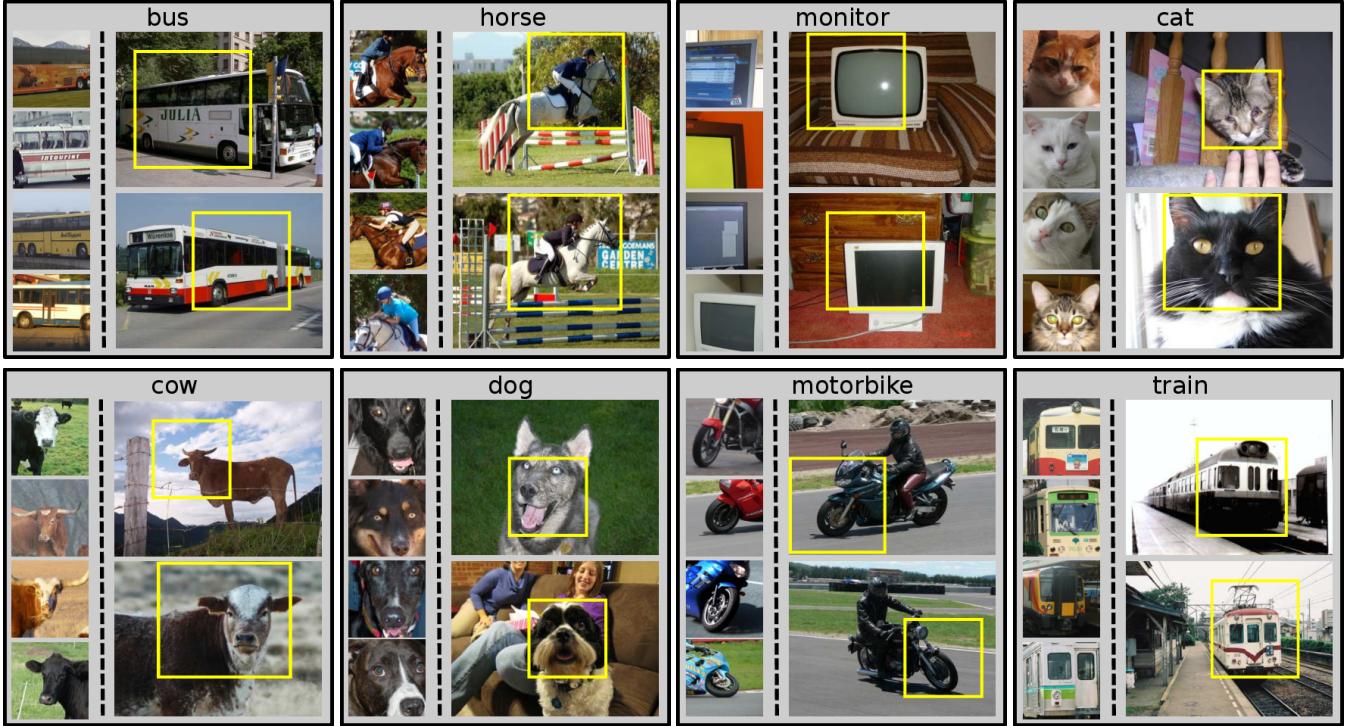


Figure 5. Discovered mid-level visual elements and their corresponding detections on test images on the PASCAL VOC 2007 dataset.

Method	mAP(%)	Comments
CNN-G	67.3	CNN for whole image
CNN-Avg	70.9	average pooling
CNN-Max	71.1	max pooling
CNN-SVM [18]	73.9	OverFeat toolbox
PRE-1000C [20]	77.7	bounding box used
CNN-jittered [40]	72.2	jittered CNN
SCFVC [42]	76.9	new Fisher encoding
CL-45C [43]	75.0	cross-layer pooling
Ours	75.2	MDPM (50 elements)
Ours+CNN-G	77.9	concatenation

Table 4. Classification results of methods using CNN activations on the PASCAL VOC 2007 dataset.

75.2% mean average precision (mAP), significantly outperforming the baseline that using CNN activations as a global feature (67.3%), as well as its average pooling and max pooling counterparts. Compared with state-of-the-arts, Oquab *et al.* [20] fine-tune the pre-trained network on the ImageNet, however, it relies on bounding box annotation which makes the task easier, so it is not surprising that it outperforms ours which does not use bounding box annotation. The best result on PASCAL VOC 2007 (77.9%) is achieved when the proposed MDPM feature and global CNN activation are concatenated, marginally outperforming fine-tuning with bounding box annotation [20]. This is

despite the fact that the bounding box annotations constitute extra information which is time-consuming to gather.

Visualization. We visualize some visual elements discovered and their firings on the test images of the VOC 2007 dataset in Fig. 5. It is clear that the discovered mid-level visual elements capture some discriminative parts of object (*e.g.*, dog faces). It is worth noting here that “parts” have been shown to be extremely important for state-of-the-art object recognition, such as Deformable Part Models [47] and Poselets [48].

7. Conclusion

We have addressed mid-level visual element discovery from the perspective of pattern mining. In the process we have shown not only that it is profitable to apply pattern mining technique to mid-level visual element discovery, but also that, from the right perspective, CNN activations are particularly well suited to the task. This is significant because CNNs underpin many current state-of-the-art methods in vision, and pattern mining underpins significant elements of the state-of-the-art in Big Data processing.

References

- [1] S. Singh, A. Gupta, and A. A. Efros, “Unsupervised discovery of mid-level discriminative patches,” in *Proc. Eur. Conf. Comp. Vis.*, 2012, pp. 73–86.

- [2] C. Doersch, A. Gupta, and A. A. Efros, “Mid-level visual element discovery as discriminative mode seeking,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2013, pp. 494–502.
- [3] M. Juneja, A. Vedaldi, C. V. Jawahar, and A. Zisserman, “Blocks that shout: Distinctive parts for scene classification,” in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2013, pp. 923–930.
- [4] J. Sun and J. Ponce, “Learning discriminative part detectors for image classification and cosegmentation,” in *Proc. IEEE Int. Conf. Comp. Vis.*, 2013, pp. 3400–3407.
- [5] Q. Li, J. Wu, and Z. Tu, “Harvesting mid-level visual concepts from large-scale internet images,” in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2013, pp. 851–858.
- [6] X. Wang, B. Wang, X. Bai, W. Liu, and Z. Tu, “Max-margin multiple-instance dictionary learning,” in *Proc. Int. Conf. Mach. Learn.*, 2013, pp. 846–854.
- [7] L. Bossard, M. Guillaumin, and L. V. Gool, “Food-101 mining discriminative components with random forests,” in *Proc. Eur. Conf. Comp. Vis.*, 2014, pp. 446–461.
- [8] A. Jain, A. Gupta, M. Rodriguez, and L. S. Davis, “Representing videos using mid-level discriminative patches,” in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2013, pp. 2571–2578.
- [9] L. Wang, Y. Qiao, and X. Tang, “Motionlets: Mid-level 3d parts for human motion recognition,” in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2013, pp. 2674–2681.
- [10] C. Doersch, S. Singh, A. Gupta, J. Sivic, and A. A. Efros, “What makes paris look like paris?,” *Ann. ACM SIGIR Conf.*, vol. 31, no. 4, pp. 101, 2012.
- [11] Y. J. Lee, A. A. Efros, and M. Hebert, “Style-aware mid-level representation for discovering visual connections in space and time,” in *Proc. IEEE Int. Conf. Comp. Vis.*, 2013, pp. 1857–1864.
- [12] D. F. Fouhey, A. Gupta, and M. Hebert, “Data-driven 3d primitives for single image understanding,” in *Proc. IEEE Int. Conf. Comp. Vis.*, 2013, pp. 3392–3399.
- [13] M. Aubry, B. C. Russell, and J. Sivic, “Painting-to-3d model alignment via discriminative visual elements,” *Ann. ACM SIGIR Conf.*, vol. 33, no. 2, pp. 14, 2014.
- [14] M. Aubry, D. Maturana, A. A. Efros, B. C. Russell, and J. Sivic, “Seeing 3d chairs: exemplar part-based 2d-3d alignment using a large dataset of cad models,” in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2014, pp. 3762–3769.
- [15] N. Dalal and B. Triggs, “Histograms of oriented gradients for human detection,” in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2005, pp. 886–893.
- [16] C. Vondrick, A. Khosla, T. Malisiewicz, and A. Torralba, “Hoggles: Visualizing object detection features,” in *Proc. IEEE Int. Conf. Comp. Vis.*, 2013, pp. 1–8.
- [17] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2012, pp. 1106–1114.
- [18] A. S. Razavian, H. Azizpour, J. Sullivan, and S. Carlsson, “Cnn features off-the-shelf: An astounding baseline for recognition,” in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn. Workshops*, 2014, pp. 512–519.
- [19] R. Girshick, J. Donahue, T. Darrell, and J. Malik, “Rich feature hierarchies for accurate object detection and semantic segmentation,” in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2014, pp. 580–587.
- [20] M. Oquab, L. Bottou, I. Laptev, and J. Sivic, “Learning and transferring mid-level image representations using convolutional neural networks,” in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2014, pp. 1717–1724.
- [21] P. Agrawal, R. Girshick, and J. Malik, “Analyzing the performance of multilayer neural networks for object recognition,” in *Proc. Eur. Conf. Comp. Vis.*, 2014, pp. 329–344.
- [22] Y. Gong, L. Wang, R. Guo, and S. Lazebnik, “Multi-scale orderless pooling of deep convolutional activation features,” in *Proc. Eur. Conf. Comp. Vis.*, 2014, pp. 392–407.
- [23] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell, “Caffe: Convolutional architecture for fast feature embedding,” *arXiv preprint arXiv:1408.5093*, 2014.
- [24] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and F.-F. Li, “Imagenet: A large-scale hierarchical image database,” in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2009, pp. 248–255.
- [25] Y. LeCun, B. E. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. E. Hubbard, and L. D. Jackel, “Handwritten digit recognition with a back-propagation network,” in *Proc. Adv. Neural Inf. Process. Syst.*, 1989, pp. 396–404.
- [26] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. S. Bernstein, A. C. Berg, and L. Fei-Fei, “Imagenet large scale visual recognition challenge,” *Int. J. Comp. Vis.*, 2015.
- [27] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell, “Decaf: A deep convolutional activation feature for generic visual recognition,” in *Proc. Int. Conf. Mach. Learn.*, 2014, pp. 647–655.
- [28] T. Quack, V. Ferrari, B. Leibe, and L. J. V. Gool, “Efficient mining of frequent and distinctive feature configurations,” in *Proc. IEEE Int. Conf. Comp. Vis.*, 2007, pp. 1–8.
- [29] J. Yuan, M. Yang, and Y. Wu, “Mining discriminative co-occurrence patterns for visual recognition,” in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2011, pp. 2777–2784.
- [30] B. Fernando, É. Fromont, and T. Tuytelaars, “Effective use of frequent itemset mining for image classification,” in *Proc. Eur. Conf. Comp. Vis.*, 2012, pp. 214–227.
- [31] W. Voravuthikunchai, B. Crémilleux, and F. Jurie, “Histograms of pattern sets for image classification and object recognition,” in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2014, pp. 224–231.
- [32] J. Wang, Z. Liu, Y. Wu, and J. Yuan, “Mining actionlet ensemble for action recognition with depth cameras,” in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2012, pp. 1290–1297.

- [33] B. Yao and F.-F. Li, “Grouplet: A structured image representation for recognizing human and object interactions,” in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2010, pp. 9–16.
- [34] A. Quattoni and A. Torralba, “Recognizing indoor scenes,” in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2009, pp. 413–420.
- [35] R. Agrawal and R. Srikant, “Fast algorithms for mining association rules in large databases,” in *Proc. Int. Conf. Very Large Databases*, 1994, pp. 487–499.
- [36] B. Hariharan, J. Malik, and D. Ramanan, “Discriminative decorrelation for clustering and classification,” in *Proc. Eur. Conf. Comp. Vis.*, 2012, pp. 459–472.
- [37] T. Malisiewicz, A. Gupta, and A. A. Efros, “Ensemble of exemplar-svms for object detection and beyond,” in *Proc. IEEE Int. Conf. Comp. Vis.*, 2011, pp. 89–96.
- [38] A. Shrivastava, T. Malisiewicz, A. Gupta, and A. A. Efros, “Data-driven visual similarity for cross-domain image matching,” *Ann. ACM SIGIR Conf.*, vol. 30, no. 6, pp. 154, 2011.
- [39] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin, “Liblinear: A library for large linear classification,” *J. Mach. Learn. Res.*, vol. 9, pp. 1871–1874, 2008.
- [40] H. Azizpour, A. S. Razavian, J. Sullivan, A. Maki, and S. Carlsson, “From generic to specific deep representations for visual recognition,” *CoRR*, vol. abs/1406.5774, 2014.
- [41] B. Zhou, A. Lapedriza, J. Xiao, A. Torralba, and A. Oliva, “Learning deep features for scene recognition using places database,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2014.
- [42] L. Liu, C. Shen, L. Wang, A. van den Hengel, and C. Wang, “Encoding high dimensional local features by sparse coding based fisher vectors,” in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 1143–1151.
- [43] L. Liu, C. Shen, A. van den Hengel, and C. Wang, “The treasure beneath convolutional layers: Cross-convolutional-layer pooling for image classification,” in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2015.
- [44] H. Jegou, M. Douze, C. Schmid, and P. Pérez, “Aggregating local descriptors into a compact image representation,” in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2010, pp. 3304–3311.
- [45] F. Perronnin, J. Sánchez, and T. Mensink, “Improving the fisher kernel for large-scale image classification,” in *Proc. Eur. Conf. Comp. Vis.*, 2010, pp. 143–156.
- [46] M. Everingham, L. J. V. Gool, C. K. I. Williams, J. M. Winn, and A. Zisserman, “The pascal visual object classes (VOC) challenge,” *Int. J. Comp. Vis.*, vol. 88, no. 2, pp. 303–338, 2010.
- [47] P. F. Felzenszwalb, R. B. Girshick, D. A. McAllester, and D. Ramanan, “Object detection with discriminatively trained part-based models,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 9, pp. 1627–1645, 2010.
- [48] L. D. Bourdev and J. Malik, “Poselets: Body part detectors trained using 3d human pose annotations,” in *Proc. IEEE Int. Conf. Comp. Vis.*, 2009, pp. 1365–1372.