

# In-N-Out: Towards Good Initialization for Inpainting and Outpainting

Changho Jo  
timegate@kaist.ac.kr

Woobin Im  
iwbn@kaist.ac.kr

Sung-Eui Yoon  
sungeui@kaist.edu

School of Computing,  
Korea Advanced Institute of Science  
and Technology (KAIST),  
Daejeon, Korea

---

## Abstract

In computer vision, recovering spatial information by filling in masked regions, e.g., inpainting, has been widely investigated for its usability and wide applicability to other various applications: image inpainting, image extrapolation, and environment map estimation. Most of them are studied separately depending on the applications. Our focus, however, is on accommodating the opposite task, e.g., image outpainting, which would benefit the target applications, e.g., image inpainting. Our self-supervision method, In-N-Out, is summarized as a training approach that leverages the knowledge of the opposite task into the target model. We empirically show that In-N-Out – which explores the complementary information – effectively takes advantage over the traditional pipelines where only task-specific learning takes place in training. In experiments, we compare our method to the traditional procedure and analyze the effectiveness of our method on different applications: image inpainting, image extrapolation, and environment map estimation. For these tasks, we demonstrate that In-N-Out consistently improves the performance of the recent works with In-N-Out self-supervision to their training procedure. Also, we show that our approach achieves better results than an existing training approach for outpainting.

## 1 Introduction

Learning to write a paper helps the learner better fill in blanks in a sentence. Learning to draw from scratch gives the ability to fill in the masks, i.e., inpainting. Often, an educational curriculum for humans gives learners such tasks, in the order we mentioned, or vice versa.

Thus, some tasks are complementary; each task is mutually cooperative with each other so that learning one task is eventually good for the other tasks. So are tasks of computer vision, which have been studied in the form of transfer learning (pretraining) [11, 52], multi-task learning [16, 27], and continual learning [42, 46].

In this work, we introduce a transfer learning strategy, In-N-Out, for both image inpainting and image outpainting tasks. In-N-Out is motivated by the complementary relationship between image inpainting and image outpainting tasks. For instance, changing an outpainting task to an inpainting task [17] and cycle consistency between both tasks [15] are

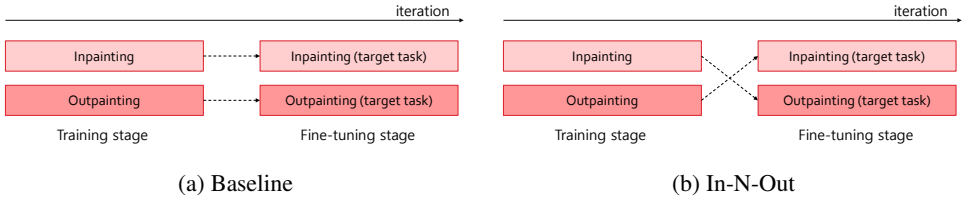


Figure 1: **(a)**: General procedure (baselines) for inpainting and outpainting, **(b)**: In-N-Out approach for inpainting and outpainting.

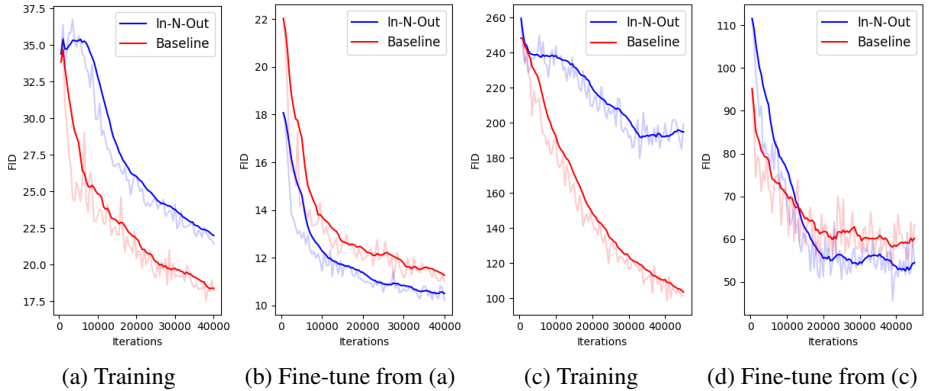


Figure 2: **(a)-(b)**: training procedure for image inpainting task ( $\mathcal{T}_{in}$ ). Each graph represents **(a)** training stage from scratch, and **(b)** fine-tuning stage. **(c)-(d)**: training procedure for image outpainting task ( $\mathcal{T}_{out}$ ). Each graph represents **(c)** training stage from scratch, and **(d)** fine-tuning stage. Each graph shows FID (Fréchet inception distance) [11] during each stage. Though our method, In-N-Out, shows worse performance in the training stage, In-N-Out converges more fastly and shows better performance in the end, where each method learns each target task by fine-tuning. Interestingly, good initialization does not require having a good metric score (FID) before fine-tuning. Also, it can be confirmed that In-N-Out (blue line) is an approach that can be well applied to various filling mask tasks such as inpainting and outpainting.

introduced. Inspired by this, we propose our method, In-N-Out, which is a simple and general training practice that can be applied to a network quickly. In short, our method In-N-Out learns an outpainting task for the initialization of a model for an inpainting task, and learns an inpainting task for the initialization of a model for an outpainting task as in Figure 1.

Figure 2 shows the test results of each training iteration. In Figure 2a and Figure 2c, although In-N-Out shows worse performance as it does not learn the target task in the training stage, it shows the faster convergence and better performance than the baseline in the fine-tuning stage as shown in Figure 2b and Figure 2d. Note that the baseline learns the target task in both the training and the fine-tuning stage, i.e., inpainting  $\rightarrow$  inpainting for inpainting task.

To summarize, our contributions are as follows:

- This work provides a simple and general self-supervised pretraining method for both inpainting and outpainting tasks.



Figure 3: Image outpainting (or image extrapolation) shares the same pipeline as image inpainting, where the masked input image becomes the completed image.

- We analyze and provide extensive experiments to validate In-N-Out procedure in both image inpainting and image outpainting tasks.
- We show the effectiveness of our method on various applications: inpainting task, image extrapolation, and environment map estimation.

## 2 Related Work

**Deep Learning-based Inpainting and Outpainting.** Image inpainting is aimed at filling holes in images. In a general image inpainting pipeline where the masked input image is filled in, generative adversarial networks (GANs) [8] are mainly employed, and various techniques have been developed in terms of network architectures [27, 56, 41] and loss functions [12, 25, 27, 52, 57] to infer the invisible masked region better from visible known regions. New layers have been proposed to transfer textures from input images, including shift-connection layer [56] and contextual attention layer [41]. In addition, loss functions have been proposed in terms of global and local consistency [12, 57], and patch distribution [25, 52].

On the other hand, image outpainting fills outer masked regions from visible inner regions. Similar to the image inpainting task, image outpainting has been developed mainly in terms of network architectures [53, 59] and loss functions [30, 53, 59]. Outer propagation of inner visible features has been proposed with a recurrent module [59] and a normalization module [53]. WGAN [9] losses are mainly used for outpainting methods [30, 53, 59] and spatially variant losses are applied [53].

Apart from the architecture and losses, we focus on the complementary relationship between both tasks: can knowledge from one side give an advantage to the other side? Regarding that both tasks do not require a specific label, i.e., self-supervised, exploiting the opposite task is feasible and desirable. It has been shown that exploring across the tasks is beneficial by solving outpainting with inpainting [17], or learning cycle-consistency between both tasks [15]. We step further by simplifying the cross-task learning and generalize it to both tasks with transfer learning, which is a well-known general strategy in deep learning [43].

Specifically, our method (In-N-Out) lets a model first experience the opposite task (e.g., inpainting), which is transferred to the target task (e.g., outpainting). We show by the experiment (Table 1c) that our simple transfer strategy gives an equivalent or a larger amount of performance gain to the task-specific (i.e., horizontal extrapolation) method [17], which also exploits the cross-task relation. Compared to common practice for the inpainting or outpainting, where pretraining on the target task is performed, our method gives no overhead, since it substitutes the pretraining task with the opposite task without additional overhead.

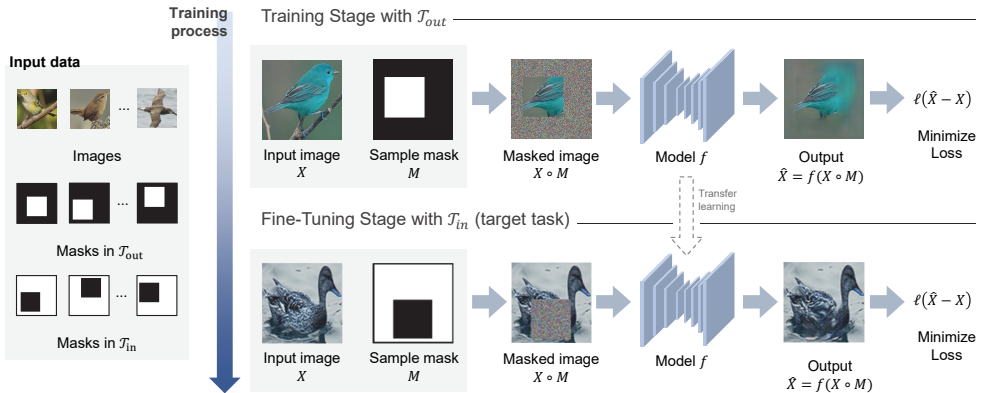


Figure 4: In-N-Out training process for inpainting task.

**Self-supervised Learning.** Self-supervised learning has been proposed to learn feature representations from annotation-free images. For instance, image transformation [5], rotation [19], position of patches [2], and color [25] can be used as surrogate labels obtained from the image itself for feature learning. While a majority of self-supervision methods is validated on classification tasks, different kinds of features fit better for different tasks [43]. Therefore, specific self-supervision methods have been proposed for better performance in different tasks: optical flow [13, 24], reflection removal [18, 38], image inpainting [27, 32], image outpainting [28, 30], and light estimation [6, 29].

Our method is regarded as a self-supervised method, which is defined by performing the opposite task (e.g., inpainting) for the target task (e.g., outpainting). Our self-supervision is devised to help the ‘filling masked region’ task, including inpainting, extrapolation, and environment map estimation.

### 3 In-N-Out: Inpainting and Outpainting

In this section, we propose a training scheme, In-N-Out, for the inpainting and outpainting tasks. The key idea of our method is to enable the knowledge from a counterpart task; the counterpart task is defined as inpainting if the main task is outpainting, and outpainting if the target task is inpainting (Fig. 1). We found that transfer learning, which has been widely applied and investigated by computer vision researchers, shows a fair amount of effectiveness in our purpose. Thus, we divide the whole training process into two parts: **training** with the counterpart and **fine-tuning** with the target task (Fig. 1).

Specifically, let  $\mathcal{T}$  denote a task, which is a data distribution for the general ‘fill masked region’ task. That is,  $\mathcal{T}$  is a joint distribution of images and masks, where we sample image  $X$  and mask  $M$  from distribution  $\mathcal{T}$  to perform a self-supervision task; the task is to learn to restore  $X$  from the input that masked by mask  $M$ . We denote inpainting task  $\mathcal{T}_{in}$  and outpainting task  $\mathcal{T}_{out}$ . Inpainting task  $(X, M) \sim \mathcal{T}_{in}$  generates sample mask  $M$  whose masked region is local, which is often called as a hole (or holes) in an image (Fig. 3a). On the other hand, outpainting mask  $M$  from distribution  $\mathcal{T}_{out}$  has the visible region locally, and the outer regions are masked (Fig. 3b). For each task, a prediction model  $f_{\theta}$  can be trained by minimizing loss:

$$\mathcal{L}(\theta) = \mathbb{E}_{(X, M) \sim \mathcal{T}} [\ell(f_{\theta}(X \circ M), X)], \quad (1)$$

where  $\circ$  is the Hadamard product operator with filling operation for the masked region; the region is filled with normally uniform value [27, 40] or random noise [9, 40]. Note that  $\ell(A, B)$  is a loss function to measure a distance between  $A$  and  $B$ , e.g., L1 loss. Our loss includes the adversarial loss, ID-MRF loss [32], etc., depending on the target tasks; the details are mentioned as implementation details in Section 4.

We transfer knowledge from opposite tasks ( $\mathcal{T}_{\text{out}}$  and  $\mathcal{T}_{\text{in}}$ ), and define our process as follows. If the target task is  $\mathcal{T}_{\text{in}}$ , we run the first  $N$  steps with  $\mathcal{T}_{\text{out}}$  with loss function  $\mathcal{L}(\theta)$  in Eq. 1. Then, we continue to fine-tune on the target task  $\mathcal{T}_{\text{in}}$  for  $K$  steps additional to the training step. This simple rule applies vice versa when  $\mathcal{T}_{\text{out}}$  is the target. In short,  $\mathcal{T}_{\text{out}} \rightarrow \mathcal{T}_{\text{in}}$  for inpainting task, and  $\mathcal{T}_{\text{in}} \rightarrow \mathcal{T}_{\text{out}}$  for outpainting task.

Intuition for transferability is based on the complementary property between inpainting and outpainting. Inpainting and outpainting have been taken to account together [13, 17]. We simplify and show the relationship in the form of transfer learning. We empirically show that this simple approach can have an equivalent or better performance than the task-specific method [17].

**Transfer learning.** It has been shown that visual tasks help each other [43], e.g., learning surface normal  $\rightarrow$  learning depth, which is effectively shown with transfer learning. Our work aims to discover the relationship between inpainting and outpainting with transfer learning, and it provides an effective yet simple methodology for an inpainting or an outpainting task.

**Complementarity.** Inpainting fills a hole region by seamlessly expanding a visible outer area into the hole and outpainting transfers a context from a small visible region to an invisible outer region. The difference between the hidden regions – a hole or an outer area – makes an inpainting model focus on low-level restoration of the hole and an outpainting network focus on deeper understanding of context, rather than low-level prediction. Thus, learning from both tasks lets the model have a sort of synergy by benefiting from both high and low level knowledge. Similarly, in literature, an outpainting task is often regarded as a more difficult one than inpainting from the perspective of a model. Regarding this, In-N-Out can be viewed as curriculum learning [10] or reverse curriculum learning [30].

## 4 Experiments & Results

We evaluate our method and show its generality on various applications in both inpainting and outpainting tasks. The evaluation tasks include image inpainting, image outpainting, and environment map estimation. Throughout all tasks, we follow the scheme described in Sec. 3 and Fig. 1. That is, for an outpainting task (e.g., environment map estimation), we use the inpainting task before fine-tuning. We start a model from scratch with the self-supervised training, followed by the fine-tuning step for each desired task. For all experiments, we use a variant of Semantic Regeneration Network [33], unless otherwise specified. We provide the architecture details in the supplementary material. Also, for inpainting and outpainting, many types of masks have been used. For instance, uniform values (to black or white), or noises are used. We use uniform noise that randomly chooses all possible color values unless otherwise stated. In the subsequent subsections, we describe the evaluation results of each task.

Method	PSNR	SSIM	FID
Baseline	21.03	0.830	11.02
<b>In-N-Out</b>	<b>21.14</b>	<b>0.833</b>	<b>10.22</b>

(a) Image inpainting task.

Method	PSNR	SSIM	FID
Baseline	<b>15.11</b>	0.552	63.45
<b>In-N-Out</b>	15.10	<b>0.554</b>	<b>54.13</b>

(b) Image outpainting task.

Method	PSNR	SSIM	FID
SRN [63]	18.22	0.513	-
PSL [14]	18.78	<b>0.716</b>	35.86
<b>In-N-Out</b>	<b>19.52</b>	0.711	<b>30.17</b>

(c) Outpainting compared to PSL [14].

Method	PSNR	SSIM	FID
Baseline	24.83	0.813	63.63
<b>In-N-Out</b>	<b>24.86</b>	<b>0.817</b>	<b>59.91</b>

(d) Inpainting using irregular masks [14], and MEDFE [14] network.

Method	PSNR	SSIM	FID
Baseline	31.64	0.950	29.63
<b>In-N-Out</b>	<b>32.26</b>	<b>0.955</b>	<b>22.46</b>

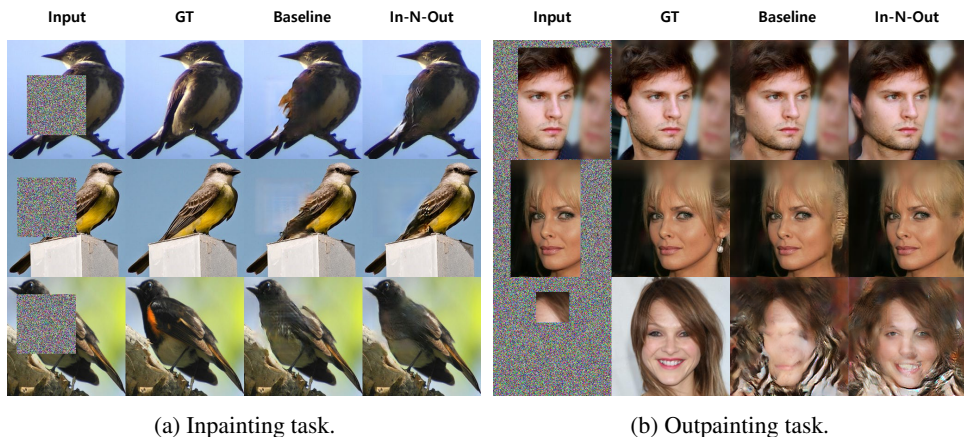
(e) Inpainting using irregular masks [22], and Shift-Net [64].

Table 1: (a) Results of In-N-Out (outpainting  $\rightarrow$  inpainting) and baseline (inpainting  $\rightarrow$  inpainting) for the inpainting task. (b) Results of In-N-Out (inpainting  $\rightarrow$  outpainting) and baseline (outpainting  $\rightarrow$  outpainting) for the outpainting task. (c) Results of In-N-Out compared to PSL [14], for the outpainting task. (d) Results of In-N-Out and baseline for the inpainting task, using irregular masks [14] on MEDFE network. (e) Results of In-N-Out and baseline for the inpainting task, using irregular masks [22] on Shift-Net. In-N-Out can be a good practice by leveraging the knowledge from the opposite task.

## 4.1 Image Inpainting

**Experimental Setting.** For the inpainting task, we tested our models on the CUB200 dataset [65], consisting of 11,788 bird images. We randomly crop the images with given bird locations and resize the images into  $256 \times 256$ . For the size of the random mask, i.e., blocked region, we use resolution  $128 \times 128$ . The training and test splits are the same as in [63]. For In-N-Out, the model goes through the schedule of outpainting  $\rightarrow$  inpainting (Fig. 4). On the other hand, the baseline for comparison has the schedule that has the same training scheme to the fine-tuning step, i.e., inpainting  $\rightarrow$  inpainting, which has been a customary process for many computer vision tasks. For each stage, we use the strategies in [63], which is training with reconstruction losses, then adding the adversarial loss and ID-MRF loss [62] in the fine-tuning stage. We run 40,000 iterations of training steps and 40,000 iterations of fine-tuning steps with batch size 8. In the training stage, baseline uses  $128 \times 128$  mask and learns to paints inside the region, while In-N-Out uses inverse mask and learns to paint outside of the  $128 \times 128$  region. Both baseline and In-N-Out learn to inpaint in the fine-tuning stage.

**Quantitative Results.** We provide image inpainting results on CUB200 dataset in Table 1a. When compared to the baseline, improvements can be observed in all three metrics: PSNR, SSIM, and Frchet Inception Distance (FID) [14]. While the amount of improvements are not significant in PSNR and SSIM, FID is shown to be improved in a clear gap ( $11.02 \rightarrow 10.22$ ). That is because FID measures the quality of generated images, which is



(a) Inpainting task.

(b) Outpainting task.

Figure 5: Qualitative results on inpainting task (CUB200 dataset [53]) and outpainting task (CelebA-HQ dataset [14]).

considered more important in image inpainting; PSNR or SSIM can possibly be worse even with higher quality images as in also noted by [53]. This performance improvement is supported by the training graph (Fig. 2a and 2b). Even though In-N-Out shows higher FID in the initial training stage (Fig. 2a), In-N-Out converges faster in the fine-tuning stage (Fig. 2b), since it gives better initialization than the baseline.

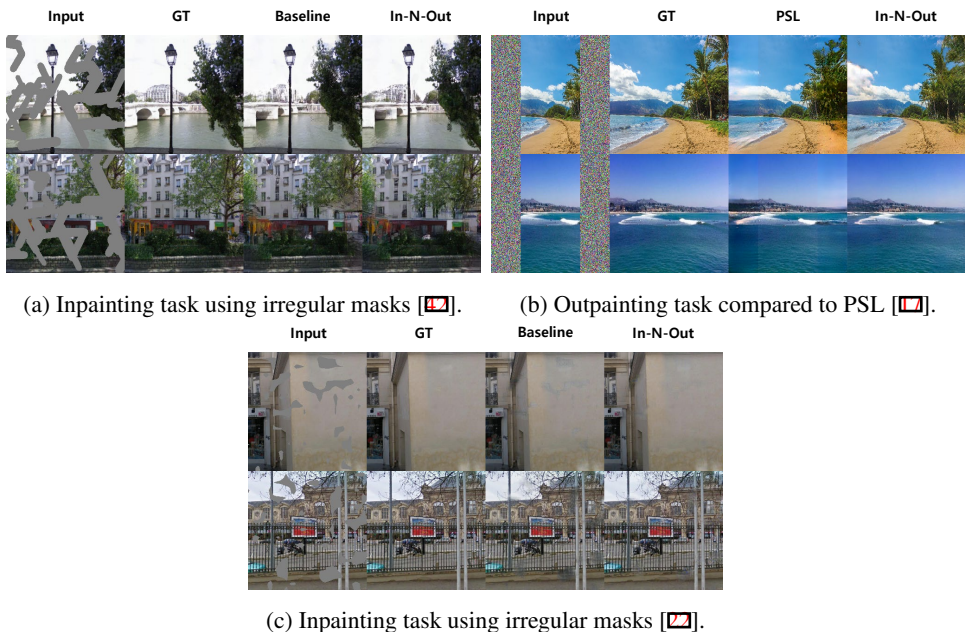
**Qualitative Results.** The visual comparisons are shown in Figure 5a. Compared to the baseline, In-N-Out shows more convincing results. In-N-Out better restores the overall shapes of the birds – including their wings and bodies – with reasonable color tones, where the baseline often fails to create continuous shape and suffers from some artifacts. More results can be found in our supplementary material.

## 4.2 Image Outpainting

**Experimental Setting.** For the outpainting task, we test the models on CelebA-HQ dataset [14] which consists of 30,000 face images. To show resiliency of our approach to the mask size, we use rectangular masks with random sizes for our testing. The details of used masks are described in the supplementary material. We use the official training and test split of the dataset, and we resize images to  $256 \times 256$  for both training and testing. We train with reconstruction losses, then adversarial loss and ID-MRF loss are added in the fine-tuning stage as in Section 4.1. With batch size 8, we run 45,000 iterations of training steps and 45,000 iterations of fine-tuning. In the training stage, the baseline learns to paint outside of a  $128 \times 128$  inner region, while In-N-Out learns to inpaint. Both baseline and In-N-Out learn to outpaint in the fine-tuning stage.

**Quantitative Results.** Test FIDs are provided in Figure 2c and 2d, and the numerical results are reported in Table 1b. As shown in Table 1b, In-N-Out shows better results in terms of FID. Interestingly, In-N-Out shows to be more effective on the outpainting task than on the inpainting task (Table 1a). This is partly because, inpainting tends to be learned easier than outpainting, since pixels to be inpainted can get more directional cues [17], so that the GAN is able to be stabilized faster. This is also observed in the graph (Fig. 2d), where the





(a) Inpainting task using irregular masks [17].

(b) Outpainting task compared to PSL [17].

Input GT Baseline In-N-Out



(c) Inpainting task using irregular masks [17].

Figure 6: Qualitative results on inpainting task (Paris Street-View dataset [3]) and outpainting task (beach dataset [28]). In-N-Out shows more seamless and semantically plausible results.

baseline FID measure doesn’t improve much in the fine-tuning stage, even with its lower error in the training stage (Fig. 2c). This is also supported by the results reported by the original SRN paper [53], where quantitative measures get worsen<sup>1</sup> after finetuning on SRN, which is our baseline network. On the other hand, In-N-Out is shown to ameliorate the issue (Fig. 2d), and shows better quantitative results after fine-tuning.

**Qualitative Results.** We also provide visual comparisons in Figure 5b. In-N-Out shows plausible restoration of hair, mouth, and ear. While baseline suffers from some artifacts to far invisible region from the mask, In-N-Out shows plausible images to the far invisible region. More results are available in our supplementary material.

### 4.3 Comparison to PSL [17]

**Experimental Setting.** We provide evaluation compared to progressive step learning (PSL) [17], which is a strategy where an outpainting task is intentionally substituted by an inpainting task. We compare our method with PSL, since PSL gives knowledge from the opposite task (i.e., inpainting) for the target task (i.e., outpainting), similar to In-N-Out. In short, PSL changes the outpainting task to the inpainting task and performs progressive inpainting. Please refer to [17] for details. Compared to PSL, our method is simple and general, since we do not constrain our method to a specific task. To compare with the approach, we train the model on the beach dataset [28], which contains 10,515 images of nature landscapes. With

<sup>1</sup>Please refer to Table 5 in [53] for more details.



batch size 8, we run 40,000 iterations of training steps and 40,000 iterations of fine-tuning steps.

**Results.** In Table 1c, we report the performance of our baseline network (SRN), which is originally reported in [14], progressive step learning (PSL) [14], and In-N-Out (ours). First, we notice that In-N-Out clearly improves over the baseline network (SRN) on PSNR and SSIM metrics. When compared to a task-specific method (PSL), In-N-Out shows clearly better results in PSNR and FID, while SSIM is in a similar range. As mentioned in the earlier sections, FID has been regarded as more important than PSNR or SSIM in generation tasks, since it better measures the quality of generated images. In Figure 6b, In-N-Out predicts seamless (e.g., color of sky) and more naturally extrapolated scenery (e.g., clouds) than PSL.

## 4.4 Image Inpainting with Irregular Masks

**Experimental Setting.** We also test our method on inpainting tasks with irregular masks, from DeepFillv2 [12] and PConv [12]. We use Paris Street-View dataset [3], which consists of 15,000 street-view images. We use masks with a uniform constant and networks designed for inpainting, i.e., MEDFE [23] and Shift-Net [36]. The details of experiment setting are available in our supplementary material.

**Results.** We provide quantitative results in Table 1d and Table 1e. In-N-Out shows improvements from the baseline in terms of FID. Also, in Figure 6a and Figure 6c, it shows to predict background structures (first row of 6a) and shape of trees (second row of 6a and 6c) better. These experiments demonstrate that our In-N-Out is widely applicable to different shapes of masks including irregular masks, and different network architectures.

## 4.5 Environment Map Estimation

**Experimental Setting.** Our approach also finds its application to environment map estimation. To get a supervision from HDR images, we added LDR to HDR conversion network from [4] and train our model in an end-to-end manner. We train and test our model on the Laval Indoor HDR Dataset [5] that consists of 2,233 indoor panoramas. The details of experiment setting and test masks are available in our supplementary material.

**Qualitative Results.** We provide results compared to Gardner *et al.*, even though it uses a different mask configuration to ours <sup>2</sup>, since the result gives an insight about the benefit of In-N-Out for the environment map task. In detail, In-N-Out successfully recovers reflection of light in the ceiling (first row), overall color tones in the panorama (second row), and the shape of the floor (third row). While Gardner *et al.* is predicting well in terms of light, but it may fail in semantic detail (e.g., overall color-tone of the ceiling and floor).

<sup>2</sup>Masks correspond to the central 60-degree fov, and output panoramas are post-processed with the spherical warping.



Figure 7: Visual comparison on environment map estimation task. The results demonstrate that our In-N-Out approach can be useful for light estimation or have more plausible tone-mapped results. These results show that our In-N-Out can help another application, environment map estimation.

## 5 Conclusion

By leveraging mutual complementarity, we have proposed In-N-Out, a training strategy for both inpainting and outpainting, which can work with various networks and loss functions. In-N-Out learns an outpainting task for the initialization of a model for an inpainting task, and vice versa. We showed they could benefit each other, and our approach shows promising results in inpainting, outpainting, and environment map estimation.

Although we showed the effectiveness of our approach, there is more interesting future work lies ahead. We tackled the two tasks: inpainting and outpainting, while there are numerous ‘fill in masks’ tasks between them. We expect we can devise a better self-supervision task from one side to others by a smoother transition with these tasks. Moreover, we would like to point out that our work can get waterfall effects from developments of network architectures because our approach is parallel to the network architectures. For example, leveraging transferability of Learning without Forgetting (LwF) [21] may improve our approach further. Also, context encoders [22] showed that pretrained features using image inpainting can help other computer vision tasks (e.g., classification, detection, and semantic segmentation); our work may be able to help other vision tasks by leveraging inpainting and outpainting, and showing this is left as future work.

**Acknowledgments** This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP-2015-0-00199) and the National Research Foundation of Korea (No. 2019R1A2C3002833), grant funded by the Korea government (MSIT).

## References

- [1] Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. Curriculum learning. In *Proceedings of the 26th annual international conference on machine learn-*

- ing, pages 41–48, 2009.
- [2] Jiankang Deng, Shiyang Cheng, Niannan Xue, Yuxiang Zhou, and Stefanos Zafeiriou. Uv-gan: Adversarial facial uv map completion for pose-invariant face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7093–7102, 2018.
  - [3] Carl Doersch, Saurabh Singh, Abhinav Gupta, Josef Sivic, and Alexei A. Efros. What makes paris look like paris? *ACM Transactions on Graphics (SIGGRAPH)*, 31(4): 101:1–101:9, 2012.
  - [4] Carl Doersch, Abhinav Gupta, and Alexei A Efros. Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE international conference on computer vision*, pages 1422–1430, 2015.
  - [5] Alexey Dosovitskiy, Philipp Fischer, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with exemplar convolutional neural networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(9):1734–1747, 2015.
  - [6] Marc-André Gardner, Kalyan Sunkavalli, Ersin Yumer, Xiaohui Shen, Emiliano Gabbareto, Christian Gagné, and Jean-François Lalonde. Learning to predict indoor illumination from a single image. *ACM Trans. Graph.*, 36(6), November 2017. ISSN 0730-0301. doi: 10.1145/3130800.3130891. URL <https://doi.org/10.1145/3130800.3130891>.
  - [7] Vasileios Gkitsas, Nikolaos Zioulis, Federico Alvarez, Dimitrios Zarpalas, and Petros Daras. Deep lighting environment map estimation from spherical panoramas. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 640–641, 2020.
  - [8] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2, NIPS'14*, page 2672–2680, Cambridge, MA, USA, 2014. MIT Press.
  - [9] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. Improved training of wasserstein gans. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, page 5769–5779, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.
  - [10] Yunhui Guo, Honghui Shi, Abhishek Kumar, Kristen Grauman, Tajana Rosing, and Rogerio Feris. Spottune: Transfer learning through adaptive fine-tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
  - [11] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, page 6629–6640, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.

- [12] Satoshi Iizuka, Edgar Simo-Serra, and Hiroshi Ishikawa. Globally and locally consistent image completion. *ACM Transactions on Graphics (ToG)*, 36(4):1–14, 2017.
- [13] Woobin Im, Tae-Kyun Kim, and Sung-Eui Yoon. Unsupervised learning of optical flow with deep feature similarity. In *European Conference on Computer Vision*, pages 172–188. Springer, 2020.
- [14] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=Hk99zCeAb>.
- [15] Sai Hemanth Kasaraneni and Abhishek Mishra. Image completion and extrapolation with contextual cycle consistency. In *2020 IEEE International Conference on Image Processing (ICIP)*, pages 1901–1905. IEEE, 2020.
- [16] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7482–7491, 2018.
- [17] Kyunghun Kim, Yeohun Yun, Keon-Woo Kang, Kyeongbo Kong, Siyeong Lee, and Suk-Ju Kang. Painting outside as inside: Edge guided image outpainting via bidirectional rearrangement with progressive step learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2122–2130, 2021.
- [18] Soomin Kim, Yuchi Huo, and Sung-Eui Yoon. Single image reflection removal with physically-based training images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5164–5173, 2020.
- [19] Nikos Komodakis and Spyros Gidaris. Unsupervised representation learning by predicting image rotations. In *International Conference on Learning Representations (ICLR)*, 2018.
- [20] Yijun Li, Sifei Liu, Jimei Yang, and Ming-Hsuan Yang. Generative face completion. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3911–3919, 2017.
- [21] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017.
- [22] Guilin Liu, Fitsum A Reda, Kevin J Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image inpainting for irregular holes using partial convolutions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 85–100, 2018.
- [23] Hongyu Liu, Bin Jiang, Yibing Song, Wei Huang, and Chao Yang. Rethinking image inpainting via a mutual encoder-decoder with feature equalizations. In *Proceedings of the European Conference on Computer Vision*, 2020.
- [24] Pengpeng Liu, Michael Lyu, Irwin King, and Jia Xu. Selflow: Self-supervised learning of optical flow. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4571–4580, 2019.

- [25] Roey Mechrez, Itamar Talmi, and Lihi Zelnik-Manor. The contextual loss for image transformation with non-aligned data. In *Proceedings of the European conference on computer vision (ECCV)*, pages 768–783, 2018.
- [26] Ishan Misra, Abhinav Shrivastava, Abhinav Gupta, and Martial Hebert. Cross-stitch networks for multi-task learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3994–4003, 2016.
- [27] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2536–2544, 2016.
- [28] Mark Sabini and Gili Rusak. Painting outside the box: Image outpainting with gans. *arXiv preprint arXiv:1808.08483*, 2018.
- [29] Shuran Song and Thomas Funkhouser. Neural illumination: Lighting prediction for indoor environments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6918–6926, 2019.
- [30] Piotr Teterwak, Aaron Sarna, Dilip Krishnan, Aaron Maschinot, David Belanger, Ce Liu, and William T Freeman. Boundless: Generative adversarial networks for image extension. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10521–10530, 2019.
- [31] Wei Wang, Isaac Caswell, and Ciprian Chelba. Dynamically composing domain-data selection with clean-data selection by “co-curricular learning” for neural machine translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1282–1292, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1123. URL <https://aclanthology.org/P19-1123>.
- [32] Yi Wang, Xin Tao, Xiaojuan Qi, Xiaoyong Shen, and Jiaya Jia. Image inpainting via generative multi-column convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 331–340, 2018.
- [33] Yi Wang, Xin Tao, Xiaoyong Shen, and Jiaya Jia. Wide-context semantic image extrapolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1399–1408, 2019.
- [34] Wei Wei, Deyu Meng, Qian Zhao, Zongben Xu, and Ying Wu. Semi-supervised transfer learning for image rain removal. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3877–3886, 2019.
- [35] Peter Welinder, Steve Branson, Takeshi Mita, Catherine Wah, Florian Schroff, Serge Belongie, and Pietro Perona. Caltech-ucsd birds 200. 2010.
- [36] Zhaoyi Yan, Xiaoming Li, Mu Li, Wangmeng Zuo, and Shiguang Shan. Shift-net: Image inpainting via deep feature rearrangement. In *Proceedings of the European conference on computer vision (ECCV)*, pages 1–17, 2018.

- [37] Chao Yang, Xin Lu, Zhe Lin, Eli Shechtman, Oliver Wang, and Hao Li. High-resolution image inpainting using multi-scale neural patch synthesis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6721–6729, 2017.
- [38] Jie Yang, Dong Gong, Lingqiao Liu, and Qinfeng Shi. Seeing deeply and bidirectionally: A deep learning approach for single image reflection removal. In *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018.
- [39] Zongxin Yang, Jian Dong, Ping Liu, Yi Yang, and Shuicheng Yan. Very long natural scenery image prediction by outpainting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10561–10570, 2019.
- [40] Raymond A Yeh, Chen Chen, Teck Yian Lim, Alexander G Schwing, Mark Hasegawa-Johnson, and Minh N Do. Semantic image inpainting with deep generative models. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5485–5493, 2017.
- [41] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Generative image inpainting with contextual attention. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5505–5514, 2018.
- [42] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Free-form image inpainting with gated convolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4471–4480, 2019.
- [43] Amir R Zamir, Alexander Sax, William Shen, Leonidas J Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3712–3722, 2018.
- [44] Mengyao Zhai, Lei Chen, Frederick Tung, Jiawei He, Megha Nawhal, and Greg Mori. Lifelong gan: Continual learning for conditional image generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2759–2768, 2019.
- [45] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *European conference on computer vision*, pages 649–666. Springer, 2016.
- [46] Man Zhou, Jie Xiao, Yifan Chang, Xueyang Fu, Aiping Liu, Jinshan Pan, and Zheng-Jun Zha. Image de-raining via continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4907–4916, June 2021.