

# Scene Classification in Indoor Environments for Robots using Context Based Word Embeddings

Bao Xin Chen, Raghavender Sahdev, Dekun Wu, Xing Zhao, Manos Papagelis and John K. Tsotsos

**Abstract**—Scene Classification has been addressed with numerous techniques in the computer vision literature. However with the increasing number of scene classes in datasets in the field, it has become difficult to achieve high accuracy in the context of robotics. In this paper, we implement an approach which combines traditional deep learning techniques with natural language processing methods to generate a word embedding based Scene Classification algorithm. We use the key idea that context (objects in the scene) of an image should be representative of the scene label meaning a group of objects could assist to predict the scene class. Objects present in the scene are represented by vectors and the images are re-classified based on the objects present in the scene to refine the initial classification by a Convolutional Neural Network (CNN). In our approach we address indoor Scene Classification task using a model trained with a reduced pre-processed version of the Places365 dataset and an empirical analysis is done on a real world dataset that we built by capturing image sequences using a GoPro camera. We also report results obtained on a subset of the Places365 dataset using our approach and additionally show a deployment of our approach on a robot operating in a real world environment.

## I. INTRODUCTION

Scene Classification has been extensively researched by the computer vision [1], [2], [3] community. Most datasets have millions of images which makes the task challenging and thereby encourages the vision community to come up with better algorithms to achieve higher accuracies. However, on some of these datasets, current state of the art algorithms have low accuracies which renders their usage in the real world impractical. For instance the top-1 accuracy on the Places365 dataset [1] is only around 56%. Such an accuracy is not of much use when deploying algorithms on a robot as nearly half of the time the robot would make erroneous predictions. In the past, there have been datasets with a limited number of different places, like [4], [5], which had only upto 17 places, however these datasets’ usage has been rendered obsolete due to already achieved near human accuracy ( $\approx 95\%$ ).

In this paper, we propose an approach for Scene Classification in real time in the context of robotics. We primarily target indoor scenes in this paper. In indoor environments, GPS cannot be relied on for metric level precision. We can only get a rough estimate of the position of the robot from GPS in indoor scenes. We leverage this inaccuracy in the GPS to reduce the search space of the places. We propose a taxonomy based approach to perform Scene Classification

All authors are with the Department of Electrical Engineering and Computer Science at York University, Toronto, Ontario, Canada {baoxchen, sahdev, jackwu, xingzhao, papaggel, tsotsos}@cse.yorku.ca

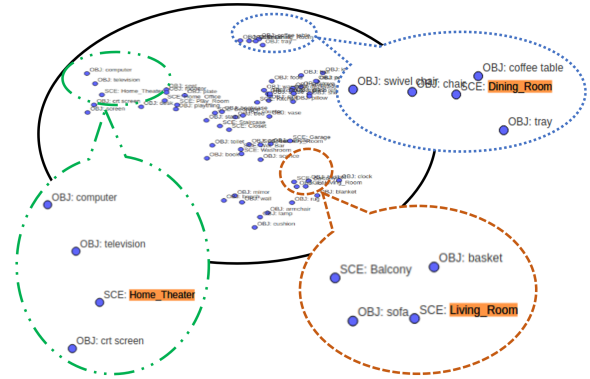


Fig. 1. Object and Scene vectors represented in a 2D space after learning from our word2vec model. Objects related to specific scenes are closer to each other, e.g., *coffee table*, *tray* in *Living Room*; *crt screen*, *television* in *Home Theatre*. Figure generated using tSNE Visualization tool <sup>2</sup>

by dividing the places according to the area type (e.g., *school*, *shopping-mall*, *home*, etc.). Once we know roughly which region the robot is in, we need to only search for a subset of places rather than searching all the classes. For example, from GPS we can know that the robot is in a school environment. This would reduce the search space to places like *corridors*, *lecture hall*, *seminar room*, *washroom*, etc. (see table I). We no longer need to consider the places that belong to any unrelated class like *living room*, *pet shop*, *jewelry store*, *forest*, *waterfall*, *mountains*, etc. After this step, many classes are pruned leaving a smaller set of candidate classes to select from. To classify the places, we propose to use different CNN models for each of the indoor areas present in the Places365 dataset [1].

After employing a CNN for Scene Classification, we obtain the top-5 predictions for a given image. To further refine the accuracy of the prediction, we make use of different objects present in the scene which are detected using a Scene Parsing Module (see Figure 2). Each object present in the image is represented using a word-embedding [6]. In most cases, these object embeddings are similar for objects belonging to a particular scene and the scene embedding itself, e.g., *tray*, *coffee table*, *chairs* in a *dinning room* (See Figure 1) should have similar embedding. Using these embeddings, we further refine the top-5 scores to improve the accuracy. Finally, we build our own test dataset to validate our approach as there are many inconsistencies in

<sup>2</sup><https://cs.stanford.edu/people/karpathy/tsnejs/csvdemo.html>

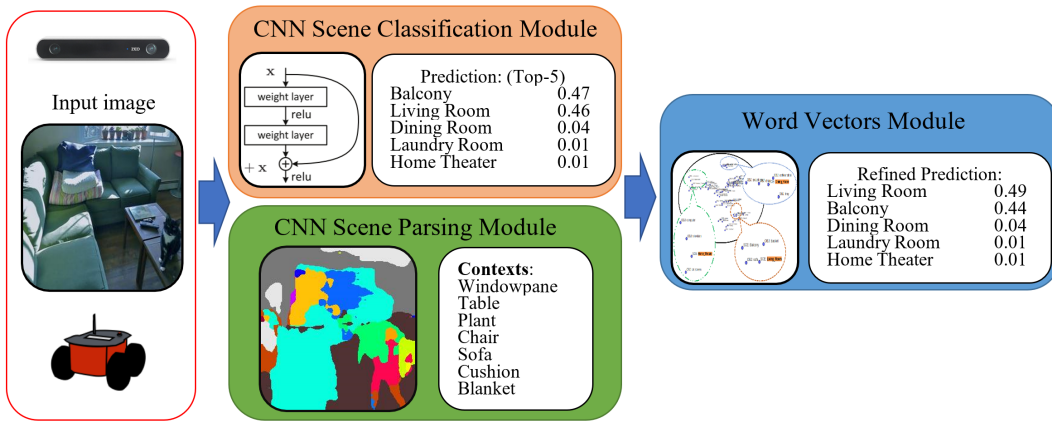


Fig. 2. Overview of our algorithm. Our approach takes an RGB image as input. Then passes it onto two CNN modules. The top module is for Scene Classification, which gives us an initial top-5 prediction. The bottom module is for Scene Parsing, which detects the scene contents including background (*window pane, plant, etc.*) and foreground (*table, chair, sofa, etc.*) objects. The Word Vectors Module computes a vector for the contents of the image, and a vector for each of the top-5 predicted labels. Then Word Vectors Module refines the ranking of these top-5 predicted labels by comparing the vector similarity.

the Places365 dataset as described in section IV. We also report the performance on a subset of the Places365 dataset.

The major contributions of this work are: (i) A taxonomy based approach to make Scene Classification work in real time with high accuracy for robots using GPS information, (ii) An empirical evaluation showing performance of our proposed approach, (iii) A real world dataset for the task of indoor Scene Classification. The paper is structured as follows: we describe the relevant work in Section II. Section III describes our proposed approach. Section IV describes our test dataset. Experimental results of our approach are described in Section V and finally we conclude our work in Section VI.

## II. RELATED WORK

In this section, we give a brief overview of the existing approaches used for Scene Classification. Li et al. [7] used object detectors as features to form an Object Bank Representation to assist Scene Classification. Our proposed approach is similar to [7], instead of using an Object Bank representation, we employed a word vector [6] based feature representation to find similarities between objects in the image and the scene category. Yang et al. [8] used key point detection to create feature vectors. Bosch et al. [9] used visual vocabulary as features to train a Support Vector Machine (SVM) classifier. Furthermore, different types of CNN architectures (AlexNet [10], VGG [11], GoogleNet [12]) were used in Scene Classification recently. These CNNs achieved decent state-of-the-art performance ( $\approx 56\%$ ) on Places365 [1] dataset. As well as, Khan et al. [13] integrated Places-VGG with Spectral Features to improve Scene Classification. However, the datasets they used are not appropriate for real world application.

The CNN generally learns a feature representation of an image but the network can not tell us the object relations in the image. However, a special CNN architecture (SegNet [14]) can be used to perform pixel level segmenta-

tion (Scene Parsing). A newly released pre-trained CNN model [15] based on a cascade segmentation module can find 150 types of contents in an image, which includes background (*wall, doors, windows, floor, etc.*) and foreground (*person, television, table, chair, etc.*) information. There are several approaches which can detect objects in an image; YOLO [16] detects most of the foreground objects. But their approach does not yield anything about the background of the image. Background information also plays an important role when a human is classifying a scene.

Word vectors [17] are very popular in Natural Language Processing (NLP). This is also known as word embedding. It improves the performance of LSTM-CRF [18] significantly on the named entity recognition systems without any language-specific knowledge [19]. Word vectors were trained to learn the vector representation of a given vocabulary. Word vectors have been used in many domains like sentiment analysis [6], detecting magnitude of events [20], named entity recognition [21], etc.

Scene Classification can also be useful for Place Recognition as it reduces the search space. Alternatively Place recognition can also be used to tackle scene classification (in case where all test images come from a set of predefined places and each place is associated with a scene type). Some of the existing Place recognition works include [3], [22], [4], [23].

**Data sets:** Many well-known data sets and benchmarks exist. For instance, Places88 [24] is the very first version of MIT Places benchmark<sup>3</sup>. It has 88 scene categories. The latest Scene Classification data set is MIT Place365 [1], which has 365 scene categories. In this paper we use Place365 as our training and evaluation data-set, as it is the largest scene classification data-set, and it contains broad categories. Some other old datasets also exist for Scene Classification. The Pascal VOC [25] data set has scene context annotations which are used for object detection and

<sup>3</sup><http://places.csail.mit.edu/index.html>

object segmentation tasks. Another data set is MS-COCO [26]. COCO is mainly used for instance detection and scene segmentation. The pre-trained CNN Scene Parsing model [15] used in this paper is trained on COCO.

### III. APPROACH

Our aim is to provide a Scene Classification approach which can perform well in a known real world environment (e.g., *school, home, shopping mall*). We build a taxonomy of different environments each containing separate scenes as shown in Table I. We pre-process the Places365 dataset to clean and reduce it to have only sensible indoor places for our use. Places365 has many inconsistencies as described in section IV.

We use two existing CNN models, one for Scene Classification [1] and the other for Scene Parsing [15]. Figure 2 shows the overview of our approach. The input is an RGB image for the two CNN modules. We call them CNN Scene Classification Module and CNN Scene Parsing Module (Figure 2). The Classification Module computes the initial raw top-5 predictions. The Scene Parsing module computes the scene contents from foreground and background. The Word Vector module computes the vector similarity between the contents present in the input image and the top-5 predicted labels. Using the computed similarity score, we output a refined re-ranked top-5 labels for the input image. We describe the details about each module in following subsections.

#### A. CNN Scene Classification Module

In this module, we train the ResNet models (one for each environment) on a reduced version of the Places365 dataset (based on our taxonomy, see table I). The code and model parameters can be downloaded from this link <sup>4</sup>. CNN Scene Classification Module computes the top-5 prediction labels which is further used by the Word Vectors Module to refine these prediction scores.

#### B. CNN Scene Parsing Module

In this module, we use a pre-trained CNN Scene Parsing model <sup>5</sup>. The CNN model was trained on ADE20K dataset [15]. It is trained to detect 150 different object categories from the given input image. The Scene Parsing module tells us the different objects/contents present in a given input image. For a kitchen input image, the parser would output the presence of *kettle, stove, oven, glasses, plates, etc.* We convert these object labels as English words and pass that onto Word Vectors module, described in the section below.

#### C. Word Vectors Module

Now we know the objects present in each of the images in the dataset. Knowing the top-5 predictions of each image and the objects present in each of the images, we need to make use of this information to refine the top-5 scores.

This section is the main contribution of this work. First, we define some of the notations being used and then describe our approach using these notations.

**Notations used:** let  $V$  be the dictionary/vocabulary containing the objects,  $O$  and the scenes,  $S$  present in the dataset.  $V = \{O, S\}$ , where  $O = \{o_1, o_2, o_3, o_4, \dots, o_{150}\}$  and  $S = \{s_1, s_2, s_3, \dots, s_n\}$ , where each  $o_i$  is a vector representing each of the objects and each  $s_j$  is a vector representing the scene classes present in the dataset. We further define a 2D weight matrix,  $W$  of scalars. The dimensionality of the weight matrix is  $|O| * |S|$ . where  $|O| = 150$ , the number of different objects we trained on and  $|S| = n$ , the number of defined scenes in a particular environment (e.g., for school,  $n = 24$ ). The matrix  $W = \{w_{1,1}, \dots, w_{i,j}, \dots, w_{150,n}\}$  where  $i \in [1, 150]$  and  $j \in [1, n]$ . Also we define  $T$  as the matrix containing the top-5 predictions for each of the images.  $T$  has dimensionality  $5 * |D|$  where  $|D|$  is the cardinality of the dataset.  $T_k$  represents the row of the top-5 predicted classes in the  $k^{th}$  Image,  $I_k$ . Similar to  $T$ ,  $ACC$  is the confidence for the predicted top-5 classes.  $IO_k$  is the set of objects present in the  $k^{th}$  image,  $I_k$ . A sample  $IO_k$  (for image,  $I_k$ ) would be like  $\{o_{14}, o_{52}, o_{78}, o_{113}, o_{143}, o_{149}\}$  implying say an input image (e.g., *home\_office*) consists of 6 objects  $\{laptop, keyboard, table, desk, chair, book\}$ .

---

#### Algorithm 1 Pseudocode of the proposed approach

---

**Input :**

A dataset,  $D$  of RGB images,  $I_k \in \mathbb{D} : D = \{I_1, I_2, \dots, I_n\}$

A vocabulary,  $V$  of object vectors,  $O$  and  $S : V = \{O, S\}$

A list of Object vectors,  $O$  in  $V : O = \{o_1, o_2, \dots, o_{150}\}$

A list of Scene vectors,  $S$  in  $V : S = \{s_1, s_2, \dots, s_n\}$

A weight matrix,  $W$  for objects and scene,  $W : \{w_{i,j}, \text{ where } i \in [1, 150]; j \in [1, n]\}$

**Output :**

Refined Top5 Prediction,  $RT_k$  for each image,  $I_k$  after Refinement

**Procedure 1, main:**

1. **for**  $I_k \in D$  **do**
2.    $T_k, ACC_k = SceneClassificationModule(I_k)$
3.    $IO_k = SceneParsingModule(I_k)$
4.    $RT_k = WordVectorModule(IO_k, T_k, ACC_k)$
5. **return**  $RT$

**Procedure 2, WordVectorModule ( $IO_k, T_k, ACC_k$ ) :**

1.  $Similarity[5] = \{0, 0, 0, 0, 0\}$
  2. **for**  $s_j \in T_k$  **do**
  3.    $IVector_k(j) = \langle 0 \rangle \text{ vector}$
  4.   **for**  $o_i \in IO_k$  **do**
  5.      $IVector_k(j) += o_i * w_{i,j}$
  6.      $Similarity[j] = cosine(IVector_k(j), s_j)$
  7.  $Similarity = Normalize(Similarity)$
  8.  $RT_k = descendingOrder(ACC_k * Similarity)$
  9. **return**  $RT_k$
- 

Now we have the initial top-5 predicted classes  $T_k$  for each image,  $I_k$  obtained from the Scene Classification module and the set of objects,  $IO_k$  present in each image obtained from the scene parser module. Our approach can be summarized in Algorithm 1. To compute the refined top-5 scene classes predictions,  $RT_k$ : first we compute a vector representation for the image with respect to each of the top-5 predicted

<sup>4</sup><http://places2.csail.mit.edu/>

<sup>5</sup><https://github.com/CSAILVision/sceneparsing>

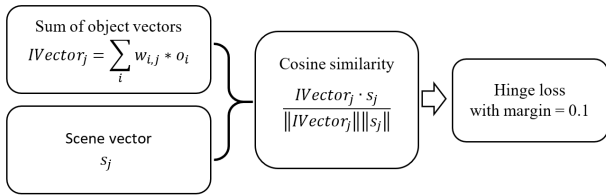


Fig. 3. Training Model for learning the new word2vector  $V$ , and the weighted table  $W$ .

classes. We compute a weighted vector sum of the detected objects in the image to represent the image vector. The vector representation for a given input image,  $I$  with respect to a scene  $s_j$  is computed as follows:

$$IVector_k(j) = \sum_i w_{i,j} * o_i \quad (1)$$

Here  $o_i \in IO_k$ . The weights,  $w_{i,j}$  and vectors,  $o_i$  are learned by the word vector module (see Figure 3).  $IVector_k(j)$  is the vector representation of the input image,  $I_k$  in terms of the objects present *w.r.t.* to scene label  $j$ . Refer to line 4,5 in Algorithm 1, Procedure 2.

**Learning the object and scene vectors,  $V$ , and the weight matrix,  $W$ :**

Vectors in  $V$  are initialized by the pre-trained word2vector<sup>6</sup>, and  $W$  are initialized by *tf-idf*<sup>7</sup>. When computing *tf-idf*, we considered each scene class is a document, and the objects in the image are the words. The object representing the scene has a higher weight and the object occurring less frequently in the scene would have a lower weight. We pre-compute the training predictions,  $T$ ,  $ACC$  and objects in images,  $IO$ . We feed this information into our cosine similarity training model (see Figure 3) to learn the weights,  $W$  and Vocabulary,  $V$ . The loss function used in our model was a hinge loss with margin equal to 0.1.

There are some objects,  $o_i \in O$  which might exist in most of the scenes. Such objects contribute less in determining the class of the scene. So a lesser weight should be assigned to such objects while computing the weighted sum representation of the image vector,  $IVector_k(j)$ . On the other hand, if an object exists only in some particular scenes, we need to increase the weight of this object. This is intuitive, for instance: the chances of finding a “sofa” is more likely in a “living-room” scene than any other place. Now we describe the dataset we built to validate our approach and provide empirical results obtained on our dataset and the Places365 dataset.

#### IV. THE DATASET

One of the current widely used datasets for Scene Classification is the Places365 dataset [1]. This dataset is the

<sup>6</sup><https://radimrehurek.com/gensim/models/word2vec.html>

<sup>7</sup><http://www.tfidf.com/>

latest and largest Scene Classification dataset with more than 1 million images. Places365 has 365 scene categories with 3000 to 5000 images per class. However in this dataset there are many confusing places that are practically the same but have been labelled distinctly. For example, the places {*bedroom, dorm room*}, {*lecture room, classroom*}, {*pharmacy, drug-store*}, etc. are some sets of places that are referring to the same but labelled differently. Moreover, after a close inspection of this dataset, it was found that many images in the dataset are incorrectly labelled. Due to these mentioned issues with the Places365 dataset, we decided to merge certain classes into one before training our models (see taxonomy file on the project page<sup>8</sup>). Additionally since in our work, we are only concerned about indoor places, we remove all outdoor places in the Places365 dataset. After this pre-processing step of removal and merging of places, we get a total of 156 different scenes for 8 different environments. We train our model using this pre-processed dataset.

For testing our proposed approach, we built a dataset in the real world of 69 different places (see table I), which we arranged in a hierarchical manner with 2 levels. First level (environment-type) lists out the major areas like school, home, shopping mall, etc. Each of the categories in the first level has sub-class scenes associated with them. For e.g., a shopping mall would have a set of places like *gift-shop, food-court, salon, drug-store, etc.* A detailed list of places in our testing dataset can be found in table I and at our project page<sup>8</sup>. We have a total of 10,000 testing images in our dataset. Out of these approximately 1000 images were extracted from videos on youtube as some places were not easily available in the real world and others were captured from a GoPro camera.

#### V. EVALUATION

In our taxonomy, we described 8 different environment types (*School, Shopping Mall, Home, Condo Buildings, Airport, Public Transit, Hospital, and Hotel*) using 156 indoor scenes. In this work, we select 3 environments to evaluate our approach, which are *School, Home, and Shopping Mall* (see Table I).

In our experiments, we trained a ResNet model for each environment with the selected scenes. The performance is listed in Table II. As one would expect, the ResNet model performs well on the reduced Places365 validation set. The performance on our test set is even higher, since our dataset is labeled with fewer human errors. Furthermore, we used a GoPro camera with wide-angle view mode, so the images perfectly represent the scenes. Overall, our method is better than a CNN model (ResNet50) in terms of accuracy. Our results proved that the “*image context (objects)*” plays a vital role in Scene Classification. Although, the refinement on Shopping Mall scenes is not noticeable, this is because the CNN Scene Parsing model we used could not find objects that can distinguish the scenes well enough. For instance, the

<sup>8</sup><http://jtl.lassonde.yorku.ca/2018/04/scene-classification-robots/>

TABLE I

SAMPLE DATA CLASSES IN OUR DATASET. FOR COMPLETE TAXONOMY SEE PROJECT PAGE <sup>8</sup>

First Level Place (Environments)	Second Level Place (Scenes)
<b>School (24-scenes)</b>	Classroom, Kindergarten, Office, MeetingRoom, ComputerLab, ChemistryLab, BiologyLab, PhysicsLab, Library, Corridor, Elevator, Escalator, Cafeteria, Washroom, Auditorium, Gymnasium, LockerRoom, IndoorSwimmingPool, BasketballCourt, VolleyballCourt, BadmintonCourt, TableTennis, DormRoom, Lobby
<b>Home (14-scenes)</b>	Kitchen, WetBar, LivingRoom, DiningRoom, Bedroom, Closet, PlayRoom, HomeTheater, HomeOffice, LaundryRoom, Washroom, Garage, Staircase, Balcony
<b>Shopping Mall (31-scenes)</b>	BookStore, CandyStore, VideoStore, MusicStore, HardwareStore, ShoeShop, DrugStore, ToyStore, ClothingStore, HatShop, FloristStore, JewelryStore, Optician, Supermarket, Bakery, Salon, PetShop, GiftShop, Foodcourt, Bar, Restaurant, CoffeeShop, TeaShop, DepartmentStore, Reception, Fountain, Elevator, Escalator, Washroom, IndoorParking, MovieTheater

TABLE II

TOP-1 ACCURACY ON THE REDUCED VALIDATION SET ON PLACES365 AND OUR TEST SET.

	School (24)		Home (14)		Shopping Mall (31)	
	Places365 val.-set	our test-set	Places365 val.-set	our test-set	Places365 val.-set	our test-set
Places365-ResNet	51.45%	58.18%	57.13%	60.68%	50.17%	62.58%
ResNet50	73.82%	90.33%	83.46%	92.03%	<b>70.47%</b>	87.31%
ResNet50+Word2Vec (Ours)	<b>74.28%</b>	<b>92.25%</b>	<b>83.67%</b>	<b>93.27%</b>	70.44%	<b>87.39%</b>

model could not detect shoes, watches, hats, etc., which are commonly found in shopping malls.

Figure 4 discusses the results of some images in our dataset. Finally, we also demonstrate our approach by deploying our algorithm on a Pioneer3AT robot in a university environment. The robot was configured with a ZED camera from which monocular images were used by our algorithm. The robot was equipped with a Razer Blade laptop with a GTX 1060 GPU. The demo video of the deployment is available on our project web-page <sup>8</sup>.

## VI. CONCLUSION

In this paper, we introduced a new method for Scene Classification using a taxonomy for different indoor environments. Using our approach, robots can recognize different indoor places with high confidence/accuracy. A hierarchical taxonomy of places allowed us to prune out many irrelevant classes, thereby reducing the complexity/training time of our approach. A word embedding based approach was implemented to refine the top-5 scores for the scenes. It was shown that context has the potential to improve the Scene Classification results to some extent. We additionally tested our approach with a real world dataset that we built to show the practical applicability of our approach. We also deployed our algorithm on a robot in a university environment. On our dataset we could get promising results for doing Scene Classification for robots.

## REFERENCES

- [1] B. Zhou, A. Lapedriza, A. Khosla, A. Oliva, and A. Torralba, "Places: A 10 million image database for scene recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- [2] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba, "Sun database: Large-scale scene recognition from abbey to zoo," in *Computer vision and pattern recognition (CVPR), 2010 IEEE conference on*. IEEE, 2010, pp. 3485–3492.
- [3] D. Kumar, "Deep learning based place recognition for challenging environments," Master's thesis, University of Waterloo, 2016.
- [4] R. Sahdev and J. K. Tsotsos, "Indoor place recognition system for localization of mobile robots," in *Computer and Robot Vision (CRV), 2016 13th Conference on*. IEEE, 2016, pp. 53–60.
- [5] S. Lazebnik, C. Schmid, and J. Ponce, "Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories," in *Computer vision and pattern recognition, 2006 IEEE computer society conference on*, vol. 2. IEEE, 2006, pp. 2169–2178.
- [6] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," *arXiv preprint arXiv:1301.3781*, 2013.
- [7] L.-J. Li, H. Su, L. Fei-Fei, and E. P. Xing, "Object bank: A high-level image representation for scene classification & semantic feature sparsification," in *Advances in neural information processing systems*, 2010, pp. 1378–1386.
- [8] J. Yang, Y.-G. Jiang, A. G. Hauptmann, and C.-W. Ngo, "Evaluating bag-of-visual-words representations in scene classification," in *Proceedings of the international workshop on Workshop on multimedia information retrieval*. ACM, 2007, pp. 197–206.
- [9] A. Bosch, A. Zisserman, and X. Muñoz, "Scene classification using a hybrid generative/discriminative approach," *IEEE transactions on pattern analysis and machine intelligence*, vol. 30, no. 4, pp. 712–727, 2008.
- [10] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [11] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.
- [12] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 1–9.
- [13] S. H. Khan, M. Hayat, and F. Porikli, "Scene categorization with spectral features," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 5638–5648.
- [14] V. Badrinarayanan, A. Kendall, and R. Cipolla, "Segnet: A deep convolutional encoder-decoder architecture for image segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- [15] B. Zhou, H. Zhao, X. Puig, S. Fidler, A. Barriuso, and A. Torralba, "Scene parsing through ade20k dataset," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017.
- [16] J. Redmon and A. Farhadi, "Yolo9000: better, faster, stronger," *arXiv preprint arXiv:1612.08242*, 2016.
- [17] J. Pennington, R. Socher, and C. Manning, "Glove: Global vectors for word representation," in *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 2014, pp. 1532–1543.
- [18] Z. Huang, W. Xu, and K. Yu, "Bidirectional lstm-crf models for sequence tagging," *arXiv preprint arXiv:1508.01991*, 2015.
- [19] G. Lample, M. Ballesteros, S. Subramanian, K. Kawakami, and C. Dyer, "Neural architectures for named entity recognition," *arXiv preprint arXiv:1603.01360*, 2016.

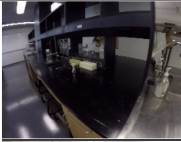






Input	ResNet50		Image contexts		Ours		Ground Truth
	Top-5	A	Objects	B	Refined Top-5	A × B	
1 	Physics Lab	0.16	wall, floor, ceiling,	0.68	Chemistry Lab	0.32	Chemistry Lab
	Chemistry Lab	0.15	windowpane, table,	0.91	Physics Lab	0.25	
	Gymnasium	0.14	sink, bottle, glass	0.63	Gymnasium	0.21	
	Office	0.11		0.53	Office	0.14	
	Computer Lab	0.08		0.41	Computer Lab	0.08	
2 	Computer Lab	0.33	wall, floor, ceiling,	0.64	Meeting Room	0.36	Meeting Room
	Meeting Room	0.29	door, table, chair,	0.93	Computer Lab	0.28	
	Office	0.16	painting, seat, light,	0.73	Classroom	0.19	
	Classroom	0.16	screen, crt screen	0.93	Office	0.15	
	Dorm Room	0.02		0.67	Dorm Room	0.02	
3 	Washroom	0.54	wall, floor,	0.19	Home Office	0.67	Home Office
	Home Office	0.34	windowpane, table,	0.96	Washroom	0.22	
	Closet	0.04	chair, sofa, shelf, desk,	0.45	Play Room	0.06	
	Play Room	0.03	lamp, cushion, book,	0.91	Closet	0.04	
	Staircase	0.02	computer	0.58	Staircase	0.02	
4 	Kitchen	0.54	wall, floor, table, chest	0.58	Laundry Room	0.55	Laundry Room
	Laundry Room	0.44	of drawers, washer	0.92	Kitchen	0.43	
	Home Office	0.00		0.83	Home Office	0.01	
	Closet	0.00		0.53	Washroom	0.00	
	Washroom	0.00		0.68	Closet	0.00	
5 	Physics Lab	0.29	wall, floor, ceiling,	0.85	Office	0.33	Physics Lab
	Office	0.29	chair, painting, shelf,	0.98	Computer Lab	0.29	
	Computer Lab	0.27	desk, box, computer,	0.93	Physics Lab	0.29	
	Biology Lab	0.05	swivel chair, light,	0.90	Biology Lab	0.06	
	Library	0.04	monitor, clock	0.69	Library	0.03	
6 	Laundry Room	0.91	wall, floor, cabinet,	0.61	Laundry Room	0.91	Kitchen
	Kitchen	0.05	table, sink, stove, bottle	0.96	Kitchen	0.08	
	Home Theater	0.01		0.60	Home Theater	0.01	
	Living Room	0.01		0.25	Bedroom	0.00	
	Bedroom	0.01		0.27	Living Room	0.00	
7 	Home Office	0.30	wall, floor, ceiling,	0.68	Washroom	0.38	Laundry Room
	Washroom	0.30	windowpane, cabinet,	0.90	Home Office	0.29	
	Kitchen	0.18	sink, toilet, light,	0.96	Kitchen	0.25	
	Bedroom	0.12	washer	0.19	Laundry Room	0.05	
	Laundry Room	0.03		0.95	Bedroom	0.03	

Fig. 4. Column A is the top-5 confidences obtained from the CNN Scene Classification module. Column B is the cosine similarity score from the Word Vectors module, Column AxB represents the refined top-5 predictions (refer to line 8 in Algorithm 1, Procedure 2). Sample results: (1)-(4) are correct examples, (5)-(7) are incorrect examples on our test set. The first 4 rows show that our model correctly predicted the scene labels. In the first example, CNN gave an unclear result by predicting the image as *physics lab*, *chemistry lab*, *gymnasium* with very similar scores. But, our word2vec model gave a very high score on *chemistry lab*. By combining the two scores, our model gave a higher predicted score on the ground truth label. On the other hand, our method also yielded some bad predictions. For row (5), our word2vec model provided a very high confidence on the false label *office* by using the context of the image. As a result, this image was incorrectly classifying as *office* instead of *physics lab*. Another example in row (6), the CNN model very confidently predicted a *kitchen* as *laundry room*. Even though word2vec said it is a *kitchen*, our model still could not refine the CNN predictions.

- [20] A. Agrawal, R. Sahdev, H. Davoudi, F. Khonsari, A. An, and S. McGrath, "Detecting the magnitude of events from news articles," in *Web Intelligence (WI), 2016 IEEE/WIC/ACM International Conference on*. IEEE, 2016, pp. 177–184.
- [21] M. Xu, H. Jiang, and S. Watcharawittayakul, "A local detection approach for named entity recognition and mention detection," in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, vol. 1, 2017, pp. 1237–1247.
- [22] Y. Hou, H. Zhang, and S. Zhou, "Bocnf: efficient image matching with bag of convnet features for scalable and robust visual place recognition," *Autonomous Robots*, pp. 1–17.
- [23] M. Shakeri and H. Zhang, "Illumination invariant representation of natural images for visual place recognition," in *Intelligent Robots and Systems (IROS), 2016 IEEE/RSJ International Conference on*. IEEE, 2016, pp. 466–472.
- [24] B. Zhou, A. Lapedriza, J. Xiao, A. Torralba, and A. Oliva, "Learning deep features for scene recognition using places database," in *Advances in neural information processing systems*, 2014, pp. 487–495.
- [25] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes (voc) challenge," *International journal of computer vision*, vol. 88, no. 2, pp. 303–338, 2010.
- [26] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *European conference on computer vision*. Springer, 2014, pp. 740–755.