

Context-Aware Image Inpainting with Learned Semantic Priors

Wendong Zhang¹, Junwei Zhu², Ying Tai², Yunbo Wang^{1*},
Wenqing Chu², Bingbing Ni^{1*}, Chengjie Wang² and Xiaokang Yang¹

¹MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University

²Youtu Lab, Tencent

{diergent, yunbow, nibingbing}@sjtu.edu.cn

Abstract

Recent advances in image inpainting have shown impressive results for generating plausible visual details on rather simple backgrounds. However, for complex scenes, it is still challenging to restore reasonable contents as the contextual information within the missing regions tends to be ambiguous. To tackle this problem, we introduce pretext tasks that are semantically meaningful to estimating the missing contents. In particular, we perform knowledge distillation on pretext models and adapt the features to image inpainting. The learned semantic priors ought to be partially invariant between the high-level pretext task and low-level image inpainting, which not only help to understand the global context but also provide structural guidance for the restoration of local textures. Based on the semantic priors, we further propose a context-aware image inpainting model, which adaptively integrates global semantics and local features in a unified image generator. The semantic learner and the image generator are trained in an end-to-end manner. We name the model SPL to highlight its ability to *learn and leverage semantic priors*. It achieves the state of the art on Places2, CelebA, and Paris StreetView datasets[†].

1 Introduction

Image inpainting aims at generating visually plausible image structures and local details for missing regions. This task has been researched for a long time and can be applied to many applications such as image restoration, object removal, and image manipulation [Criminisi *et al.*, 2004]. In summary, the main challenge for restoring realistic and reasonable contents lies in how to resolve the context ambiguities in missing regions. Although some approaches exploit stochastic models to learn the multi-modal distributions of the missing contents [Zheng *et al.*, 2019; Zhao *et al.*, 2020], in this paper, we focus on learning deterministic models that can leverage global context cues to reduce the context ambiguities.

*Corresponding Authors

[†]Code available at <https://github.com/WendongZh/SPL>

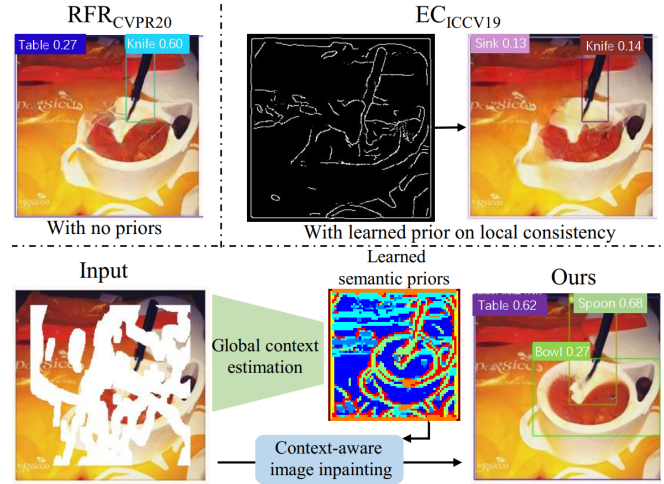


Figure 1: Illustration of SPL. Unlike existing methods (*e.g.*, RFR [Li *et al.*, 2020] and EC [Nazeri *et al.*, 2019]) that mainly focus on local textures and are hard to restore reasonable contents especially when there exist context ambiguities, our model tackles this problem from a more semantic view. The detection results on the generated images show the benefits of learning semantic priors.

In the realm of deterministic image inpainting, previous methods generally focus more on local texture consistency. The early work in this field usually uses the local texture similarities to perform pixel interpolations, such as texture diffusion or patch-based methods [Barnes *et al.*, 2009; Levin *et al.*, 2003]. However, these methods may fail to restore non-repetitive structures. Recently, many learning-based approaches are proposed in form of deep neural networks and directly reformulate the image inpainting as a conditional image generation problem [Yu *et al.*, 2020; Li *et al.*, 2020; Xie *et al.*, 2019; Nazeri *et al.*, 2019; Liu *et al.*, 2020]. Among these methods, some approaches further use particular structural supervisions to enhance the structure completion in final results, such as edge image [Nazeri *et al.*, 2019] and smooth image [Liu *et al.*, 2020]. Although these learning-based methods produce more realistic local details compared with early work, restoring reasonable image contents for more complex scenes is still hard as shown in Figure 1. The main reason lies in that these methods fail to understand the global context but mainly focus on local texture consistency, which results in the distorted global structure.

In contrast to the above literature, we suggest that the high-level semantics can be considered as common knowledge between specific pretext tasks and image inpainting and that neural networks pre-learned for the pretext tasks can provide strong global relational priors that can be used to infer low-level missing structures. To properly model the global relations between the missing contents and the remainder, we take the high-level feature maps of a pretrained multi-label classification (detection) model as the semantic supervisions*. These feature maps contain discriminative pre-learned semantic information of multiple objects which can help the global context understanding. Besides, as the pretext tasks mainly focus on the extractions of high-level features, they improve the robustness of our model to local textures and structural defects. The learned contextual features are referred to as *semantic priors*, and we visualize an example via the k-means algorithm in Figure 1. Following this line, a novel context-aware image inpainting model is proposed which considers both local texture consistency and global semantic consistency. The whole pipeline of our method can be summarized into two phases. In the first phase, we exploit two different encoders to separately extract the low-level image features and learn high-level semantic priors. By performing knowledge distillation on networks from pretext tasks†, we obtain the global context representations. In the second phase, the learned semantic priors are spatially injected into image features to help image completion. By learning and leveraging semantic priors, our approach (SPL) not only strengthens the local consistency of the generated image but also facilitates the structural inference of the missing contents from the remainder.

Our contributions to image inpainting can be summarized as follows:

- We show that distilling high-level knowledge from specific pretext tasks can benefit the understanding of global contextual relations and is thus semantically meaningful to low-level image inpainting. It requires no additional human annotations on the image inpainting datasets.
- We propose SPL, a new model for context-aware image inpainting, which adaptively incorporates the learned semantic priors and the local image features in a unified generative network. It is trained with the semantic prior learner in an end-to-end fashion.
- SPL achieves state-of-the-art performance on three standard image inpainting datasets.

2 Related Work

Rule-based image inpainting. Early work for image inpainting usually designs particular rules on low-level image features to establish the pixel-level mappings. More specifically, exemplar-based approaches [Criminisi *et al.*, 2004;

Barnes *et al.*, 2009] exploit the patch or pixel-level similarities to directly copy textures in visible regions into missing regions. On the other hand, fusion-based methods [Ballester *et al.*, 2001; Levin *et al.*, 2003] try to smoothly propagate proper image contents into missing regions based on neighboring visible pixels. Although these methods produce vivid textures for simple background completion, they usually generate non-realistic images containing repetitive patterns due to the lack of high-level semantics.

Learning-based image inpainting. Recent learning-based methods are mostly built upon generative adversarial networks [Goodfellow *et al.*, 2014], and formulate the image inpainting problem as conditional image generation. The work of context encoder [Pathak *et al.*, 2016] first utilizes an encoder-decoder architecture and proposes a channel-wise fully connected layer to estimate the missing contents. Subsequently, the global and local discriminators are introduced by [Iizuka *et al.*, 2017] to further enhance the constraints of local consistency. After that, a two-stage coarse-to-fine generation framework is proposed by [Yu *et al.*, 2018] with a context attention module. To exclude the influence of invalid pixels during the feature extraction process, the gated convolution layer [Yu *et al.*, 2019] and the region normalization block [Yu *et al.*, 2020] are proposed successively. Besides, the attention mechanism [Xie *et al.*, 2019] and multi-scale learning framework [Wang *et al.*, 2019; Li *et al.*, 2020] are also introduced. Although these approaches use different network architectures for learning image features, they usually fail to restore reasonable structures in complex scenes due to the unresolved context ambiguities.

Structure-enhanced image inpainting. To resolve the structural distortion of the missing contents, some two-stage models [Nazeri *et al.*, 2019; Ren *et al.*, 2019] first generate edges or smooth images and then use them to enhance the structural details. Other approaches [Liu *et al.*, 2020; Liao *et al.*, 2020] also use edge-preserving images or semantic maps as extra supervisions and require dense human annotations on the image inpainting dataset. In comparison, our model is trained on the image inpainting dataset in a fully unsupervised manner, and only exploits pre-learned features from the pretext tasks, where data annotations are more accessible. Another difference is that, in our model, the semantic priors are learned and spatially injected into the image features in an end-to-end manner, which provides a more adaptive way for the restoration of the global and local structures.

3 Approach

The overall architecture of the proposed approach is shown in Figure 2. We name it SPL to highlight its ability to learn and leverage semantic priors. It contains three main components: an image encoder, a semantic priors learner, and an image decoder. In summary, the whole pipeline can be divided into two phases. In the encoding phase, we learn image features and semantic priors from the remainder of corrupted images. In the generation phase, the learned semantic priors are spatially incorporated into the image features to help the restoration of the global and local structural contents.

*Various pretext tasks can be used to provide semantic supervisions. In the rest of the paper, we mainly use multi-label classification as a typical pretext task without loss of generality.

†We do not train or fine-tune these models directly on the image inpainting datasets, but only use existing models that were pretrained on the Open Images dataset [Kuznetsova *et al.*, 2018].

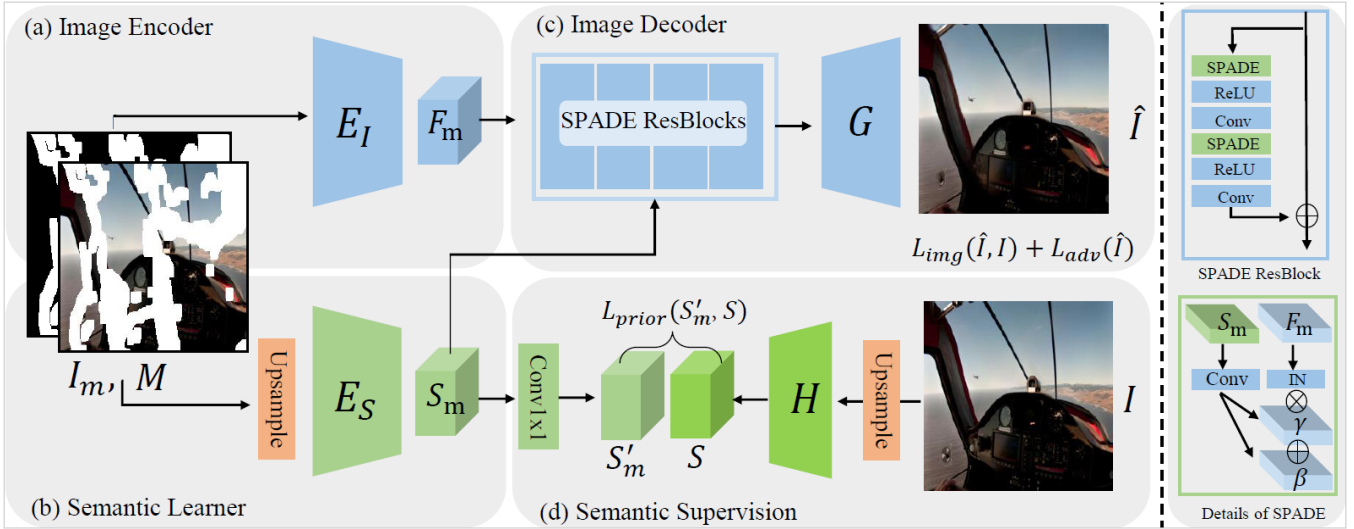


Figure 2: Overview of SPL. In the encoding phase, it learns image features F_m and semantic priors S_m from input via the image encoder E_I and the semantic learner E_S . In the generation phase, the semantic priors are spatially injected into the image features via spatially-adaptive normalization module (SPADE) [Park *et al.*, 2019] to help context-aware image inpainting. We exploit features S of a multi-label classification model (H) [Ben-Baruch *et al.*, 2020] as the supervisions for learning semantic priors, whose parameters are pretrained on the Open Image dataset and fixed during training. Modules (a-c) are trained in an end-to-end manner.

3.1 Image Encoding

Given corrupted image I_m and corresponding mask $M \in \{0, 1\}$ where 1 represents the locations for invalid pixels and 0 represents those for valid pixels, we first utilize two different encoders E_I and E_S to learn image feature F_m and semantic priors S_m . As these two features contain different information, we exploit different strategies to design the architectures of these two encoders, and they do not share any parameters.

For image encoder E_I shown in Figure 2(a), we aim to extract rich low-level features from visible contents to form the conditional input. These features play important roles in restoring vivid local textures. Therefore we exploit two down-sample layers to embed corrupted input into image feature:

$$F_m = E_I(I_m, M), \quad (1)$$

where $I_m \in \mathbb{R}^{3 \times H \times W}$, $M \in \mathbb{R}^{1 \times H \times W}$, $F_m \in \mathbb{R}^{c \times \frac{H}{4} \times \frac{W}{4}}$. The image feature F_m will be further fed to the generation phase to restore plausible visual details.

3.2 Semantic Priors Learning

For the semantic learner E_S shown in Figure 2(b), we aim to learn complete semantic priors for corrupted visual elements with supervisions from pretrained deep neural networks. To understand the global context which usually contains multiple different objects, as shown in Figure 2(d), we exploit the multi-label classification model H pretrained on the Open Image dataset [Kuznetsova *et al.*, 2018] with asymmetric loss (ASL) [Ben-Baruch *et al.*, 2020] to provide supervisions, and we do not fine-tune this model on our image inpainting datasets. For this encoder, the texture details are less important and the encoder needs to pay more attention to understanding the global context in the remainder of the corrupted images. To this end, we make two modifications.

First, as we aim to learn semantic information from higher layers of deep neural networks, we up-sample the input images to preserve more local structures. More specifically, we up-sample the full image $I \in \mathbb{R}^{3 \times H \times W}$ into a large one $I' \in \mathbb{R}^{3 \times 2H \times 2W}$, and the corrupted image I_m and corresponding mask M are also up-sampled in the same manner to form the input I'_m and M' for learning semantic priors. After that, we utilize I' as the input to the pretrained multi-label classification model H to extract feature maps $S \in \mathbb{R}^{d \times \frac{H}{4} \times \frac{W}{4}}$ as the supervisions of the semantic priors:

$$S = H(I'). \quad (2)$$

Second, to establish the mappings between visible contents and semantic priors, we exploit three down-sampling layers followed by five residual blocks to form the semantic learner E_S . The semantic learner E_S utilizes enlarged corrupted image and mask as input to generate the semantic priors S_m :

$$S_m = E_S(I'_m, M'), \quad (3)$$

where $I'_m \in \mathbb{R}^{3 \times 2H \times 2W}$, $M' \in \mathbb{R}^{1 \times 2H \times 2W}$, $S_m \in \mathbb{R}^{c \times \frac{H}{4} \times \frac{W}{4}}$. We then exploit a 1×1 convolutional layer to adapt S_m to the representations from the pretext task:

$$S'_m = \text{Conv}_{1 \times 1}(S_m), \quad (4)$$

where $S'_m \in \mathbb{R}^{d \times \frac{H}{4} \times \frac{W}{4}}$. Finally, an ℓ_1 reconstruction loss with extra constraints on missing regions is proposed to supervise the semantic learner:

$$\mathcal{L}_{\text{prior}} = \|(S - S'_m) \odot (1 + \alpha M_s)\|_1, \quad (5)$$

where \odot represents Hadamard product operator, α represents the extra weight for missing regions and M_s represents resized mask which has the same spatial size as S . With Eq. 4 and Eq. 5, we can distill the knowledge that are useful for image inpainting and filter out the task-agnostic components in S .

3.3 Context-Aware Image Inpainting

In the generation phase, we exploit the image decoder shown in Figure 2(c) to tackle another key problem that how to utilize the learned semantic priors \mathbf{S}_m to help the image restoration process. As the image feature \mathbf{F}_m and learned semantic priors \mathbf{S}_m focus on different aspects of image contents, directly concatenating these features for feature fusion will not only disturbs local textures in visible regions but also affects the learning processes of the corresponding encoders. To this end, we exploit 8 ResBlocks with the *spatially-adaptive normalization module* (SPADE) [Park *et al.*, 2019] to spatially inject the semantic priors into the decoding phase. In this case, our model provides adaptive modifications on image features and decreases the unnecessary influences. Instead of utilizing images or semantic maps as conditional inputs, in our model, the affine transformation parameters in SPADE are learned from semantic priors and we provide the architecture of this module in Figure 2. In more details, the SPADE modules in our model first normalize the input image feature \mathbf{F}_m with non-parametric instance normalization. Then, two different sets of parameters are learned from semantic priors \mathbf{S}_m to perform the spatially pixel-wise affine transformations on image feature \mathbf{F}_m :

$$[\gamma, \beta] = \text{SPADE}(\mathbf{S}_m), \quad (6)$$

$$\mathbf{F}'_m = \gamma \cdot \text{IN}(\mathbf{F}_m) + \beta, \quad (7)$$

where $\gamma, \beta, \mathbf{F}_m, \mathbf{F}'_m \in \mathbb{R}^{c \times \frac{H}{4} \times \frac{W}{4}}$ and IN represents non-parametric instance normalization. Finally, the output image is generated from the modified image feature \mathbf{F}'_m with generator G which contains two upsampling layers:

$$\hat{\mathbf{I}} = G(\mathbf{F}'_m). \quad (8)$$

Novelty of the use of SPADE. Although the SPADE ResBlocks are the core modules of SPL, they are used for different purposes and in different ways from the original SPADE model. Unlike those used for image-to-image translation, here they are jointly trained with two separate encoders and integrate the learned low-level features (instead of random variables) with the distilled high-level priors (instead of the raw layout data). The contribution of SPL is to show that extracting high-level knowledge from specific pretext tasks can be helpful to low-level image inpainting, which has not been thoroughly studied by previous work.

3.4 Loss Function

We exploit the reconstruction and adversarial losses to train the whole model of SPL. For the reconstruction loss, we use ℓ_1 loss for the reconstructed image and focus more on the content of missing regions:

$$\mathcal{L}_{\text{img}} = \|(\mathbf{I} - \hat{\mathbf{I}}) \odot (\mathbf{1} + \delta \mathbf{M})\|_1, \quad (9)$$

where \odot represents the Hadamard product operator and δ represents the weight for missing regions. To generate more plausible local details, we also utilize the adversarial loss to train our model, along with a discriminator in an iterative training scheme:

$$\mathcal{L}_{\text{adv}} = -\mathbb{E}_{\hat{\mathbf{I}}}[\log(1 - D(\hat{\mathbf{I}}))], \quad (10)$$

$$\mathcal{L}_{\text{adv}}^D = \mathbb{E}_I[\log D(\mathbf{I})] + \mathbb{E}_{\hat{\mathbf{I}}}[\log(1 - D(\hat{\mathbf{I}}))]. \quad (11)$$

The full objective function of SPL can be written as

$$\mathcal{L} = \lambda_1 \mathcal{L}_{\text{img}} + \lambda_2 \mathcal{L}_{\text{adv}} + \lambda_3 \mathcal{L}_{\text{prior}}. \quad (12)$$

We tune the hyper-parameters with grid search, and finally obtain $\lambda_1 = 10$, $\lambda_2 = \lambda_3 = 1$, $\alpha = 3$ and $\delta = 5$.

4 Experiments

4.1 Experimental Setups

Datasets. We evaluate our approach on three datasets, including Places2 [Zhou *et al.*, 2017], CelebA [Liu *et al.*, 2015], and Paris StreetView [Doersch *et al.*, 2015]. We use the standard training set of Places2 with over 4 million images and evaluate all models on its validation set with 36,500 images. The CelebA human face dataset contains over 160,000 training images and about 19,900 testing images. The StreetView dataset contains 15,900 training samples and 100 testing samples. We use the mask set of [Liu *et al.*, 2018] with 12,000 irregular masks that are pre-grouped into six intervals according to the size of the masks (0%-10%, 10%-20%, ..., 50%-60%). All images in Places2 and StreetView are resized to 256×256 . As for CelebA, we first apply center cropping and then resize the images to 256×256 , which is a common practice of previous work [Li *et al.*, 2020; Yu *et al.*, 2020].

Compared methods. We compare our method (SPL) with state-of-the-art methods including the *region normalization model* (RN) [Yu *et al.*, 2020], the *recurrent feature reasoning model* (RFR) [Li *et al.*, 2020] and the *mutual feature equalization mode* (MFE) [Liu *et al.*, 2020]. Among these methods, RN and RFR mainly focus on the consistency of local textures, and MFE exploits smooth images as extra supervisions to enhance structural completion. For methods that provide pretrained models, we directly evaluate them in our settings. Otherwise, we retrain the models based on the provided code.

Implementation and training details. To provide supervisions for semantic priors, we use multi-label classification as the pretext task, in which the model was pretrained on the Open Image dataset [Kuznetsova *et al.*, 2018] with the asymmetric loss (ASL) [Ben-Baruch *et al.*, 2020]. We do not fine-tune the model on any image inpainting datasets. We adopt the patch-based discriminator from previous work [Yu *et al.*, 2020; Nazeri *et al.*, 2019]. We apply the Adam solver with $\beta_1 = 0.0$ and $\beta_2 = 0.9$ for model optimization. The initial learning rate is $1e^{-4}$ for all experiments and decayed to $1e^{-5}$ at different epochs for different datasets. For Places2 and CelebA, we decay the learning rate at 30 epochs and further fine-tune the model for another 10 epochs. For Paris StreetView, we decay the learning rate at 50 epochs and fine-tune the model for 20 epochs. The batch size is set to 8 for all datasets and we train our model with two V100 GPUs.

4.2 Qualitative Comparisons

We provide qualitative results on three different datasets in Figure 3. The Places2 dataset is the most challenging dataset as it contains different complex scenes and various objects.




Figure 3: Qualitative comparisons of our approach (SPL) with RN [Yu *et al.*, 2020], RFR [Li *et al.*, 2020], and MFE [Liu *et al.*, 2020]. Notably, as shown by the detection results on the generated images, SPL provides a better estimation of object-level contents and structures.

For this dataset, we can observe from the first four columns that results from our methods contain more complete global and local structures. On the contrary, results from other methods contain obvious distorted structures and suffer from blur effects. These results suggest that our method successfully reduces the context ambiguities in missing regions and obtains more complete visual concepts for foreground objects.

For CelebA dataset, results are shown in the fifth column. In this case, all methods produce complete human faces as this dataset only contains various aligned human faces. However, as shown in the fifth column, previous approaches generate meaningless textures in local details. In other words, these methods fail to globally understand the context information and pay more attention to fill the missing regions with textures based on local distributions. Compared with them, our results contain more clear local details, which suggests that the learned semantic priors can also help local texture restoration. Similar observations can also be found in Paris dataset.

We further provide detection results of the output images in the last column in Figure 3. From these results, we can observe that our result contains more complete objects with higher detection confidences than other methods.

4.3 Quantitative Comparisons

We further provide the quantitative comparison results in Table 1. To quantitatively evaluate our performance, we select peak signal-to-noise ratio (PSNR), structural similarity (SSIM), and mean ℓ_1 error (MAE) as our evaluation metrics which are commonly used in previous work. The evaluation process is as follows. We first randomly sample three different mask sets from the whole irregular mask dataset. These mask sets are assigned to different image datasets to form the mask-image pairs. For each image dataset, the mask-image mappings are held for different methods to obtain fair comparisons. The results of the whole testing sets are shown in columns with mask ratio being “All”, and we gather those images with masks in corresponding mask ratios to calculate other results in Table 1. We can see that our method obtain obvious improvements on all three metrics, which further demonstrates the effectiveness of our method.

4.4 Ablation Studies

Alternatives of the semantic supervisions. We then conduct two different experiments to evaluate the effects of the proposed semantic supervisions. First, we directly remove

Dataset		Places2			CelebA			Paris StreetView		
Mask Ratio		20%-40%	40%-60%	All	20%-40%	40%-60%	All	20%-40%	40%-60%	All
SSIM \uparrow	MFE	0.820	0.655	0.809	0.916	0.806	0.900	0.872	0.707	0.834
	RFR	0.858	0.706	0.842	0.931	0.837	0.917	0.893	0.763	0.863
	RN	0.877	0.728	0.858	0.941	0.850	0.923	0.892	0.757	0.861
	Ours	0.895	0.762	0.877	0.950	0.868	0.935	0.911	0.790	0.882
PSNR \uparrow	MFE	23.07	19.22	24.24	29.01	23.78	30.14	27.22	22.07	27.07
	RFR	24.52	20.49	25.72	30.30	25.02	31.39	28.32	23.71	28.36
	RN	25.55	21.16	26.58	31.15	25.43	31.94	28.44	23.53	28.37
	Ours	26.58	22.14	27.71	32.14	26.34	33.27	29.34	24.47	29.38
MAE \downarrow	MFE	0.0260	0.0531	0.0287	0.0120	0.0289	0.0146	0.0159	0.0412	0.0226
	RFR	0.0209	0.0435	0.0233	0.0099	0.0240	0.0122	0.0134	0.0324	0.0182
	RN	0.0184	0.0403	0.0212	0.0091	0.0232	0.0119	0.0136	0.0337	0.0187
	Ours	0.0161	0.0355	0.0187	0.0080	0.0205	0.0102	0.0115	0.0292	0.0161

Table 1: Quantitative comparisons of the generated images by our approach (SPL) and the state of the art on three image inpainting datasets.





Figure 4: Comparisons on the StreetView dataset of (1) different ways to integrate the learned semantic priors into the image decoder, and (2) using different semantic supervisions (vs. w/ YoloV5).


 Figure 5: Visualization of the semantic priors using the k-means algorithm, compared with *edgeconnect* (EC) [Nazeri *et al.*, 2019].

the semantic supervisions from our model, named as “w/o S”. Then, we replace the pretrained multi-label classification network with a pretrained detection network YoloV5 [Jocher and others, 2020] to provide semantic supervisions. In this case, we aim to explore whether other pretext models can also provide sufficient information for global context understanding. From Figure 4, we can obtain two observations: 1) The performance of our method will suffer from significant degeneration without explicit supervision from pretrained model. 2) Supervisions from the detection model can also provide helpful semantic information for global context understanding.

Alternatives of SPADE in SPL. Instead of using the SPADE modules, in this experiment, we directly concatenate the learned semantic priors and the features from the image

encoder to form the input of ResBlocks, which is commonly used by previous tow-stages generation frameworks [Nazeri *et al.*, 2019; Ren *et al.*, 2019]. Results in Figure 4 show that the straightforward operation of concatenation leads to decreases in all three evaluation metrics.

4.5 Visualization of Learned Semantic Priors

We further visualize the learned semantic priors and compare the generation results with the edge pre-generation method (EC) [Nazeri *et al.*, 2019] in Figure 5. Specifically, we directly conduct the k-means algorithm on the learned semantic priors and use different colors to represent different clusters. From this figure, we can see that the EC model cannot restore reasonable structures for complex scenes. Besides, the distorted structures in edge images will further mislead the final results. Instead, the semantic priors learned by our method reduce the context ambiguities and provide explicit guidance for both global and local structure restoration.

5 Conclusion

In this paper, we propose a new solution to image inpainting named SPL. Instead of focusing on the consistency of local textures, we show that the high-level features from deep neural networks, which were well-trained under specific pretext tasks (*e.g.*, multi-label classification and object detection), contain rich semantic information and can benefit global context understanding. We transfer the high-level knowledge from the pretext models to low-level feature space and use the learned priors as the structural guidance of image inpainting. Finally, we propose the context-aware image inpainting model that adaptively incorporates global semantics with local textures in a unified image decoding module. SPL significantly outperforms the state-of-the-art approaches on three datasets.

Acknowledgments

This work was supported by NSFC (U19B2035, U20B2072, 61976137). It was also supported by the Shanghai Municipal Science & Technology Major Project (2021SHZDZX0102).

References

- [Ballester *et al.*, 2001] Coloma Ballester, Marcelo Bertalmio, Vicent Caselles, Guillermo Sapiro, and Joan Verdera. Filling-in by joint interpolation of vector fields and gray levels. *IEEE transactions on image processing*, 10(8):1200–1211, 2001.
- [Barnes *et al.*, 2009] Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. Patchmatch: A randomized correspondence algorithm for structural image editing. *ACM Transactions on Graphics*, 28(3):24, 2009.
- [Ben-Baruch *et al.*, 2020] Emanuel Ben-Baruch, Tal Ridnik, Nadav Zamir, Asaf Noy, Itamar Friedman, Matan Protter, and Lihi Zelnik-Manor. Asymmetric loss for multi-label classification. *arXiv preprint arXiv:2009.14119*, 2020.
- [Criminisi *et al.*, 2004] Antonio Criminisi, Patrick Pérez, and Kentaro Toyama. Region filling and object removal by exemplar-based image inpainting. *IEEE Transactions on Image Processing*, 13(9):1200–1212, 2004.
- [Doersch *et al.*, 2015] Carl Doersch, Saurabh Singh, Abhinav Gupta, Josef Sivic, and Alexei A Efros. What makes paris look like paris? *Communications of the ACM*, 58(12):103–110, 2015.
- [Goodfellow *et al.*, 2014] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *NeurIPS*, 27:2672–2680, 2014.
- [Iizuka *et al.*, 2017] Satoshi Iizuka, Edgar Simo-Serra, and Hiroshi Ishikawa. Globally and locally consistent image completion. *ACM Transactions on Graphics*, 36(4):1–14, 2017.
- [Jocher and others, 2020] Glenn Jocher et al. Yolov5. <https://github.com/ultralytics/yolov5>, 2020. Accessed: 2020-12-07.
- [Kuznetsova *et al.*, 2018] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, et al. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *arXiv preprint arXiv:1811.00982*, 2018.
- [Levin *et al.*, 2003] Anat Levin, Assaf Zomet, and Yair Weiss. Learning how to inpaint from global image statistics. In *ICCV*, pages 305–312, 2003.
- [Li *et al.*, 2020] Jingyuan Li, Ning Wang, Lefei Zhang, Bo Du, and Dacheng Tao. Recurrent feature reasoning for image inpainting. In *CVPR*, pages 7760–7768, 2020.
- [Liao *et al.*, 2020] Liang Liao, Jing Xiao, Zheng Wang, Chia-Wen Lin, and Shin’ichi Satoh. Guidance and evaluation: Semantic-aware image inpainting for mixed scenes. In *ECCV*, pages 683–700, 2020.
- [Liu *et al.*, 2015] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *ICCV*, pages 3730–3738, 2015.
- [Liu *et al.*, 2018] Guilin Liu, Fitsum A Reda, Kevin J Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image inpainting for irregular holes using partial convolutions. In *ECCV*, pages 85–100, 2018.
- [Liu *et al.*, 2020] Hongyu Liu, Bin Jiang, Yibing Song, Wei Huang, and Chao Yang. Rethinking image inpainting via a mutual encoder-decoder with feature equalizations. In *ECCV*, pages 725–741, 2020.
- [Nazeri *et al.*, 2019] Kamyar Nazeri, Eric Ng, Tony Joseph, Faisal Z Qureshi, and Mehran Ebrahimi. Edgeconnect: Generative image inpainting with adversarial edge learning. *arXiv preprint arXiv:1901.00212*, 2019.
- [Park *et al.*, 2019] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *CVPR*, pages 2337–2346, 2019.
- [Pathak *et al.*, 2016] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *CVPR*, pages 2536–2544, 2016.
- [Ren *et al.*, 2019] Yurui Ren, Xiaoming Yu, Ruonan Zhang, Thomas H Li, Shan Liu, and Ge Li. Structurflow: Image inpainting via structure-aware appearance flow. In *CVPR*, pages 181–190, 2019.
- [Wang *et al.*, 2019] Ning Wang, Jingyuan Li, Lefei Zhang, and Bo Du. Musical: Multi-scale image contextual attention learning for inpainting. In *IJCAI*, 2019.
- [Xie *et al.*, 2019] Chaohao Xie, Shaohui Liu, Chao Li, Ming-Ming Cheng, Wangmeng Zuo, Xiao Liu, Shilei Wen, and Errui Ding. Image inpainting with learnable bidirectional attention maps. In *ICCV*, 2019.
- [Yu *et al.*, 2018] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Generative image inpainting with contextual attention. In *CVPR*, pages 5505–5514, 2018.
- [Yu *et al.*, 2019] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Free-form image inpainting with gated convolution. In *CVPR*, 2019.
- [Yu *et al.*, 2020] Tao Yu, Zongyu Guo, Xin Jin, Shilin Wu, Zhibo Chen, Weiping Li, Zhizheng Zhang, and Sen Liu. Region normalization for image inpainting. In *AAAI*, pages 12733–12740, 2020.
- [Zhao *et al.*, 2020] Lei Zhao, Qihang Mo, Sihuan Lin, Zhizhong Wang, Zhiwen Zuo, Haibo Chen, Wei Xing, and Dongming Lu. Uctgan: Diverse image inpainting based on unsupervised cross-space translation. In *CVPR*, pages 5741–5750, 2020.
- [Zheng *et al.*, 2019] Chuanxia Zheng, Tat-Jen Cham, and Jianfei Cai. Pluralistic image completion. In *CVPR*, pages 1438–1447, 2019.
- [Zhou *et al.*, 2017] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(6):1452–1464, 2017.