

Scene recognition with CNNs: objects, scales and dataset bias

Luis Herranz, Shuqiang Jiang, Xiangyang Li

Key Laboratory of Intelligent Information Processing of Chinese Academy of Sciences (CAS)
Institute of Computer Technology, CAS, Beijing, 100190, China

{luis.herranz, shuqiang.jiang, xiangyang.li}@vip1.ict.ac.cn

Abstract

Since scenes are composed in part of objects, accurate recognition of scenes requires knowledge about both scenes and objects. In this paper we address two related problems: 1) scale induced dataset bias in multi-scale convolutional neural network (CNN) architectures, and 2) how to combine effectively scene-centric and object-centric knowledge (i.e. Places and ImageNet) in CNNs. An earlier attempt, Hybrid-CNN[23], showed that incorporating ImageNet did not help much. Here we propose an alternative method taking the scale into account, resulting in significant recognition gains. By analyzing the response of ImageNet-CNNs and Places-CNNs at different scales we find that both operate in different scale ranges, so using the same network for all the scales induces dataset bias resulting in limited performance. Thus, adapting the feature extractor to each particular scale (i.e. scale-specific CNNs) is crucial to improve recognition, since the objects in the scenes have their specific range of scales. Experimental results show that the recognition accuracy highly depends on the scale, and that simple yet carefully chosen multi-scale combinations of ImageNet-CNNs and Places-CNNs, can push the state-of-the-art recognition accuracy in SUN397 up to 66.26% (and even 70.17% with deeper architectures, comparable to human performance).

1. Introduction

State-of-the-art in visual recognition is based on the successful combination of deep representations and massive datasets. Deep convolutional neural networks (CNNs) trained on ImageNet (i.e. ImageNet-CNNs) achieve impressive performance in object recognition, while CNNs trained on Places (Places-CNNs) do in scene recognition[23, 2]. However, CNNs also have limitations, such as the lack of invariance to significant scaling. This problem is particularly important in scene recognition, due to a wider range of scales and a larger amount of objects per image.

As an alternative to Places-CNN holistic representation,

some recent works[4, 1, 21, 16] have shown that CNN features extracted locally in patches can be also aggregated into effective scene representations. Often, these approaches combine multiple scales, that are pooled using VLAD[4] or Fisher vector (FV)[21] encoding. Dixit et al[1] suggested applying the pooling directly on the semantic representation, arguing that semantic representations are more invariant. Recently, Wu et al[16] proposed an architecture in which dense sampling of patches is replaced by region proposals and discriminative patch mining. In general, these works use ImageNet-CNN to extract the local activations instead of Places-CNN, since local patches are closer to objects than to scenes. However, a largely overlooked aspect in this multi-scale scenario is the role of the scale and its relation with the feature extractor (i.e. CNN). One limitation of current multi-scale approaches is the naive use of CNNs by simply considering CNNs as general purpose feature extractors[10, 2]. Using the same fixed CNN model for all the scales inevitably leads to dataset bias[13], since the properties of the data vary at different scales, while the feature extractor remains fixed.

Since objects are main components of scenes, knowledge about objects may be helpful in scene recognition. Although, Places-CNN itself develops suitable object models at intermediate layers[22], the information in ImageNet might be valuable. However, in a previous attempt, a network trained with the combined dataset ImageNet+Places (Hybrid-CNN[23]) show that including ImageNet, far from helpful was harmful. We will see how this problem is also connected to scale-related dataset bias.

In this paper we will study these two problems (i.e. dataset bias in patch-based CNNs under different scaling conditions, and how to effectively combine Places and ImageNet) and will see that they are related. Torralba and Efros[13] studied the dataset bias as a cross-dataset generalization problem, in which the same classes may have slightly different feature distributions in different datasets. In our particular case, this bias in the feature distribution is induced by scaling the image. If the scaling operation is considerable, the characteristics of the data may change

completely, switching from scene data to object data. Understanding and quantifying this bias can help us to design better multi-scale architectures, and even better ways to combine object and scene knowledge. In particular, we propose multi-scale architectures with *scale-specific networks* as a principled way to address scale-related dataset bias and combine scene and object knowledge (i.e. Places and ImageNet). In the next sections:

- We show that using a single CNN as a generic feature extractor from patches is quite limited, due to the dataset bias induced by scale changes. We show how networks trained on different datasets are suitable for different scale ranges. In particular, ImageNet-CNN and Places-CNN have very different optimal ranges, due to their object-centric and scene-centric natures.
- We evaluate two strategies to alleviate the dataset bias by using scale-specific networks: hybrid Places/ImageNet architectures and fine tuning. By combining after reducing the dataset bias, our method is also a more effective way to combine Places and ImageNet. Extensive experiments with different scale combinations and hybrid variations (optionally fine tuned) lead to some variations achieving state-of-the-art performance in scene recognition.

2. Objects and scenes

2.1. Objects in object datasets and scene datasets

The knowledge learned by CNNs lies in the data seen during training, and will be of limited use if tested in a different type of data. Thus, CNNs trained with ImageNet are limited when used for scene recognition due to this training/test bias, while Places-CNNs perform better in this task. While this is essentially true, objects and scenes are closely related, so knowledge about objects can be still helpful to recognize scenes, if used properly.

Understanding the characteristics of the datasets involved is essential to better explain the causes of dataset bias. In our case, we want to analyze the properties of objects found in scene and object datasets. We focus on two aspects related with the objects: *scales* and *density*.

To evaluate the dataset bias we use SUN397[18, 17] as target dataset. Since Places contains scene data, with 205 scene categories overlapping with SUN397, and significantly more data, we can expect a low dataset bias. Thus we focus on ImageNet (in particular ILSVRC2012), which contains mostly object data. Fortunately, both ImageNet and SUN have a fraction of images with region annotations and labels, so we can collect some relevant statistics and compare their distributions (we used the LabelMe toolbox[14]). Since we will use this information to interpret the variations in recognition accuracy in next experiments,

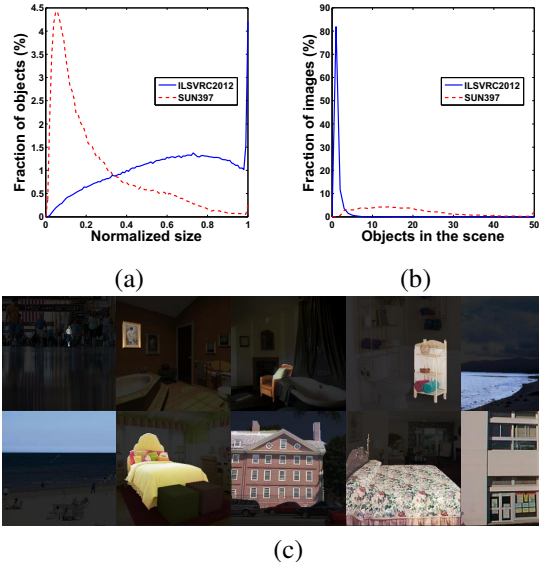


Figure 1. Characteristics of objects in ILSVRC2012 (object data) and SUN397 (scene data): (a) distribution of objects sizes (normalized), (b) number of objects per scene, and (c) examples of objects by increasing normalized size.

we focus on the 397 categories of the SUN397 benchmark (rather than the 908 categories of the full SUN database).

Scale. Fig. 1a shows the distribution of object sizes, and Fig. 1c some examples of objects of different normalized sizes. We normalized the size of the object relative to the equivalent training crop. While objects in ImageNet are mostly large, often covering the whole image, objects in SUN397 are much smaller, corresponding to the real distribution in scenes. Thus Fig. 1a shows an obvious mismatch between both datasets.

Density. Fig. 1b shows the distribution of object annotations per scene image. We can observe that images in ImageNet usually contain just one big object, while images in SUN397 typically contain many small objects.

2.2. Dataset bias in object recognition

In order to study the behaviour of ImageNet-CNNs and Places-CNNs in object recognition, we need object data extracted from scenes datasets. We selected 100 images per category from the 75 most frequent object categories in SUN397, so we can have enough images to train SVM classifiers. We took some precautions to avoid selecting too small objects.

In contrast to most object and scene datasets, in this case we have the segmentation of the object within the scene, so we can use it to create variations over the same objects. Thus we defined two scales:

- Original scale: the scale of the object in original scene.

- Canonical scale: the object is centered and rescaled to fill the crop (keeping the aspect ratio). So in this case its normalized size is 1.

Then we created four variations (see Fig. 2): *original masked*, *original with background*, *canonical masked* and *canonical with background*. In particular, to study the response to different scaling, the canonical variant is scaled in the range 10%-100%. Note how scaling the variant with background shifts progressively the content of the crop from object to scene.

2.3. Scale sensitivity and object density

We trained a SVM classifier with 50 images per class, and tested on the remaining 50 images. The input feature was the output of the fc7 activation. The results are shown in Fig. 3. We use two variants: *objects masked* and *objects with background* (see Fig. 2). Regarding *objects masked*, where the background is removed, we can see that in general the performance is optimal when the object is near full size, above 70-80%. This is actually the most interesting region, with ImageNet-CNN performing slightly better than Places-CNN. This is interesting, since Places-CNN was trained with scenes containing more similar objects to the ones in the test set, while ImageNet-CNN was trained with the less related categories found in ILSVRC2012 (e.g. dogs, cats). However, as we saw in Fig. 1a, objects in ILSVRC2012 cover a large portion of the image in contrast to smaller objects in SUN397, suggesting that a more similar scale in the training data may be more important than more similar object categories. As the object becomes smaller, the performance of both models degrades similarly, again showing a limited robustness to scale changes.

Focusing now on the *objects with background* variant, the performance is worse than when the object is isolated from the background. This behaviour suggests that the background may introduce some noise in the feature and lead to poorer performance. In the range close to full object size, both ImageNet-CNN and Places-CNN have similar performance. However, as the object becomes smaller, and the content is more similar to scenes, Places-CNN has much better performance than ImageNet-CNN, arguably due to the fact it has learned contextual relations between objects and global scene properties. In any case, scales with low accuracy are probably too noisy and not suitable for our purpose.

3. Multi-scale architecture with scale-specific networks

3.1. Overview

For scene recognition we introduce our multi-scale architecture, which combines several networks that operate in parallel over patches extracted from increasingly larger

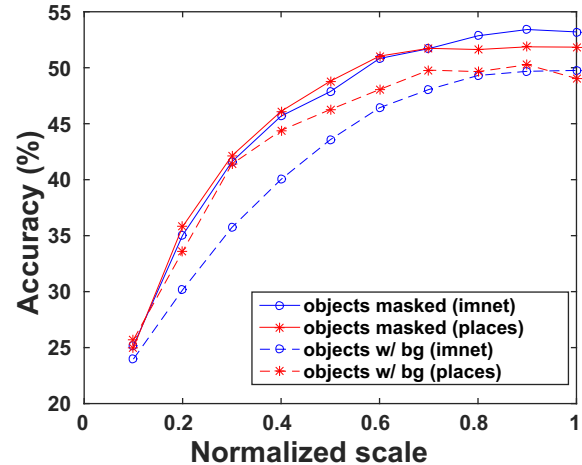


Figure 3. Object recognition accuracy on SUN397 (75 categories).

versions of the input image. We use a standard multi-scale architecture combining several AlexNet CNNs (CaffeNet[5] in practice) where 227x227 patches are extracted from each full image. For faster processing, instead of extracting patches independently we use a fully convolutional network. In contrast to recent works[4, 3, 21, 1], we adopt simple max pooling to aggregate patch features into image features.

The previous analysis and experimental results on object recognition evidence the limitations of using either ImageNet-CNNs or Places-CNNs to deal with such a broad range of scales, and will be confirmed in the next sections by the experiments on scene data. For these reasons, we propose a hybrid architecture introducing two simple, yet crucial modifications in the architecture (discussed previously in Section 2.2).

- Instead of using naively the same CNN model for all the scales, we select the most suitable one for each (ImageNet-CNN, Places-CNN or fine tuned).
- Optionally we fine tune each CNN model to further adapt it to the range of each scale. This requires resizing the image to target size and extracting patches for training.

3.2. Differences with previous works

Our architecture is similar to others proposed in previous multi-scale approaches[4, 21, 1], with the subtle difference of using scale-specific networks in a principled way to alleviate the dataset bias induced by scaling. The main emphasis in these works is on the way multi-scale features are combined, implemented as either VLAD or FV encoding, while leaving the CNN model fixed. While adding a BOW encoding layer can help to alleviate somewhat the dataset bias, the main problem is still the rigid CNN model. In contrast, our method addresses better the dataset bias related




Figure 2. The two variants used in the object recognition experiments: object masked (top row) and object with background (bottom row) with two examples of (a) *armchair* and (b) *streetlight*. Left crops show the object in the original scale in the scene. Right crops show the object scaled progressively from the canonical size (100%) down to 10%. All the images are centered in the object of interest.

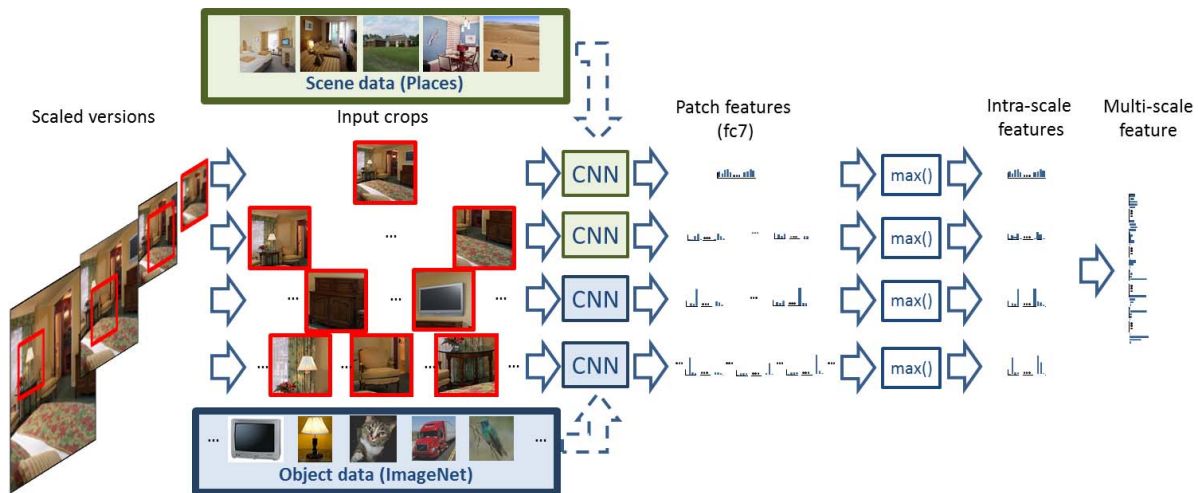


Figure 4. Multi-scale architecture combining scale-specific networks (spliced architecture). ImageNet-CNNs and Places-CNNs are combined according to the scale of the input patches. This can effectively alleviate the dataset bias by adapting test data to the underlying training data. Intra-scale features are obtained using max pooling within each scale, and then concatenated into a single multi-scale feature.

with scale and achieves significantly better performance, by simply adapting the CNN model to the target scale, even without relying to sophisticated pooling methods.

We can also regard our approach as a way to combine the training data available in Places and ImageNet. This was explored previously by Zhou et al[23], who trained a Hybrid-CNN using the AlexNet architecture and the combined Places+ImageNet dataset. However, Hybrid-CNN performs just slightly better than Places-CNN on MIT Indoor 67 and worse on SUN397. We believe that the main reason was that this way of combining data from ImageNet and Places ignores the fact that objects found in both datasets in two different scale ranges (as shown in Fig. 1). In contrast, our architecture combines the knowledge in a scale-adaptive way via either ImageNet-CNN or Places-CNN. Wu et al[16] use Hybrid-CNN on patches at different scales. Again, the main

limitation is that the CNN model is fixed, not adapting to the scale-dependent distributions of patches.

4. Experiments on scene recognition

In this section we perform experiments directly over scene data, to evaluate the relation between scale, training dataset and dataset bias by analyzing the scene recognition performance. Then we combine and evaluate multi-scale architectures.

4.1. Datasets

We evaluate the proposed architectures with three widely used scene benchmarks. 15 scenes[6] is a small yet popular dataset with 15 natural and indoor categories. Models are trained with 100 images per category. MIT Indoor 67[9] contains 67 categories of indoor images, with 80 images

per category available for training. Indoor scenes tend to be rich in objects, which in general makes the task more challenging, but also more amenable to architectures using ImageNet-CNNs on patches. SUN397[18, 19] is a larger scene benchmark (at least considered as such before Places) containing 397 categories, including indoor, man-made and natural categories. This dataset is very challenging, not only because of the large number of categories, but also because of the more limited amount of training data (50 images per category) and a much larger variability in objects and layout properties. It is widely accepted as the reference benchmark for scene recognition. We consider seven scales in our experiments, obtained by scaling images between 227x227 and 1827x1827 pixels.

4.2. Single scale

4.2.1 Accuracy

Average accuracy is a reasonable metric to evaluate a deep representation in the context of a classification task. For the different scales, we extracted *fc7* activations locally in patches as features, and then trained SVMs. In addition to the seven scales evaluated, we included 256x256 pixels as a baseline, since off-the-shelf ImageNet-CNN and Places-CNN are trained on this scale. The results for the three datasets are shown in Fig. 5, with similar patterns. Places-CNN achieves the best performance when is applied globally at scene level (227x227 or 256x256), while rapidly degrades for more local scales. ImageNet-CNN exhibits a very different behaviour, with a more modest performance at global scales, and achieving optimal performance on patches at intermediate scales, and outperforming Places-CNN at most local scales. These curves somewhat represent the operational curve of CNNs and the scale. In particular, the performance of ImageNet-CNN can be increases notably just by using an appropriate scale.

An interesting observation is that there is one point (around 643 or scale 0.35) that splits the range into two parts, one dominated by ImageNet-CNN and another one dominated by Places-CNN, which we can loosely identify as object range and scene range. We will use this observation later in Section 4.4 to design spliced architectures.

4.2.2 Effect of fine tuning

A popular way to deal with dataset bias in CNNs is fine tuning, which basically continues training on a pretrained model with the target dataset data. Similarly in our case, we expect that fine tuning can modify somehow the weights and thus adapt to the objects or at least the scales in the target dataset. However, in practice that is often not possible because of the limited data, overfitting and difficulty of setting the training process itself. In our case, fine tuning on scales where patches are very local is very difficult

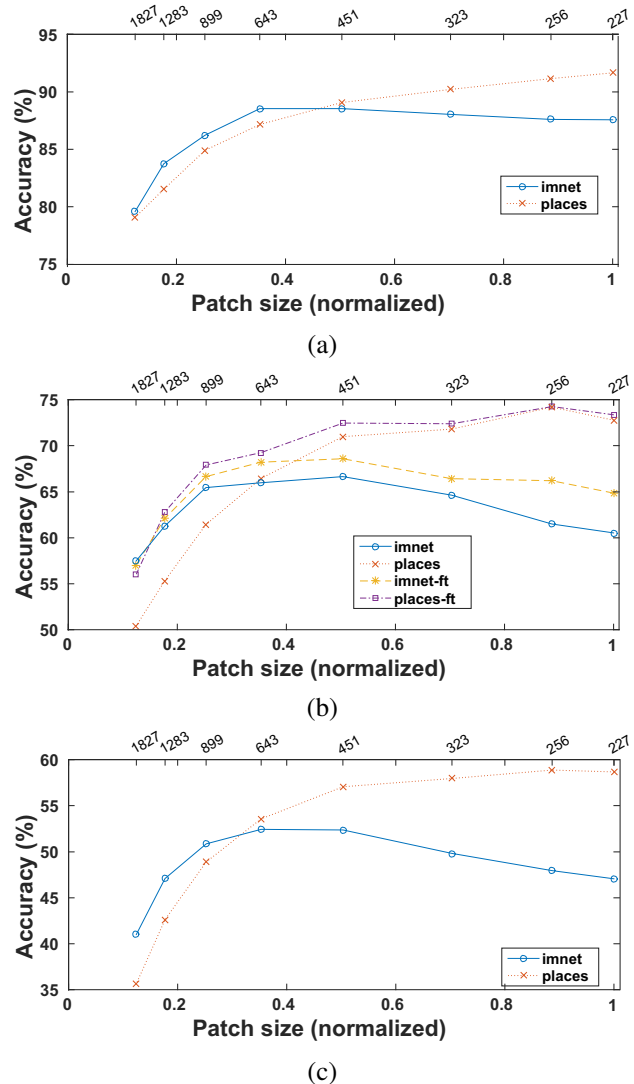


Figure 5. Scene recognition accuracy for different scales: (a) 15 scenes, (b) MIT Indoor 67, and (c) SUN397.

due since the patch often contains objects or parts while labels indicates scenes categories. In addition, the number of patches is huge, so only a tiny fraction of them can be used in practice, rendering fine tuning not very effective.

We evaluated fine tuning on MIT Indoor 67. For scales with few patches, and thus limited training data, we only fine tune the fully connected layers. For larger images we can collect more patches, up to 500K patches (randomly selected). Fig. 5b shows the results. Interestingly, there is a moderate gain in those range of scales where the original CNNs perform poorly, i.e. global scales for ImageNet-CNN and local scales for Places-CNN, while marginal or no gain in ranges where they have already good performance. Thus, fine tuning has certain “equalizing” effect over the accuracy vs scale curve. but limited overall improvement. In particular the gain is such that now Places-CNN (fine tuned) has the best performance in the whole range of scales.

Fine tuning has impact mostly on the top layers, obtaining a similar effect to adding a BOW pooling layer. However, the effectiveness is limited, since intermediate layers remain biased to the (pre-)training data.

4.2.3 Discriminability and redundancy

Accuracy provides a good indirect measure of the utility of the feature for a given target task (e.g. scene recognition) via a classifier (e.g. SVM). Here we also consider two information theoretic metrics measuring directly the discriminability and redundancy of the deep feature[8]. We define the *discriminability* of a feature $\mathbf{x} = (x_1, \dots, x_{4096})$ with respect to a set of classes $C = \{1, \dots, M\}$

$$D(\mathbf{x}, C) = \frac{1}{|C||S|} \sum_{c \in C} \sum_{x_i \in \mathbf{x}} I(x_i; c)$$

where $I(x_i; c)$ is the filter x_i and the class c . In order to evaluate how redundant is the feature (compared with other filters), we use the *redundancy* of a feature \mathbf{x} , defined as

$$R(\mathbf{x}) = \frac{1}{|S|^2} \sum_{x_j \in \mathbf{x}} \sum_{x_i \in \mathbf{x}} I(x_i; x_j)$$

In the next experiment we compute $D(\mathbf{x}, C)$ and $R(\mathbf{x})$ of the fc7 activation for ImageNet-CNN and Places-CNN in MIT Indoor 67. While we can find similarities with the accuracy curve, a direct comparison is not easy, since more discriminability not always means higher accuracy. If we observe the discriminability of ImageNet-CNN (see Fig. 6a), the curve follows a similar pattern to the accuracy, with a peak around the scales where the accuracy was best, and bad discriminability at global scales. Places-CNN extracts the most discriminative features at more global scales. Comparing ImageNet-CNN and Places-CNN, the former obtains more discriminative features yet also more redundant. Too local scales (e.g. 1827x1827) increase significantly the redundancy of the feature and the noise

4.3. Two scales

In the next experiment we evaluated pairwise combinations of CNNs used at different scales. This dual architecture consists simply of two CNNs processing images at different scales. We then concatenate the two resulting fc7 activations into a 8192-dim feature and then train the SVM. The results in Fig. 7 show that the dual architectures with best performance are hybrid combinations of Places-CNN extracting features at global scales (typically 227x227) with ImageNet-CNN extracting features from patches at more local scales. The result is a considerable boost in the performance, achieving a remarkable accuracy of 64.10% on SUN397 using only two AlexNet CNNs.

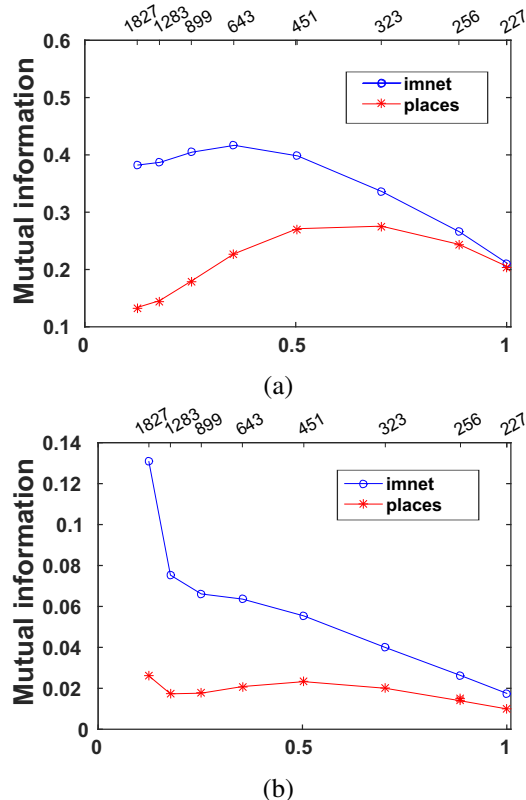


Figure 6. Discriminability (a) and redundancy (b) of fc7 feature in MIT Indoor 67.

Note that this way of combining ImageNet and Places data is much more effective than Hybrid-CNN[23] (see Table 1). Our dual architecture does not mix object and scene knowledge (obtained from ImageNet and Places, respectively) and adapts the learned models to scales with similar properties. Dixit et al[1] combine Places-CNN with a four-scales architecture built on top of ImageNet-CNN. Similarly to our framework, Places-CNN operates at scene scale while ImageNet-CNN operates at object scales. Note, however, that we obtain comparable performance on MIT Indoor 67 and significantly better on SUN397, using just two networks instead of five.

4.4. Multiple scales

Finally, we consider the combination of all the scales to see whether more complex architectures could be helpful in terms of accuracy. In this experiment we evaluate the concatenation of all the fc7 features of each of the seven scale-specific networks. In this case we use PCA to reduce the dimension of each features vector so the combined dimension is approximately 4096. We achieve 74.33% (all scales using ImageNet-CNN) and 78.21% (all scales using Places-CNN) accuracy for MIT Indoor 67, and 58.71% and 63.81% for SUN397, respectively. Note that both are better than the corresponding dual architecture, yet below the corresponding dual hybrids (78.28% and 64.10%). This suggests than

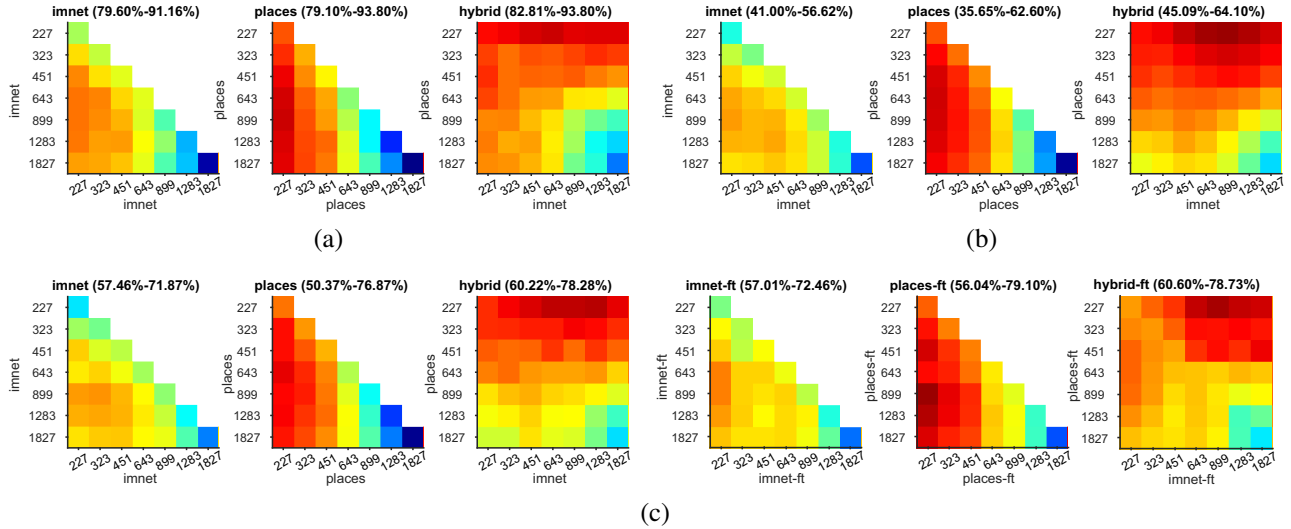


Figure 7. Accuracy in dual architectures combining two networks (only ImageNet-CNNs, only ImageNet-CNNs and hybrid combinations): (a) 15 scenes, (b) SUN397 and, (c) MIT Indoor 67 (and fine tuned versions). Diagonals of only ImageNet-CNNs and only Places-CNNs variations show single scale accuracy.

including more scales while keeping the same CNN model is marginally helpful and increases significantly the extraction cost and the noise in the representation.

So the key is to find an appropriate combination of Places-CNNs and ImageNet-CNNs. While in dual architectures evaluating all the combinations is very costly, with seven networks the combinations is impractical. Since the optimal ranges of both are complementary, we can design the full hybrid architecture as global scales using Places-CNN and local scales using ImageNet-CNN, just as shown in Fig. 4. We can consider only one free parameter which is the splicing point. The results for SUN397 are shown in Fig. 8. As expected, we can reach slightly better performance (80.97% and 65.38%) than in dual architectures. The performance of hybrid spliced is also significantly better than a 7 network architecture with a fixed CNN model.

Finally we also evaluate a double full architecture, in which both full ImageNet-CNN and full Places-CNN are combined in a complex 14 CNNs architecture by concatenating the previous features. This combination does not help in MIT Indoor 67, and slightly in SUN397, reaching an accuracy of 66.26%.

4.5. Deeper networks and other works

The experiments presented so far are based on the AlexNet architecture. Deeper architectures such as GoogLeNet[12] and VGG-net[11] have demonstrated superior performance by exploiting deeper models. We repeated some of the experiments using the 16 layer VGG architecture, obtaining state-of-the-art results in the three datasets. The experiments with VGG in dual architectures

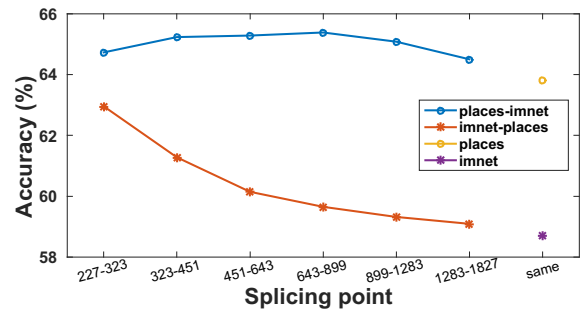


Figure 8. Accuracy on SUN397 of full and hybrid spliced architectures (7 AlexNet networks). The combination *same* indicates that the 7 networks share the same CNN model (i.e. trained with the same dataset).

are consistent with those with AlexNet, but with a more moderate gain. However, experiments combining more networks were surprisingly disappointing, performing even worse than single network baselines. VGG applied on small patches tends to be very noisy with poor performance. We tried an intermediate hybrid architecture, including a total of three scales, achieving slightly better performance than with dual architectures.

Overall, for the small 15 scenes dataset, it seems that the performance is somewhat saturated, with a best performance of 95.18% (94.51% with AlexNet). The best performance in MIT Indoor 67 is 86.04% (compared with 80.97% with AlexNet) and in SUN397 is 70.17% (compared with 66.26% with AlexNet). This performance is better than human recognition by “good workers” (68.5%), and close to human expert performance (70.6%), as reported in [18].

Table 1. Accuracy for different multi-scale variations and architectures.

Architecture	Pretraining dataset	#scales	15 scenes		MIT Indoor 67 (w/ FT)		SUN 397	
			Alex	VGG	Alex	VGG	Alex	VGG ³
Baseline ¹	IN	1	87.60	90.69	61.49	72.31	47.93	55.19
	PL	1	91.16	92.90	74.18	80.45	58.87	66.50
Best single ²	IN	1	88.54	91.86	66.64 (68.21)	76.42	52.42	59.71
	PL	1	91.65	93.73	72.76 (73.35)	80.90	58.88	66.23
Dual	IN	2	91.16	93.84	71.87 (72.46)	79.04	56.62	61.07
	PL	2	93.80	95.18	76.87 (79.40)	83.43	62.60	68.49
Dual hybrid	IN/PL	1+1	93.80	95.18	78.28 (78.81)	85.59	64.10	69.20
Three ⁴	IN/PL	1+2	93.37	95.14	78.28	86.04	63.03	70.17
Full	IN	7	91.66	92.86	74.33 (75.97)	70.22	58.71	55.18
	PL	7	93.77	94.51	78.21 (79.70)	77.81	63.81	58.80
Full hybrid (spliced)	IN/PL	7	93.90	94.08	80.97 (80.75)	80.22	65.38	63.19
Double full hybrid	IN/PL	2x7	94.51	94.84	80.97 (79.85)	80.7	66.26	62.01
Hybrid-CNN[23]	IN+PL	1	53.86	-	70.80	-	53.86	-
MOP-CNN[4]	IN	3	-	-	68.88	-	51.98	-
MPP[21]	IN	7	-	-	75.67	-	-	-
MPP+DSFL[21]	IN	7+DSFL	-	-	80.78	-	-	-
SFV[1]	IN	4	-	-	72.86	-	54.4	-
SFV+Places[1]	IN/PL	4+1	-	-	79.0	-	61.72	-
MetaObject-CNN[16]	Hybrid (IN+PL)[23]	1 (variable)	-	-	78.90	-	58.11	-
DAG-CNN[20]	IN	1	-	92.9	-	77.5	-	56.2
DSP[3, 15]	IN	1	-	91.78	-	78.28	-	59.78
Human (good)[17]				68.5%				
Human (expert)[17]				70.6%				

¹ 256x256 central crop (conventional settings for single crop).

² Excluding 256x256.

³ Six scales (1827x1827 was not included).

⁴ Only evaluated the combination Places-CNN 227x227, Places-CNN 451x451, ImageNet-CNN 899x899.

5. Conclusions

In contrast to previous works, in this paper we analyzed multi-scale CNN architectures focusing on the local CNN model, rather than on the pooling method. In particular, we showed that scaling images induces a bias between training and test data, which has a significant impact on the recognition performance. We also showed how ImageNet-CNN and Places-CNN in this context are implicitly tuned for different scale ranges (object scales and scene scales). Based on these findings, we suggest that addressing this bias is critical to improve scene recognition, and propose including scale-specific networks in the multi-scale architecture. The proposed method is also a more principled way to combine scene-centric knowledge (Places) and object-centric knowledge (ImageNet) than previous attempts (e.g. Hybrid-CNN).

In fact, recent scene recognition approaches fall into two apparently opposite directions: global holistic recognition (Places-CNN) versus local object recognition and pooling (multi-scale CNNs). In this paper we describe them as two

particular cases in a more general view of how multi-scale features can be combined for scene recognition. They are not incompatible, and actually when combined properly to reduce the dataset bias the results can be excellent, even reaching human recognition performance simply with just two or three networks carefully chosen. Our hybrid parallel architecture also suggests some similarities with perceptual and cognitive models, where object recognition and global scene features follow two distinct yet complementary neural pathways which are later integrated to accurately recognize the visual scene[7].

Acknowledgements. This work was supported in part by the National Basic Research Program of China (973 Program): 2012CB316400, the National Natural Science Foundation of China: 61532018, 61322212, 51550110505 and 61550110505, the National Hi-Tech Development Program (863 Program) of China: 2014AA015202, and the CAS President’s International Fellowship Initiative: 2011Y1GB05. This work is also funded by Lenovo Outstanding Young Scientists Program (LOYS).

References

- [1] M. Dixit, S. Chen, D. Gao, N. Rasiwasia, and N. Vasconcelos. Scene classification with semantic fisher vectors. In *Computer Vision and Pattern Recognition (CVPR)*, June 2015. 1, 3, 6, 8
- [2] 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 *International Conference on Machine Learning (ICML)*, 2014. 1
- [3] B. Gao, X. Wei, J. Wu, and W. Lin. Deep spatial pyramid: The devil is once again in the details. *arXiv*, 1504.05277, 2015. 3, 8
- [4] Y. Gong, L. Wang, R. Guo, and S. Lazebnik. Multi-scale orderless pooling of deep convolutional activation features. In *European Conference on Computer Vision (ECCV)*, pages 392–407, 2014. 1, 3, 8
- [5] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *ACM International Conference on Multimedia*, pages 675–678. ACM, 2014. 3
- [6] 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 (CVPR)*, 2006. 4
- [7] A. Oliva. Scene perception. In J. S. Werner and L. M. Chalupa, editors, *The New Visual Neurosciences*, pages 725–732. The MIT Press, 2014. 8
- [8] H. Peng, F. Long, and C. Ding. Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(8):1226–1238, Aug 2005. 6
- [9] A. Quattoni and A. Torralba. Recognizing indoor scenes. In *Computer Vision and Pattern Recognition (CVPR)*, pages 413–420, 2009. 4
- [10] A. S. Razavian, H. Azizpour, J. Sullivan, and S. Carlsson. CNN features off-the-shelf: An astounding baseline for recognition. In *Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 512–519, June 2014. 1
- [11] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations (ICLR)*, 2015. 7
- [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 *Computer Vision and Pattern Recognition (CVPR)*, pages 1–9, June 2015. 7
- [13] A. Torralba and A. A. Efros. Unbiased look at dataset bias. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1521–1528, June 2011. 1
- [14] A. Torralba, B. C. Russell, and J. Yuen. Labelme: Online image annotation and applications. *Proceedings of the IEEE*, 98(8):1467–1484, Aug. 2010. 2
- [15] X. S. Wei, B. B. Gao, and J. Wu. Deep spatial pyramid ensemble for cultural event recognition. In *International Conference on Computer Vision (ICCV) Workshops*, pages 280–286, Dec. 2015. 8
- [16] R. Wu, B. Wang, W. Wang, and Y. Yu. Harvesting discriminative meta objects with deep CNN features for scene classification. In *International Conference on Computer Vision (ICCV)*, 2015. 1, 4, 8
- [17] J. Xiao, K. Ehinger, J. Hays, A. Torralba, and A. Oliva. Sun database: Exploring a large collection of scene categories. *International Journal of Computer Vision*, pages 1–20, 2014. 2, 8
- [18] J. Xiao, J. Hayes, K. Ehringer, A. Olivia, and A. Torralba. SUN database: Largescale scene recognition from abbey to zoo. In *Computer Vision and Pattern Recognition (CVPR)*, 2010. 2, 5, 7
- [19] Y. Xiao, J. Wu, and J. Yuan. mCENTRIST: A multi-channel feature generation mechanism for scene categorization. *IEEE Transactions on Image Processing*, 23(2):823–836, Feb 2014. 5
- [20] S. Yang and D. Ramanan. Multi-scale recognition with DAG-CNNs. In *International Conference on Computer Vision (ICCV)*, pages 1215–1223, Dec. 2015. 8
- [21] D. Yoo, S. Park, J.-Y. Lee, and I. So Kweon. Multi-scale pyramid pooling for deep convolutional representation. In *Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2015. 1, 3, 8
- [22] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. Object detectors emerge in deep scene cnns. In *International Conference on Learning Representations (ICLR)*, 2015. 1
- [23] B. Zhou, A. Lapedriza, J. Xiao, A. Torralba, and A. Oliva. Learning deep features for scene recognition using places database. In *Neural Information Processing Systems (NIPS)*, pages 487–495, 2014. 1, 4, 6, 8