

# Learning Foveated Reconstruction to Preserve Perceived Image Statistics

LUCA SURACE, Università della Svizzera italiana, Switzerland

MAREK WERNIKOWSKI, West Pomeranian University of Technology, Poland

OKAN TURSUN, Università della Svizzera italiana, Switzerland

KAROL MYSZKOWSKI, Max Planck Institute for Informatics, Germany

RADOSŁAW MANTIUK, West Pomeranian University of Technology, Poland

PIOTR DIDYK, Università della Svizzera italiana, Switzerland

Foveated image reconstruction recovers full image from a sparse set of samples distributed according to the human visual system’s retinal sensitivity that rapidly drops with eccentricity. Recently, the use of Generative Adversarial Networks was shown to be a promising solution for such a task as they can successfully hallucinate missing image information. Like for other supervised learning approaches, also for this one, the definition of the loss function and training strategy heavily influences the output quality. In this work, we pose the question of how to efficiently guide the training of foveated reconstruction techniques such that they are fully aware of the human visual system’s capabilities and limitations, and therefore, reconstruct visually important image features. Due to the nature of GAN-based solutions, we concentrate on the human’s sensitivity to hallucination for different input sample densities. We present new psychophysical experiments, a dataset, and a procedure for training foveated image reconstruction. The strategy provides flexibility to the generator network by penalizing only perceptually important deviations in the output. As a result, the method aims to preserve perceived image statistics rather than natural image statistics. We evaluate our strategy and compare it to alternative solutions using a newly trained objective metric and user experiments.

## 1 INTRODUCTION

Wide-field-of-view displays, such as virtual and augmented reality headsets, require efficient methods for generating and transmitting high-resolution images. The novel, foveated techniques try to overcome the problem by leveraging the humans’ non-uniform sensitivity to spatial distortions across a wide field of view and present high image quality only around gaze location indicated by an eye-tracking device. Such foveated systems usually consist of two main steps [22, 47, 54]. First, an image is generated or transmitted in a form of a sparse set of samples generated according to the gaze location. Second, the image is reconstructed from the sparse information before it is shown to the observer. A particular example of such a technique is foveated rendering [14] where fewer image samples are computed for the peripheral vision to save computation during rendering (Figure 1).

Our work focuses on the second step, i.e., reconstructing an image from sparse samples. While simple techniques, such as interpolation [47] can be used, it has been demonstrated that generative adversarial networks (GANs) can provide superior results [22] thanks to their ability to hallucinate missing content based on learned image or video statistics. More precisely, even though such reconstruction requires additional computation, it can provide similar reconstruction quality with lower number of input samples. However, like other supervised learning solutions, the performance of such methods is heavily influenced by the choice of architecture and training loss and strategy. And even though it is acknowledged that for any task where the perceived quality is critical the loss function must capitalize on visually-important image features, it is unclear how to fully exploit the properties of the human visual perception to guide foveated reconstructions. While many foveated systems have previously

---

Authors’ addresses: Luca Surace, luca.surace@usi.ch, Università della Svizzera italiana, Switzerland; Marek Wernikowski, marek.wernikowski@zut.edu.pl, West Pomeranian University of Technology, Poland; Okan Tursun, okan.tarhan.tursun@usi.ch, Università della Svizzera italiana, Switzerland; Karol Myszkowski, karol@mpi-inf.mpg.de, Max Planck Institute for Informatics, Germany; Radosław Mantiuk, rmantiuk@wi.zut.edu.pl, West Pomeranian University of Technology, Poland; Piotr Didyk, piotr.didyk@usi.ch, Università della Svizzera italiana, Switzerland.

exploited the loss of sensitivity to high-frequency information [14, 54], the effect does not fully explain the visibility of missing information in the peripheral vision. For example, previous psychophysical experiments suggest that even though a perfect reconstruction of fine image details in the periphery is not critical, a complete lack of high frequencies is detectable [51]. Another effect that is not fully explained by the reduced sensitivity to higher spatial frequencies in the periphery is increased positional uncertainty [27]. Building upon these findings, we pose the question of how to efficiently guide the training of foveated reconstruction techniques such that they are fully aware of the human visual system (HVS) capabilities and limitations.

To address the above question, we establish a connection between perceived quality degradation and machine learning-based foveated reconstruction by introducing a new training scheme. First, we design and conduct psychophysical experiments to study the human’s sensitivity to content hallucination across a wide visual field. We employ a texture synthesis technique guided by image statistics of the original images to generate stimuli with varying amounts of hallucinated content. We argue that this type of distortions resembles the content synthesized using GAN-based image reconstruction, and our experiments quantify their visibility. Next, we demonstrate how to incorporate the experimental results into training. A known and applied strategy for including perceptual findings into the training GAN is to use a perceptual loss function, e.g., LPIPS [64], as a generator loss. Our main hypothesis is that this is insufficient, and also the discriminator training should be altered to better comply with perception. Consequently, instead of providing the discriminator with original images during training, we propose to use a dataset with imperceptible distortions derived from our perceptual experiment. This way the discriminator network inherits limitations of the HVS represented in the data. The key idea here is to make the discriminator penalize the generator network only for the deviations in the output with visible distortions and not for perceptually-plausible hallucination. Such a strategy, by focusing on perceptually-important image features, can provide the same perceived quality using fewer samples, leading to savings in bandwidth or rendering time. We argue that this is possible because our foveated reconstruction method aims to preserve perceived image statistics rather than natural image statistics as done by standard, GAN-based image reconstruction. We refer here to perceived image statistic as a term which describes statistics computed on images that encode relevant for the human visual system information. This is in contrast to natural image statistics which characterize natural images [44] and do not account for human perception. The new dataset also allows us to calibrate application-specific objective metrics that predict image quality. We use the new metric as well as perceptual experiments to evaluate our new training strategy and compare it with alternative solutions.

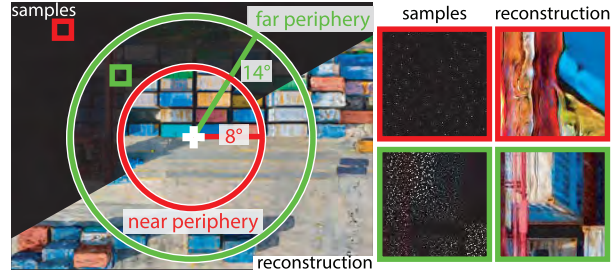


Fig. 1. The input to the foveated image reconstruction are sparse samples (magnify the top-left part of the image or see the insets), based on which, the technique reconstructs the images (bottom-right part of the image). In this example, the image consists of three regions: fovea (100% samples), near periphery (12% samples), and far periphery (0.7% samples). The white cross in the center indicates the gaze position. The reconstructed regions are combined using linear blending to create a smooth transition between them.

## 2 RELATED WORK

Our work takes inspiration from and bridges the expertise in visual perception, computer graphics, and machine learning. Here, we provide an overview of relevant works from these fields.

## 2.1 Foveal vs. peripheral vision

*Retina.* Perceptual capabilities of the HVS has been studied extensively in the past for different positions of the visual stimuli in the visual field. The perception is not uniform across the visual field because of optical and physiological limitations. Studies on the retina revealed that the human photoreceptor density is highly heterogeneous [7, 62]. The central region of the retina called *fovea centralis* (or *fovea*) is characterized by relatively high density of the cone photoreceptors and retinal ganglion cells (RGCs). This provides the foveal vision with a superior perceptual capability compared to the non-foveal (or the *peripheral*) vision. Although fovea provides a sharp central vision, it is relatively small and corresponds to approximately  $2^\circ$  of the visual field, which spans up to  $160\text{--}170^\circ$  [25]. On the other hand, the peripheral vision corresponds to more than 99 % of our visual field.

*Peripheral contrast sensitivity.* In order to study the differences between foveal and peripheral vision, earlier psychophysical studies focused on the measurements of the contrast sensitivity function (CSF) which is the sensitivity to changes in the contrast at different spatial frequencies [3, 24, 28]. In the fovea, studies showed that the human CSF curve has a peak around 4-8 cycles per degree (cpd) with its tail reaching up to 50-60 cpd. Later, Peli *et al.* [33] and more recently, Chwesiuk *et al.* [6] extended these measurements to the peripheral vision and observed that the decline in the contrast sensitivity is characterized by a smaller peak that shifts towards lower spatial frequencies as the eccentricity increases. This implies a loss of sensitivity to the high spatial frequency content in the peripheral vision.

*Foveated rendering.* The differences between foveal and peripheral vision which are mentioned above led to gaze-contingent techniques that process and display images depending on the gaze position of the observer. In this domain, foveated rendering is an actively studied gaze-contingent technique, which uses the gaze position from an eye tracker for low-resolution image reconstruction in the periphery [5, 14, 23, 30, 32, 47]. These works significantly improve the computational cost of rendering because they effectively reduce the number of pixels to be rendered [63]. However, their reconstruction methods are mostly based on simple interpolation of subsampled image and postprocessing steps such as temporal antialiasing and contrast enhancement. Unfortunately, such a simple reconstruction approach does not aim to replace the high-frequency spatial details lost due to undersampling of the underlying content, which leads to visible quality degradation.

*Hallucinating image details.* Psychophysical measurements show that the peripheral vision requires a more sophisticated model than a simple boundary between perceptible and imperceptible regions of contrast guided by the shape of the CSF [40, 42]. Thibos *et al.* [51] revealed that the resolution threshold declines from 14 cpd to 2.6 cpd in the eccentricity range  $5^\circ$  to  $35^\circ$ , whereas the detection threshold drops from 46 cpd to 28 cpd in the same eccentricity range. As a result of the faster drop-off in the resolution threshold, for each specific eccentricity there exists a band of spatial frequencies that can be detected but not accurately resolved. Additional studies show that the spatial phase discrimination performance also decreases with increasing eccentricity and a higher positional uncertainty is observed in visual perception [27, 31, 35]. Rosenholtz [39] claimed that the HVS encodes image statistics rather than the precise location information in the peripheral vision, which leads to a performance decline in resolving the stimulus position. These studies have important implications on the design of foveated image reconstruction methods because they clearly show that the HVS models driving the reconstruction must be comprehensive enough to take multiple aspects of the visual perception into account. In contrast to standard reconstruction techniques mentioned above, we address this missing piece in the foveated image reconstruction pipeline.

## 2.2 Metamers

The visual stimuli which are physically different but perceptually identical are called *metamers*. The limitations of the HVS perception has inspired several important studies on metamerism. Initial works aimed towards

texture synthesis, where the goal is to obtain perceptually identical image patches from an exemplar. To this end, Portilla and Simoncelli [34] used an iterative optimization that is run until a randomly initialized image patch converges to the same summary statistics as the target texture. Balas *et al.* [2] revealed that the summary statistics representation could explain the crowding effects observed in the periphery. Inspired by this perceptual phenomenon, Rosenholtz *et al.* [41, 42] introduced texture tiling model (TTM) based on summary statistics to model the performance of visual search in the periphery. Later, Fridman *et al.* [12] proposed a convolutional network to reproduce the outputs of TTM, and Deza *et al.* [9] introduced a generative model for creating metamers using neural networks. Instead of using hand-crafted summary statistics as previously introduced by Portilla and Simoncelli [34], Deza *et al.* used channel autocorrelation statistics computed from the pretrained VGG network features [45]. Although those works showed promising results, their main goal is to study foveal texture perception or to provide a reference model for studying the properties of the peripheral vision (e.g. for visual crowding).

Recently, Kaplanyan *et al.* [22] introduced a foveated image compression solution using a GAN model where they reconstruct perceptually plausible image sequences from significantly sparse samples while maintaining temporal coherence. Our goal is also attaining the best perceived quality while minimizing the required number of samples from the underlying content. Different from Kaplanyan *et al.*’s work, we focus on reconstructing by hallucination, capitalizing on positional uncertainty and distortions as suggested by Thibos [51]. We aim to find out how much flexibility we have in such a training scheme using the perceptual data collected through a series of systematic experiments.

### 2.3 Image metrics and perceptual loss

One possibility to guide image reconstruction is to use image metrics. Existing foveal quality metrics use the properties of central vision and provide inaccurate predictions for the periphery. The growing research on peripheral vision and its applications to foveated image reconstruction suggest a need for new foveated metrics [16, 19, 37, 43, 49, 52, 53, 56, 59]. These metrics are promising candidates for guiding the loss function in learning-based approaches. However, their complex implementations, costly computations, and in some cases, non-differentiable operations pose challenges for training image reconstruction models using them. An alternative to this is to use a training loss defined on the feature maps from a pre-trained deep network. This has become one of the most common approaches for learning-based image reconstruction, especially for super-resolution techniques [10, 21, 64]. Compared to simpler loss functions like mean squared error (MSE), the hierarchical architecture of deep networks more closely resembles how the HVS processes visual information. But still, there may be significant differences between a deep network representation and the human visual perception [11]. In addition, some of the most commonly used pre-trained networks are shown to have redundancy in their feature representations when reconstructing for the best-perceived quality [50]. The losses defined on those feature representations improve the perceived quality in the fovea, but they are not specifically optimized for the peripheral vision. In this work, we take an orthogonal approach in the context of GAN training. Apart from using a perceptual loss to train a generator, our main contribution is a modification to the discriminator training such that it better reflects the discriminative power of a human observer.

## 3 PERCEPTUAL EXPERIMENTS

There is an important connection between the studies on metamers and foveated image reconstructions using GANs discussed above. While the former postulates the importance of preserving image statistics for the peripheral vision, the latter reconstructs the content according to the discriminator trained on natural images and videos. However, we argue that training the discriminator using natural images does not sufficiently reflect the HVS’ lack of sensitivity to spatial distortions and leads to overconstraining the generator in hallucinating content. Therefore,

we propose to train the discriminator on a dataset composed of images that include distortions imperceptible for observers. It is important to note that the primary goal of this procedure is not to make the GAN produce distortions but rather make it insensitive to the distortions that humans cannot detect and focus on penalizing perceptually important artifacts. To derive this dataset, we rely on a texture synthesis technique where we control the strength of image distortions, and consequently, their visibility. This allows us to perform perceptual experiments to quantify the visibility of the introduced distortions and construct a training set with imperceptible distortions.

### 3.1 Stimuli generation

Our stimuli generation model is based on the texture synthesis method proposed by Gatys *et al.* [13]. Their method is formulated as an optimization procedure on the feature maps of pretrained VGG-19 [45] network which optimizes  $\hat{x}$  for an input exemplar  $\vec{x}$  by minimizing the loss function:

$$\mathcal{L}(\vec{x}, \hat{x}) = \sum_{l=0}^L w_l \sum_{i,j} \frac{1}{4N_l^2 M_l^2} (\hat{G}_{ij}^l - G_{ij}^l), \quad (1)$$

where  $\hat{G}_{ij}^l$  and  $G_{ij}^l$  are the Gram matrices for feature maps  $i$  and  $j$  in layer  $l$ ,  $N_l$  is the number of feature maps,  $M_l$  is the total number of neurons in a layer, and  $w_l$  is an additional weight associated with layer  $l$ . To synthesize images for our experiment, we use the same procedure, but we additionally constrain a portion of randomly chosen pixel values in  $\hat{x}$  to have equal RGB intensities as the corresponding pixels in  $\vec{x}$ . We refer to these pixels as *guiding samples*. We enforce the quality constraint by projecting the solution to the feasible space in each iteration of a gradient descent optimization.

We observed that running the constrained optimization results in subtle artifacts in the form of checkerboard patterns due to the backpropagation of gradients through convolutional layers (Figure 2, first row – middle). We identified two solutions to address this issue. The first one consists of running the constrained optimization until convergence, then removing the checkerboard artifacts characterized by a high spatial frequency by applying a low-pass Gaussian filter ( $\sigma = 1$ ) and running a second round of the optimization without the constraint. We also perceptually verified that similar results may be achieved when the constrained optimization is initialized with the guiding samples filtered by a Gaussian filter (Figure 2, second row).

Using the above procedure, we compute images  $\hat{x}_p$ , where  $p\%$  is the percentage of the guiding samples (Figure 3, left). For  $p = 0$ , our synthesis is equivalent to the original technique presented by Gatys *et al.*

The loss of high spatial frequencies (as commonly observed when reducing input resolution) is an important factor on the perceived quality. In order to improve the sensitivity of the trained metrics to a visible decline of resolution, we also created a separate dataset of image with different amount of Gaussian blur, i.e., different  $\sigma$  parameter of the Gaussian kernel (Figure 3, right). The results of the perceptual experiment obtained with these stimuli were used to expand the dataset of images used for training image metrics in Section 5.2.

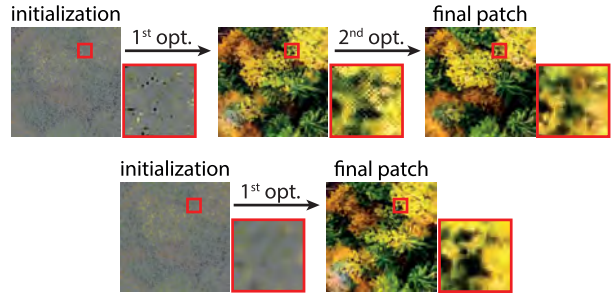


Fig. 2. The figure presents two strategies for synthesizing stimuli with guiding samples. Top: a two-step procedure where second unconstrained optimization is performed to remove subtle checkerboard artifacts after applying Gaussian filter with  $\sigma = 1$ . Bottom: a constrained optimization initialized with a blurred version of guiding samples.

### 3.2 Experiment protocol

The amount of guiding samples  $p$  and  $\sigma$  value of the Gaussian kernel provide a parametrization for our investigation of the HVS' sensitivity to the deviations from the original images. More precisely, in our main experiment with stimuli generated using texture synthesis, we seek a direct relationship between the amount of guiding samples and the probability of detecting the distortions by a human observer at a particular eccentricity. We later use this relation to generate a much larger dataset of images with imperceptible distortions required for training GAN-based foveated reconstruction. In contrast, the additional experiment with Gaussian blur only seeks pairs of images and the corresponding probability of blur detection, as a smaller dataset is sufficient for training our image metrics.

*Stimuli.* We prepared 20 image patches of size  $256 \times 256$  each from a different 4K image. The images are grouped into two main categories: nature and architecture, to create a mixture of images with and without a clear structure, as we expected that reconstructing images with clear structure may be more challenging. For each patch, we generated a corresponding set of distorted patches for  $p \in \{0, 3, 5, 7.5, 10\}$ . The set of  $p$  values was determined in a preliminary experiment where we found that  $p > 10\%$  leads to images that are almost always perceptually indistinguishable from the originals. The same set of 20 ground truth patches was used to generate a set of blurred patches for  $\sigma \in \{0.25, 1.25, 2.25, 3.25, 4.25\}$ . The range was chosen to uniformly span the range of visible blur across the considered field of view [54]. The sample of our stimuli is presented in Figure 3.

*Task.* In each trial, three patches were presented on the screen: (1) the original patch at the fixation point, (2) a synthesized stimulus on either right or left side at a given eccentricity, and (3) the original patch on the opposite side at the same eccentricity. The stimuli were visible to both eyes, and the participants were asked to select the patch that was more similar to the reference by pressing the left or right arrow keys. With this procedure we tested the visibility of distortions at the following eccentricities:  $8^\circ$  (the end of parafovea [57]),  $14^\circ$  (the center of periphery [48, 57]), and  $20^\circ$ , for which the stimuli spanned  $3.21^\circ$ ,  $3.08^\circ$  and  $2.89^\circ$ . The participants were instructed to maintain their gaze on the central image. At the same time, we kept the position of the stimuli relative to the gaze position obtained from an eye tracker fixed to maintain the desired localization of the stimuli in the periphery. In total, each participant performed 1500 trials. The order of the images, eccentricities, and the sides

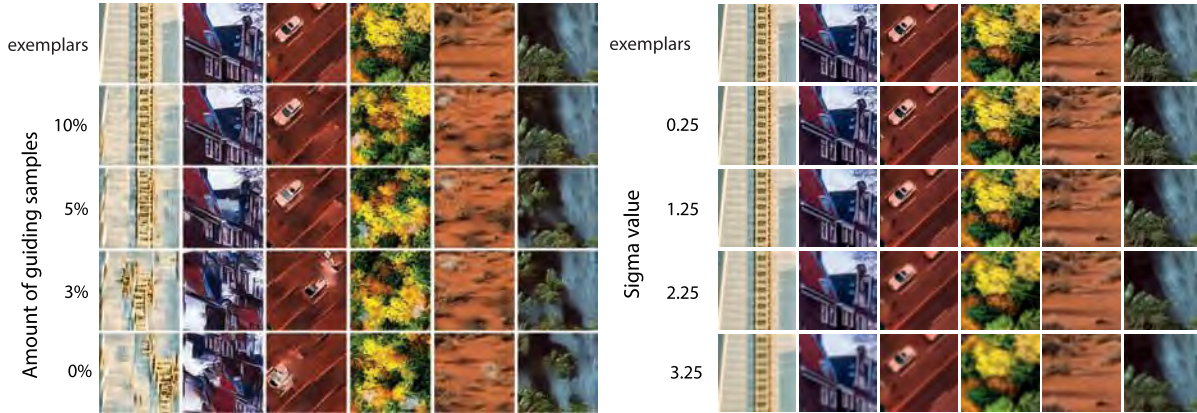


Fig. 3. Results of our stimuli generation for different amount of guiding samples (left) and  $\sigma$  values (right) of the Gaussian filter. Note the increased distortions with respect to the exemplar when the number of guiding samples decreases (left) and when  $\sigma$  increases (right).

on which the test stimuli were presented was randomized. The same experiment was performed with images containing artifacts from texture synthesis and Gaussian blur.

*Hardware.* We used a setup consisting of a 27" Acer Predator display operating at a resolution of 3840×2160 px at 120 Hz and a peak luminance of 170 cd/m<sup>2</sup>. For tracking the gaze position, we used a Tobii Pro Spectrum eye tracker with 600 Hz sampling rate.

*Participants.* 5 participants between ages 25–36 with normal or corrected-to-normal vision took part in the experiment. 4 of them were the authors and the remaining participant was completely naive to the study. During a short initial warm-up phase, the participants were given instructions about the task.

### 3.3 Results

The results for the Gaussian blur were directly used for training image metrics (Section 5.2). The rest of the data, i.e., for texture synthesis stimuli, was processed by aggregating the results across all images for each eccentricity. Figure 4 presents the probability of identifying the synthesized stimuli as a function of the amount of guiding samples and eccentricity at which the stimuli were presented. Using the inverse function and interpolating the values by a cubic function, we found the sampling rates corresponding to a detection rate of 75%, usually considered as a visibility threshold: 9.09%, 6.89% and 4.71%, for eccentricities 8°, 14° and 20°, respectively. In other words, we found threshold sampling rates required to generate images with artifacts that remain invisible to human observers. These values were later used as the amount of guiding samples to prepare the inputs for the discriminator in GAN training (Section 4.1).

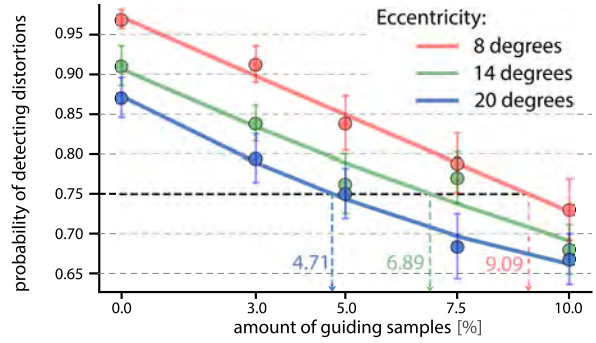


Fig. 4. The probability of detecting differences between original and synthesized images as a function of amount of guiding samples and eccentricity. The error bars visualize standard errors of the means.

## 4 METHOD

The results of the perceptual experiments described in the previous section provide the measured thresholds for structural distortions for a standard observer. We use these data to control the learned manifold of target images in foveated image reconstruction. To this end, we present an improved training scheme for GAN where the training data consists of a set of both natural and synthesized images.

Our foveated reconstruction network uses Wasserstein GAN [1] training scheme to produce perceptually optimized reconstructions from subsampled images. The network topology is based on the encoder-decoder structure of UNet with skip connections [38] (Figure 5). This network design is similar to the model previously used by Kaplanyan *et al.* [22]. The encoder part of the generator network (G) consists of downsampling residual blocks using average pooling layers [17]. Each residual block of the encoder consists of two convolutional layers with a filter size of  $5 \times 5$ , except for the main branch, where we use a  $1 \times 1$  filter to adjust the dimensionality. The numbers of filters are 16-32-64-128-128 in each block, respectively. The decoder part is a mirrored version of the first four encoder blocks with upsampling blocks using bilinear interpolation instead of average pooling. The encoder and the decoder are connected by an additional bilinear upsampling layer. We use LeakyReLU activation

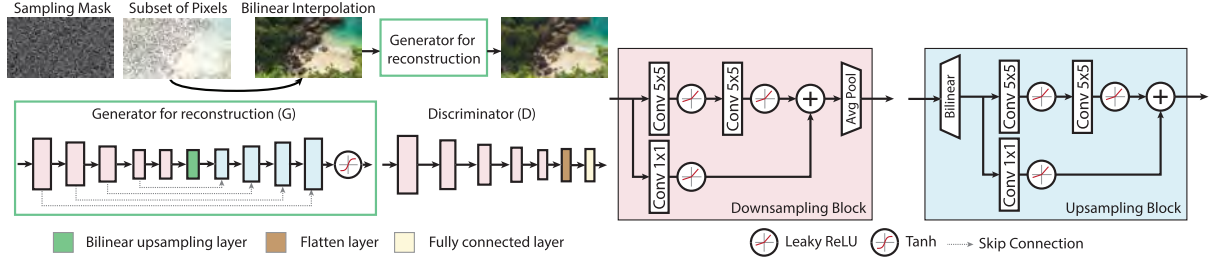


Fig. 5. The architecture of our proposed network.

with a negative slope coefficient of  $\alpha = 0.2$  throughout the network, except for the final layer of the generator, which uses tanh activation.

The discriminator network or *critic* ( $D$ ) is based on PatchGAN [20] with a patch size of  $64 \times 64$ . The discriminator consists of downsampling blocks similar to the encoder part of the generator (the number of filters: 16-32-64-128-128). The output of downsampling blocks is flattened and passed to a fully connected layer which produces a scalar. Compared to the previous model of Kaplanyan *et al.*, we use a more compact generator with the number of filters reduced by half in the first and last three blocks. Using such a compact generator is made possible by our more permissive training scheme, which allows for imperceptible deviations from the target image statistics. In contrast, the discriminator loss used by Kaplanyan *et al.* aims to match the target image statistics as closely as possible. This important difference in our training scheme makes it possible to retain the perceptual quality with a more compact network.

#### 4.1 Dataset

We used two datasets for the inputs of the generator and for the set of images provided to the discriminator, respectively. The input dataset of the generator consists of patches from natural images with a size of  $256 \times 256$ . The patches are generated by cropping images with random offsets. In order to maintain a balanced data representation, 50 images are randomly selected from each one of the 1000 classes in the ImageNet dataset [8], which provided us with a total of 50K patches. These images are later subsampled using the void-and-cluster algorithm [55] with a sampling rate 12% for the near periphery and 0.7% for the far periphery. This choice is done according to a content-aware foveated rendering method proposed by Tursun *et al.* [54]. The subsampling is followed by bilinear interpolation before passing the images to the generator as inputs.

As for the inputs of the discriminator, we used the texture synthesis method described in Section 3.1 using the results of the perceptual experiment described in Section 3.2. This dataset consists of 50K patches synthesized using 9.09 % and 6.89 % of the pixels as guiding samples respectively for near and far peripheral regions in addition to the full-resolution ground truth images.

#### 4.2 Discriminator loss

The training of the discriminator,  $D$ , uses the same loss function as the original WGAN design [1]:

$$\mathcal{L}_{adv} = D(x) - D(G(z)), \quad (2)$$

where  $z$  represents the inputs to the generator network  $G$ ,  $D(x)$  is the output of the discriminator to real samples (natural images), and  $D(G(z))$  is the output of the discriminator to the reconstructions from  $G$ . This training is equivalent to the optimizations performed in previous works when  $x \in \mathbb{I}$ , where  $\mathbb{I}$  is the set of images from ImageNet.



Fig. 6. Training with (left) and without (right) input sampling mask.

In our experiments, we update this formulation by using  $x^* \in \mathbb{I}^*$  where  $\mathbb{I}^*$  is the set of images with visually imperceptible structural distortions, as we described in Section 3.2. We denote the discriminator loss operating on this manifold of synthesized images with structural distortions as:

$$\mathcal{L}_{adv}^* = D(x^*) - D(G(z)). \quad (3)$$

For training stability of WGAN, we imposed a soft Lipschitz constraint using gradient penalty [15].

#### 4.3 Generator loss

Our optimization trains generator networks with a weighted sum of different types of losses. To have a comprehensive evaluation of our new training scheme, we focused our analysis on three types of generators trained with standard and perceptual losses. The first generator,  $G_{L2}$ , was trained with a combination of the standard MSE loss and the adversarial loss. The second generator,  $G_{LPIPS}$ , was trained with the learned perceptual image patch similarity loss term. We used the learned linear weights on top of the VGG network as provided by the authors in their work [64]. Additionally, inspired by [18], we add the generator  $G_{Lapl}$  that used Laplacian-based loss. It is defined as a weighted sum of mean squared error between the corresponding levels of the Laplacian pyramid for the reconstruction and the ground truth. We assigned the weights to each level according to a Gaussian with  $\sigma = 1.0$ . By centering the Gaussian around different levels we can put more emphasis on the reconstruction fidelity of different spatial frequencies in the pyramid decomposition. The main motivation behind the loss was that by putting more weight on lower spatial frequencies, the network will be given more freedom in hallucinating high spatial frequencies, which might be desired in the periphery. All the generator losses used in our experiments are listed in Table 1.

#### 4.4 Training

Similar to the previous work of Guenter *et al.* [14], we divide the image into three regions with different levels of structural distortions controlled by the amount of guiding samples used while training the generators. In order to provide conservative approximations, we take our estimations from Section 3 at the eccentricities corresponding to the inner boundary of near and far peripheral regions, which are at  $8^\circ$  and  $14^\circ$ , respectively (represented as red and green circles in Figure 1). The contents of the foveal region are directly transferred from a full resolution image. We train two distinct generator networks, each of which is responsible for the reconstruction of near and far peripheral regions. For benchmarking purposes, we also train separate networks for each one of the two different discriminator losses (using our  $\mathcal{L}_{adv}^*$  and the standard loss  $\mathcal{L}_{adv}$ ) and three different generator losses

L2	$\mathcal{L}_G^{L2} = w_{L2} \cdot \mathcal{L}_{L2} + w_{adv} \cdot \mathcal{L}_{adv}$
L2 ours	$\mathcal{L}_{G^*}^{L2} = w_{L2} \cdot \mathcal{L}_{L2} + w_{adv} \cdot \mathcal{L}_{adv}^*$
LPIPS	$\mathcal{L}_G^{LPIPS} = w_{LPIPS} \cdot \mathcal{L}_{LPIPS} + w_{adv} \cdot \mathcal{L}_{adv}$
LPIPS ours	$\mathcal{L}_{G^*}^{LPIPS} = w_{LPIPS} \cdot \mathcal{L}_{LPIPS} + w_{adv} \cdot \mathcal{L}_{adv}^*$
Laplacian	$\mathcal{L}_G^{Lapl} = w_{Lapl} \cdot \mathcal{L}_{Lapl} + w_{adv} \cdot \mathcal{L}_{adv}$
Laplacian ours	$\mathcal{L}_{G^*}^{Lapl} = w_{Lapl} \cdot \mathcal{L}_{Lapl} + w_{adv} \cdot \mathcal{L}_{adv}^*$

Table 1. The loss functions that we used for training the generator in our evaluations.

(using  $\mathcal{L}_G^{L2}$ ,  $\mathcal{L}_G^{LPIPS}$ , and  $\mathcal{L}_G^{Lapl}$ ). Relative weights of loss terms are set to  $w_{L2} = 2000$ ,  $w_{LPIPS} = 100$ ,  $w_{Lapl} = 100$ , and  $w_{adv} = 1$ .

We used Adam optimizer with a learning rate of  $2 \times 10^{-5}$  ( $\beta_1 = 0.5$ ,  $\beta_2 = 0.999$ ,  $\epsilon = 10^{-8}$ ). The training lasts for 20–30 epochs until convergence which takes approximately one day on an Nvidia 2080 Ti GPU. We assume convergence when the training loss reaches a plateau. The sample reconstructions from the converged network are also visually checked against potential instabilities during training.

#### 4.5 Sampling mask

Capitalizing on the potential correlation between subsequent frames, the network introduced by Kaplanyan *et al.* uses recurrent connections as an important part of their design which are able to retain information from the previously subsampled frames. This high-level temporal reprojection provides the network with additional information when the underlying content is only partially observed due to sparse subsampling. In order to observe the effects of different training schemes more clearly, we decided to use the information from only one frame and isolate the reconstruction from the effects of this temporal information flow. In our initial experiments, we observed that such a design decision made the network more sensitive to the sampling mask used in the inputs due to the absence of temporal information flow, which would otherwise compensate for the lack of information about the true values of missing samples. In order to address this issue, as a first attempt we filled in the missing information by interpolating the sampled pixels while keeping the sampling mask as a channel of the input. However, visual inspection revealed visual artifacts collocated with the sampled pixels (Figure 6) whose visibility was dependent on the weights assigned to the loss terms. The effect was the most pronounced when we used  $\mathcal{L}_{L2}$  in training and less visible with  $\mathcal{L}_{LPIPS}$ . As a remedy, we removed the sampling mask from the training input, decided to provide the generator with the bilinearly interpolated input consisting of RGB channels, as shown in Figure 5. This solution seemed to be effective in removing visual artifacts from the reconstruction (please refer to Figure 6 for a visual comparison).

## 5 RESULTS AND EVALUATION

We evaluated our strategy of training foveated image reconstruction using trained image metric and perceptual experiment. We compared all six methods for training the network. They are combinations of using LPIPS, L2, and Laplacian pyramid as the generator loss function, as well as using original and our new patch dataset for training the discriminator network (Table 1).

### 5.1 Visual inspection

Figure 7 presents reconstruction results obtained using differently trained architectures for four different images. For reference, we also include original high-resolution and standard foveated reconstruction using an interpolation with Gaussian weights. For the results of training using Laplacian loss we introduce a notation consisting of



Fig. 7. Sample reconstructions of all evaluated methods for nature and architecture images performed for far periphery. *Lapl MM* refers to the highest weight placed on the medium spatial frequencies, while *Lapl HH* puts the highest weight on the high spatial frequencies.

two letters, *Lapl XY*, where  $X, Y \in \{H, M\}$  encode the position of the Gaussian peak at far and near periphery, respectively. The letter H represents the position of the peak located at the first level of the pyramid (characterized by placing more emphasis on high spatial frequencies) whereas M represents the position of the peak located at the fourth level of the pyramid (medium spatial frequencies). For example, the method denoted as *Lapl HM*, refers to a reconstruction using a network trained with the Laplacian pyramid-based loss where for the far periphery the high spatial frequencies are weighted more, and for near-periphery, the medium frequencies are given higher importance.

The first observation is that all eight GAN-based reconstruction results exhibit clear hallucinated results, and the reconstruction of very fine details is not exact. While this is visible with direct visual inspection, such deviations are less visible when shown in the periphery. Furthermore, all reconstructions introduce high spatial frequencies and strong edges, but training with L2 loss makes them sparser and more exaggerated. The visual comparison (Figure 7) between the discriminator trained with and without our synthesized dataset (i.e., L2 vs. L2 ours, LPIPS vs. LPIPS ours, Lapl MM vs. Lapl MM ours, Lapl HH vs. Lapl HH ours) show that our results include more high spatial frequencies. We argue that this is due to the flexibility of the discriminator, which penalizes hallucinations of high spatial frequencies less harshly. This is a desired effect because while the HVS is sensitive to removing some of the high spatial frequencies in the periphery, it is less sensitive to the changes in their positions (Section 2).

To further investigate the spatial frequency distribution characteristics of our reconstructions, we visualize the frequency band decomposition outputs from the Laplacian Pyramid and compute the differences of two layers

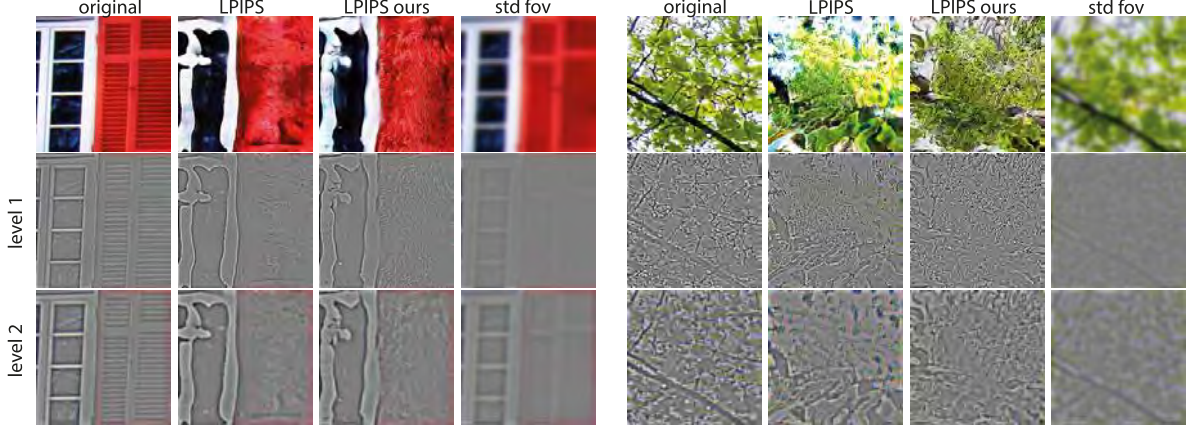


Fig. 8. High-frequency content hallucination using LPIPS, LPIPS ours, and standard Gaussian blur.

from the bottom of the pyramid<sup>1</sup>, which encode the highest frequency band and one octave below it (Figure 8). We observe that our reconstructions provide additional hallucinated high-frequency details that do not exist in the traditional foveated image reconstruction. Please refer to supplementary material for an interactive demo with more results.

## 5.2 Objective image metric

Since there is no established foveated image metric, we propose to use our perceptual data to calibrate existing metrics: L2, SSIM [60], MS-SSIM [61], and LPIPS [64], separately for different eccentricities. The calibration was performed by fitting the following logistic function [36]:

$$y(t) = a + (k - a) / (c + q \cdot e^{-b \cdot t})^{\frac{1}{v}}. \quad (4)$$

to reflect the nonlinear relation between the distortion magnitude and detection probability, with  $a, b, c, k, q, v$  being free parameters. Inspired by LPIPS [64], we also considered reweighing the contribution of the individual convolution and pooling layers of VGG-19 for each eccentricity separately. We refer to this metric as Cal. VGG.

For all metrics, the free parameters (i.e., parameters of logistic functions, as well as weights and bias for VGG-19 layers) were obtained by minimizing mean squared error in predicting the detection probability:

$$\sum_{(\vec{x}, \hat{\vec{x}}) \in S_r} \left\| y(M(\vec{x}, \hat{\vec{x}})) - P(\vec{x}, \hat{\vec{x}}) \right\|^2, \quad (5)$$

where  $M$  is one of the original metrics,  $S_r$  is the set of distorted and undistorted pairs of images for eccentricity  $r \in 8, 14, 20$ , and  $P$  is the probability of detecting the difference. The minimization was performed using nonlinear curve-fitting through *trust-region-reflective* and *Levenberg-Marquardt* optimizations [4, 26] with multiple random initializations. Additionally, we constrained the VGG weights to be nonnegative to keep the positive correlation between image dissimilarity and the magnitude of differences in VGG features, as motivated in [64]. To make our dataset more comprehensive, we added stimuli from an additional experiment that analyzed the visibility of blur. The procedure followed the one described in Section 3.2.

<sup>1</sup>Please note the difference between frequency decomposition using Laplacian Pyramid and the Laplacian pyramid-based loss function.

To validate our calibration, we performed five-fold cross-validation. Figure 9 presents Pearson’s correlation coefficients for all trained metrics and eccentricities computed as an average across all the folds. Each bar shows the correlation for uncalibrated (bright part) and calibrated (dark part) metrics. For uncalibrated case, we use the standard sigmoid logistic function:  $y(t) = 1/(1 + e^{-t})$ . We also present aggregated results where the correlation is analyzed across all eccentricities. The individual scores show that increasing the eccentricity reduces the performance of the original metrics. The additional calibration improves prediction for all metrics significantly. An interesting observation is that LPIPS performs very well for small eccentricities (8°). For larger ones (14° and 20°), as the metric was not trained for the periphery, its performance is significantly reduced even with the optimized logistic function. However, when the weights for the deep layers of VGG-19 are optimized (Cal. VGG), the performance significantly increases. This suggests that such metrics are promising, but depending on the eccentricity the contribution of the individual layers to the prediction must change. Since our Cal. VGG performed best in the tests, we use it to benchmark the foveated image reconstruction techniques listed in Table 1.

To this end, we first extend the Cal. VGG to handle full images. To support continuous values of eccentricities, we linearly interpolate the results of the metric for 8 and 20 degrees. Additionally, we used one logistic function trained according to the procedure mentioned above using experimental data for all eccentricities. Since such metric operates on patches, we compute the average value across all the patches. For the benchmark of the different reconstruction methods, we randomly selected 10 publicly available 4K images with both architectural and natural features. Before applying different reconstruction techniques, we split the images into three regions: fovea, near periphery, and far periphery and draw sparse samples according as visualized in (Figure 1). We later reconstruct the images and compute the error with respect to the ground truth images using adopted Cal. VGG. Additionally, to test the quality provided for different sample distributions, we analyze the methods for different radius of far periphery region.

Figure 10 presents the results. Training the reconstruction using LPIPS provides reconstructions that are the least likely to be distinguished from the original images. Additionally, with L2 and LPIPS, better results are achieved when the discriminator was trained using our new strategy, which demonstrates its effectiveness. We also computed the predictions by separating the images into two groups: nature and architecture (with orderly structured features). The results show that the preference of our method for nature images increases but decreases for architecture (Table 3).

### 5.3 Subjective experiments

The psychovisual experiment used to derive the data for training our reconstruction methods was performed by five participants. While it is common to use a smaller number of participants in such experiments due to their complexity and the fact that they should capture the general properties of the HVS, such experiments do not investigate potential differences in population. Also, it is unclear whether the method derived from the perceptual data is effective or not. Therefore, to further validate our claims regarding the new training strategy as well as verify the significance of the improvements revealed using calibrated metrics, we conducted an additional user subjective experiment, in which naive participants directly compared different reconstruction methods.

*Stimuli and task.* 10 4K resolution images from two classes (nature and architecture) were used in the experiment. They were sub-sampled and reconstructed following Figure 1, using L2 and LPIPS methods with and without our

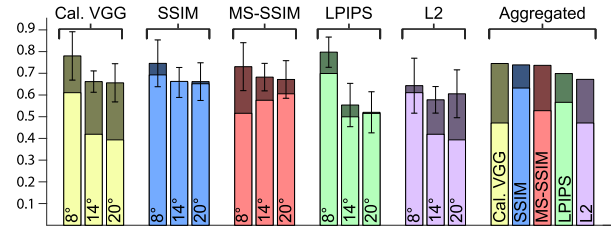


Fig. 9. Pearson’s correlation coefficient for analyzed methods. Bright bars: uncalibrated metrics, dark bars: calibrated metrics.

Table 2. Preference rates for the methods L2 ours, LPIPS ours over L2, LPIPS.

	Nature	Architecture	All
L2 ours vs L2	0.75	0.37	0.56
LPIPS ours vs LPIPS	0.75	0.43	0.59
All ours vs All	0.75	0.40	0.57

modification to the discriminator training. In each trial, the participants were shown the original image on the left and one of the two reconstructions on the right. They were asked to select the reconstruction, which was more similar to the reference by pressing a keyboard key. During the experiment, the eye-tracker was used to align the images to the participants’ gaze.

*Participants.* 15 participants with normal or corrected-to-normal vision took part in the experiment. All were naive to the purpose of the study and given instructions at its beginning. Each participant compared all pairs of techniques for each image (60 comparisons per participant).

*Results.* Table 2 contains the preference rates of the results produced using networks trained using our procedure (LPIPS ours, L2 ours) when compared to standard procedure (LPIPS, L2). We report the results for all images (last column) but also when split into nature and architecture images. For statistical significance, we performed a series of binomial tests and mentioned here all significant results. The reconstructions obtained using a network trained with our strategy are preferred in 57% of cases ( $p = 0.013$ ). The difference is significant when LPIPS is considered separately ( $p = 0.04$ ). Regarding different image classes, our method provides significant benefits for nature images, either when we consider L2 and LPIPS separately or jointly (for each case preferred in 75 % of cases,  $p < 0.001$ ). For architecture we observe the preference for L2 (63%,  $p = 0.037$ ). For all techniques, our method is preferred in 40 % of cases ( $p = 0.018$ ). This follows the results from VGG metric calibration, where our method had lower probability of detection for nature images and similar probability for architecture. We hypothesize that there might be several reasons for lower performance on architecture images. First, architecture images contain clear structure which might have not been represented well in  $256 \times 256$  patches used for the experiment. We limited the size in such a way as larger patches spread across a wide eccentricity range, posing challenges to the experiment and perceptual modeling. This, however, limits the scale of considered distortions and the consequent detection rate. The patches might be too small to represent larger structures typically presented on architecture images. Another possible cause is that we performed the perceptual experiment together for architecture and nature images and did not make any distinction when modeling the perception of artifacts for these two groups. We believe that the performance can be improved by tailoring both the experiments and training to such content.

In Table 3, we report the preference of the individual methods when compared to all other training strategies, including different loss functions, i.e., perceptual LPIPS, and simple L2. LPIPS ours gained the highest preference of 38% ( $p < 0.001$ ) while L2 the lowest (24%,  $p < 0.001$ ). When divided into classes, LPIPS ours and L2 ours were

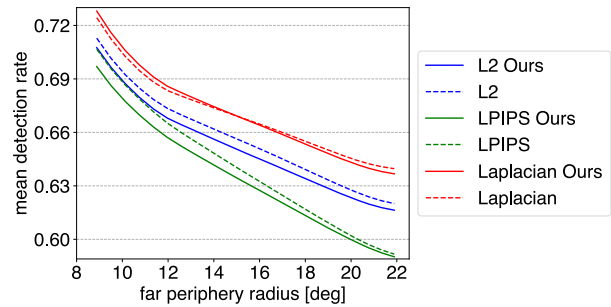


Fig. 10. Detection rate averaged across all image patches as predicted by Cal. VGG for increasing radius of far periphery. The lower values indicates a higher quality of reconstruction.

Table 3. Preference rate for analyzed methods. Overall, the stimuli generated with our injection methods were preferred over the standard ones. The trend is more emphasized for the nature images, despite the slight worsening on the architecture images.

	Nature	Architecture	All
LPIPS ours	<b>0.41</b>	0.35	<b>0.38</b>
LPIPS	0.22	0.37	0.29
L2 ours	<b>0.37</b>	0.21	<b>0.29</b>
L2	0.20	0.27	0.24

preferred the most: 41% ( $p < 0.001$ ) and 37% ( $p = 0.003$ ), respectively. Other methods reached lower preference levels - 22% for LPIPS ( $p < 0.001$ ) and 20% for L2 ( $p < 0.001$ ). For architecture images, the results show the favour for LPIPS (37%,  $p = 0.005$ ) and LPIPS ours (35%,  $p = 0.032$ ). The L2 ours was selected the fewest number of times (21%,  $p < 0.001$ ). All other results did not show statistical significance.

## 6 CONCLUSIONS AND FUTURE WORK

Current techniques use a perceptual loss to guide a network training to capitalize on perceptually-important image features. The goal of this work is to inject perceptual information into the discriminator network. To this end, during training, we provide the discriminator with images containing distortions imperceptible by a human observer. This allows the discriminator to inherit the properties of the HVS encoded in the training dataset. Our new dataset includes images with invisible spatial distortions coming from texture synthesis. We argue that such distortions are much closer to artifacts introduced by GAN-based reconstruction than previously considered blur. Additionally, the new dataset allows us to train several image metrics to improve their predictions for stimuli presented in the periphery.

We studied the suitability of the new training strategy for foveated image reconstruction. In the future, it is essential to extend the investigation to video content with potential benefits when the sensitivity of the HVS to temporal artifacts is incorporated. We trained separate networks for the near and far periphery. While this makes the training procedure easier, a more practical solution is to train one network to handle spatially-varying density. One alternative solution to reach this goal is to use a fully convolutional network in log-polar domain [46, 58]. We also did not focus on computational performance. Currently, our unoptimized inference takes 3 seconds on our hardware. While previous work [22] demonstrates the feasibility of using GAN in such scenarios, the computational efficiency is still an important concern. We believe that making the networks and their training aware of the limitations of human perception will be important for closing the gaps.

Another exciting research direction is designing a foveated image metric that accounts for a wide range of effects. While the concurrent work by Mantiuk et al. [29], takes this path; in contrast to our focus, they target image quality instead of distortion visibility. The challenge here lies in collecting large-scale perceptual data with eye-tracking information containing information about the visibility of distortions. And even though our data set collects such information, it is not enough to train a general-purpose visibility metric for both foveal and peripheral vision.

## REFERENCES

- [1] Martin Arjovsky, Soumith Chintala, and Léon Bottou. 2017. Wasserstein GAN. *arXiv:1701.07875 [cs, stat]* (Jan. 2017). arXiv: 1701.07875.
- [2] Benjamin Balas, Lisa Nakano, and Ruth Rosenholtz. 2009. A summary-statistic representation in peripheral vision explains visual crowding. *Journal of vision* 9, 12 (2009), 13–13.
- [3] Peter GJ Barten. 1999. *Contrast sensitivity of the human eye and its effects on image quality*. SPIE press.

- [4] Mary Ann Branch, Thomas F Coleman, and Yuying Li. 1999. A subspace, interior, and conjugate gradient method for large-scale bound-constrained minimization problems. *SIAM Journal on Scientific Computing* 21, 1 (1999), 1–23.
- [5] Valentin Bruder, Christoph Schulz, Ruben Bauer, Steffen Frey, Daniel Weiskopf, and Thomas Ertl. 2019. Voronoi-Based Foveated Volume Rendering. In *EuroVis (Short Papers)*. The Eurographics Association, 5.
- [6] Michał Chwesiuk and Radosław Mantiuk. 2019. Measurements of Contrast Sensitivity for Peripheral Vision. In *ACM Symposium on Applied Perception 2019 (Barcelona, Spain) (SAP '19)*. Association for Computing Machinery, New York, NY, USA, Article 20, 9 pages. <https://doi.org/10.1145/3343036.3343123>
- [7] Christine A. Curcio, Kenneth R. Sloan, Robert E. Kalina, and Anita E. Hendrickson. 1990. Human photoreceptor topography. *Journal of comparative neurology* 292, 4 (1990), 497–523.
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*. IEEE, 248–255.
- [9] Arturo Deza, Aditya Jonnalagadda, and Miguel Eckstein. 2017. Towards Metamerism via Foveated Style Transfer. *arXiv:1705.10041 [cs]* (May 2017). arXiv: 1705.10041.
- [10] Alexey Dosovitskiy and Thomas Brox. 2016. Generating images with perceptual similarity metrics based on deep networks. In *Advances in neural information processing systems*. 658–666.
- [11] Jenelle Feather, Alex Durango, Ray Gonzalez, and Josh McDermott. 2019. Metamers of neural networks reveal divergence from human perceptual systems. In *Advances in Neural Information Processing Systems*. 10078–10089.
- [12] Lex Fridman, Benedikt Jenik, Shaiyan Keshvari, Bryan Reimer, Christoph Zetzsche, and Ruth Rosenholtz. 2017. A Fast Foveated Fully Convolutional Network Model for Human Peripheral Vision. *arXiv:1706.04568 [cs]* (2017). arXiv:1706.04568 [cs.NE]
- [13] Leon Gatys, Alexander S Ecker, and Matthias Bethge. 2015. Texture synthesis using convolutional neural networks. In *Advances in neural information processing systems*. 262–270.
- [14] Brian Guenter, Mark Finch, Steven Drucker, Desney Tan, and John Snyder. 2012. Foveated 3D Graphics. *ACM Trans. Graph.* 31, 6 (Nov. 2012), 164:1–164:10.
- [15] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. 2017. Improved training of Wasserstein GANs. In *Advances in neural information processing systems*. 5767–5777.
- [16] Peiyao Guo, Qiu Shen, Zhan Ma, David J. Brady, and Yao Wang. 2018. Perceptual Quality Assessment of Immersive Images Considering Peripheral Vision Impact. *arXiv:1802.09065 [cs]* (Feb. 2018). arXiv: 1802.09065.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [18] Alexander Hepburn, Valero Laparra, Ryan McConville, and Raul Santos-Rodriguez. 2019. Enforcing perceptual consistency on generative adversarial networks by using the normalised laplacian pyramid distance. *arXiv preprint arXiv:1908.04347* (2019).
- [19] Chih-Fan Hsu, Anthony Chen, Cheng-Hsin Hsu, Chun-Ying Huang, Chin-Laung Lei, and Kuan-Ta Chen. 2017. Is Foveated Rendering Perceivable in Virtual Reality?: Exploring the Efficiency and Consistency of Quality Assessment Methods. In *Proceedings of the 25th ACM International Conference on Multimedia (MM '17)*. ACM, New York, NY, USA, 55–63. <https://doi.org/10.1145/3123266.3123434> event-place: Mountain View, California, USA.
- [20] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2017. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1125–1134.
- [21] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. 2016. Perceptual Losses for Real-Time Style Transfer and super-Resolution. In *Computer Vision – ECCV 2016 (Lecture Notes in Computer Science)*, Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling (Eds.). Springer International Publishing, 694–711.
- [22] Anton Kaplanyan, Anton Sochenov, Thomas Leimkühler, Mikhail Okunev, Todd Goodall, and Gizem Rufo. 2019. Deepfovea: Neural Reconstruction for Foveated Rendering and Video Compression Using Learned Natural Video Statistics. In *ACM SIGGRAPH 2019 Talks (Los Angeles, California) (SIGGRAPH '19)*. ACM, New York, NY, USA, Article 58, 2 pages.
- [23] Jonghyun Kim, Zander Majercik, Peter Shirley, Josef Spjut, Morgan McGuire, David Luebke, Youngmo Jeong, Michael Stengel, Kaan Akşit, Rachel Albert, Ben Boudaoud, Trey Greer, Joohwan Kim, and Ward Lopes. 2019. Foveated AR: dynamically-foveated augmented reality display. *ACM Trans. Graph.* 38, 4 (July 2019), 1–15.
- [24] Kil Joong Kim, Rafal Mantiuk, and Kyoung Ho Lee. 2013. Measurements of achromatic and chromatic contrast sensitivity functions for an extended range of adaptation luminance. In *Human vision and electronic imaging XVIII*, Vol. 8651. International Society for Optics and Photonics, 86511A.
- [25] George F Koob, Michel Le Moal, and Richard F Thompson. 2010. *Encyclopedia of behavioral neuroscience*. Elsevier.
- [26] Kenneth Levenberg. 1944. A method for the solution of certain non-linear problems in least squares. *Quarterly of applied mathematics* 2, 2 (1944), 164–168.
- [27] Dennis M Levi, Stanley A Klein, and Yen Lee Yap. 1987. Positional uncertainty in peripheral and amblyopic vision. *Vision research* 27, 4 (1987), 581–597.

- [28] James Mannos and David Sakrison. 1974. The effects of a visual fidelity criterion of the encoding of images. *IEEE Transactions on Information Theory* 20, 4 (1974), 525–536.
- [29] Rafal K. Mantiuk, Gyorgy Denes, Alexandre Chapiro, Anton Kaplanyan, Gizem Rufo, Romain Bachy, Trisha Lian, and Anjul Patney. 2021. FovVideoVDP: A visible difference predictor for wide field-of-view video. *ACM Transactions on Graphics (TOG)* XX (2021). Issue X. (to appear).
- [30] Xiaoxu Meng, Ruofei Du, Matthias Zwicker, and Amitabh Varshney. 2018. Kernel Foveated Rendering. *Proc. ACM Comput. Graph. Interact. Tech.* 1, 1 (July 2018), 1–20.
- [31] M Concetta Morrone, David C Burr, and Donatella Spinelli. 1989. Discrimination of spatial phase in central and peripheral vision. *Vision research* 29, 4 (1989), 433–445.
- [32] Anjul Patney, Marco Salvi, Joohwan Kim, Anton Kaplanyan, Chris Wyman, Nir Benty, David Luebke, and Aaron Lefohn. 2016. Towards Foveated Rendering for Gaze-tracked Virtual Reality. *ACM Trans. Graph.* 35, 6 (Nov. 2016), 179:1–179:12.
- [33] Eli Peli, Jian Yang, and Robert B Goldstein. 1991. Image invariance with changes in size: The role of peripheral contrast thresholds. *JOSA A* 8, 11 (1991), 1762–1774.
- [34] Javier Portilla and Eero P Simoncelli. 2000. A parametric texture model based on joint statistics of complex wavelet coefficients. *International journal of computer vision* 40, 1 (2000), 49–70.
- [35] Ingo Rentschler and Bernhard Treutwein. 1985. Loss of spatial phase relationships in extrafoveal vision. *Nature* 313, 6000 (1985), 308–310.
- [36] FJ Richards. 1959. A flexible growth function for empirical use. *Journal of experimental Botany* 10, 2 (1959), 290–301.
- [37] Snježana Rimac-Drlje, Mario Vranješ, and Drago Žagar. 2010. Foveated mean squared error—a novel video quality metric. *Multimedia Tools and Applications* 49, 3 (Sept. 2010), 425–445. <https://doi.org/10.1007/s11042-009-0442-1>
- [38] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-Net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*. Springer, 234–241.
- [39] Ruth Rosenholtz. 2011. What your visual system sees where you are not looking. In *SPIE: Human Vision and Electronic Imaging XVI*, Bernice E. Rogowitz and Thrasyvoulos N. Pappas (Eds.). San Francisco Airport, California, USA. <https://doi.org/10.1117/12.876659>
- [40] Ruth Rosenholtz. 2016. Capabilities and Limitations of Peripheral Vision. *Annual review of vision science* 2 (2016), 437–457. <https://doi.org/10.1146/annurev-vision-082114-035733>
- [41] Ruth Rosenholtz, Jie Huang, and Krista A Ehinger. 2012. Rethinking the role of top-down attention in vision: Effects attributable to a lossy representation in peripheral vision. *Frontiers in psychology* 3 (2012), 13.
- [42] Ruth Rosenholtz, Jie Huang, Alvin Raj, Benjamin J. Balas, and Livia Ilie. 2012. A summary statistic representation in peripheral vision explains visual search. *Journal of Vision* 12, 4 (April 2012), 14–14. <https://doi.org/10.1167/12.4.14>
- [43] Sanghoon Lee, M. S. Pattichis, and A. C. Bovik. 2002. Foveated video quality assessment. *IEEE Transactions on Multimedia* 4, 1 (March 2002), 129–132. <https://doi.org/10.1109/6046.985561>
- [44] Eero P Simoncelli and Bruno A Olshausen. 2001. Natural image statistics and neural representation. *Annual review of neuroscience* 24, 1 (2001), 1193–1216.
- [45] Karen Simonyan and Andrew Zisserman. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *ICLR*. arXiv:1409.1556 [cs.CV]
- [46] Fabio Solari, Manuela Chessa, and Silvio P Sabatini. 2012. Design strategies for direct multi-scale and multi-orientation feature extraction in the log-polar domain. *Pattern Recognition Letters* 33, 1 (2012), 41–51.
- [47] Michael Stengel, Steve Grogorick, Martin Eisemann, and Marcus Magnor. 2016. Adaptive Image-Space Sampling for Gaze-Contingent Real-time Rendering. *Computer Graphics Forum* 35, 4 (2016), 129–139.
- [48] Hans Strasburger, Ingo Rentschler, and Martin Jüttner. 2011. Peripheral vision and pattern recognition: A review. *Journal of vision* 11, 5 (2011), 13–13.
- [49] Nicholas T. Swafford, José A. Iglesias-Guitian, Charalampos Koniaris, Bochang Moon, Darren Cosker, and Kenny Mitchell. 2016. User, Metric, and Computational Evaluation of Foveated Rendering Methods. In *Proceedings of the ACM Symposium on Applied Perception (SAP '16)*. ACM, New York, NY, USA, 7–14. <https://doi.org/10.1145/2931002.2931011>
- [50] Taimoor Tariq, Okan Tarhan Tursun, Munchurl Kim, and Piotr Didyk. 2020. Why Are Deep Representations Good Perceptual Quality Features?. In *The European Conference on Computer Vision (ECCV)*.
- [51] LN Thibos, FE Cheney, and DJ Walsh. 1987. Retinal limits to the detection and resolution of gratings. *JOSA A* 4, 8 (1987), 1524–1529.
- [52] Huyen T. T. Tran, Duc V. Nguyen, Nam Pham Ngoc, Trang H. Hoang, Truong Thu Huong, and Truong Cong Thang. 2019. Impacts of Retina-related Zones on Quality Perception of Omnidirectional Image. *arXiv:1908.06239 [cs, eess]* (Aug. 2019). arXiv: 1908.06239.
- [53] W. Tsai and Y. Liu. 2014. Foveation-based image quality assessment. In *2014 IEEE Visual Communications and Image Processing Conference*. 25–28. <https://doi.org/10.1109/VCIP.2014.7051495>
- [54] Okan Tarhan Tursun, Elena Arabadzhyska-Koleva, Marek Wernikowski, Radosław Mantiuk, Hans-Peter Seidel, Karol Myszkowski, and Piotr Didyk. 2019. Luminance-contrast-aware foveated rendering. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 1–14.

- [55] Robert A Ulichney. 1993. Void-and-cluster method for dither array generation. In *Human Vision, Visual Processing, and Digital Display IV*, Vol. 1913. International Society for Optics and Photonics, 332–343.
- [56] Mario Vranješ, Snježana Rimac-Drlje, and Denis Vranješ. 2018. Foveation-based content adaptive root mean squared error for video quality assessment. *Multimedia Tools and Applications* 77, 16 (Aug. 2018), 21053–21082. <https://doi.org/10.1007/s11042-017-5544-6>
- [57] Brian Wandell and Stephen Thomas. 1997. Foundations of vision. *Psychocritiques* 42, 7 (1997).
- [58] Panqu Wang and Garrison W Cottrell. 2017. Central and peripheral vision for scene recognition: A neurocomputational modeling exploration. *Journal of vision* 17, 4 (2017), 9–9.
- [59] Zhou Wang, Alan C. Bovik, Ligang Lu, and Jack L. Kouloheris. 2001. Foveated wavelet image quality index. In *Applications of Digital Image Processing XXIV*, Andrew G. Tescher (Ed.). San Diego, CA, 42–52. <https://doi.org/10.1117/12.449797>
- [60] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. 2004. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* 13, 4 (2004), 600–612.
- [61] Zhou Wang, Eero P Simoncelli, and Alan C Bovik. 2003. Multiscale structural similarity for image quality assessment. In *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, Vol. 2. Ieee, 1398–1402.
- [62] Andrew B. Watson. 2014. A formula for human retinal ganglion cell receptive field density as a function of visual field location. *Journal of Vision* 14, 7 (2014), 15–15.
- [63] Martin Weier, Michael Stengel, Thorsten Roth, Piotr Didyk, Elmar Eisemann, Martin Eisemann, Steve Grogorick, André Hinkenjann, Ernst Kruijff, Marcus Magnor, et al. 2017. Perception-driven Accelerated Rendering. *Computer Graphics Forum* 36, 2 (2017), 611–643.
- [64] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 586–595.

# Learning Foveated Reconstruction to Preserve Perceived Image Statistics

Supplementary material

LUCA SURACE, Università della Svizzera italiana, Switzerland

MAREK WERNIKOWSKI, West Pomeranian University of Technology, Poland

OKAN TURSUN, Università della Svizzera italiana, Switzerland

KAROL MYSZKOWSKI, Max Planck Institute for Informatics, Germany

RADOSŁAW MANTIUK, West Pomeranian University of Technology, Poland

PIOTR DIDYK, Università della Svizzera italiana, Switzerland

Foveated image reconstruction recovers full image from a sparse set of samples distributed according to the human visual system’s retinal sensitivity that rapidly drops with eccentricity. Recently, the use of Generative Adversarial Networks was shown to be a promising solution for such a task as they can successfully hallucinate missing image information. Like for other supervised learning approaches, also for this one, the definition of the loss function and training strategy heavily influences the output quality. In this work, we pose the question of how to efficiently guide the training of foveated reconstruction techniques such that they are fully aware of the human visual system’s capabilities and limitations, and therefore, reconstruct visually important image features. Due to the nature of GAN-based solutions, we concentrate on the human’s sensitivity to hallucination for different input sample densities. We present new psychophysical experiments, a dataset, and a procedure for training foveated image reconstruction. The strategy provides flexibility to the generator network by penalizing only perceptually important deviations in the output. As a result, the method aims to preserve perceived image statistics rather than natural image statistics. We evaluate our strategy and compare it to alternative solutions using a newly trained objective metric and user experiments.

## A INTRODUCTION

This document provides supplementary information for ICCV 2021 paper submission 9929: *Learning Foveated Reconstruction to Preserve Perceived Image Statistics*. The material provides additional information regarding the stimuli used in the experiments and the calibration of the image metrics and the obtained parameters.

The anonymous GIT repository <sup>1</sup> containing the source code, also includes an offline web-based demo which provides an interactive version of the results obtained using different training strategies. The interface allows for easy navigation through the foveated images from center to far-periphery and a comparison between different methods. The demo serves as an extension to Figures 1 and 7 from the main paper. To run the demo, the reader is encouraged to open the file “reconstructions.html” in the “demo” folder.

## B ADDITIONAL STIMULI

In the main perceptual experiment (Section 3, main paper), we investigated the sensitivity of the HVS to hallucination. The full set of stimuli used in the experiment is shown in Figure 11 and Figure 12.

---

<sup>1</sup><https://github.com/LFRPPIS/Learning-Foveated-Reconstruction>

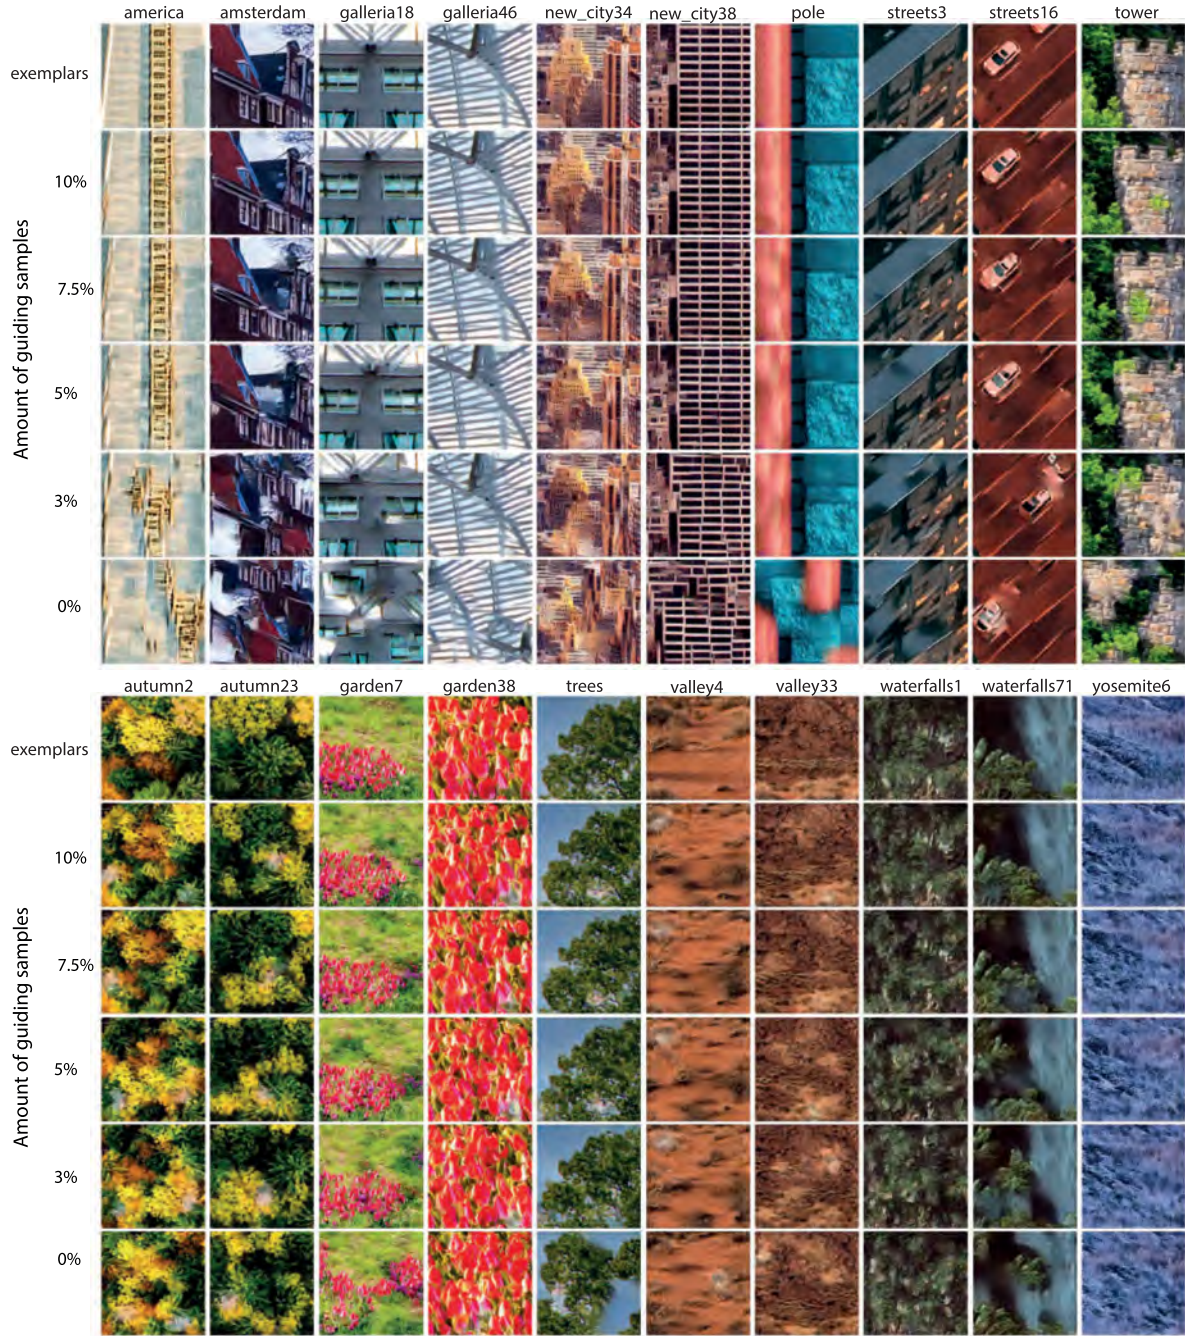


Fig. 11. Results of our stimuli generation for different  $\sigma$  values of the Gaussian filter. Note the increased blur with respect to the exemplar as  $\sigma$  increases.

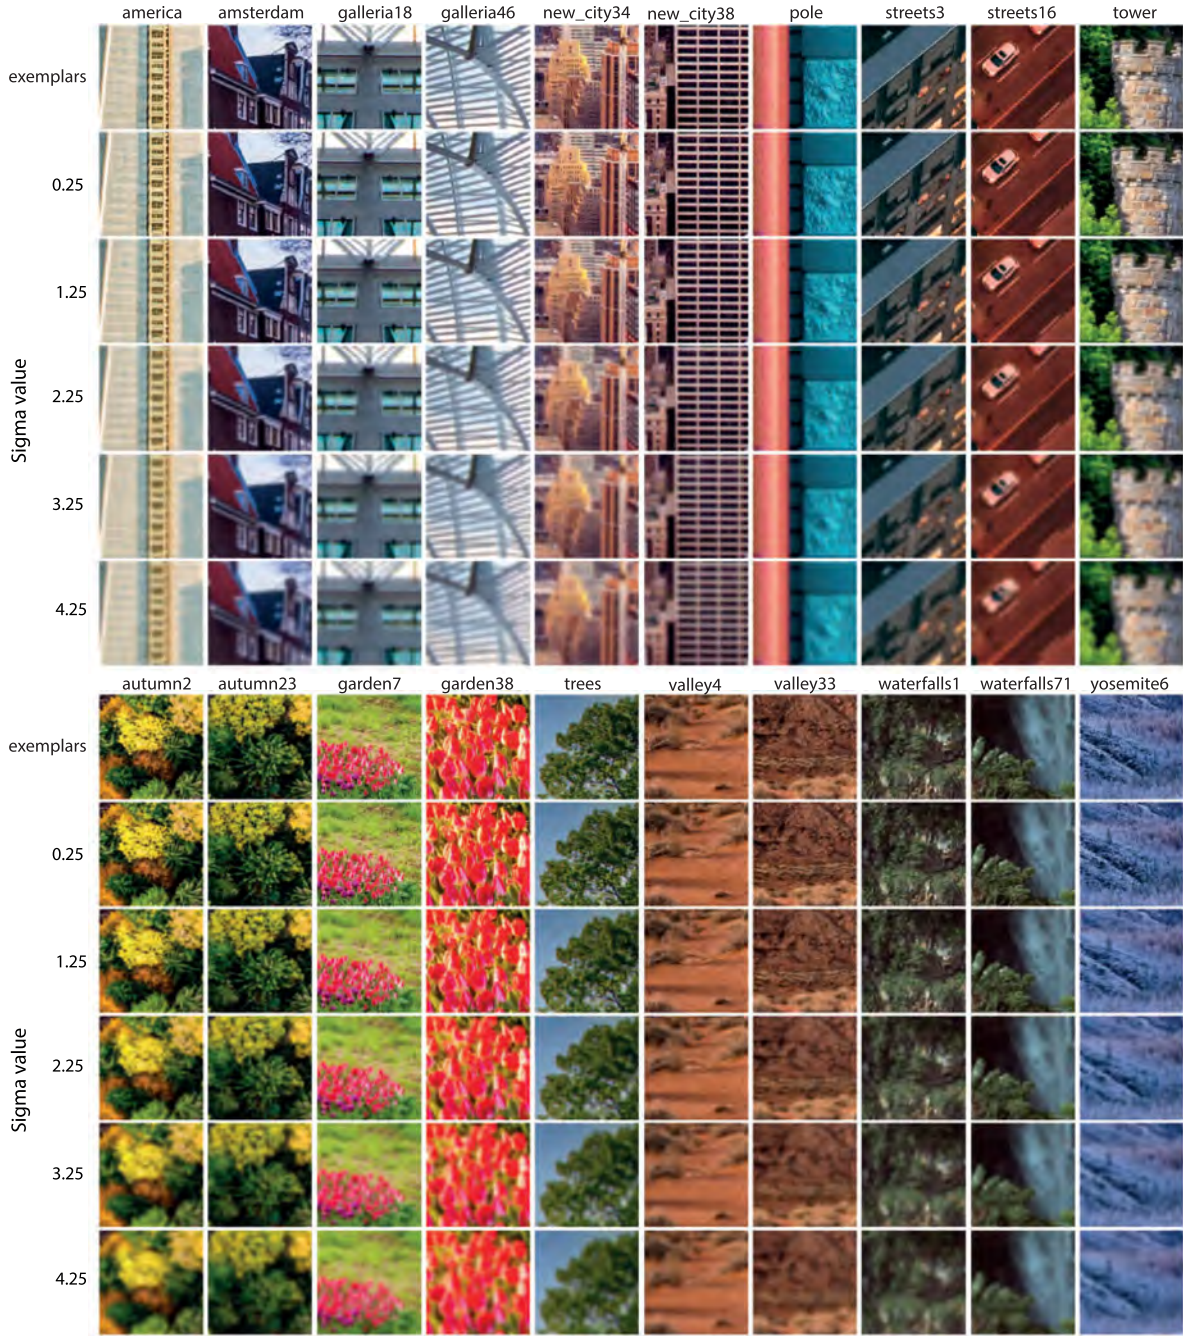


Fig. 12. Results of our stimuli generation for different  $\sigma$  values of the Gaussian filter. Note the increased blur with respect to the exemplar as  $\sigma$  increases.

## C OBJECTIVE IMAGE METRIC

In Section 5.2 of the main paper, we described how we trained multiple objective image metrics to evaluate different training strategies. To this end, we fit the following logistic psychometric function to compute the probability of detection by a human observer from image metrics scores:

$$y(t) = a + \frac{k - a}{(c + q \cdot e^{-b \cdot t})^{\frac{1}{\sigma}}} \quad (6)$$

In Figure 15, we present the fitting result for all of the metrics in 45 subplots. Each column shows the results of fitting the metric. In each column, every three rows correspond to different datasets used for fitting. The first three show the fitting for the images with all types of distortions, the middle three were fitted using only blurred images, and the last three using only images with structural distortions. In those triplets, every individual subplot corresponds to the data gathered at the eccentricity of 8°, 14° and 20°. Every subplot shows the samples used during the training.  $x$ -axis represents the metric’s value whereas  $y$ -axis is the real detection rate of the difference between given images. The red line is the fitted logistic function that has to be applied to the metric in order to obtain the predicted detection rate value.

As described in Section 5.2 of the main paper, the Calibrated VGG metric performed the best in our tests. Consequently, we extended the metric to introduce a predictor for the probability of a human observer detecting the foveation and used it to evaluate different reconstruction methods. We compute the prediction on 256x256 patches separately and average across them. In addition to Figure 10 in the paper, which analyzes the detection rate of different reconstruction methods as a function of far-periphery radius, Figures 14a and 14b present the same data but separately for nature and architectural images. In the figures, it can be observed that the strategy proposed in the paper provides a more significant improvement for nature images. This is manifested by a larger gap between solid lines (our strategy) and dashed lines (standard strategy). This observation is in line with the results of our validation in subjective experiments. Please see the discussion in Section 5.3 of the main paper.

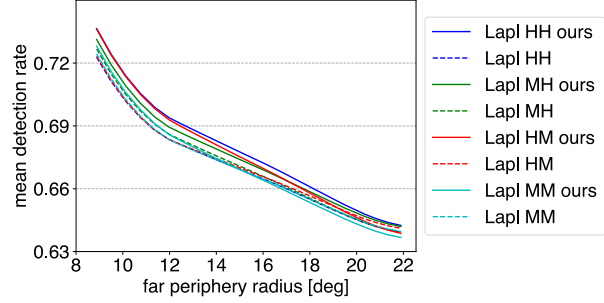


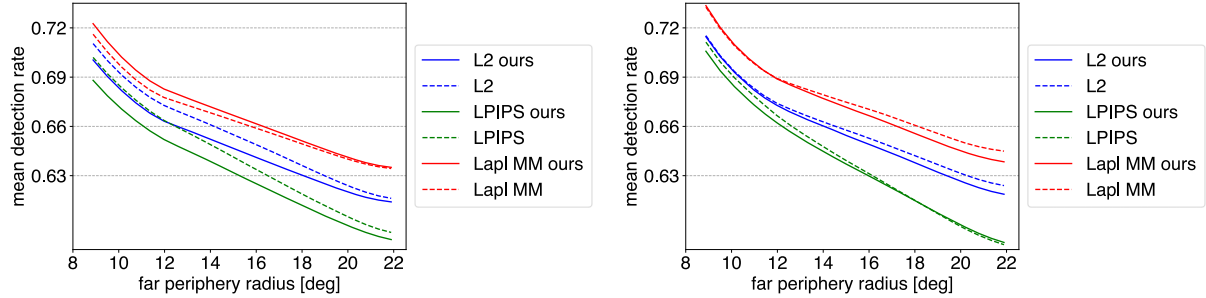
Fig. 13. Detection rate averaged across all images using Calibrated VGG metric, computed for results of different weighting schemes with Laplacian pyramid loss function.

Method name	$a$	$k$	$c$	$q$	$b$	$v$
L2, fitted for blurred images at 8 deg	-0.8275	1.1566	1.0839	0.4307	0.0580	0.9784
L2, fitted for blurred images at 14 deg	0.0603	0.5351	0.0002	1.4677	0.2120	18.0421
L2, fitted for blurred images at 20 deg	0.1296	0.5012	0.0014	1.4559	0.1324	11.1431
L2, fitted for synthesized images at 8 deg	1.0000	0.8578	1.0000	1.0000	1.0000	1.0000
L2, fitted for synthesized images at 14 deg	1.0000	0.7928	1.0000	1.0000	1.0000	1.0000
L2, fitted for synthesized images at 20 deg	1.0000	0.7470	1.0000	1.0000	1.0000	1.0000
L2, fitted for all images at 8 deg	-1.5724	1.1916	1.0003	0.0005	0.0610	0.0024
L2, fitted for all images at 14 deg	-1.5250	1.1680	1.0017	0.0023	0.0240	0.0147
L2, fitted for all images at 20 deg	-1.5994	1.1420	1.0017	0.0031	0.0153	0.0176
LPIPS, fitted for blurred images at 8 deg	0.4115	0.5541	0.1556	2.7971	8.1785	1.3117
LPIPS, fitted for blurred images at 14 deg	0.5261	1.3795	1.0222	0.4079	6.7496	0.0243
LPIPS, fitted for blurred images at 20 deg	0.5038	1.1974	1.0089	1.8666	11.7229	0.0092
LPIPS, fitted for synthesized images at 8 deg	-1.8664	2.0275	1.0035	0.0025	5.4190	0.0113
LPIPS, fitted for synthesized images at 14 deg	-1.7855	1.7645	1.0029	0.0018	3.7249	0.0109
LPIPS, fitted for synthesized images at 20 deg	-1.9407	1.8200	1.0279	0.0304	1.7105	0.1318
LPIPS, fitted for all images at 8 deg	-1.9437	1.9609	1.0067	0.0049	3.7454	0.0240
LPIPS, fitted for all images at 14 deg	-1.6496	1.5661	1.0551	0.0353	3.1857	0.2225
LPIPS, fitted for all images at 20 deg	-1.2599	1.4337	1.2268	0.1641	3.5770	0.7646
MS-SSIM, fitted for blurred images at 8 deg	1.0065	-0.8641	-0.6884	2.0076	0.1594	0.0185
MS-SSIM, fitted for blurred images at 14 deg	1.1278	-0.6995	-0.4912	2.2396	0.3064	0.1335
MS-SSIM, fitted for blurred images at 20 deg	0.8689	-0.5834	0.4690	1.5626	0.8504	0.0981
MS-SSIM, fitted for synthesized images at 8 deg	1.1339	0.9955	-1.3392	2.2630	0.4319	1.1751
MS-SSIM, fitted for synthesized images at 14 deg	0.8920	0.2792	0.5939	1.3321	1.0262	0.1648
MS-SSIM, fitted for synthesized images at 20 deg	0.8476	0.1738	0.8122	1.0888	1.3951	0.1464
MS-SSIM, fitted for all images at 8 deg	0.9162	0.3365	0.1974	1.4729	0.5961	0.0314
MS-SSIM, fitted for all images at 14 deg	0.8490	0.4375	0.5523	1.9660	1.4316	0.1052
MS-SSIM, fitted for all images at 20 deg	0.8198	0.7478	-0.6954	2.4228	0.5294	0.2014
SSIM, fitted for blurred images at 8 deg	0.9719	-0.2663	1.0032	0.8654	3.6542	0.0278
SSIM, fitted for blurred images at 14 deg	0.8564	-0.0452	1.0009	0.0208	4.1712	0.0012
SSIM, fitted for blurred images at 20 deg	0.7681	0.0005	1.0001	0.0037	6.4049	0.0001
SSIM, fitted for synthesized images at 8 deg	0.9441	0.8724	0.4893	3.2029	4.2918	0.2933
SSIM, fitted for synthesized images at 14 deg	0.8616	0.4078	0.9651	2.1697	3.7867	0.1680
SSIM, fitted for synthesized images at 20 deg	0.8086	-0.0623	1.0219	0.5260	4.6636	0.0318
SSIM, fitted for all images at 8 deg	0.9422	0.3296	0.8555	2.3633	2.5179	0.1648
SSIM, fitted for all images at 14 deg	0.8518	0.0553	1.0404	1.0015	4.2294	0.0610
SSIM, fitted for all images at 20 deg	0.7950	0.0058	1.0049	0.1659	5.7480	0.0057
Calibrated VGG, fitted for blurred images at 8 deg	-0.0049	0.1984	0.1912	1.2225	0.9965	0.9955
Calibrated VGG, fitted for blurred images at 14 deg	-0.0043	0.0996	0.0915	1.0736	0.9574	0.9959
Calibrated VGG, fitted for blurred images at 20 deg	-0.0008	1.1439	1.2118	0.4494	0.8092	0.9977
Calibrated VGG, fitted for synthesized images at 8 deg	-0.0135	1.6502	1.7198	4.7628	0.5564	0.9832
Calibrated VGG, fitted for synthesized images at 14 deg	-0.0076	0.8010	0.8718	0.4328	0.9957	0.9935
Calibrated VGG, fitted for synthesized images at 20 deg	-0.0130	0.7177	0.7731	3.2664	0.8968	0.9884
Calibrated VGG, fitted for all images at 8 deg	0.0060	1.6617	1.6213	3.0235	0.9019	0.9989
Calibrated VGG, fitted for all images at 14 deg	0.0038	0.4219	0.4057	1.1894	0.8611	0.9976
Calibrated VGG, fitted for all images at 20 deg	-0.0236	0.9341	0.9850	0.8796	0.9291	0.9805

Table 4. Generalized logistic function parameters for all trained metrics.

Layer	Blurred images			Synthesized images			All images		
	8 deg	14 deg	20 deg	8 deg	14 deg	20 deg	8 deg	14 deg	20 deg
<i>bias</i>	1.587	2.312	-1.246	-7.164e-1	-6.493e-1	1.100	3.713e-1	1.120	-1.006e-1
<b>conv1-1</b>	5.115e-5	3.909e-5	7.193e-4	9.443e-4	5.728e-4	5.235e-5	4.233e-4	5.029e-4	6.661e-4
<b>conv1-2</b>	3.190e-4	1.112e-4	5.738e-4	5.887e-4	2.881e-4	1.807e-4	1.134e-4	1.441e-4	1.126e-5
<b>pool1</b>	3.144e-4	1.044e-4	2.909e-4	3.950e-4	4.605e-4	2.882e-4	1.169e-4	6.371e-5	2.600e-4
<b>conv2-1</b>	1.756e-4	1.875e-5	3.040e-4	2.324e-4	7.695e-6	2.836e-5	1.549e-5	1.375e-5	2.895e-5
<b>conv2-2</b>	2.709e-4	4.976e-5	1.174e-4	1.050e-4	6.867e-5	1.387e-5	1.398e-7	3.212e-5	5.933e-5
<b>pool2</b>	8.101e-6	5.861e-6	2.164e-4	2.938e-4	2.036e-4	1.122e-5	2.769e-5	4.773e-5	4.259e-5
<b>conv3-1</b>	2.220e-5	8.808e-7	4.938e-5	1.586e-4	4.253e-5	2.148e-5	4.129e-6	6.549e-7	9.656e-6
<b>conv3-2</b>	8.622e-5	3.225e-6	1.120e-4	2.631e-4	5.301e-5	6.497e-5	1.524e-5	1.140e-6	3.508e-6
<b>conv3-3</b>	4.218e-4	7.714e-5	1.203e-4	3.956e-4	1.406e-4	1.433e-4	7.521e-5	1.397e-5	4.584e-5
<b>conv3-4</b>	5.386e-4	1.156e-4	1.916e-4	4.445e-4	1.197e-5	2.207e-5	2.353e-5	4.176e-5	1.857e-5
<b>pool3</b>	1.265e-4	2.265e-5	2.915e-4	2.620e-4	4.622e-5	1.164e-5	3.413e-5	1.353e-6	1.388e-5
<b>conv4-1</b>	6.064e-6	2.436e-6	4.221e-6	4.415e-4	2.184e-5	2.459e-5	1.168e-4	9.106e-13	1.469e-6
<b>conv4-2</b>	1.163e-4	1.994e-5	6.841e-5	5.248e-3	8.723e-5	1.057e-4	1.011e-3	1.339e-5	2.908e-5
<b>conv4-3</b>	3.092e-4	9.000e-5	2.400e-4	2.119e-3	1.621e-4	9.065e-5	2.025e-4	2.260e-4	1.295e-4
<b>conv4-4</b>	2.883e-4	1.505e-4	3.017e-4	1.203e-3	1.290e-4	2.036e-4	1.449e-4	4.483e-4	1.398e-4
<b>pool4</b>	6.958e-4	4.035e-4	3.000e-3	1.318e-3	1.411e-4	6.272e-4	5.519e-4	2.970e-3	9.817e-3
<b>conv5-1</b>	3.772e-3	7.210e-4	9.182e-4	2.443e-3	8.311e-4	1.433e-3	5.852e-4	2.003e-3	6.144e-4
<b>conv5-2</b>	2.356e-3	3.301e-3	2.616e-4	5.079e-4	1.040e-2	7.866e-3	3.365e-3	1.694e-3	7.141e-4
<b>conv5-3</b>	8.696e-4	8.612e-4	8.410e-4	1.352e-3	8.421e-3	1.015e-2	6.328e-3	8.783e-4	5.821e-4
<b>conv5-4</b>	1.626e-2	1.805e-2	5.045e-4	2.617e-3	1.583e-2	4.425e-2	4.423e-2	2.114e-2	2.848e-3
<b>pool5</b>	7.985e-3	9.232e-3	7.596e-4	5.953e-4	2.428e-3	1.960e-3	6.607e-4	1.357e-2	1.039e-3

Table 5. Weights for all layers of calibrated VGG network.



(a) Detection rate averaged across nature images using Calibrated VGG metric.

(b) Detection rate averaged across architectural images using Calibrated VGG metric.

Fig. 14. A discussion about the visible (a) and small (b) difference between dashed and solid lines is provided in Section 5.3 of the main paper.

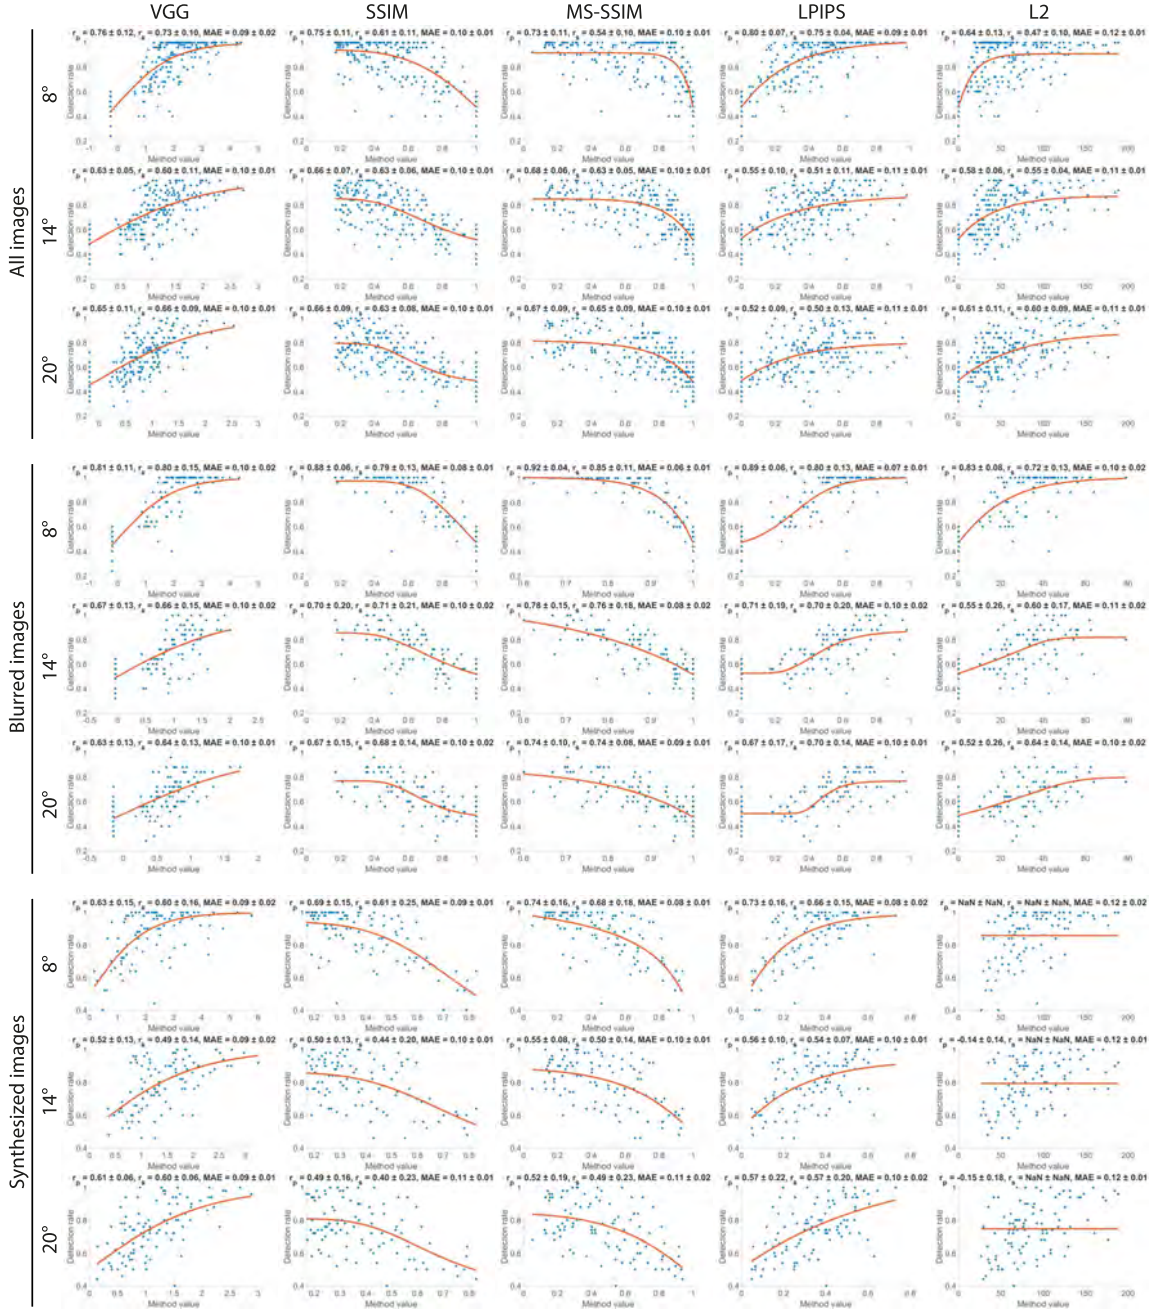


Fig. 15. Plots showing fitting and correlation coefficients of all trained metrics.  $r_p$ ,  $r_s$  and MAE stand for Pearson correlation coefficient, Spearman correlation coefficient, and mean absolute error, respectively.

1. fold				
trees, s. 0%	waterfalls1, s. 0%	america, s. 10%	autumn2, s. 10%	galleria18, s. 10%
new_city38, s. 10%	trees, s. 10%	waterfalls71, s. 10%	america, s. 3%	garden38, s. 3%
new_city34, s. 3%	pole, s. 3%	streets16, s. 3%	tower, s. 3%	amsterdam, s. 5%
garden7, s. 5%	valley33, s. 5%	galleria46, s. 7.5%	tower, s. 7.5%	waterfalls1, s. 7.5%
yosemite6, s. 7.5%	autumn23, b. 0.25 $\sigma$	waterfalls1, b. 0.25 $\sigma$	amsterdam, b. 1.25 $\sigma$	galleria46, b. 1.25 $\sigma$
new_city38, b. 1.25 $\sigma$	tower, b. 1.25 $\sigma$	trees, b. 1.25 $\sigma$	new_city38, b. 2.25 $\sigma$	pole, b. 2.25 $\sigma$
america, b. 3.25 $\sigma$	amsterdam, b. 3.25 $\sigma$	tower, b. 3.25 $\sigma$	valley4, b. 3.25 $\sigma$	streets16, b. 4.25 $\sigma$
streets3, b. 4.25 $\sigma$	tower, b. 4.25 $\sigma$	valley4, b. 4.25 $\sigma$		
2. fold				
america, s. 0%	pole, s. 0%	amsterdam, s. 10%	garden38, s. 10%	tower, s. 10%
valley4, s. 10%	waterfalls1, s. 10%	autumn23, s. 3%	valley33, s. 3%	valley4, s. 3%
yosemite6, s. 3%	america, s. 5%	galleria18, s. 5%	garden38, s. 5%	new_city34, s. 5%
tower, s. 5%	autumn2, s. 7.5%	garden38, s. 7.5%	america, b. 0.25 $\sigma$	galleria18, b. 0.25 $\sigma$
streets16, b. 0.25 $\sigma$	valley4, b. 0.25 $\sigma$	waterfalls71, b. 0.25 $\sigma$	autumn2, b. 1.25 $\sigma$	galleria18, b. 1.25 $\sigma$
streets16, b. 1.25 $\sigma$	streets3, b. 1.25 $\sigma$	valley33, b. 1.25 $\sigma$	amsterdam, b. 2.25 $\sigma$	autumn2, b. 2.25 $\sigma$
galleria18, b. 2.25 $\sigma$	garden38, b. 2.25 $\sigma$	waterfalls71, b. 2.25 $\sigma$	autumn23, b. 3.25 $\sigma$	galleria18, b. 3.25 $\sigma$
galleria46, b. 3.25 $\sigma$	new_city34, b. 3.25 $\sigma$	streets3, b. 3.25 $\sigma$	trees, b. 3.25 $\sigma$	autumn2, b. 4.25 $\sigma$
pole, b. 4.25 $\sigma$	waterfalls1, b. 4.25 $\sigma$			
3. fold				
autumn23, s. 0%	new_city34, s. 0%	new_city38, s. 0%	valley4, s. 0%	waterfalls71, s. 0%
galleria46, s. 10%	new_city34, s. 10%	pole, s. 10%	streets16, s. 10%	streets3, s. 10%
valley33, s. 10%	amsterdam, s. 3%	autumn2, s. 3%	waterfalls1, s. 3%	galleria46, s. 5%
trees, s. 5%	waterfalls1, s. 5%	autumn23, s. 7.5%	trees, s. 7.5%	garden7, b. 0.25 $\sigma$
new_city38, b. 0.25 $\sigma$	pole, b. 0.25 $\sigma$	tower, b. 0.25 $\sigma$	autumn23, b. 1.25 $\sigma$	garden7, b. 1.25 $\sigma$
galleria46, b. 2.25 $\sigma$	garden7, b. 2.25 $\sigma$	streets3, b. 2.25 $\sigma$	tower, b. 2.25 $\sigma$	trees, b. 2.25 $\sigma$
valley33, b. 2.25 $\sigma$	valley4, b. 2.25 $\sigma$	yosemite6, b. 2.25 $\sigma$	yosemite6, b. 3.25 $\sigma$	america, b. 4.25 $\sigma$
amsterdam, b. 4.25 $\sigma$	autumn23, b. 4.25 $\sigma$	garden38, b. 4.25 $\sigma$	valley33, b. 4.25 $\sigma$	waterfalls71, b. 4.25 $\sigma$
yosemite6, b. 4.25 $\sigma$				
4. fold				
amsterdam, s. 0%	autumn2, s. 0%	galleria46, s. 0%	garden7, s. 0%	streets16, s. 0%
tower, s. 0%	garden7, s. 10%	galleria18, s. 3%	galleria46, s. 3%	new_city38, s. 3%
trees, s. 3%	autumn23, s. 5%	new_city38, s. 5%	pole, s. 5%	streets16, s. 5%
streets3, s. 5%	waterfalls71, s. 5%	yosemite6, s. 5%	amsterdam, s. 7.5%	galleria18, s. 7.5%
garden7, s. 7.5%	new_city34, s. 7.5%	pole, s. 7.5%	valley33, s. 7.5%	waterfalls71, s. 7.5%
autumn2, b. 0.25 $\sigma$	galleria46, b. 0.25 $\sigma$	garden38, b. 0.25 $\sigma$	new_city34, b. 0.25 $\sigma$	streets3, b. 0.25 $\sigma$
yosemite6, b. 0.25 $\sigma$	america, b. 1.25 $\sigma$	garden38, b. 1.25 $\sigma$	pole, b. 1.25 $\sigma$	new_city34, b. 2.25 $\sigma$
streets16, b. 2.25 $\sigma$	waterfalls1, b. 2.25 $\sigma$	autumn2, b. 3.25 $\sigma$	garden7, b. 3.25 $\sigma$	pole, b. 3.25 $\sigma$
waterfalls1, b. 3.25 $\sigma$	galleria46, b. 4.25 $\sigma$	trees, b. 4.25 $\sigma$		
5. fold				
galleria18, s. 0%	garden38, s. 0%	streets3, s. 0%	valley33, s. 0%	yosemite6, s. 0%
autumn23, s. 10%	yosemite6, s. 10%	garden7, s. 3%	streets3, s. 3%	waterfalls71, s. 3%
autumn2, s. 5%	valley4, s. 5%	america, s. 7.5%	new_city38, s. 7.5%	streets16, s. 7.5%
streets3, s. 7.5%	valley4, s. 7.5%	amsterdam, b. 0.25 $\sigma$	trees, b. 0.25 $\sigma$	valley33, b. 0.25 $\sigma$
new_city34, b. 1.25 $\sigma$	valley4, b. 1.25 $\sigma$	waterfalls1, b. 1.25 $\sigma$	waterfalls71, b. 1.25 $\sigma$	yosemite6, b. 1.25 $\sigma$
america, b. 2.25 $\sigma$	autumn23, b. 2.25 $\sigma$	garden38, b. 3.25 $\sigma$	new_city38, b. 3.25 $\sigma$	streets16, b. 3.25 $\sigma$
valley33, b. 3.25 $\sigma$	waterfalls71, b. 3.25 $\sigma$	galleria18, b. 4.25 $\sigma$	garden7, b. 4.25 $\sigma$	new_city34, b. 4.25 $\sigma$
new_city38, b. 4.25 $\sigma$				

Table 6. Cross validation folds used for the fitting. **b.** is an abbreviation for *blurred*, **s.** for *guided samples*.