

Selective Feature Connection Mechanism: Concatenating Multi-layer CNN Features with a Feature Selector

Chen Du^{1,2}, Chunheng Wang¹, Yanna Wang¹, Cunzhao Shi¹, Baihua Xiao¹

¹Institute of Automation, Chinese Academy of Sciences(CASIA)

²University of Chinese Academy of Sciences(UCAS)

{duchen2016, chunheng.wang, wangyanna2013, cunzhao.shi, baihua.xiao}@ia.ac.cn

Abstract

Different layers of deep convolutional neural networks(CNNs) can encode different-level information. High-layer features always contain more semantic information, and low-layer features contain more detail information. However, low-layer features suffer from the background clutter and semantic ambiguity. During visual recognition, the feature combination of the low-layer and high-level features plays an important role in context modulation. If directly combining the high-layer and low-layer features, the background clutter and semantic ambiguity may be caused due to the introduction of detailed information. In this paper, we propose a general network architecture to concatenate CNN features of different layers in a simple and effective way, called Selective Feature Connection Mechanism (SFCM). Low-level features are selectively linked to high-level features with a feature selector which is generated by high-level features. The proposed connection mechanism can effectively overcome the above-mentioned drawbacks. We demonstrate the effectiveness, superiority, and universal applicability of this method on multiple challenging computer vision tasks, including image classification, scene text detection, and image-to-image translation.

1. Introduction

Deep Convolutional Neural Networks (CNNs) have achieved great success on a variety computer vision tasks, such as image classification [18], semantic segmentation [23, 9], and object detection [7, 25, 21, 24, 5]. In order to understand why the CNNs frameworks perform so well, Zeiler *et al.*[35] introduce a novel visualization technique that gives insight into the function of intermediate feature layers and the operation of the classifier. Actually, feature maps of different layers extract information of different levels from input image. The low-layer features tend to response the patches with similar simple patterns and with

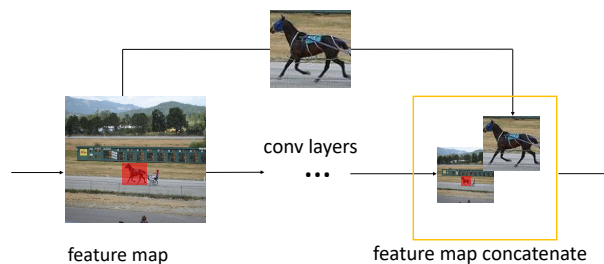


Figure 1. During visual recognition, the feature combination of the low-layer and high-level features plays an important role in context modulation.

more ambiguity. The feature maps of higher layers care more about semantic information but less detail information about an image, since higher layers are closer to the last layer with category labels. Combining high-level information with low-level information effectively can improve the performance of CNNs in many computer vision tasks. The U-Net [26], which combines the low-layer features and high-level features directly, works with very few training images and yields more precise segmentation. The EAST [36] adopt the idea from U-shape, merging feature maps gradually, to detect word regions of which the sizes vary tremendously. For many image translation problems, there is a great deal of low-level information shared between the input and output, and it would be desirable to shuttle this information directly across the net. Pix2pix [12] is proposed to add skip connections for the generator to circumvent the information bottleneck. DenseNets [11] improve the flow of information and gradients throughout the network, which can also be seen as combining the low-layer features and high-level features. The FPN(Feature Pyramid Network) [20] develops a top-down architecture with lateral connections to build high-level semantic feature maps at all scales. This architecture shows significant improvement as a generic feature extractor in several applications.

As shown in Fig.1, during visual recognition, the feature combination of the low-layer and high-level features plays

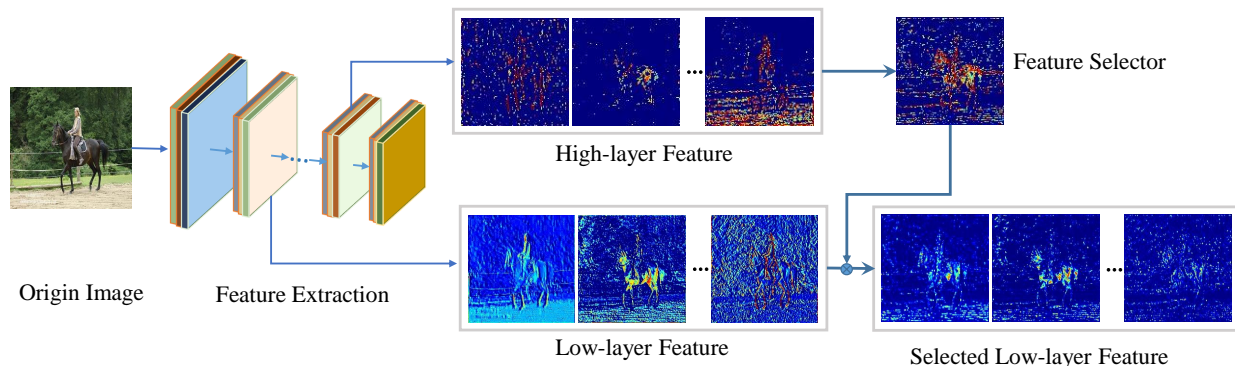


Figure 2. Different layers of deep convolutional neural networks(CNN) can encode different-level information. High-layer features always contain more semantic information, and low-layer features contain more detail information. In this paper, low-layer features are selected for further use by a feature selector which is generated by high-level features.

an important role in context modulation. When observing a horse in a big image, we will pay more attention to the area that represents the detail of the horse. And in the human visual recognition system, high-level features play an important role to generate an attention map for selectively combining the low-layer features. But all these features combining methods [26, 11, 20] just combine the low-layer features and high-level features directly, which is not in accordance with the biological fact.

In this paper, we propose a general network architecture to concatenate CNN features of different layers in a simple and effective way, called **Selective Feature Connection Mechanism (SFCM)**. Inspired by the human visual recognition mechanism, low-level features are selectively linked to high-level features with a feature selector which is generated by high-level features. Fig.2 shows an example of low-layer features selected for further use by a feature selector which is generated by high-level features. We demonstrate the effectiveness, superiority, and universal applicability of this method on many challenging computer vision tasks, such as image classification, scene text detection, and image-to-image translation.

Our main contributions are as follows:

- We propose the SFCM, a general network architecture to concatenate CNN features of different layers in a simple and effective way. Different layers of deep ConvNet can encode different-level information. High-layer features always contain more semantic information, and low-layer features contain more detail information. However, low-layer features suffer from the background clutter and semantic ambiguity. The SFCM can be used in many existing frameworks to achieve better performance as it can combine high-resolution maps with low-level features without harm the semantic representation capacity of high-layer fea-

tures.

- We give two connection modes (direct connection and residual connection) to implement our SFCM. Both modes are effective and light weights. They add a small number of parameters and computation for the feature selector learning. They can be embedded into any existing deep CNNs based methods and the feature selector can be easily trained end-to-end with standard backpropagation.
- The proposed SFCM was tested on many challenging computer vision tasks, such as image classification, scene text detection, and image-to-image translation. Experimental results show that our SFCM can be used in many existing frameworks to achieve better performance.

2. Related Work

Network Architecture for improving NN performance.

There are many notable network architecture innovations which have yielded competitive results. The Network in Network[19] structure includes micro multi-layer perceptrons into the filters of convolutional layers to extract more complicated features. Spatial Transformer Networks [13] is proposed to give neural networks the ability to actively spatially transform feature maps. [6]introduce deformable convolution and deformable RoI pooling to enhance the transformation modeling capability of CNNs. In [31], Deeply-Fused Nets (DFNs) were proposed to improve information flow by combining intermediate layers of different base networks.

Methods using multiple layers. A number of recent approaches improve detection and segmentation by using different layers in a ConvNet. The U-Net [26], which combines the low-layer features and high-level features directly,

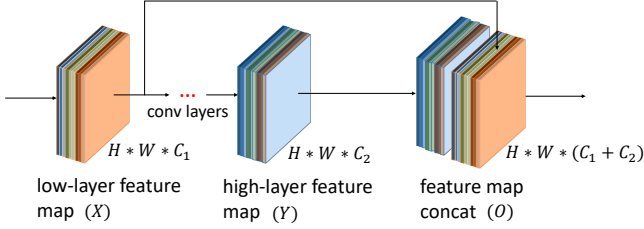


Figure 3. Low-layer and high-layer features in CNN are combined directly.

works with very few training images and yields more precise segmentation. FCN [23] sums partial scores for each category over multiple scales to compute semantic segmentation. HyperNet [16], ParseNet [22], and ION [1] concatenate features of multiple layers before computing predictions, which is equivalent to summing transformed features. SSD [21] and MS-CNN [2] predict objects at multiple layers of the feature hierarchy without combining features or scores. The FPN (Feature Pyramid Network) [20] develops a top-down architecture with lateral connections to build high-level semantic feature maps at all scales. But all these features combining methods just combine or use the low-layer features and high-level features directly, they have not considered the selectivity of high-level features to low-level features when the low-layer features are combined to the high-layer features.

Attention mechanism. Recently, attention mechanisms have been widely applied in different areas [8, 33, 34]. Similar to human visual processing, attention-based models tend to selectively focus on the meaningful object. Vaswani et.al. [28] propose the self-attention method for machine translation. A self-attention module computes the response at a position in a sequence by attending to all positions and taking their weighted average in an embedding space. Based on the recent work of soft attention [3, 13], the Residual Attention Network [30] is composed of multiple attention modules which generate attention-aware features. The attention aware features from different modules change adaptively as layers go deeper.

3. Method

3.1. Selective Feature Connection Mechanism

In the existing features concatenation methods, feature maps obtained from different layers are combined directly. As shown in Fig.3, the combined features:

$$O = [X, Y] \quad (1)$$

where $X \in \mathbf{R}^{H \times W \times C_1}$ denotes the low-layer feature map. C_1 is the number of channels, W and H denote the width and height. $Y \in \mathbf{R}^{H \times W \times C_2}$ is the high-level features of a ConvNet. $O \in \mathbf{R}^{H \times W \times (C_1 + C_2)}$ is the combined features.

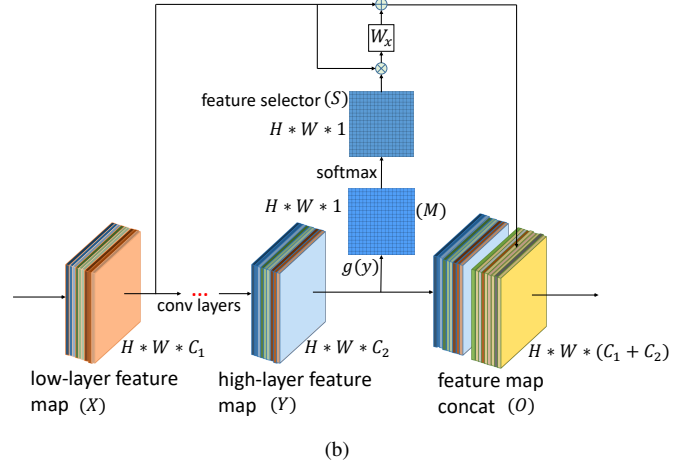
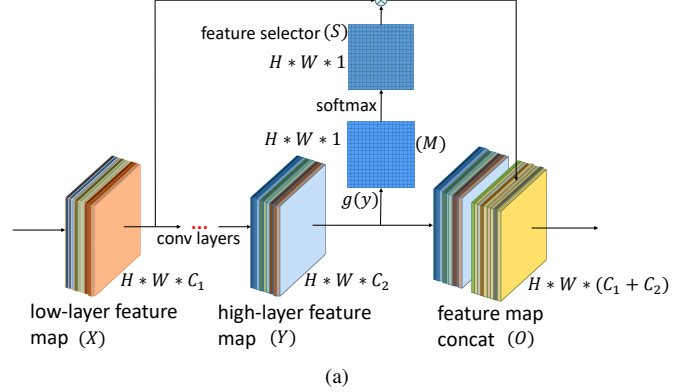


Figure 4. The proposed SFCM. "⊗" denotes matrix multiplication, "⊕" denotes element-wise sum. (a)Direct connection. (b)Residual connection.

However, direct concatenation can not sufficiently apply the complementary information of high-layer and low-layer features. High-layer features can represent semantic information of an image and low-layer features contain more detail information. When directly combining the high-layer and low-layer features, the background clutter and semantic ambiguity may be caused due to the introduction of detailed information.

Inspired by the attention mechanism and recent advances in the deep neural network, we propose the SFCM. We assign attention scores for each local position on the low-layer feature map which indicates the importance of the low-layer features. As shown in Fig.4, the attention scores M are learned by $C_2 \times 1 \times 1 \times 1$ convolutional filters.

$$M = W_g * Y \quad (2)$$

where "*" denotes a convolution operation. W_g is the learned weight matrices, which is implemented as 1×1 convolution.

In order to ensure the non-negative of the feature selec-

tor, we subsequently use the softmax normalization on M to get feature selector S :

$$S_{i,j} = \frac{\exp(M_{i,j})}{\sum_{i=1}^W \sum_{j=1}^H \exp(M_{i,j})} \quad (3)$$

In this equation, $S \in \mathbf{R}^{H \times W \times 1}$, and $S_{i,j}$ is the score at position (i, j) .

The feature selector is learned to indicate the importance of the low-layer features. Thus, more noteworthy features in lower-layer can be screened out by multiplying the feature selector with the low-layer features. We denote the value in X corresponding to channel c and spatial location (i, j) as $X_{i,j,c}$. Thus, the new lower-layer features X^s is given by

$$X_{i,j,c}^s = S_{i,j} * X_{i,j,c} \quad (4)$$

Then the new lower-layer features are used to connect with higher-layer features.

3.2. Adding SFCM To Deep ConvNets

We give two connection modes to add our method to a deep ConvNets, direct connection and residual connection [10, 32]. Fig.4 shows the building block that constructs our SFCM. **Direct connection:** In the direct connection, as shown in Fig.4 (a), X^s , given in Eq.(3), is concatenated to the higher level features Y :

$$O = [X^s, Y] \quad (5)$$

Residual connection: As exemplify in Fig.4 (b), the selective feature connection method can also be incorporated into many existing architectures in the form of residuals.

We further multiply X^s by a scale parameter W_x and add back the lower-layer features X :

$$X' = W_x * X^s + X \quad (6)$$

where X^s is given in Eq.(3) and "+ X " denotes a residual connection.

The residual connection allows us to insert the Selective Feature Connection block into any pre-trained model, without breaking its initial behavior. In the residual connection, X' is concatenated to the higher level features Y :

$$O = [X', Y] \quad (7)$$

3.3. Discussion

We discuss why the SFCM is efficient by visualizing the feature layer of the network and the feature selector in SFCM, and then we discuss what exactly learned of the feature selector and how the SFCM works on improving state-of-art model performance. Fig.2 shows an example of low-layer features selected for further use by a feature selector which is generated by high-level features. As shown in

Fig.2, the "Feature selector" can learn to enhance the area of foreground and suppress the area of background. With the SFCM, most of the pixels in the low-level feature map are suppressed, so it can achieve better performance as it can combine high-resolution maps with low-level features without harm the semantic representation capacity of high-layer features and add more detail features of the foreground to further increase the representation ability of features.

4. Experiments

We perform experiments on a variety of tasks, like image classification, scene text detection, and image-to-image translation, to explore the effectiveness and universality of our SFCM. The output forms of these tasks, *Label-Output* in classification, *Coordinate-Regression* in detection and *Image-Output* in image-to-image translation, contain all output types in computer vision tasks, so the experiments can prove the generality of our method.

4.1. Image classification:

Datasets: The effect of SFCM in image classification is tested on the CIFAR-10 and CIFAR-100 datasets [17]. CIFAR-10 (C10) and CIFAR-100 (C100) consist of images in 10 classes and 100 classes respectively and all images in the two CIFAR datasets are colored natural images with 3232 pixels. The training sets contain 50,000 images and the test sets contain 10,000 images. In the training progress we hold out 5,000 training images as a validation set. A standard data augmentation scheme (mirroring/shifting)[10, 19] is used for these two datasets and we denote this data augmentation scheme by a + mark at the end of the dataset name (e.g., C10+).

Base Network: We adopt the DenseNet [11] as the base model and change the dense block to include the SFCM. Fig.5 (a) shows a dense block in DenseNet. A dense block with SFCM is shown in Fig.5 (b). Our focus is on the behaviors of a deep network with SFCM, but not on pushing the state-of-the-art results, so we intentionally use the simplest architectures in the DenseNet (three dense blocks, $k = 12$, $depth = 40$).

All the networks are trained end-to-end using stochastic gradient descent(SGD) and batch size 64 for 300 epochs. The initial learning rate is set to 0.1 and is divided by 10 at 50% and 75% of the total number of training epochs.

Result Discussion: Table4 evaluates the effect of SFCM on DenseNet. When using the proposed SFCM in the first dense block, the error rates of 6.28% with direct connection and 6.21% with residual connection are both lower than that of 7.00% achieved by the DenseNet without SFCM achieves on CIFAR-10 dataset. When using the proposed SFCM in the first two dense blocks, the error rates of 6.10% with direct connection and 6.21% with residual connection are both lower than that of 7.00% achieved by the DenseNet

Table 1. Error rates(%) on CIFAR datasets. ”+” indicates standard data augmentation (translation and/or mirroring).The ”Direct Connection” refers to use the proposed SFCM in the Direct Connection mode. The error rates of direct and residual connections on CIFAR-10 and CIFAR-100 are both significantly lower than that of the DenseNet. Residual connection is more effective on existing architectures, as the selective feature connection method does not change the existing architectures in the form of residuals.

| Method | Direct Connection | | | | Residual Connection | | | |
|--------------------------------------|-------------------|-------------|--------------|--------------|---------------------|-------------|--------------|--------------|
| | C10 | C10+ | C100 | C100+ | C10 | C10+ | C100 | C100+ |
| DenseNet(baseline) | 7.00 | 5.24 | 27.55 | 24.42 | 7.00 | 5.24 | 27.55 | 24.42 |
| The first dense block with SFCM | 6.28 | 4.26 | 25.72 | 22.17 | 6.21 | 4.21 | 24.94 | 23.26 |
| The first two dense blocks with SFCM | 6.10 | 4.02 | 24.37 | 20.79 | 5.93 | 3.94 | 23.76 | 20.51 |
| All dense blocks with SFCM | 5.92 | 3.81 | 22.96 | 19.31 | 5.62 | 3.73 | 22.18 | 19.27 |

Table 2. Results on ICDAR 2015 Challenge 4 Incidental Scene Text Localization task and COCO-Text. * This baseline is provided by the published model by EAST.

| Method | ICDAR2015 | | | COCO-Text | | |
|-----------------------|---------------|---------------|---------------|---------------|---------------|---------------|
| | Recall | Precision | F-score | Recall | Precision | F-score |
| EAST-ResNet(baseline) | 0.7722 | 0.8464 | 0.8076 | 0.3240 | 0.5039 | 0.3945 |
| EAST with SFCM | 0.7916 | 0.8687 | 0.8254 | 0.3928 | 0.5416 | 0.4554 |

without SFCM achieves on CIFAR-10 dataset. And the experimental results clearly show that:

- The accuracy is steadily improved when more layers use the proposed SFCM. When changing all the three dense blocks with the SFCM, the error drops to 5.62% which is close to 20% lower than the baseline.
- Both the direct connection and residual connection significantly improve the accuracy of classification. The error rates of direct and residual connections on both CIFAR-10 and CIFAR-100 are both significantly lower than those of the DenseNet. This suggests that the SFCM can improve the representation learning ability of CNN models.
- Residual connection is more effective on existing architectures, as the selective feature connection method does not change the existing architectures in the form of residuals.

4.2. Scene text detection

We further verify the effectiveness of our method in scene text detection. Scene text detection is challenging due to text may exist in natural images with arbitrary size and orientation. The core of text detection is the design of features to distinguish text from backgrounds. Merging feature maps of different layers should be possible to improve the performance of detecting text in various size.

Datasets: We use the ICDAR 2015 Incidental Text datasets [14] and the COCO-Text [29] to test the effect of SFCM on Scene text detection. The ICDAR 2015 Incidental Text datasets issues from Challenge 4 of the ICDAR 2015

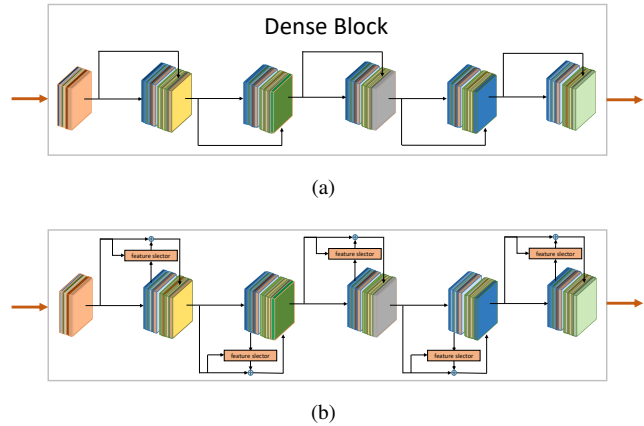


Figure 5. The dense block for image classification. (a) A dense block in DenseNet. (b) A dense block with SFCM (shown in orange).

Robust Reading Competition. It includes 1000 training images and 500 testing images. These images are taken by Google Glass in an incidental way. Therefore text in the scene can be in arbitrary orientations, or suffer from motion blur and low resolution. Each image may contain multi-oriented text. Annotations are provided in terms of word bounding boxes. COCO-Text is a large dataset which contains 63686 images, where 43,686 of the images is used for training, 10,000 for validation, and 10,000 for testing. It is one of the challenges of ICDAR 2017 Robust Reading Competition.

Base Network: We use the EAST [36] as the base model and only change the feature-merging branch with SFCM(Fig.6). We take the residual connection of the SFCM because it is more effective on existing architectures.

Table 3. Segmentation accuracy for the Cityscapes labels \leftrightarrow photos task. The cGAN with SFCM outperform the original cGAN [12].

| Method | Per-pixel acc. | Per-class acc. | Class IOU |
|-----------------|----------------|----------------|-----------|
| cGAN (baseline) | 0.66 | 0.23 | 0.17 |
| cGAN with SFCM | 0.71 | 0.24 | 0.17 |
| Ground truth | 0.80 | 0.26 | 0.21 |

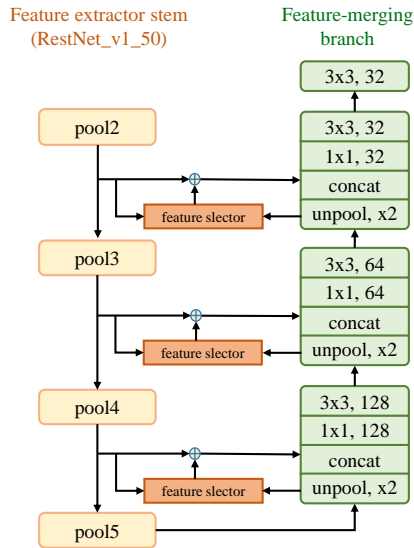


Figure 6. The feature-merging branch of the EAST is modified with SFCM.

The network is trained end-to-end using ADAM [15] optimizer. To speed up learning, we uniformly sample 512x512 crops from images to form a minibatch of size 14 on one GPU. Learning rate of ADAM starts from 1e-3 and is decayed to its 0.94 every 10000 minibatches. The network is trained until performance stops improving.

Result Discussion: Table 4 evaluates the effect of SFCM on EAST. When changing the feature-merging branch with the SFCM, the EAST achieves an F-score of 0.8254 on IC-DAR 2015 dataset and 0.4554 on COCO-Text dataset, and both accuracy and recall rate have been greatly improved on these two datasets. This is because the model with the proposed connection mechanism reduces the misjudgment of small targets and makes the detection of larger targets more accurate, as shown in Fig.7. This further verifies that our SFCM can concatenate the low-level features and high-level features more effectively.

4.3. Image-to-image translation

Image-to-image translation is one of the image generation tasks aiming to translate an input image into a corresponding output image. A defining feature of image-to-image translation problems is that they map a high-



(a)



(b)

Figure 7. Some results on images in ICDAR-15 dataset. (a)Result of EAST-ResNet(baseline). (b)Result of EAST-ResNet with SFCM.

resolution input grid to a high-resolution output grid. Combining high-level semantic information with low-level detail information effectively can significantly improve the performance of the translation tasks. We use the pix2pix (cGAN) [12] as the base model to test the effect of SFCM. This model solves the image generation task that the input and output differ in surface appearance, but both are renderings of the same underlying structure. We test the method on two tasks and datasets:

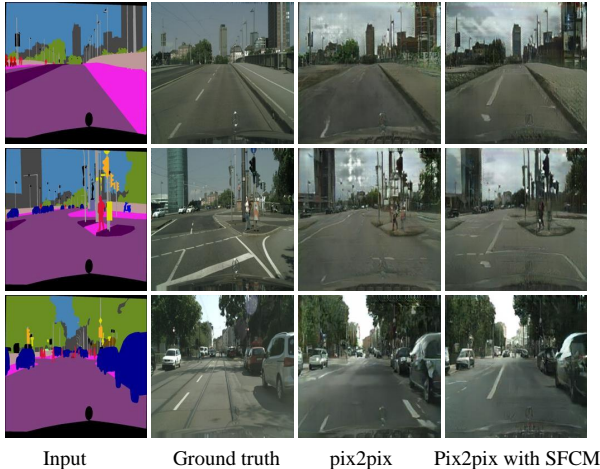
1. Semantic labels to photo, trained on the Cityscapes dataset [4]
2. Architectural labels to photo, trained on CMP Facades [27]

Quantitative evaluation of generative models is known to be challenging. In this paper, we use the FCN-score [12] to evaluate the performance of our SFCM. We adopt the popular FCN-8s [23] architecture for semantic segmentation, and train it on the cityscapes dataset. We then score synthesized photos by the classification accuracy against the labels these photos were synthesized from.

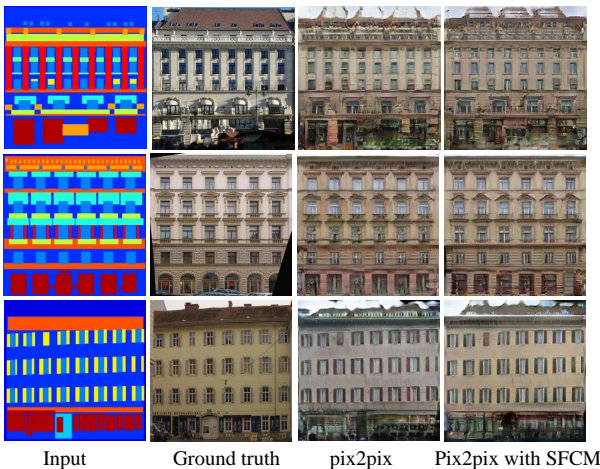
Table 4.2 evaluates the effect of SFCM on cGAN. The cGAN with SFCM outperforms the original cGAN. Some results are shown in Fig.8. The resolution of the images translated by the pix2pix is low and the image edges are illegible. With the SFCM, images translated by the pix2pix show much higher quality especially in details. This suggests that the SFCM has learned a more effective way to combine high-level semantic information with low-level detail information in the training of a neural network model.

5. Conclusion and Future Work

We have presented the SFCM, a general network architecture to concatenate CNN features of different layers, and



(a)



(b)

Figure 8. Some results on the image-to-image translation tasks. The pix2pix model with the SFCM results in much higher quality results, especially in details. (a) Example results on the Cityscapes dataset. (b) Example results on CMP Facades.

given two connection modes, direct connection and residual connection, to use this architecture. With our SFCM, low-level features are selectively linked to high-level features with an attention map which is generated by high-level features. Experiments show that our method is effective and generic on a variety of tasks, including image classification, scene text detection, and image-to-image translation. Thus, it provides a practical solution for research and applications to use multi-layer CNN Features more effectively.

References

[1] S. Bell, C. Lawrence Zitnick, K. Bala, and R. Girshick. Inside-outside net: Detecting objects in context with skip pooling and recurrent neural networks. In *Proceedings of*

the IEEE conference on computer vision and pattern recognition, pages 2874–2883, 2016.

[2] Z. Cai, Q. Fan, R. S. Feris, and N. Vasconcelos. A unified multi-scale deep convolutional neural network for fast object detection. In *European Conference on Computer Vision*, pages 354–370. Springer, 2016.

[3] L.-C. Chen, Y. Yang, J. Wang, W. Xu, and A. L. Yuille. Attention to scale: Scale-aware semantic image segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3640–3649, 2016.

[4] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016.

[5] J. Dai, Y. Li, K. He, and J. Sun. R-fcn: Object detection via region-based fully convolutional networks. In *Advances in neural information processing systems*, pages 379–387, 2016.

[6] J. Dai, H. Qi, Y. Xiong, Y. Li, G. Zhang, H. Hu, and Y. Wei. Deformable convolutional networks. *CoRR*, abs/1703.06211, 1(2):3, 2017.

[7] R. Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015.

[8] K. Gregor, I. Danihelka, A. Graves, D. J. Rezende, and D. Wierstra. Draw: A recurrent neural network for image generation. *arXiv preprint arXiv:1502.04623*, 2015.

[9] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask r-cnn. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, pages 2980–2988. IEEE, 2017.

[10] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.

[11] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger. Densely connected convolutional networks. In *CVPR*, volume 1, page 3, 2017.

[12] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. *arXiv preprint*, 2017.

[13] M. Jaderberg, K. Simonyan, A. Zisserman, et al. Spatial transformer networks. In *Advances in neural information processing systems*, pages 2017–2025, 2015.

[14] D. Karatzas, L. Gomez-Bigorda, A. Nicolaou, S. Ghosh, A. Bagdanov, M. Iwamura, J. Matas, L. Neumann, V. R. Chandrasekhar, S. Lu, et al. Icdar 2015 competition on robust reading. In *Document Analysis and Recognition (ICDAR), 2015 13th International Conference on*, pages 1156–1160. IEEE, 2015.

[15] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[16] T. Kong, A. Yao, Y. Chen, and F. Sun. Hypernet: Towards accurate region proposal generation and joint object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 845–853, 2016.

[17] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.

- [18] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- [19] M. Lin, Q. Chen, and S. Yan. Network in network. *arXiv preprint arXiv:1312.4400*, 2013.
- [20] T.-Y. Lin, P. Dollár, R. B. Girshick, K. He, B. Hariharan, and S. J. Belongie. Feature pyramid networks for object detection. In *CVPR*, volume 1, page 4, 2017.
- [21] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016.
- [22] W. Liu, A. Rabinovich, and A. C. Berg. Parsenet: Looking wider to see better. *arXiv preprint arXiv:1506.04579*, 2015.
- [23] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
- [24] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.
- [25] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [26] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [27] R. Tyleček and R. Šára. Spatial pattern templates for recognition of objects with regular structure. In *German Conference on Pattern Recognition*, pages 364–374. Springer, 2013.
- [28] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008, 2017.
- [29] A. Veit, T. Matera, L. Neumann, J. Matas, and S. Belongie. Coco-text: Dataset and benchmark for text detection and recognition in natural images. *arXiv preprint arXiv:1601.07140*, 2016.
- [30] F. Wang, M. Jiang, C. Qian, S. Yang, C. Li, H. Zhang, X. Wang, and X. Tang. Residual attention network for image classification. *arXiv preprint arXiv:1704.06904*, 2017.
- [31] J. Wang, Z. Wei, T. Zhang, and W. Zeng. Deeply-fused nets. *arXiv preprint arXiv:1605.07716*, 2016.
- [32] X. Wang, R. Girshick, A. Gupta, and K. He. Non-local neural networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [33] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pages 2048–2057, 2015.
- [34] Z. Yang, X. He, J. Gao, L. Deng, and A. Smola. Stacked attention networks for image question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 21–29, 2016.
- [35] M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *European conference on computer vision*, pages 818–833. Springer, 2014.
- [36] X. Zhou, C. Yao, H. Wen, Y. Wang, S. Zhou, W. He, and J. Liang. East: an efficient and accurate scene text detector. In *Proc. CVPR*, pages 2642–2651, 2017.