

Stay Positive: Non-Negative Image Synthesis for Augmented Reality

Katie Luo* Guandao Yang* Wenqi Xian Harald Haraldsson
 Bharath Hariharan Serge Belongie
 Cornell University

Abstract

In applications such as optical see-through and projector augmented reality, producing images amounts to solving non-negative image generation, where one can only add light to an existing image. Most image generation methods, however, are ill-suited to this problem setting, as they make the assumption that one can assign arbitrary color to each pixel. In fact, naive application of existing methods fails even in simple domains such as MNIST digits, since one cannot create darker pixels by adding light. We know, however, that the human visual system can be fooled by optical illusions involving certain spatial configurations of brightness and contrast. Our key insight is that one can leverage this behavior to produce high quality images with negligible artifacts. For example, we can create the illusion of darker patches by brightening surrounding pixels. We propose a novel optimization procedure to produce images that satisfy both semantic and non-negativity constraints. Our approach can incorporate existing state-of-the-art methods, and exhibits strong performance in a variety of tasks including image-to-image translation and style transfer.

1. Introduction

The design of images that combine views of the real world with aligned, overlaid content remains a key challenge on the path toward widespread adoption of augmented reality devices. Consider the problem of image generation in an optical see-through (OST) or projector AR device, where the image is formed by combining light from the real world with light generated by a projector; see Fig. 1. Despite the success of modern image generation methods in applications including deep image generation [39, 40, 37], image synthesis [62, 83], image-to-image translation [87, 51, 35], and style transfer [33, 74], all of the aforementioned methods assume full control over each pixel’s color. This assumption, however, is not realistic in the OST or projector AR setting where one can add, but




Figure 1. Illustration of the image formation process in augmented reality settings. Left: optical see-through setting. Right: projector AR setting. In both settings, the image is created by adding light to existing light sources from the real world, which motivates the non-negative image generation problem.

not subtract, light. In fact, naively applying any of the current state-of-the-art image generation methods will result in low-quality images in conflict with the underlying image, often riddled with ghosting artifacts.

In this work, we focus on the task of image synthesis under a non-negativity constraint, which limits the types of images that can be generated. Since we restrict ourselves to adding a non-negative residual to the input image, we would in theory only be able to generate images in our half-space. Intuitively, one cannot add black-framed glasses onto Brad Pitt’s face, since doing so would require us to attenuate the light in the desired area, *i.e.*, adding negative light, which is not physically possible. With such limitations, current hardware solutions compromise by relying on dimming the incoming light instead, using methods such as a sunglasses-like effect for OST and a darkened room for projector AR. In fact, current solutions use very dark tints, admitting only 15–20% [22, 58] of the light from the original scene. While this delivers increased control over the pixel dynamic range, blocking too much light in the scene limits the kind of settings in which OST or projector AR can be applied.

Without the ability to subtract light or dim the scene, is it still possible to succeed in image-to-image translation tasks such as turning a day image into a night image? If humans perceived luminance and chrominance in an objective manner, then this would be virtually impossible. Fortunately, the human visual system is tolerant to a wide range of photometric deviations from a “true” image [70, 78], as illus-

*Denotes equal contribution.




Figure 2. Examples of lightness illusions. Left: the Checker-shadow illusion [1]. Areas A and B have the same color, yet B appears brighter. Right: Koffka ring illusion [45]. Both half-rings have the same color, but the right one looks darker. These two examples illustrated that human do not perceive brightness in an absolute way.

trated by popular optical illusions that trick us into seeing colors or relative brightness orderings that aren't there (e.g. Figure 2). These quirks offer us a path to a solution for the challenging task of non-negative image generation: we can leverage properties of the human visual system to enlarge the solution space. For example, *lightness constancy* [1] suggests that humans do not perceive light in an absolute way, but instead take into account many factors such as geometry, texture, and surrounding neighborhoods. We may therefore no longer need to block light to create the perception of a darker pixel – we could instead add light to its peripheral locations. Another quirk of human vision is that we are relatively insensitive to visual glitches when we are not paying attention to them [60]. This allows us to introduce errors in some non-salient locations without being noticed.

In this paper, we approach the novel task of non-negative image generation by designing a framework that can take advantage of both existing image-to-image translation methods and human visual perception quirks. Specifically, we first incorporate existing image-to-image translation models to generate a proposal image that satisfies the task, but is not necessarily feasible in the non-negative generation setting. We then finetune the proposal image so that it can be produced by adding non-negative light to the input image while remaining perceptually unchanged. Our method is compatible with almost all existing image-to-image translation methods, thereby permitting its application to a wide variety of tasks. While we mainly appeal to *lightness constancy* in this paper, our framework can be extended to model other types of perceptual illusions, which we leave to future explorations. We empirically show that our method outperforms all strong baselines that naively adapt state-of-the-art models on this task, both quantitatively and qualitatively, even in settings when much more light is allowed in the optical combiner. Finally, we provide detailed analysis of how – and when – our method works.

2. Non-negative Image Generation

We now specify the non-negative image generation problem, in which the model can only add light to the input image when producing the output image. We narrow our scope to the specific problem of producing an image that fulfills a certain semantic or style category, as opposed to any arbitrary image. This is motivated by the idea that AR users are more likely to specify a category instead of a detailed image.

During training, one is provided with samples from P_X and P_Y , the distributions of the images of the input domain and target domain, respectively. The goal is to learn a model that can take an input image $x \sim P_X$ and produce a non-negative residual $R_\theta(x)$ that can be combined with the input to produce an image in P_Y . Formally, non-negative image generation amounts to solving the following constrained optimization problem:

$$\begin{aligned} \arg \max_{\theta} \quad & \mathbb{E}_{x \sim P_X} [\log P_Y(\alpha x + \beta R_\theta(x))] \\ \text{s.t.} \quad & \forall i, j, 0 \leq R_\theta(x)_{i,j} \leq 1, \end{aligned} \quad (1)$$

where $\alpha, \beta \in [0, 1]$ are parameters controlling influence from the input image and the predicted residual, respectively, and (with a slight abuse of notation) $R_\theta(x)_{i,j}$ denotes the i - by j -th pixel of residual $R_\theta(x)$.

While the problem formulation in Equation 1 can be used in a number of AR settings, we focus on the optical see-through (OST) case. Here, α represents the amount of light transmitted through the semi-transparent mirror. At the same time, the amount of light reflected from the semi-transparent mirror is $1 - \alpha$, so the upper bound is $\beta = 1 - \alpha$. We refer to this simplified optical formulation for the rest of the work. Note that under this formulation, when $\alpha = 0$, the problem reduces to the video see-through (VST) setting, i.e., the traditional unconstrained version. The higher the value of α , the harder the problem becomes, since the amount of light at the disposal of the generator R_θ becomes more limited.

3. Method

In developing our proposed method, we aim to draw upon the strengths of state-of-the-art image generation methods while availing ourselves of degrees of freedom provided by the human visual system.

3.1. A two stage approach

To make a model compatible with existing image generators and adaptable to quirks of human visual perception, we propose to break down the non-negative image generation problem into two steps. The first step is to solve the image generation problem without the non-negativity constraint. In the second step, we finetune the image produced by the

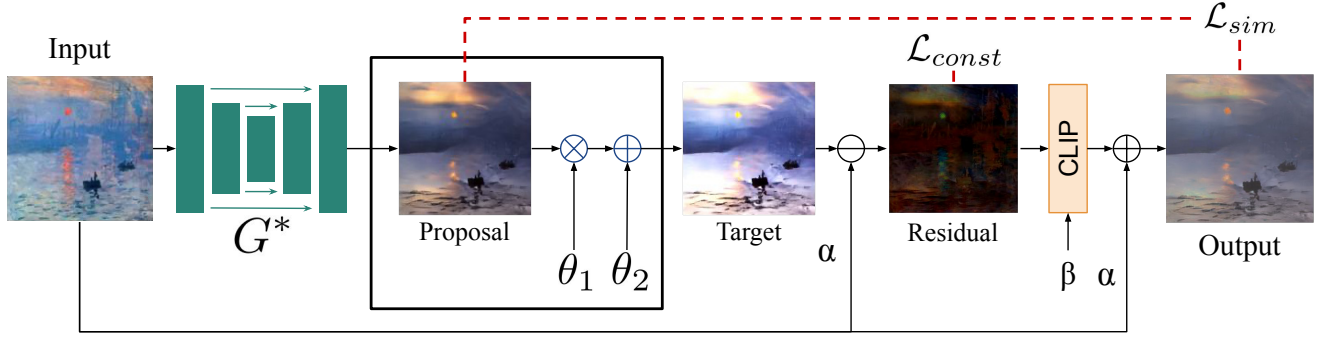


Figure 3. **Illustration of our two-stage training pipeline.** Our framework’s overall schematic consists of the image proposal step and semantic-preserving generator step for the residual. The final output is an optical combination between the input and residual, clipped to be physically feasible. We rely on very few parameters for efficient training.

first step with the aim of making it perceptually unchanged under the non-negativity constraint.

Formally, we assume that we can produce a model G^* that takes an image from domain X and outputs one that is in domain Y while preserving the structure of the input images. Specifically, we take a model G^* that minimizes the following objective:

$$\min_G \mathbb{E}_{x \sim \mathcal{X}} [\mathcal{S}(G(x), x) - \log P_Y(G(x))], \quad (2)$$

where $\mathcal{S}(\cdot, \cdot)$ returns the structural similarity between two images. Most image-to-image translation works such as Pix2Pix [36] and CycleGAN [87] can achieve this step.

At this point, we can produce an image proposal $y = G^*(x)$ for any input image x . The image proposal y is a good solution as long as it can be feasibly produced under the non-negativity constraints imposed by Equation 1. The end goal, however, does not require the output image to be exactly the same as y — in fact, any image that is perceptually indistinguishable from y should suffice. Therefore, the second step can be formulated as finding a target image such that it is perceptually similar to input image x while being physically realizable in the non-negative image generation regime. Formally, for each pair of input image x and image proposal y , we want to find an image $F_\theta(x, y)$ that optimizes the following objective and constraints:

$$\begin{aligned} \min_{\theta} \quad & \mathcal{L}_{sim}(F_\theta(x, y), y) \\ \text{s.t.} \quad & \forall i, j, 0 \leq F_\theta(x, y)_{i,j} - \alpha x_{i,j} \leq \beta, \end{aligned} \quad (3)$$

where $\mathcal{L}_{sim}(\cdot, \cdot)$ measures the perceptual similarity between two images. For example, if y' represents y altered by perturbations undetectable to the human eye, then $\mathcal{L}_{sim}(y', y) = 0$. Prior works have studied measuring perceptual similarity for a variety of applications including lossy image compression [57, 66] and image generation [85]. In the following sections, we will show how to design \mathcal{L}_{sim} to encode lightness constancy, and how to solve

this problem by optimizing a few parameters using stochastic gradient descent.

3.2. Perceptual similarity

The simplest (and strictest) formulation for \mathcal{L}_{sim} is that two images are the same only if the pixel values match everywhere (*i.e.*, $\mathcal{L}_{sim}(a, b) = \|a - b\|^2$). We instead propose to relax this by incorporating the notion of lightness constancy. Formally, we define the similarity metric as:

$$\mathcal{L}_{sim}(a, b) = \|N(a) - N(b)\|^2, \quad (4)$$

where $N(\cdot)$ is a normalization function that brings two perceptually similar images closer together. For example, to capture the idea that humans are visually insensitive to global lightness changes, which in part explains lightness constancy, we normalize the dynamic range the image by using Color Histogram Stretching [81]:

$$N(x) = \frac{x - x_{min}}{x_{max} - x_{min}}, \quad (5)$$

where x_{min} and x_{max} are the smallest and largest pixel values in image x , respectively. For RGB images, we apply such normalization independently to all three channels. With this operation in place, adding a constant to every pixel in the image will not affect the output since $N(ax + b) = N(x)$. While Equation 4 is not restricted to just one normalization method or just one notion of perceptual quality, we will focus on the lightness constancy property and will assume the use of Equation 5 unless specified otherwise. Figure 4 shows how loss with normalization leads to feasible solutions and outputs with less artifacts.

3.3. Soft Constraint Loss

The above expression of perceptual similarity with normalization makes the constrained optimization problem in Equation 3 non-convex, and thereby a more challenging problem. As mentioned in Section 1, we are not sensitive to minor changes in color, especially in high texture



Figure 4. Visualizations of our Semantic-Preserving Generator outputs (Left). By varying the θ parameters, we can control the brightness and contrast of the final output image. We visualize the Perceptual Similarity loss with and without normalization (Middle and Right respectively), and outline the region of parameters corresponding to physically realizable residuals in red. Darker color indicates larger loss. The output image of the minimum of the two losses are shown to the right of the loss contour. Observe that normalizing the images accomplishes two things: 1. the loss contour changes to include the physically feasible region, and 2. the perceptual quality of the image is improved and the output exhibits less “ghosting” artifacts from the input image.

Algorithm 1: Semantics-Preserving Generator

Result: Output image $O_{\theta}(x, y)$

Input: Device parameters α, β , image x , proposal y , learning rate lr ;

Initialize $\theta = [1, 0]$;

while not converged **do**

$F_{\theta}(x, y) = \theta_1 y + \theta_2$;
 $r = F_{\theta}(x, y) - \alpha x$;
 $O_{\theta}(x, y) = \min(\max(r, 0), \beta) + \alpha x$;
 $\mathcal{L}(\theta) \leftarrow \mathcal{L}_{sim}(O_{\theta}(x, y), y) + \mathcal{L}_{const}(r, 0, \beta)$;
 $\theta \leftarrow Adam(\nabla_{\theta} \mathcal{L}(\theta), lr)$;

end

regions of a natural image. This allows our system to make some mistakes by making the target $F_{\theta}(x, y)$ not *completely* physically realizable at all pixels, as long as the error between the actual output and the target output is relatively small. Thus, we relax the hard constraint on the residual r into minimizing the following soft constraint:

$$\mathcal{L}_{const}(r, a, b) = \gamma \sum_{i,j} |\max(\min(r_{i,j}, b), a) - r_{i,j}|, \quad (6)$$

where $\gamma > 0$ is a hyper-parameter that controls the trade-off between perceptual similarity and residual constraint satisfaction. This soft constraint loss gives a continuous penalty to places where the pixel is not physically realizable, and it is differentiable everywhere except at a or b . By using the soft constraint as a loss function, we get the added benefit of penalizing the values depending on how infeasible they are. This allows the model to balance between producing a strictly physically realizable image and producing an image that is perceptually similar to the image proposal, albeit containing a few errors.

3.4. Semantics-Preserving Generator

The final piece of the puzzle is the design of the generator $F_{\theta}(x, y)$. F_{θ} takes the input image x and the semantic prototype y and aims to predict a target image that looks similar to y and is also physically realizable by adding light to αx . The naive solution for $F_{\theta}(x, y)$ is to return y directly, but such solution is likely to break the constraints since the brightness of the X domain could be drastically different from that of the Y domain. In the other extreme, we can use a universal approximator like a neural network to approximate the function (*i.e.*, predicting colors of each pixel). Since such a generator imposes almost no domain specific knowledge, it contains lots of parameters and can be empirically difficult to train. In this paper, we propose a semantics-preserving generator that is flexible enough to include a physically feasible solution while building upon y as a very good starting point.

Specifically, we want the generator to be able to explore images that preserve the structure of the semantic image y in order to find an image that best respects the physical constraints. Note an affine transformation on an image y into $\theta_1 y + \theta_2$ does not change the perceptual loss \mathcal{L}_{sim} since $N(\theta_1 y + \theta_2) = N(y)$ for any scalar θ_1 and θ_2 . This reflects the assumption that humans have the ability to do white balancing and adjust to different contrasts. One can potentially find a physically realizable image by choosing $\theta = (\theta_1, \theta_2)$. Formally, we define the target image $F_{\theta}(x, y)$ and the physically realizable output image $O_{\theta}(x, y)$ as follow:

$$\begin{aligned} F_{\theta}(x, y) &= \theta_1 y + \theta_2, \\ O_{\theta}(x, y) &= \max(\min(F_{\theta}(x, y) - \alpha x, \beta), 0) + \alpha x. \end{aligned} \quad (7)$$

We adjust the final objective to also minimize the perceptual similarity between the actual output image O_{θ} and the image proposal y . This allows \mathcal{L}_{sim} to also provide training signal, which empirically leads to better performance.

Method	Satellite → Map [36]			Map → Satellite [36]			Day → Night [46]		
	FID(↓)	KID(↓)	N-PSNR(↑)	FID(↓)	KID(↓)	N-PSNR(↑)	FID(↓)	KID(↓)	N-PSNR(↑)
Heuristic	252.33	19.03	20.16	200.96	18.08	33.07	134.97	4.53	39.00
From scratch	268.79	19.38	26.19	305.82	33.29	30.63	218.09	11.75	27.67
Finetuning	254.86	16.22	30.77	265.87	24.97	34.63	195.14	11.86	29.00
Ours	238.93	17.90	23.78	125.75	12.76	68.90	110.83	3.13	47.42

Table 1. Method performances on Map↔Satellite and Day→Night datasets.

Thus, the final training objective is:

$$\mathcal{L}(\theta, x, y) = \mathcal{L}_{sim}(O_{\theta}(x, y), y) + \mathcal{L}_{const}(F_{\theta}(x, y) - \alpha x, 0, \beta). \quad (8)$$

Note that this objective is differentiable almost everywhere and can be optimized efficiently using stochastic gradient descent. Our method is summarized in Algorithm 1, with hyper-parameters and implementation details provided in the supplementary material.

4. Experiments

In this section, we begin by empirically evaluating our proposed method in comparison to baselines in a variety of tasks in Section 4.1 and Section 4.2. We also conduct an ablation study on major design choices for our model (Section 4.3), and provide an analysis of several interesting properties of our method (Section 4.4). We provide our code at <https://github.com/katieluo88/StayPositive>.

Datasets. Most of our experiments are conducted on datasets from the image translation and image editing literature, which include semantic label↔photo [36, 87], image→style [34, 87], dog→cat [35], and CelebA-HQ dataset [71, 74, 38]. For unaligned datasets, we employ the provided pretrained models to generate a set of targets y . If the dataset is paired, we will then use the ground truth images as the proposed image.

Metrics. Following prior works [41, 52, 16, 36], we will use Frechet Inception Distance (FID) [29], Kernel Inception Distance (KID) [2], and Peak signal-to-noise ratio (PSNR). KID computes the squared Maximum Mean Discrepancy of features extracted from Inception Network [76] while FID computes the Frechet distance from those features. At the same time, we also report the PSNR to quantify the similarity between the proposed images and the final output image, after normalization. All reported numbers are computed by averaging across multiple α values.

Datasets	Baseline	Ours
Zebra → Horse	12.63	11.53
Winter → Summer	1.16	1.03
Photo → Monet	1.59	1.38
Photo → Von Gogh	3.27	3.26
Photo → Cezanne	4.70	4.51
Photo → Uyeoko	10.33	9.71

Table 2. Kernel Inception Distance×100 metric on [87]. Our model outperforms the heuristic baseline on a wide range of tasks.

4.1. Comparisons

We consider the following three baselines. The first one clips the difference required to reach the proposed image before adding to the input images to ensure physical feasibility (Heuristic). The second one trains a Pix2pix [36] model from scratch to predict non-negative residuals used to predict the output image (From Scratch). The third baseline uses our proposed loss to fine-tune a Pix2Pix generator (Finetune). We train these three baselines and our methods on three tasks, including generating synthetic images (satellite→map), generating photorealistic images (map→satellite), and generating images with less light (day→night). The results are shown in Table 1 and Figure 5. Our models out-perform baselines in FID in all tasks, and all but one task in KID and PSNR. Even though our model does not lead in these metrics in the task of Satellite→map, we can see from Figure 5 that our methods produce images with fewer ghosting artifacts for all three tasks. In addition to the above mentioned metrics, we also conducted a user studies using Amazon Mechanical Turk. The results (Figure 6) suggest that most users chose images by our method over those from the baselines by a large margin.

4.2. Applications

As mentioned in Section 3.1, our method can be combined with different image translation methods. In this section, we quantitatively compare our method with the heuristic baseline on a set of unpaired image translation datasets listed in Table 2, adapted from CycleGAN [87]. The table suggests that our model achieves comparable or bet-

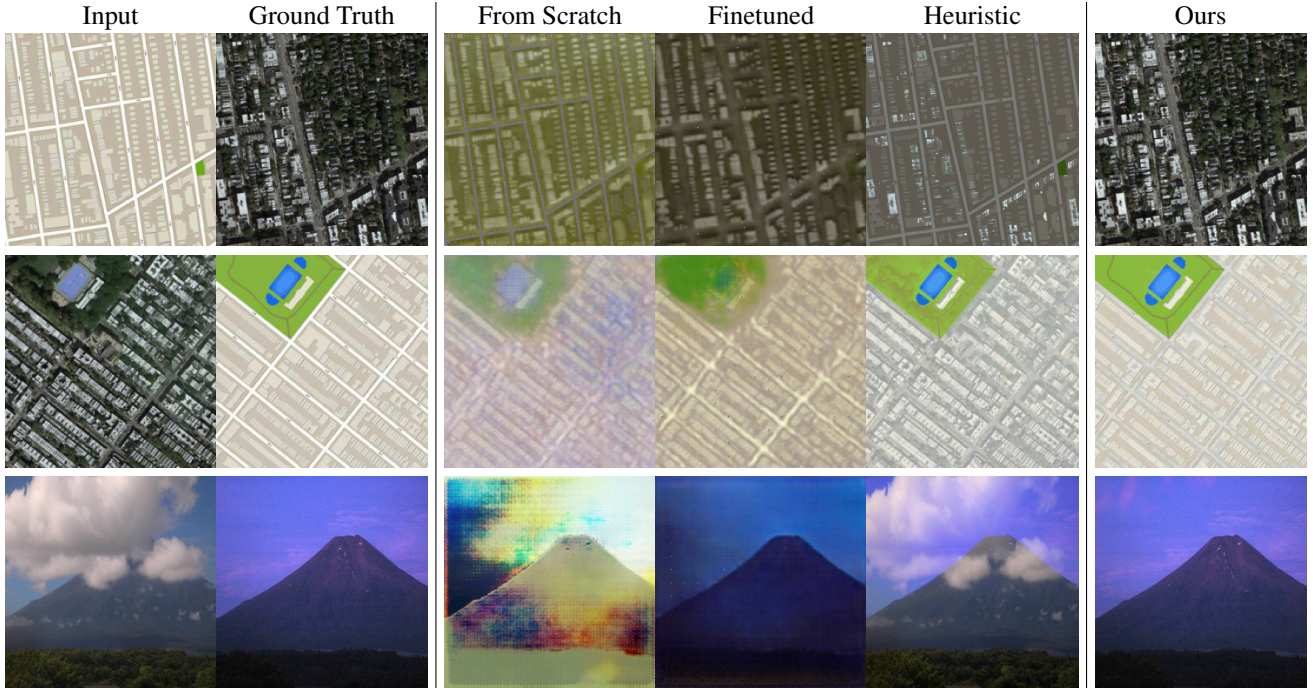


Figure 5. Visual comparison to the baseline methods. Rows from top to bottom: Map→Satellite, Satellite→Map, Day→Night




Figure 6. Results from user study. We use AMT to recruit users to answer three questions asking for their preference on “satellite image”, “photo-realistic”, and “detail”. Most users prefer the output of our model over other baselines by a large margin.

ter KID than the heuristic baseline in all tasks. We also apply our method in a variety of more difficult datasets such as high resolution day→night [4], style transfer [34], sketch to photo [34], dog→cat [35], and face attribute editing [71, 73]. The qualitative examples of aforementioned tasks are presented in Figure 7. While the images produced by the heuristic based method tend to exhibit ghosting artifacts arising from enforcing non-negativity, our method is able to avoid such shortcomings and produce high quality images that can resemble the target.

4.3. Analysis

In this section, we will first conduct an ablation study to examine our design decisions. Then we will present an

Method	FID(↓)	KID(↓)	N-PSNR(↑)
w/o normalization	146.95	13.23	45.86
w/o \mathcal{L}_{const}	139.95	13.74	69.09
w/o \mathcal{L}_{sim}	145.21	14.57	65.71
Per-pixel prediction	132.82	13.90	78.17
Ours	125.75	12.76	68.90

Table 3. Ablation study on the Map→Satellite task. We explore in this order: normalization methods, loss function importance, and our method’s transformation groups.

analysis of how our methods perform according to different α values. Finally, we will use qualitative examples to illustrate how our proposed approach takes advantage of human perceptual quirks.

Ablations. We evaluate three design decisions: normalizing the dynamic range before computing the perceptual similarity \mathcal{L}_{sim} (Section 3.2), minimizing both perceptual similarity \mathcal{L}_{sim} and constraint violation \mathcal{L}_{const} (Section 3.3), and predicting global modification instead of per-pixel modification (Section 3.4). We run ablations of our method on the task map→satellite and reported results in Table 3. It’s interesting to note that predicting per-pixel modification achieves better PSNR score, which is expected due to the large amount of parameters, but produces worse FID and KID, which suggests that the method creates im-

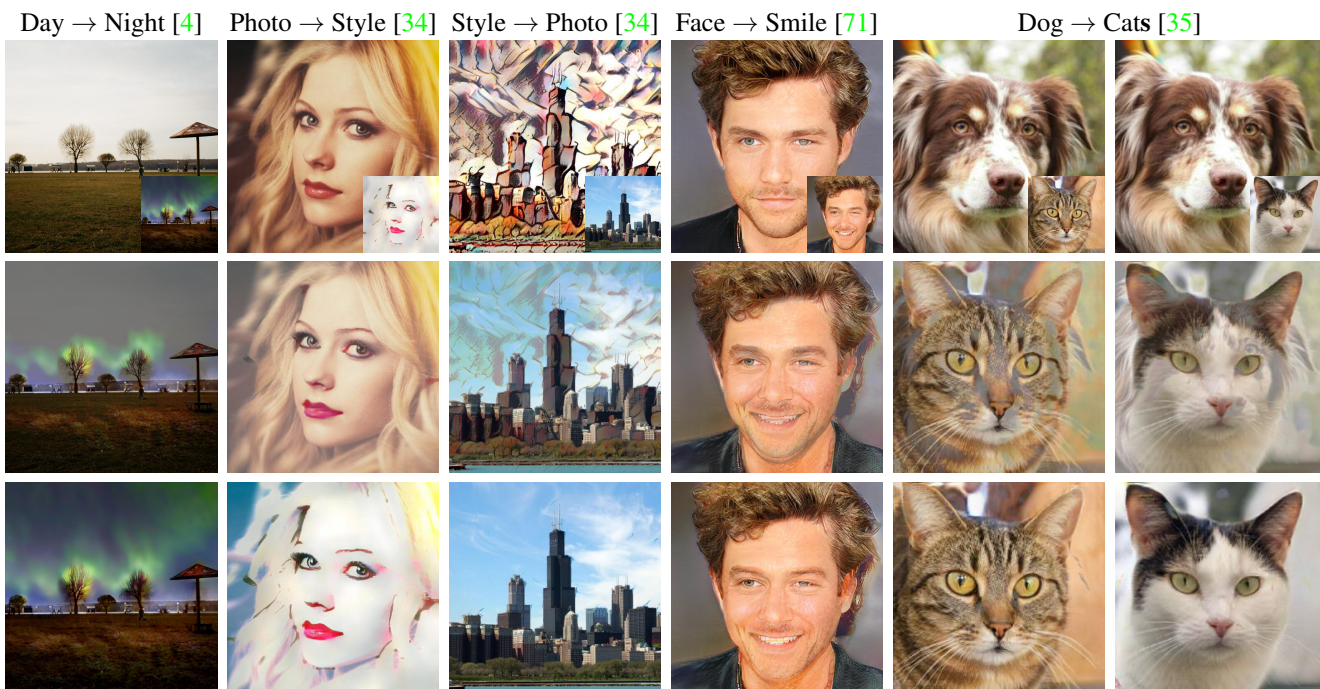


Figure 7. We compare our results across a variety of domains. From top to bottom: input (with proposal overlaid), heuristic baseline, our method. Observe that with our framework, we are able to obtain multimodal outputs from the generative model.

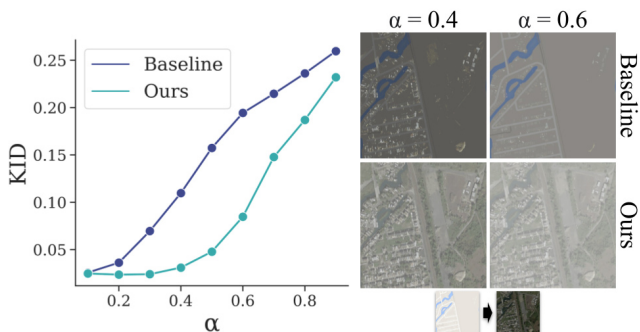


Figure 8. Results on a range of α values for Map→Satellite. We plot KID metrics on our method vs. heuristic baseline. Our method outperforms baselines across almost all α . Visually, our method also produces less artifacts comparing to the baseline.

ages that are semantically far away from the target domain. On the other hand, our methods achieve the best FID and KID among all variants.

Effect of varying α . While previous experiments show that our methods outperforms the baselines when averaging across different values of α , it's thus far unclear how our method compares to the baselines at different α values. In Figure 8, we present the change of KID across 10 different α values between our methods and the heuristic baseline. When α is small, both the baseline and our method perform well, as expected, since the task is very similar to the one

without non-negativity constraint. On the other hand, when α is large, both our method and the baseline tend to perform badly (though our method still out-performs the baseline), which is expected since the task is almost impossible. It's interesting to note that our method is able to maintain good performance until after $\alpha > 0.5$, while the baseline's performance degrades drastically. This suggests that our method can produce images in harmony with more ambient light. Visually, the ghosting artifacts from the input image does not show up in our method until higher α values.

4.4. Leveraging quirks of human perception

In this section we present qualitative examples to demonstrate how our model benefits from quirks of human visual perception involving.

Lightness constancy. In Figure 9, we show how our model turns a daytime image into a nighttime image, along with the associated color histogram of the input and target images. Note that such a task is infeasible, at face value, since the nighttime image contains more pixels with a lower amount of blue or red light compared to the input image, as shown in the histogram. Our model is nonetheless capable of producing an image that looks similar to the target image without subtracting any light. The model achieves this result by padding the amount of light around the points in need of darkening, which shifts the whole dynamic range higher and brighter (i.e., the histogram of the output image

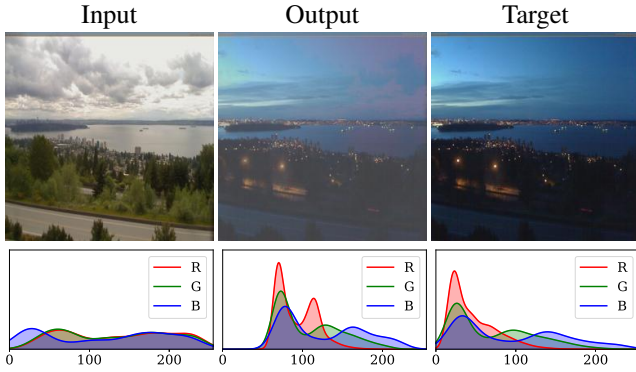


Figure 9. Demonstration that our model can leverage the *lightness constancy* produce the appearance of dark pixels through the addition of light.

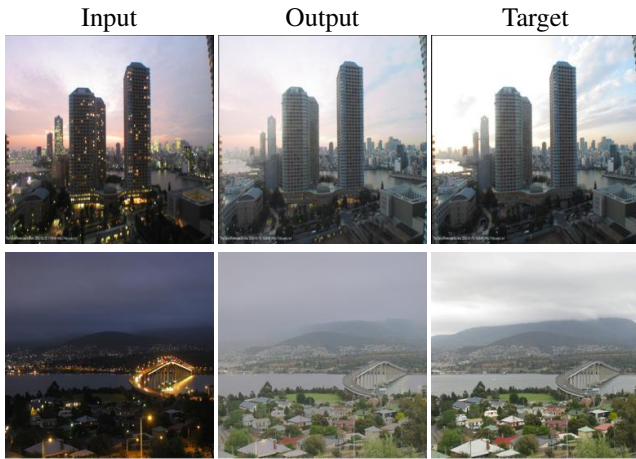


Figure 10. Our model can hide mistakes and non-salient image details. At first glance, the output image is a normal day photo. The building lights that are left on are details ignored by our method.

has few pixels with value less than 50). Though the “black” point in the middle of the image has RGB values greater than 50, such a change is generally hard to detect without direct, side-by-side comparison to the target image.

Insensitivity to errors in textured regions. Figure 10 shows examples of our method turning nighttime images to daytime images of city scenes. The challenge in such tasks lies in “turning off” streetlamps or office window lights when subtraction is not allowed. Interestingly, our model tends not to worry too much about turning off all the lights since the associated discrepancy is very small and usually does not attract users’ attention. While all of our output images appear natural at first glance, upon closer examination, these simulated daytime images retain evidence of turned-on lights. However, our visual system is capable of handling these minor errors.

5. Related Work

Recent deep generative models such as flow-based model [67, 42, 17], auto-regressive models [59, 69, 63], VAEs [44, 43, 55, 80], and GANs [25, 65, 5, 26, 39, 40, 13, 38] have contributed to improvements in unconditional image generation. The literature on conditional image generation explores related tasks including style transfer [23, 79, 14, 34, 19, 4, 8, 61], supervised image translation [36, 88, 24, 48, 49, 47, 53, 62, 77, 82, 86], unsupervised image translation [51, 87, 35, 10, 16, 52, 68], and semantic image editing [72, 9]. Please refer to Liu *et al* [50] for more details. Despite this progress, the above mentioned works assume the ability to output any color at any pixel location, which is ill-suited to our problem setting. In fact, our method can be combined with all of the above image translation models. A related constrained image generation problem is studied by Heim [28]. This work, however, focuses on semantic constraints, *e.g.*, that the generated image should be closer to one attribute group than the other. Most closely related to our present work is the recent literature on projector compensation [32, 30, 31, 11, 6, 27, 75, 3, 7], where one updates the projected images to account for the potentially non-planar and non-uniformly colored projector surface. Our work focuses instead on the OST setting, in which the image is formed via semi-transparent glass instead of at the projector surface. In the sense that our approach plays with the effective dynamic range and the contrast of the image, our work is related to *tone mapping*, a process of mapping real world colors to a restricted color space for displaying image while preserving details. Traditional tone mapping algorithms employ variants of histogram equalization [84, 15, 18, 12, 64] to enhance the contrast of LDR images. Many recent works [56, 21, 20, 54] leverage deep learning to recover the missing details in the over-exposed image regions by expanding the dynamic range of single LDR images. In our paper, we found that pairing very simple contrast and dynamic range adjustment with SGD is sufficient to obtain our desired results.

6. Conclusion

Through this work, we explored the problem of non-negative image generation and proposed an effective framework that achieves consistently better results than the baselines. Possible extensions to this work include exploring additional properties of the human visual system and removing the overly simplistic modeling assumptions. Thus far, our studies have not included implementation using a real-world AR rig. We are pursuing the testing of our method in a physical system in an ongoing work.

Acknowledgements This work was supported in part by grants from Magic Leap and Facebook AI, and a donation from NVIDIA.

References

- [1] Edward H Adelson. 24 lightness perception and lightness illusions. 2000. 2
- [2] Youssef Alami Mejjati, Christian Richardt, James Tompkin, Darren Cosker, and Kwang In Kim. Unsupervised attention-guided image-to-image translation. *Advances in Neural Information Processing Systems*, 31:3693–3703, 2018. 5
- [3] D. Aliaga, Y. H. Yeung, Alvin J. Law, Behzad Sajadi, and A. Majumder. Fast high-resolution appearance editing using superimposed projections. *ACM Trans. Graph.*, 31:13:1–13:13, 2012. 8
- [4] Ivan Anokhin, Pavel Solovev, Denis Korzhenkov, Alexey Kharlamov, Taras Khakhulin, Alexey Silvestrov, Sergey Nikolenko, Victor Lempitsky, and Gleb Sterkin. High-resolution daytime translation without domain labels. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 6, 7, 8
- [5] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017. 8
- [6] Hirotaka Asayama, D. Iwai, and K. Sato. Fabricating diminishable visual markers for geometric registration in projection mapping. *IEEE Transactions on Visualization and Computer Graphics*, 24:1091–1102, 2018. 8
- [7] M. Ashdown, T. Okabe, I. Sato, and Y. Sato. Robust content-dependent photometric projector compensation. *2006 Conference on Computer Vision and Pattern Recognition Workshop (CVPRW'06)*, pages 6–6, 2006. 8
- [8] Hyojin Bahng, Sunghyo Chung, Seungjoo Yoo, and J. Choo. Exploring unlabeled faces for novel attribute discovery. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5820–5829, 2020. 8
- [9] David Bau, Jun-Yan Zhu, Hendrik Strobelt, Bolei Zhou, Joshua B. Tenenbaum, William T. Freeman, and Antonio Torralba. Gan dissection: Visualizing and understanding generative adversarial networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019. 8
- [10] Sagie Benaim and Lior Wolf. One-sided unsupervised domain mapping. *ArXiv*, abs/1706.00826, 2017. 8
- [11] A. Boroomand, Hicham Sekkati, M. Lamm, David A Clausi, and A. Wong. Saliency-guided projection geometric correction using a projector-camera system. *2016 IEEE International Conference on Image Processing (ICIP)*, pages 2951–2955, 2016. 8
- [12] Alberto Boschetti, Nicola Adami, Riccardo Leonardi, and Masahiro Okuda. High dynamic range image tone mapping based on local histogram equalization. In *2010 IEEE International Conference on Multimedia and Expo*, pages 1130–1135. IEEE, 2010. 8
- [13] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018. 8
- [14] Tian Qi Chen and Mark Schmidt. Fast patch-based style transfer of arbitrary style. *arXiv preprint arXiv:1612.04337*, 2016. 8
- [15] Ho-Hyoung Choi and Byoung-Ju Yun. Contrast enhancement for the captured image by estimating local illumination. *Optical review*, 18(5):389–393, 2011. 8
- [16] Yunjey Choi, Min-Je Choi, Munyoung Kim, Jung-Woo Ha, S. Kim, and J. Choo. Stargan: Unified generative adversarial networks for multi-domain image-to-image translation. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8789–8797, 2018. 5, 8
- [17] Laurent Dinh, Jascha Sohl-Dickstein, and S. Bengio. Density estimation using real nvp. *ArXiv*, abs/1605.08803, 2017. 8
- [18] Jiang Duan, Marco Bressan, Chris Dance, and Guoping Qiu. Tone-mapping high dynamic range images by novel histogram adjustment. *Pattern Recognition*, 43(5):1847–1862, 2010. 8
- [19] Vincent Dumoulin, Jonathon Shlens, and M. Kudlur. A learned representation for artistic style. *ArXiv*, abs/1610.07629, 2017. 8
- [20] Gabriel Eilertsen, Joel Kronander, Gyorgy Denes, Rafał K Mantiuk, and Jonas Unger. Hdr image reconstruction from a single exposure using deep cnns. *ACM transactions on graphics (TOG)*, 36(6):1–15, 2017. 8
- [21] Yuki Endo, Yoshihiro Kanamori, and Jun Mitani. Deep reverse tone mapping. *ACM Transactions on Graphics*, 36:1–10, 11 2017. 8
- [22] Austin Erickson, Kangsoo Kim, Gerd Bruder, and Gregory F. Welch. Exploring the limitations of environment lighting on optical see-through head-mounted displays. In *Symposium on Spatial User Interaction, SUI '20*, New York, NY, USA, 2020. Association for Computing Machinery. 1
- [23] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2414–2423, 2016. 8
- [24] A. Ghosh, Viveka Kulharia, Vinay P. Namboodiri, P. Torr, and P. Dokania. Multi-agent diverse generative adversarial networks. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8513–8521, 2018. 8
- [25] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. 8
- [26] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *Advances in neural information processing systems*, pages 5767–5777, 2017. 8
- [27] M. Harville, W. Culbertson, I. Sobel, D. Gelb, A. Fitzhugh, and D. Tanguay. Practical methods for geometric and photometric correction of tiled projector. *2006 Conference on Computer Vision and Pattern Recognition Workshop (CVPRW'06)*, pages 5–5, 2006. 8
- [28] Eric Heim. Constrained generative adversarial networks for interactive image generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10753–10761, 2019. 8
- [29] 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 *Advances in neural information processing systems*, pages 6626–6637, 2017. 5
- [30] Bingyao Huang and Haibin Ling. Compennet++: End-to-end full projector compensation. *CoRR*, abs/1908.06246, 2019. 8
- [31] Bingyao Huang and Haibin Ling. End-to-end projector photometric compensation. *CoRR*, abs/1904.04335, 2019. 8
- [32] Bingyao Huang and H. Ling. Deltra: Deep light transport for projector-camera systems. *ArXiv*, abs/2003.03040, 2020. 8
- [33] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1501–1510, 2017. 1
- [34] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *ICCV*, 2017. 5, 6, 7, 8
- [35] Xun Huang, Ming-Yu Liu, Serge Belongie, and Jan Kautz. Multimodal unsupervised image-to-image translation. In *The European Conference on Computer Vision (ECCV)*, September 2018. 1, 5, 6, 7, 8
- [36] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017. 3, 5, 8
- [37] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017. 1
- [38] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017. 5, 8
- [39] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4401–4410, 2019. 1, 8
- [40] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. *arXiv preprint arXiv:1912.04958*, 2019. 1, 8
- [41] Junho Kim, Minjae Kim, Hyeonwoo Kang, and Kwanghee Lee. U-gat-it: unsupervised generative attentional networks with adaptive layer-instance normalization for image-to-image translation. *arXiv preprint arXiv:1907.10830*, 2019. 5
- [42] Diederik P. Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *ArXiv*, abs/1807.03039, 2018. 8
- [43] Diederik P. Kingma, Tim Salimans, and M. Welling. Improved variational inference with inverse autoregressive flow. *ArXiv*, abs/1606.04934, 2017. 8
- [44] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 8
- [45] Kurt Koffka. *Principles of Gestalt psychology*, volume 44. Routledge, 2013. 2
- [46] Pierre-Yves Laffont, Zhile Ren, Xiaofeng Tao, Chao Qian, and James Hays. Transient attributes for high-level understanding and editing of outdoor scenes. *ACM Transactions on Graphics (proceedings of SIGGRAPH)*, 33(4), 2014. 5
- [47] Cheng-Han Lee, Z. Liu, Lingyun Wu, and Ping Luo. Maskgan: Towards diverse and interactive facial image manipulation. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5548–5557, 2020. 8
- [48] Soochan Lee, J. Ha, and Gunhee Kim. Harmonizing maximum likelihood with gans for multimodal conditional generation. *ArXiv*, abs/1902.09225, 2019. 8
- [49] K. Li, T. Zhang, and Jitendra Malik. Diverse image synthesis from semantic layouts via conditional imle. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4219–4228, 2019. 8
- [50] M. Liu, Xun Huang, Jiahui Yu, T. Wang, and Arun Mallya. Generative adversarial networks for image and video synthesis: Algorithms and applications. *ArXiv*, abs/2008.02793, 2020. 8
- [51] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *Advances in neural information processing systems*, pages 700–708, 2017. 1, 8
- [52] Ming-Yu Liu, Xun Huang, Arun Mallya, Tero Karras, Timo Aila, J. Lehtinen, and J. Kautz. Few-shot unsupervised image-to-image translation. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10550–10559, 2019. 5, 8
- [53] Xihui Liu, Guojun Yin, J. Shao, X. Wang, and Hongsheng Li. Learning to predict layout-to-image conditional convolutions for semantic image synthesis. In *NeurIPS*, 2019. 8
- [54] Yu-Lun Liu, Wei-Sheng Lai, Yu-Sheng Chen, Yi-Lung Kao, Ming-Hsuan Yang, Yung-Yu Chuang, and Jia-Bin Huang. Single-image hdr reconstruction by learning to reverse the camera pipeline. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1651–1660, 2020. 8
- [55] Lars Maaløe, M. Fraccaro, Valentin Liévin, and O. Winther. Biva: A very deep hierarchy of latent variables for generative modeling. *ArXiv*, abs/1902.02102, 2019. 8
- [56] Demetris Marnerides, Thomas Bashford-Rogers, Jonathan Hatchett, and Kurt Debattista. Expandnet: A deep convolutional neural network for high dynamic range expansion from low dynamic range content. In *Computer Graphics Forum*, volume 37, pages 37–49. Wiley Online Library, 2018. 8
- [57] Juan Carlos Mier, Eddie Huang, Hossein Talebi, Feng Yang, and Peyman Milanfar. Deep perceptual image quality assessment for compression, 2021. 3
- [58] S. Mori, S. Ikeda, A. Plopski, and C. Sandor. Brightview: Increasing perceived brightness of optical see-through head-mounted displays through unnoticeable incident light reduction. In *2018 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pages 251–258, 2018. 1
- [59] Aaron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. *arXiv preprint arXiv:1601.06759*, 2016. 8
- [60] W. Osberger. Perceptual vision models for picture quality assessment and compression applications. 1999. 2

- [61] D. Y. Park and K. H. Lee. Arbitrary style transfer with style-attentional networks. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5873–5881, 2019. 8
- [62] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 1, 8
- [63] Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, L. Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. Image transformer. *ArXiv*, abs/1802.05751, 2018. 8
- [64] Stelios Ploumis, Ronan Boitard, Mahsa T Pourazad, and Panos Nasiopoulos. Perception-based histogram equalization for tone mapping applications. In *2016 Digital Media Industry & Academic Forum (DMI AF)*, pages 11–16. IEEE, 2016. 8
- [65] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015. 8
- [66] Rafael Reisenhofer, Sebastian Bosse, Gitta Kutyniok, and Thomas Wiegand. A haar wavelet-based perceptual similarity index for image quality assessment. *Signal Processing: Image Communication*, 61:33–43, 2018. 3
- [67] Danilo Jimenez Rezende and Shakir Mohamed. Variational inference with normalizing flows. *arXiv preprint arXiv:1505.05770*, 2015. 8
- [68] K. Saito, Kate Saenko, and M. Liu. Coco-funit: Few-shot unsupervised image translation with a content conditioned style encoder. *ArXiv*, abs/2007.07431, 2020. 8
- [69] Tim Salimans, A. Karpathy, Xi Chen, and Diederik P. Kingma. Pixelcnn++: Improving the pixelcnn with discretized logistic mixture likelihood and other modifications. *ArXiv*, abs/1701.05517, 2017. 8
- [70] Arthur Shapiro, Laysa Hedjar, Erica Dixon, and Akiyoshi Kitaoka. Kitaoka’s tomato: Two simple explanations based on information in the stimulus. *i-Perception*, 9(1):2041669517749601, 2018. 1
- [71] Yujun Shen, Jinjin Gu, Xiaou Tang, and Bolei Zhou. Interpreting the latent space of gans for semantic face editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9243–9252, 2020. 5, 6, 7
- [72] Yujun Shen, Jinjin Gu, X. Tang, and B. Zhou. Interpreting the latent space of gans for semantic face editing. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9240–9249, 2020. 8
- [73] Yujun Shen, Ceyuan Yang, Xiaou Tang, and Bolei Zhou. Interfacegan: Interpreting the disentangled face representation learned by gans. *TPAMI*, 2020. 6
- [74] Lu Sheng, Ziyi Lin, Jing Shao, and Xiaogang Wang. Avatar-net: Multi-scale zero-shot style transfer by feature decoration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8242–8250, 2018. 1, 5
- [75] C. Siegl, Matteo Colaianni, Lucas Thies, Justus Thies, M. Zollhöfer, S. Izadi, M. Stamminger, and F. Bauer. Real-time pixel luminance optimization for dynamic multi-projection mapping. *ACM Trans. Graph.*, 34:237:1–237:11, 2015. 8
- [76] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016. 5
- [77] Hao Tang, Xiaojuan Qi, D. Xu, P. H. S. Torr, and N. Sebe. Edge guided gans with semantic preserving for semantic image synthesis. *ArXiv*, abs/2003.13898, 2020. 8
- [78] Maiko Uesaki, Hiroshi Ashida, Akiyoshi Kitaoka, and Achille Pasqualotto. Cross-modal size-contrast illusion: Acoustic increases in intensity and bandwidth modulate haptic representation of object size. *Scientific Reports*, 9(1), 2019. 1
- [79] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Improved texture networks: Maximizing quality and diversity in feed-forward stylization and texture synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6924–6932, 2017. 8
- [80] Arash Vahdat and J. Kautz. Nvae: A deep hierarchical variational autoencoder. *ArXiv*, abs/2007.03898, 2020. 8
- [81] S. Wang, Yewei Zhang, Peng Deng, and Fuqiang Zhou. Fast automatic white balancing method by color histogram stretching. *2011 4th International Congress on Image and Signal Processing*, 2:979–983, 2011. 3
- [82] T. Wang, Ming-Yu Liu, Jun-Yan Zhu, A. Tao, J. Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8798–8807, 2018. 8
- [83] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018. 1
- [84] Byoung-Woo Yoon and Woo-Jin Song. Image contrast enhancement based on the generalized histogram. *Journal of electronic imaging*, 16(3):033005, 2007. 8
- [85] Richard Zhang, Phillip Isola, Alexei A. Efros, E. Shechtman, and O. Wang. The unreasonable effectiveness of deep features as a perceptual metric. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 586–595, 2018. 3
- [86] Haitian Zheng, Haofu Liao, Lele Chen, W. Xiong, T. Chen, and Jiebo Luo. Example-guided image synthesis across arbitrary scenes using masked spatial-channel attention and self-supervision. *ArXiv*, abs/2004.10024, 2020. 8
- [87] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017. 1, 3, 5, 8
- [88] Jun-Yan Zhu, Richard Zhang, Deepak Pathak, Trevor Darrell, Alexei A. Efros, O. Wang, and E. Shechtman. Toward multimodal image-to-image translation. *ArXiv*, abs/1711.11586, 2017. 8