

Learning Continuous Implicit Representation for Near-Periodic Patterns

Bowei Chen¹, Tiancheng Zhi¹, Martial Hebert¹, and Srinivasa G. Narasimhan¹

Carnegie Mellon University, Pittsburgh PA 15213, USA
{boweiche, tzhi, mhebert, srinivas}@andrew.cmu.edu

Abstract. Near-Periodic Patterns (NPP) are ubiquitous in man-made scenes and are composed of tiled motifs with appearance differences caused by lighting, defects, or design elements. A good NPP representation is useful for many applications including image completion, segmentation, and geometric remapping. But representing NPP is challenging because it needs to maintain global consistency (tiled motifs layout) while preserving local variations (appearance differences). Methods trained on general scenes using a large dataset or single-image optimization struggle to satisfy these constraints, while methods that explicitly model periodicity are not robust to periodicity detection errors. To address these challenges, we learn a neural implicit representation using a coordinate-based MLP with single image optimization. We design an input feature warping module and a periodicity-guided patch loss to handle both global consistency and local variations. To further improve the robustness, we introduce a periodicity proposal module to search and use multiple candidate periodicities in our pipeline. We demonstrate the effectiveness of our method on more than 500 images of building facades, friezes, wallpapers, ground, and Mondrian patterns in single and multi-planar scenes.

Keywords: Near-Periodic Patterns, Neural Implicit Representation, Single Image Optimization.

1 Introduction

Patterns are all around us and help us understand our visual world. In the 1990s, a human preattentive vision experiment [43] showed that periodicity is a crucial factor in high-level pattern perception. But most real patterns are not composed of perfectly periodic (tiled) motifs. Consider the commonly occurring real-world building facade scene in Figure 1 (a). While the windows are laid out periodically, they vary in their individual appearances. There are several design elements (borders, texture), shading variations, or obstructions (tree, car, street lamp) that are not periodic. These factors make it challenging to create a good computational representation for such “Near-Periodic Patterns” (NPP).

A good NPP representation must preserve both *global consistency* (similar motifs layout) and *local variations* (different appearances). For global consistency, the distance (periods) and orientations between adjacent motifs should

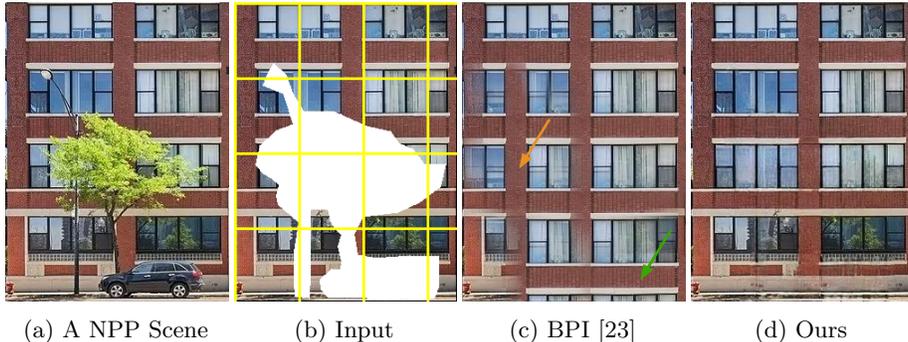


Fig. 1: Inpainting to remove the tree, street lamp, and car from a near-periodic patterned (NPP) scene in (a). Input image (b) visualizes the mask (white unknown region) and detected (but inaccurate) NPP representation (yellow lattice). Guided by this inaccurate representation, the state of the art method BPI [23] (c) fails to generate windows occluded by the tree (orange arrow) and the white strip across the bottom (green arrow). Our NPP-Net (d) maintains global consistency and local variations, while preserving the known regions.

be accurate (*e.g.*, window layout). At the same time, the local details in the scene should be fully encoded (*e.g.*, appearance variations in windows or the horizontal design strips). In this paper, we present a novel method to learn such a representation that can be used for applications such as image completion (our main focus), segmentation of periodic parts, and resolution enhanced scene remapping, *e.g.*, transforming to a fronto-parallel view (see supplementary).

Existing image completion works applicable for NPP can be classified into two categories. The first category does not explicitly consider knowledge of periodicity. They complete images by training on large datasets [36,58,22,46,57] or by exploring single image statistics [44,50,2]. However, these methods fail to generate good global consistency, especially with a large unknown mask inside (interpolation) or outside (extrapolation) the image border, or severe perspective effect. The other category [16,24,23,32] models periodicity as prior for image completion. These works extract explicit NPP representations (*e.g.*, displacement vectors) and use them to guide image completion. These methods can generate good global periodic structure *if* the estimated periodicity is accurate. However, this is hard to achieve in the presence of strong local variations.

Our work is inspired by the recent progress on implicit neural representations [35] that map image coordinates to RGB values using coordinate-based multi-layer perceptrons (MLP). But, naively using this method fails on our task due to the lack of a good periodicity prior. Thus, we present a periodicity-aware coordinate-based MLP to learn a continuous implicit neural representation, which we call NPP-Net for short. The key idea is to extract periodicity information from a partially observed NPP scene and inject it into both the MLP input and the loss function to help optimize NPP representation.

Three novel steps are proposed for the above idea: (1) The *Periodicity Proposal* step extracts periodicity in the form of a set of candidate periods and orientations that are used together to handle inaccurate detections; (2) The *Periodicity-Aware Input Warping* injects periodicity into the MLP input by warping input coordinates according to the proposed periodicities. This step preserves global consistency and the MLP converges to a good periodic pattern easily; (3) Finally, the *Periodicity-Guided Patch Loss* samples observable patches according to periodicity to optimize the representation. This step preserves local variations, improves extrapolation ability, and removes high-frequency artifacts.

Our approach only requires a single image for optimization. This is important since there are no large dedicated NPP datasets. Thus, we evaluate our approach on a total of 532 NPP subclasses chosen from three datasets [1,10,48]. The scenes include building facades, friezes, ground patterns, wallpapers, and Mondrian patterns that are tiled on one or more geometric planes and perspective warped. Our dataset is larger than those (157 at most) used in previous works [16,23,24,32] that are designed for NPP. We mainly apply NPP-Net for the image completion task, but extend it to resolution enhanced NPP remapping and NPP segmentation in the supplementary. We compare NPP-Net with four traditional [11,2,16,23] and five deep learning-based methods [50,44,58,57,46], and eight variants of NPP-Net. Experiments show that NPP-Net can interpolate and extrapolate images, in-paint large and arbitrarily shaped regions, recover blurry regions when images are remapped, segment periodic and non-periodic regions, in planar and multi-planar scenes. Figure 1 shows the effectiveness of NPP-Net, inpainting a complex NPP scene, compared to the state of the art BPI [23]. While our method is not designed for general scenes, it is a useful tool to understand a large class of man-made scenes with near-periodic patterns.

2 Related Work

Near-Periodic Patterns Completion: There are two types of image completion methods that can be applied to NPP. The first type of methods do not explicitly consider periodicity as prior for completion [36,58,22,46,57,51,6,63]. The second stream of methods takes advantage of periodicity to guide the completion. We focus on reviewing the second type of methods.

The first stage for these methods is to obtain an NPP representation to guide image completion. Existing methods aim to represent NPP by detecting the global periodicity despite local variations. The types of NPP arrangements vary [41,42,25,16,54,38] but commonly, periodic patterns are assumed to form a 2D lattice [27,26,14,21,39,37,24]. The first lattice-based work [14] for periodicity detection without human interaction finds correspondences using visual similarity and geometric consistency. Liu *et al.* [26] improve this process by incorporating generalized PatchMatch [3] and Markov Random Field. Furthermore, Lettry *et al.* [21] detect a repeated pattern model by searching in the feature space of a pre-trained CNN. Recently, Li *et al.* [24] design a compact strategy by searching on deep feature space without any implicit models. But it requires

hyperparameter tuning to achieve competitive results. All existing methods describe periodicity using an explicit representation such as keypoints [14,42,49,26], feature-based motifs [41] or displacement vectors [21,24]. But they do not preserve both global consistency and local variations well.

The second stage is to generate or inpaint an NPP image guided by the NPP representation [28,29,30,16,32,24,23,13]. One common assumption is that the NPP lie on a single plane. Liu *et al.* [29] synthesize an NPP image through multi-model deformation fields given an input NPP patch and its representation. Mao *et al.* [32] propose GAN-based NPP generation. Huang *et al.* [16] and BPI [23] extend image completion to the multi-plane case. They detect periodicities in these planes [15,31,24] and use them to guide image completion, based on [53] and [2]. Unlike earlier work Huang *et al.*, BPI uses the periodicity detection method based on feature maps extracted from a pre-trained network. Also, the state of the art BPI’s image completion step does not use their prior GAN-based method [32].

In summary, the above methods assume that NPP representation is good enough for guidance, which is not guaranteed. By contrast, we merge the two stages by optimizing the implicit representation using image reconstruction error.

Implicit Neural Representations: Recently, coordinate-based multi-layer perceptron (MLP) has been used to obtain implicit neural representation (INR). It maps coordinates to various signals such as shapes [12,9,44], scenes [35,33] and images [8,5,44]. Mildenhall *et al.* [35] represent a 3D scene from a sparse set of views for novel view synthesis. Siren [44] replaces ReLU by a periodic activation function and designs an initialization scheme for modeling finer details. Chen *et al.* [8] present a Local Implicit Image Function for the generation of arbitrary resolutions. Skorokhodov *et al.* [45] design a decoder based on INR with GAN training, for high-quality image generation.

NPP-Net differs from previous methods in two ways: (1) Directly using MLP [35,44] fails to learn accurate NPP representation without high-level structural understanding. We propose a periodicity-aware MLP. (2) Many works require a large dataset for training, while we optimize on a single image.

3 NPP-Net

We aim to build an MLP that maps image coordinates to pixel values, given a partial observation of an NPP image. We will describe NPP-Net using the image completion task. The unknown (masked) region is completed (or inpainted) by training on the remainder of the NPP image. For clarity, we first describe the method for single planar NPP scene and pre-warp the image to be fronto-parallel [60]. Then we will extend NPP-Net to handle multi-planar scenes.

Our key idea is to extract periodicity information from the known NPP region and inject it into the MLP input and the loss function. The initial pipeline of NPP-Net (Figure 2) consists of three modules: (1) Periodicity-Aware Input Warping transforms image coordinates using the detected periodicity. (2) Coordinate-based MLP maps the transformed coordinates to the corresponding

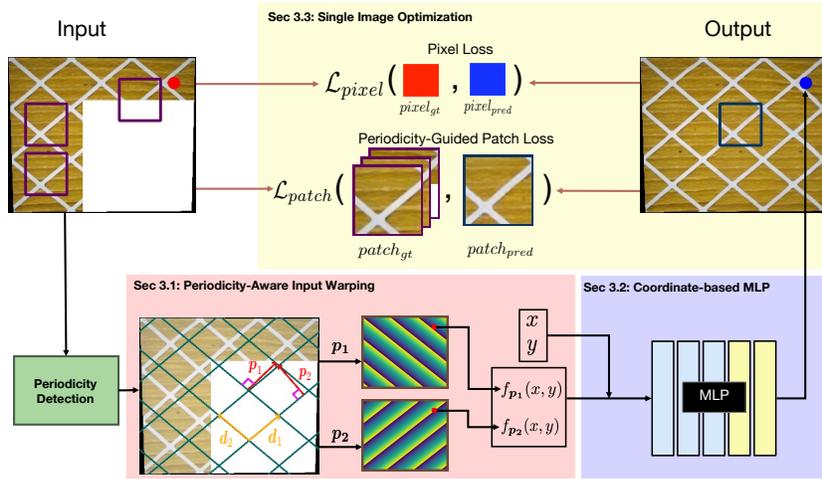


Fig. 2: Initial pipeline of NPP-Net consists of three modules. (1) *Periodicity-Aware Input Warping* (pink) warps input coordinates using detected periodicity. (2) *Coordinate-based MLP* (blue) maps warped and input coordinate features to an RGB value. (3) *Single Image Optimization* (yellow) uses pixel loss and periodicity-guided patch loss on a single NPP image. Final pipeline in section 3.4 shows how multiple periodicities are automatically detected and utilized.

RGB value. (3) Single Image Optimization provides a periodicity-guided loss function for optimizing the MLP on a single image.

3.1 Periodicity-Aware Input Warping

A traditional MLP is not good at capturing global periodic structure without additional priors. In fact, previous works [55,64] have shown that a traditional MLP is unable to extrapolate a 1D periodic signal even with many training samples. The Periodicity-Aware Input Warping module thus explicitly injects periodicity information into the MLP by warping image coordinates (x, y) .

Assuming a 2D lattice arrangement, the periodicity is represented as two displacement vectors \mathbf{d}_1 and \mathbf{d}_2 (orange arrows in Figure 2). A perfect infinite periodic pattern is invariant if shifted by $\alpha\mathbf{d}_1 + \beta\mathbf{d}_2$ ($\alpha, \beta \in \mathbb{Z}$). This representation can be transformed into periods and orientations, visualized as the magnitudes and orientations of the red arrows (\mathbf{p}_1 and \mathbf{p}_2 , called **periodicity vectors**). Mathematically, the transformation is obtained by solving $\mathbf{p}_1 \cdot \mathbf{d}_2 = \mathbf{d}_1 \cdot \mathbf{p}_2 = 0$ and $\mathbf{p}_1 \times \mathbf{d}_2 = \mathbf{d}_1 \times \mathbf{p}_2 = \mathbf{d}_1 \times \mathbf{d}_2$, where the cross product is defined using the corresponding 3D vectors on the plane $z = 0$. A **periodicity** is then defined as a vector pair $(\mathbf{p}_1, \mathbf{p}_2)$. Extension to circular patterns is available in supplementary.

One way to obtain the periodicity for an NPP image is to treat it as a learnable parameter and jointly optimize it with NPP-Net [7,17]. However, this is hard for two reasons. (1) Good periodicity is not unique (any multiple works). (2)

Many real-world NPP scenes contain strong local variations, leading to a complicated cost function. Thus we adopt an existing periodicity detection method [24] for input warping, which extracts feature maps from a pretrained CNN and performs brute force search to obtain a periodicity. Then, for each periodicity vector $\mathbf{p} = (p \cos \theta, p \sin \theta) \in \{\mathbf{p}_1, \mathbf{p}_2\}$, we define a warp as a bivariate function:

$$f_{\mathbf{p}}(x, y) = (x \cos \theta + y \sin \theta) \bmod p. \quad (1)$$

This function generates a warped coordinate value sampled from a periodic pattern with period $|\mathbf{p}|$ along direction $\frac{\mathbf{p}}{|\mathbf{p}|}$, as shown in Figure 2. Through this feature engineering, the warped coordinates explicitly encode the periodicity information. The warped coordinates $f_{\mathbf{p}_1}(x, y)$ and $f_{\mathbf{p}_2}(x, y)$, together with the original coordinates x and y , are further normalized to $[-1, 1]$ and passed through positional encoding [35] to allow the network to model high frequency signals [47]. The encoded coordinates are then input to the MLP. We keep the notations of coordinates before and after positional encoding the same for simplicity. The dimension of the features is $4d$, including d frequencies in the positional encoding, and a set of four values for each frequency: $x, y, f_{\mathbf{p}_1}(x, y)$, and $f_{\mathbf{p}_2}(x, y)$.

3.2 Coordinate-Based MLP

We adopt coordinate-based MLP to represent NPP images. It is more effective and compact than a CNN to model periodic signals since coordinates are naturally suited for encoding positional (periodic) information. Specifically, we input the warped coordinate features to enforce global consistency, and also input the original coordinate features without warping to help preserve local variations. The output of the MLP is an RGB value corresponding to the input image coordinate. Since ReLU activation function has been proven to be ineffective to extrapolate periodic signals [55], we use the more suitable SNAKE function [64].

3.3 Single Image Optimization

Pixel Loss: Pixel loss is the most intuitive way to optimize coordinate-based MLP [35], which compares predicted and ground truth pixel values. For image coordinate $\mathbf{x} = (x, y)$, we adopt the robust loss function \mathcal{L}_{rob} [4], given by:

$$\mathcal{L}_{pixel}(\mathbf{x}) = \mathcal{L}_{rob}(\hat{C}(\mathbf{x}), C(\mathbf{x})), \quad (2)$$

where $\hat{C}(\mathbf{x})$ and $C(\mathbf{x})$ are the output RGB values of the MLP and the ground truth RGB values from the input image at position \mathbf{x} , respectively. This loss is applied only to the known regions.

But merely adopting pixel loss like NeRF [35] fails to generate a good NPP for two reasons: (1) The high-dimensional input features result in the generation of high-frequency artifacts (Figure 3 (b)). See [47,61] for details about this problem. (2) Pixel loss does not enforce explicit constraints to model the correlation between a coordinate’s features and its neighbors. This constraint is critical for

preserving local variations since it helps capture local patch statistics. Thus, pixel loss fails to preserve local variations. For example, in Figure 3 (b), pixel loss generates some periodic artifacts in the top non-periodic region¹.

Periodicity-Guided Patch Loss: To address the limitations of pixel loss, we force the network to learn patch internal statistics by incorporating patch loss, which compares predicted and ground truth patches. The ground truth patches can be sampled at the same position as the predicted patch (for known regions), or sampled according to periodicity (for any region).

GT Patch at the Same Position: For a predicted patch in the known region, the ground truth patch at the same position is available. Specifically, for a square patch with size s centered at position \mathbf{x} , we input all the pixel coordinates in the patch into MLP to obtain a predicted RGB patch $\hat{I}_s(\mathbf{x})$. Let $I_s(\mathbf{x})$ be the corresponding ground truth at the same position and $M_s(\mathbf{x})$ be the mask of known pixels. We apply perceptual loss [59] on masked patches:

$$\mathcal{L}_p(\mathbf{x}) = \mathcal{L}_{pct}(\hat{I}_s(\mathbf{x}) \odot M_s(\mathbf{x}), I_s(\mathbf{x}) \odot M_s(\mathbf{x})), \quad (3)$$

where \odot is the element-wise product.

GT Patches Sampled Based on Periodicity: To train on unknown regions, we propose to sample ground truth patches based on periodicity. This is an effective way to handle the MLP extrapolation problem, which cannot be solved by merely using input warping and SNAKE activation function [64]. The input and output images in Figure 2 illustrate this sampling strategy.

Specifically, we sample multiple nearby ground truth patches for supervision by shifting position \mathbf{x} based on the estimated periodicity. The shifted patch center is defined as $\mathbf{x}_{\alpha\beta} = \mathbf{x} + \alpha\mathbf{d}_1 + \beta\mathbf{d}_2$ ($\alpha, \beta \in \mathbb{Z}$), where \mathbf{d}_1 and \mathbf{d}_2 are the displacement vectors. Because the predicted and ground truth patches are not necessarily aligned, we adopt contextual loss \mathcal{L}_{ctx} [34]:

$$\mathcal{L}_c(\mathbf{x}) = \frac{1}{N} \sum_{(\alpha, \beta) \in T_N} \mathcal{L}_{ctx}(\hat{I}_s(\mathbf{x}) \odot M_s(\mathbf{x}_{\alpha\beta}), I_s(\mathbf{x}_{\alpha\beta}) \odot M_s(\mathbf{x}_{\alpha\beta})), \quad (4)$$

where T_N is a set of (α, β) pairs corresponding to the N nearest ground truth patches, since local variations are preserved using nearby patches for supervision.

Patch Loss: Our patch loss combines the two sampling strategies: $\mathcal{L}_{patch}(\mathbf{x}) = \lambda_p \gamma(\mathbf{x}) \mathcal{L}_p(\mathbf{x}) + \lambda_c \mathcal{L}_c(\mathbf{x})$, where λ_p and λ_c are constant weights. $\gamma(\mathbf{x})$ is a binary function: 1 for \mathbf{x} in known regions and 0 for \mathbf{x} in unknown regions.

Total Loss: Our final loss is the combination of patch loss and pixel loss:

$$\mathcal{L} = \frac{\lambda_1}{|B_1|} \sum_{\mathbf{x} \in B_1} \mathcal{L}_{pixel}(\mathbf{x}) + \frac{\lambda_2}{|B_2|} \sum_{\mathbf{x} \in B_2} \mathcal{L}_{patch}(\mathbf{x}), \quad (5)$$

where λ_1 and λ_2 are constant weights. B_1 contains pixel coordinates that are randomly sampled in known areas, and B_2 contains the center coordinates sampled in both known and unknown areas in proportion.

¹ Figure 3 is generated based on the final pipeline explained in Section 3.4.

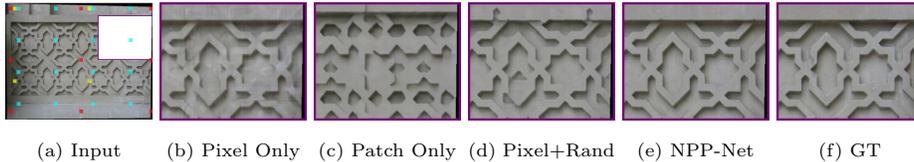


Fig. 3: Comparing different losses based on the final pipeline. The red, yellow and cyan dots in (a) visualize the Top-3 periodicities. Zoom-ins of the unknown area (white rectangle) are in (b)-(f). Merely using pixel loss (b) generates high-frequency artifacts across the image and periodic artifacts in the top part. Adopting only patch loss (c) removes the artifacts but has poor global structure. Using pixel loss and patch loss with random sampling (d) cannot preserve global consistency and local variations well since the ground truth patches are not sampled according to periodicity and might be far from the predicted patch. With pixel loss and periodicity-guided patch loss, NPP-Net (e) solves these issues.

Training with this loss preserves both global consistency and local variations, as shown in Figure 3 (e). In fact, only using patch loss cannot ensure global consistency if the detected periodicity is not accurate enough. In Figure 3 (c), the pattern structure is poorly reconstructed because it only focuses on the local structure. We also show the result for the combination of pixel loss and patch loss with patches that are randomly sampled in the known regions (we call it random sampling strategy) in Figure 3 (d). This fails to generate correct periodic patterns and good local details because the output and sampled patches have a large misalignment and are far away from each other.

3.4 Periodicity Proposal

Although the above initial pipeline shows good performance, it still fails to handle very inaccurate periodicity detection. To improve the robustness of NPP-Net, we design a Periodicity Proposal module to provide additional periodicity information. As shown in Figure 4, we first search multiple candidate periodicities and then augment the input to MLP to handle inaccurate periodicity detection.

Periodicity Searching: Our searching strategy is based on the same periodicity detection method [24] we adopt in the initial pipeline. But the authors’ original implementation requires manual hyperparameter tuning. Instead, we design an automatic tuning method, which evaluates each candidate periodicity (obtained from various hyperparameters) in the context of image completion. Specifically, we first generate M pseudo masks in the known regions and treat them as unknown masks for image completion. Then we execute the initial pipeline for each candidate periodicity, and compute its reconstruction error in pseudo mask regions for periodicity ranking. Since we focus on reconstructing a coarse global structure, we use a lightweight initial NPP-Net without patch loss for efficiency, which takes around 10 seconds for each periodicity in a Titan Xp GPU.

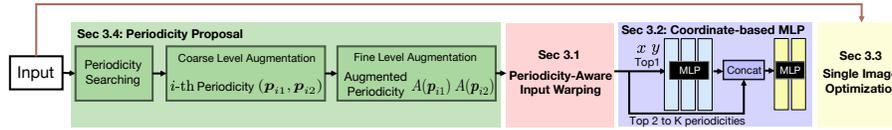


Fig. 4: Final pipeline of NPP-Net modifies two modules of the initial pipeline. (1) *Periodicity Proposal* (green) automatically searches and augments the input periodicity to handle inaccurate periodicity detection and encourage the global consistency. (2) *Coordinate-based MLP* (lavender blue) has two branches: (a) for Top-1 periodicity and original coordinates, and (b) for the rest.

Periodicity Augmentation: Prior methods [16,32,24,23] also use one periodicity to guide completion as in our initial pipeline. This cannot guarantee global consistency if the periodicity is inaccurate, especially when the unknown mask is large (see experiments). So, we augment the pattern periodicity at two levels to improve robustness. At the coarse level, instead of searching the best periodicity, we keep Top-K candidates $\{(\mathbf{p}_{i1}, \mathbf{p}_{i2}) \mid i \in \mathbb{Z}^+, i \leq K\}$, to cover multiple possible solutions. This coarse-level augmentation encourages NPP-Net to move towards the most reasonable candidate periodicity. At the fine level, we augment periodicities with small offsets to better handle smaller errors. Specifically, a periodicity vector \mathbf{p} is augmented to be $A(\mathbf{p}) = \{\mathbf{p} + \delta \frac{\mathbf{p}}{|\mathbf{p}|} \mid \delta \in \Delta\}$. We empirically define $\Delta = \{0, \pm 0.5, \pm 1\}$ (in pixels). Finally, we merge all the augmented periodicity vectors as $P = \bigcup_{i \in \mathbb{Z}^+, i \leq K} A(\mathbf{p}_{i1}) \cup A(\mathbf{p}_{i2})$. Note that $|P| = 2K|\Delta|$.

In our final pipeline, P contains $K|\Delta|$ periodicities. We perform input warping (Section 3.1) for each periodicity and input the transformed coordinate features into the MLP. We add an additional branch to the MLP, as shown in Figure 4. Since the Top-1 periodicity is likely the most accurate (see experiments), we input the coordinate features warped using the Top-1 periodicity (including fine-level augmentation) and original coordinate features, to the first branch. The coordinate features warped using the Top-2 to K periodicities are sent to the second branch. For optimization, we sample patches according to the Top-1 periodicity. All other parts remain the same in our final pipeline. We evaluate these changes in our ablation study. See supplementary for implementation details including hyperparameters, network architecture, and runtime.

3.5 Extensions

Non-NPP region segmentation: Parts of a scene may not be near-periodic (*e.g.*, trees in front of a building facade). We thus segment the non-periodic regions in an NPP image in an unsupervised manner. We use a traditional segmentation method [18] to provide an initial guess for the non-periodic regions, treated as the unknown mask in image completion. After training NPP-Net, we relabel the initial non-periodic regions with low reconstruction error as periodic regions. *Similar strategy is adopted to serve as a pre-filtering step before applying our method to any arbitrary scene (see supplementary).*

NPP remapping: NPP scenes captured from a tilted angle can result in blurry motifs after rectification. To enhance resolution, we detect blurry regions and treat them as the unknown mask in image completion. The difference is that we compute the pixel loss in the blurry regions with a smaller weight.

Multi-Plane NPP completion: Given an image with different NPPs on different planes, we first adopt a pre-trained plane segmentation network [56] to obtain a coarse plane segmentation. Then we select a bounding box in each plane as a reference to rectify the plane using TILT [60]. Note that, we do not require accurate segmentation since it is only used for bounding box selection. For each rectified plane, we first use our NPP segmentation method to segment the non-periodic regions (mainly from other planes) and treat them as invalid pixels. Then we perform NPP completion on each plane, transform it back to the original image coordinate system, and recombine the image. Figure 6 shows qualitative results for this extension. *Detailed implementation and experiments for these extensions are in supplementary.*

4 Experiments

Dataset: We evaluate NPP-Net on 532 images selected from three relevant datasets for NPP completion: PSU Near-Regular Texture Database (NRTDB) [1], Describable Textures Dataset (DTD) [10], and Facade Dataset [48]. There are 165 NPP images in the NRTDB dataset including facades, friezes, bricks, fences, grounds, Mondrian images, wallpapers, and carpets. Similarly, there are 258 NPP images in the DTD Dataset including honeycombs, grids, meshes and dots. The Facade Dataset has 109 rectified images of facades. Some of these facades are strictly *not* NPP because often the windows are not arranged periodically. But nonetheless we include these to evaluate our approach when the NPP assumption is not strictly satisfied. Finally, we also collect a small dataset with 11 NPP images for real-world applications (*e.g.*, removing trees in the scene). In general, scenes in NRTDB are more challenging than DTD since they contain more non-periodic regions (boundaries, trees, sky, etc.), complex illuminations and backgrounds, and multiple periodicities across an image (Figure 5 row 3). We use TILT [60] to rectify all the images to be fronto-parallel if needed. See supplementary for more details including sampled images and mask generation.

Metrics: No single metric can evaluate NPP image completion comprehensively. So we adopt three metrics to cover different scales, including LPIPS (perceptual distance) [59], SSIM [52], and PSNR. Lower LPIPS, higher SSIM, and higher PSNR mean better performance. A known limitation for SSIM and PSNR is that blurry images also tend to receive high scores in these metrics [20], while LPIPS handles this issue better. See supplementary for FID [40] and RMSE metrics.

Ablation Study: We perform three studies. First, we compare to a “No Periodicity” variant, which uses a standard coordinated-based MLP without a periodicity prior. Results in Table 1 show that it fails to understand the arrangement of tiled motifs. Note that the Facade dataset may have different performances because it has some non-NPP images. Second, to study the loss functions, we

Category	Method	NRTDB [1]			DTD [10]			Facade [48]		
		LPIPS ↓	SSIM ↑	PSNR ↑	LPIPS ↓	SSIM ↑	PSNR ↑	LPIPS ↓	SSIM ↑	PSNR ↑
Large Datasets	PEN-Net [57]	0.497	0.452	17.97	0.473	0.365	15.81	0.426	0.444	<u>15.78</u>
	ProFill [58]	0.401	0.300	16.35	0.443	0.249	14.30	0.374	0.391	14.73
	Lama [46]	0.196	0.551	18.64	0.274	0.479	16.39	0.207	0.468	15.24
Single Image	Image Quilting [11]	0.428	0.074	13.25	0.415	0.077	12.18	0.550	0.002	10.28
	PatchMatch [2]	0.263	0.542	18.14	0.361	0.383	15.47	0.369	0.341	14.22
	DIP [50]	0.554	0.292	16.46	0.659	0.181	13.15	0.582	0.258	15.22
	Siren [44]	0.636	0.084	14.38	0.762	0.080	13.11	0.780	0.052	12.00
	Huang <i>et al.</i> [16]	0.287	0.410	16.99	0.302	0.320	14.88	0.387	0.279	13.75
	BPI [23]	0.254	0.442	16.86	0.303	0.305	14.82	0.458	0.173	12.20
NPP-Net Variants	No Periodicity	0.429	0.449	18.22	0.468	0.350	16.38	0.379	0.443	15.54
	Pixel Only	0.308	0.618	20.20	0.397	0.473	17.86	0.427	0.458	15.54
	Patch Only	0.322	0.340	15.83	0.395	0.275	13.22	0.412	0.164	11.99
	Pixel + Random	0.216	0.670	20.74	0.264	0.501	18.03	0.316	0.426	15.43
	Initial Pipeline	0.213	0.647	20.38	0.293	0.462	17.26	0.277	0.438	15.11
	Top1 + Offsets	0.205	0.656	20.60	0.285	0.477	17.70	0.289	0.412	14.80
	Top5 + Offsets	0.197	0.661	20.73	0.259	0.492	18.03	0.265	0.483	15.51
	Top3 w/o Offsets	0.210	0.648	20.44	0.275	0.474	17.36	0.269	0.460	15.33
NPP-Net	Top3 + Offsets	0.188	0.679	21.01	0.249	0.504	18.32	<u>0.263</u>	0.485	15.93

Table 1: Comparison with baselines and NPP-Net variants for NPP completion and the metrics are evaluated only in unknown regions. The best and second-best results (excluding variants) are highlighted in bold and underline respectively. NPP-Net outperforms all the baselines on NRTDB and DTD. While Facade has some non-NPP images, NPP-Net can still outperform all other baselines except for Lama. See the supplementary for the results tested on the full images.

design three variants: (1) Only pixel loss, (2) Only patch loss, (3) Pixel loss and patch loss with random patch sampling. As discussed in Section 3.3, the results in Table 1 and Figure 3 show that NPP-Net outperforms other variants.

Third, we study the effect of the periodicity augmentation (coarse level Top-K candidate periodicities and fine level offsets) by testing four variants: (1) Initial pipeline (No augmentation), (2) Top-1 with offsets, (3) Top-5 with offsets, (4) Top-3 without offsets. Table 1 shows that the initial pipeline performs the worst. Larger K (Top-5) hurts the performance as more inaccurate periodicities may be included. But, smaller K (Top-1) also performs badly because the correct periodicity may not be included. A suitable K (Top-3) without offsets performs worse since offsets can help better handle smaller errors. With the appropriate K and offsets, NPP-Net generates the best results. See more studies in supplementary. **Baselines:** We compare against non-periodicity-guided and periodicity-guided baselines. For the former, we select two traditional methods, Image Quilting [11] and PatchMatch [2], that can handle pattern structure locally for some scenes with properly selected patch size. We then consider two learning-based methods DIP [50] (CNN-based) and Siren [44] (MLP-based) trained on a single image for inpainting. We also choose several learning-based methods: PEN-Net [57], ProFill [58] and Lama [46] that are trained on large real-world datasets [62,10,18,19] since they show competitive NPP completion examples in their work.

For periodicity-guided methods, we choose two baselines - Huang *et al.* [16] and BPI [23]. Both works were designed for multi-planar scenes, but can be

used for single-plane completion as well. BPI first segments and rectifies planes, then performs periodicity detection [24] on each plane, and inpaints each plane independently. For fair comparison in single-plane NPP image, we only compare with BPI’s completion step to remove potential inconsistency introduced from other steps (*e.g.*, plane rectification). For Huang *et al.*, we use their pipeline without modification as their method works directly for a single plane and the completion step cannot be easily separated out. We will also compare with these two methods in multi-plane NPP images.

Comparison with Baselines: Table 1 shows the quantitative results for all methods. For NRTDB and DTD datasets, among the single-image baselines, BPI obtains the almost best LPIPS because it generates a more reasonable global structure guided by periodicity. PatchMatch obtains better SSIM and PSNR even if it generates blurred results for large masks. Lama achieves the best results among the baselines since it adopts fast Fourier convolution for a larger image receptive field, which allows it to implicitly learn the underlying periodicity from large datasets. Our NPP-Net outperforms all baselines on these two datasets by optimizing only on a single image. For the Facade dataset, even if some non-NPP images are included, NPP-Net performs better than other baselines (except for Lama). Lama effectively learns scene prior from large datasets and thus works well for non-NPP images, leading to the best performance in this dataset.

Qualitative results are shown in Figure 5. Large rectangle masks are challenging since there is less information from which to estimate the representation. Perceptually, PatchMatch works well when the motifs are small (row 3) but results in blur with large masks (row 2 and 6). Although BPI and Huang *et al.* perform better than non-periodicity-guided baselines, they generate artifacts since the NPP representation (periodicity) has poor global consistency (row 1-6) or lacks local variations (row 6-9). Note that the Top-1 periodicity in row 1 (red dots) is inaccurate - the actual periods are half of the one shown. We show that NPP-Net can extrapolate NPP images well (row 5), generalize to irregular masks (row 6), and work for scenes that contain non-periodic regions (row 4). We show that NPP-Net can be extended to handle multi-planar scenes in Figure 6. Among the baselines, Lama (trained on large datasets) can better handle local variations (row 2 purple box). Although it captures some global structure when the mask is not large, Lama performs worse than BPI when the mask is outside the image border for extrapolation (all rows) and perspective effect is severe (row 2 cyan box). Learning from the Top-K periodicities, NPP-Net produces the best images, maintaining global consistency and local variations. Finally, the output of NPP-Net can in turn improve periodicity detection, leading to better image completion for both BPI and NPP-Net (see supplementary).

Influence of Mask Size: We study the influence of different mask sizes for image completion. All figures are shown in the supplementary due to the space limitation. First, for each image, if the K-th periodicity has the smallest error among Top-3 periodicities, we assign the image to the K-th periodicity. In the supplementary, we show the number of images assigned to each periodicity with different mask sizes in Figure 19, and illustrate how to compute periodicity

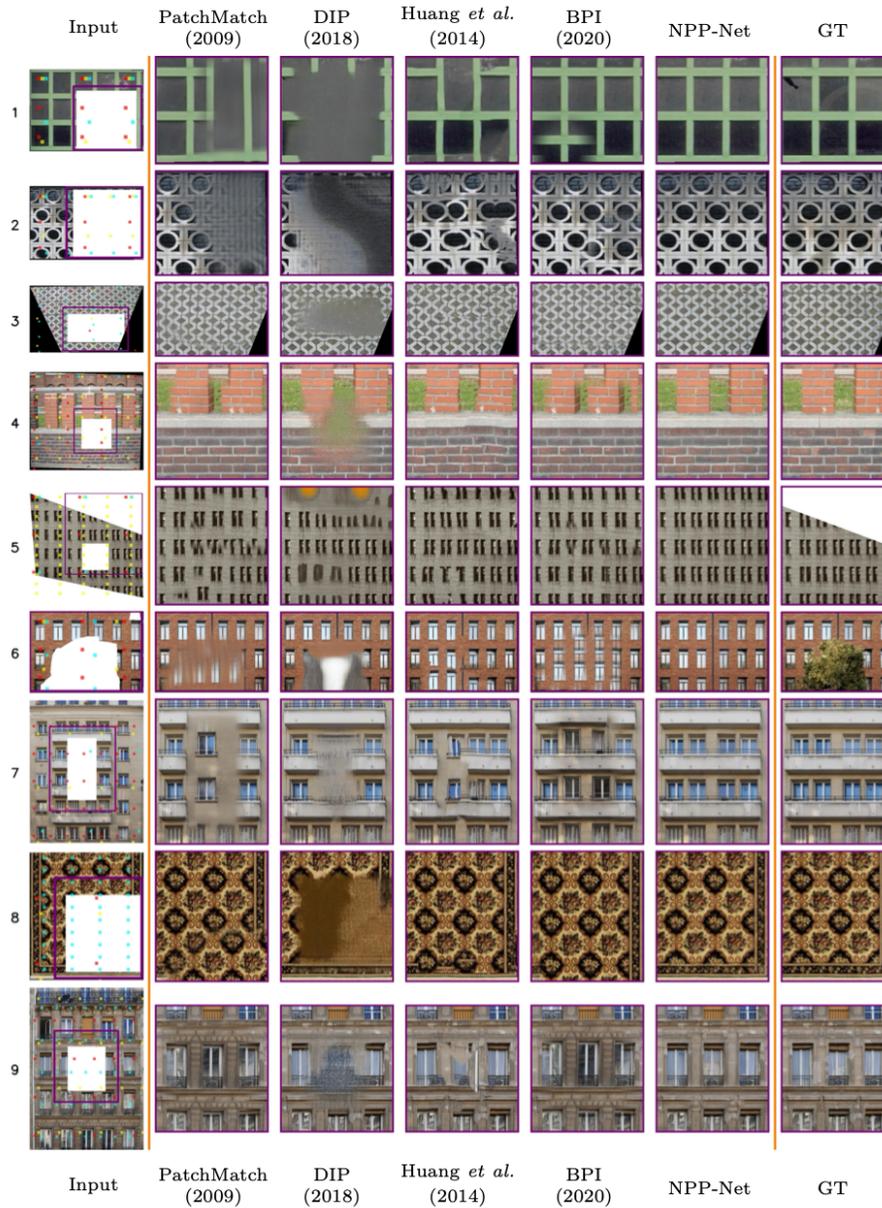


Fig. 5: Qualitative results for NPP completion. We show four baselines that operate on a single image. The red, yellow, and cyan dots in input images show the first, second, and third periodicity from periodicity searching module, respectively. For visualization, all periods are scaled by 2. Our NPP-Net outperforms all baselines for global consistency (rows 1-6) and local variations (rows 6-9).

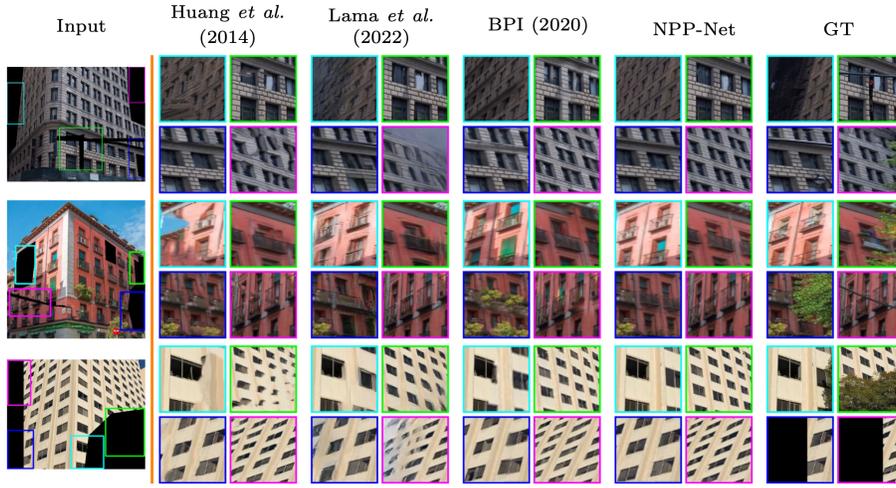


Fig. 6: Qualitative comparison for multi-plane NPP completion. We show three baselines, which are either designed for multi-plane NPP scenes (Huang *et al.* and BPI) or trained on large datasets (Lama). Some zoom-in boxes are resized for visualization. Full results are in supplementary.

error in detail. From the figure, while the Top 1st periodicity is the best one for most of the images with small mask size, the number decreases when the mask is large (64% of the image). This illustrates that the other two periodicities contain better periodicity and leveraging them by our periodicity augmentation strategy can be helpful for learning NPP representation, especially when the mask is large. Second, we show the performances for different mask sizes in Figure 20 and 21 (supplementary). In particular, we filter out images that contain large non-periodic regions. When the mask area is small (4% of the image), PatchMatch slightly outperforms NPP-Net because the unknown regions may not contain pattern structure, and simply sampling nearby patches is sufficient to produce good results. Among single-image methods, Huang *et al.*, BPI and NPP-Net perform better when the mask size increases since they are guided by periodicity. Taking better advantage of periodicity information, NPP-Net is more robust to various mask sizes, especially for larger masks.

Limitations: (1) The periodicity proposal cannot be too erroneous, allowing tolerance of about 10%. (2) It assumes a multi-planar scene with translated, circular, and potentially other types of symmetrical NPP that can be modeled.

Conclusion: We show how to learn an effective neural representation of Near-Periodic Patterns. We design the periodicity proposal, periodicity-aware input warping, periodicity-guided loss to maintain global consistency and local variations. NPP-Net provides a strong tool to study a large class of man-made scenes.

Acknowledgement: This work was supported by a gift from Zillow Group, USA, and NSF Grants #CNS-2038612, #IIS-1900821.

References

1. PSU Near-Regular Texture Database. <http://vivid.cse.psu.edu/>
2. Barnes, C., Shechtman, E., Finkelstein, A., Goldman, D.B.: Patchmatch: A randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.* p. 24 (2009)
3. Barnes, C., Shechtman, E., Goldman, D.B., Finkelstein, A.: The generalized patchmatch correspondence algorithm. In: *European Conference on Computer Vision (ECCV)*. pp. 29–43. Springer (2010)
4. Barron, J.T.: A general and adaptive robust loss function. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 4331–4339 (2019)
5. Bemana, M., Myszkowski, K., Seidel, H.P., Ritschel, T.: X-fields: Implicit neural view-, light-and time-image interpolation. *ACM Transactions on Graphics (TOG)* pp. 1–15 (2020)
6. Cao, C., Fu, Y.: Learning a sketch tensor space for image inpainting of man-made scenes. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 14509–14518 (2021)
7. Chen, H., Liu, J., Chen, W., Liu, S., Zhao, Y.: Exemplar-based pattern synthesis with implicit periodic field network. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 3708–3717 (2022)
8. Chen, Y., Liu, S., Wang, X.: Learning continuous image representation with local implicit image function. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 8628–8638. IEEE (2021)
9. Chen, Z., Zhang, H.: Learning implicit fields for generative shape modeling. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 5939–5948. IEEE (2019)
10. Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., Vedaldi, A.: Describing textures in the wild. In: *Proceedings of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)* (2014)
11. Efros, A.A., Freeman, W.T.: Image quilting for texture synthesis and transfer. In: *International Conference on Computer Graphics and Interactive Techniques (SIGGRAPH)*. p. 341–346. Association for Computing Machinery (2001)
12. Genova, K., Cole, F., Vlasic, D., Sarna, A., Freeman, W.T., Funkhouser, T.: Learning shape templates with structured implicit functions. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 7154–7164. IEEE (2019)
13. Halperin, T., Hakim, H., Vantzos, O., Hochman, G., Benaim, N., Sassy, L., Kupchik, M., Bibi, O., Fried, O.: Endless loops: Detecting and animating periodic patterns in still images. *ACM Transactions on graphics (TOG)* (2021)
14. Hays, J., Leordeanu, M., Efros, A.A., Liu, Y.: Discovering texture regularity as a higher-order correspondence problem. In: *European Conference on Computer Vision (ECCV)*. pp. 522–535. Springer (2006)
15. He, K., Sun, J.: Statistics of patch offsets for image completion. In: *European conference on computer vision*. pp. 16–29. Springer (2012)
16. Huang, J.B., Kang, S.B., Ahuja, N., Kopf, J.: Image completion using planar structure guidance. *ACM Transactions on graphics (TOG)* pp. 1–10 (2014)
17. Jetchev, N., Bergmann, U., Vollgraf, R.: Texture synthesis with spatial generative adversarial networks. *arXiv preprint arXiv:1611.08207* (2016)
18. Jiri, B., Jan, S., Jan, K., Habart, D.: Supervised and unsupervised segmentation using superpixels, model estimation, and graph cut. *Journal of Electronic Imaging* (2017)

19. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of gans for improved quality, stability, and variation. arXiv preprint arXiv:1710.10196 (2017)
20. Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., et al.: Photo-realistic single image super-resolution using a generative adversarial network. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4681–4690 (2017)
21. Lettry, L., Perdoch, M., Vanhoey, K., Van Gool, L.: Repeated pattern detection using cnn activations. In: IEEE Winter Conference on Applications of Computer Vision (WACV). pp. 47–55. IEEE (2017)
22. Li, J., Wang, N., Zhang, L., Du, B., Tao, D.: Recurrent feature reasoning for image inpainting. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE (2020)
23. Li, Y., Mao, J., Zhang, X., Freeman, W.T., Tenenbaum, J.B., Snavely, N., Wu, J.: Multi-plane program induction with 3d box priors. In: Neural Information Processing Systems (NeurIPS) (2020)
24. Li, Y., Mao, J., Zhang, X., Freeman, W.T., Tenenbaum, J.B., Wu, J.: Perspective plane program induction from a single image. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4434–4443. IEEE (2020)
25. Liu, J., Liu, Y.: Grasp recurring patterns from a single view. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2003–2010. IEEE (2013)
26. Liu, S., Ng, T.T., Sunkavalli, K., Do, M.N., Shechtman, E., Carr, N.: Patchmatch-based automatic lattice detection for near-regular textures. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 181–189. IEEE (2015)
27. Liu, Y., Collins, R.T., Tsin, Y.: A computational model for periodic pattern perception based on frieze and wallpaper groups. IEEE transactions on pattern analysis and machine intelligence (TPAMI) pp. 354–371 (2004)
28. Liu, Y., Lin, W.C.: Deformable texture: the irregular-regular-irregular cycle. Carnegie Mellon University, the Robotics Institute (2003)
29. Liu, Y., Lin, W.C., Hays, J.: Near-regular texture analysis and manipulation. ACM Transactions on Graphics (TOG) pp. 368–376 (2004)
30. Liu, Y., Tsin, Y., Lin, W.C.: The promise and perils of near-regular texture. International Journal of Computer Vision (IJCV) pp. 145–159 (2005)
31. Lowe, D.G.: Object recognition from local scale-invariant features. In: Proceedings of the IEEE international conference on computer vision (ICCV). pp. 1150–1157. IEEE (1999)
32. Mao, J., Zhang, X., Li, Y., Freeman, W.T., Tenenbaum, J.B., Wu, J.: Program-guided image manipulators. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4030–4039. IEEE (2019)
33. Martin-Brualla, R., Radwan, N., Sajjadi, M.S., Barron, J.T., Dosovitskiy, A., Duckworth, D.: Nerf in the wild: Neural radiance fields for unconstrained photo collections. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7210–7219. IEEE (2021)
34. Mechrez, R., Talmi, I., Zelnik-Manor, L.: The contextual loss for image transformation with non-aligned data. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 768–783 (2018)
35. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: European conference on computer vision (ECCV). pp. 405–421. Springer (2020)

36. Nazeri, K., Ng, E., Joseph, T., Qureshi, F., Ebrahimi, M.: Edgeconnect: Structure guided image inpainting using edge prediction. In: The IEEE International Conference on Computer Vision (ICCV) Workshops. IEEE (2019)
37. Park, M., Brocklehurst, K., Collins, R.T., Liu, Y.: Deformed lattice detection in real-world images using mean-shift belief propagation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* pp. 1804–1816 (2009)
38. Park, M., Brocklehurst, K., Collins, R.T., Liu, Y.: Translation-symmetry-based perceptual grouping with applications to urban scenes. In: Asian conference on computer vision (ACCV). pp. 329–342. Springer (2010)
39. Park, M., Collins, R.T., Liu, Y.: Deformed lattice discovery via efficient mean-shift belief propagation. In: European Conference on Computer Vision (ECCV). pp. 474–485. Springer (2008)
40. Parmar, G., Zhang, R., Zhu, J.Y.: On aliased resizing and surprising subtleties in gan evaluation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2022)
41. Pritts, J., Chum, O., Matas, J.: Detection, rectification and segmentation of coplanar repeated patterns. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2973–2980. IEEE (2014)
42. Pritts, J., Rozumnyi, D., Kumar, M.P., Chum, O.: Coplanar repeats by energy minimization. In: Proceedings of the British Machine Vision Conference (BMVC). pp. 107.1–107.12. IEEE (2016)
43. Rao, A., Lohse, G.: Identifying high level features of texture perception. *CVGIP: Graphical Models and Image Processing* pp. 218–233 (1993)
44. Sitzmann, V., Martel, J., Bergman, A., Lindell, D., Wetzstein, G.: Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems* (2020)
45. Skorokhodov, I., Ignatyev, S., Elhoseiny, M.: Adversarial generation of continuous images. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10753–10764. IEEE (2021)
46. Suvorov, R., Logacheva, E., Mashikhin, A., Remizova, A., Ashukha, A., Silvestrov, A., Kong, N., Goka, H., Park, K., Lempitsky, V.: Resolution-robust large mask inpainting with fourier convolutions. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 2149–2159 (2022)
47. Tancik, M., Srinivasan, P.P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J.T., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains. In: Neural Information Processing Systems (NeurIPS) (2020)
48. Teboul, O., Simon, L., Koutsourakis, P., Paragios, N.: Segmentation of building facades using procedural shape priors. In: 2010 IEEE computer society conference on computer vision and pattern recognition. pp. 3105–3112. IEEE (2010)
49. Torii, A., Sivic, J., Pajdla, T., Okutomi, M.: Visual place recognition with repetitive structures. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 883–890. IEEE (2013)
50. Ulyanov, D., Vedaldi, A., Lempitsky, V.: Deep image prior. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9446–9454. IEEE (2018)
51. Wang, T., Ouyang, H., Chen, Q.: Image inpainting with external-internal learning and monochromic bottleneck. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5120–5129 (2021)
52. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* pp. 600–612 (2004)

53. Wexler, Y., Shechtman, E., Irani, M.: Space-time completion of video. *IEEE Transactions on pattern analysis and machine intelligence* pp. 463–476 (2007)
54. Wu, C., Frahm, J.M., Pollefeys, M.: Detecting large repetitive structures with salient boundaries. In: *European conference on computer vision (ECCV)*. pp. 142–155. Springer (2010)
55. Xu, K., Zhang, M., Li, J., Du, S.S., Kawarabayashi, K.I., Jegelka, S.: How neural networks extrapolate: From feedforward to graph neural networks. In: *International Conference on Learning Representations (ICLR)* (2021)
56. Yu, Z., Zheng, J., Lian, D., Zhou, Z., Gao, S.: Single-image piece-wise planar 3d reconstruction via associative embedding. In: *CVPR*. pp. 1029–1037 (2019)
57. Zeng, Y., Fu, J., Chao, H., Guo, B.: Learning pyramid-context encoder network for high-quality image inpainting. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 1486–1494 (2019)
58. Zeng, Y., Lin, Z., Yang, J., Zhang, J., Shechtman, E., Lu, H.: High-resolution image inpainting with iterative confidence feedback and guided upsampling. In: *European Conference on Computer Vision (ECCV)*. pp. 1–17. Springer (2020)
59. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2018)
60. Zhang, Z., Ganesh, A., Liang, X., Ma, Y.: Tilt: Transform invariant low-rank textures. *International journal of computer vision (IJCV)* pp. 1–24 (2012)
61. Zheng, J., Ramasinghe, S., Lucey, S.: Rethinking positional encoding. *arXiv preprint arXiv:2107.02561* (2021)
62. Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., Torralba, A.: Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence* pp. 1452–1464 (2017)
63. Zhou, Y., Barnes, C., Shechtman, E., Amirghodsi, S.: Transfill: Reference-guided image inpainting by merging multiple color and spatial transformations. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2266–2276 (2021)
64. Ziyin, L., Hartwig, T., Ueda, M.: Neural networks fail to learn periodic functions and how to fix it. In: *Neural Information Processing Systems (NeurIPS)*. pp. 1583–1594 (2020)