

Transforming and Projecting Images into Class-conditional Generative Networks

Minyoung Huh^{12*} Richard Zhang² Jun-Yan Zhu²
Sylvain Paris² Aaron Hertzmann²

¹MIT CSAIL ²Adobe Research

Abstract. We present a method for projecting an input image into the space of a class-conditional generative neural network. We propose a method that optimizes for transformation to counteract the model biases in generative neural networks. Specifically, we demonstrate that one can solve for image translation, scale, and global color transformation, during the projection optimization to address the object-center bias and color bias of a Generative Adversarial Network. This projection process poses a difficult optimization problem, and purely gradient-based optimizations fail to find good solutions. We describe a hybrid optimization strategy that finds good projections by estimating transformations and class parameters. We show the effectiveness of our method on real images and further demonstrate how the corresponding projections lead to better editability of these images. The project page and the code is available at <https://minyoung.github.io/GAN-Transform-and-Project/>.

1 Introduction

Deep generative models, particularly Generative Adversarial Networks (GANs) [24], can create a diverse set of realistic images, with a number of controls for transforming the output, e.g., [48, 6, 29, 32, 48]. However, most of these methods apply only to synthetic images that are generated by GANs in the first place. In many real-world cases, a user would like to edit their own image. One approach is to train a network for each separate image transformation. However, this would require a combinatorial explosion of training time and model parameters.

Instead, a user could “project” their image to the manifold of images produced by the GAN, by searching for an appropriate latent code [60]. Then, any transformations available within the GAN could be applied to the user’s image. This could allow a powerful range of editing operations within a relatively compact representation. However, projection is a challenging problem. Previous methods have focused on class-specific models, for example, for objects [60], faces [46, 9], or specific scenes such as bedrooms and churches [5, 7]. With the challenges in both optimization and generative model’s limited capacity, we wish to find a generic method that can fit *real* images from diverse categories into the same generative model.

*Work started during an internship at Adobe Research.




Fig. 1. Given a pre-trained BigGAN [8] and a target image (left), our method uses gradient-free BasinCMA to transform the image and find a latent vector to closely reconstruct the image. Our method (top) can better fit the input image, compared to the baseline (bottom), which does not model image transformation and uses gradient-based ADAM optimization. Finding an accurate solution to the inversion problem allows us to further fine-tune the model weights to match the target image without losing downstream editing capabilities. For example, our method allows for changing the class of the object (top row), compared to the baseline (bottom).

This paper proposes the first method for projecting images into class-conditional models. In particular, we focus on BigGAN [8]. We address the main problems with these tasks, mainly, the challenges of optimization, object alignment, and class label estimation:

- To help avoid local minima during the optimization process, we systematically study choices of both gradient-based and gradient-free optimizers and show Covariance Matrix Adaptation (CMA) [26] to be more effective than stand-alone gradient-based optimizers, such as L-BFGS [40] and Adam [33].
- To better fit a real image into the latent space, we account for the model’s center bias by simultaneously estimating both spatial image transformation (translation, scale, and color) and latent variable. Such a transformation can then be inverted back to the input image frame. Our simultaneous transformation and projection method largely expands the scope and diversity of the images that a GAN can reconstruct.
- Finally, we show that estimating and jointly optimizing the continuous embedding of the class variable leads to better projections. This ultimately leads to more expressive editing by harnessing the representation of the class-conditional generative model.

We evaluate our method against various baselines on projecting real images from ImageNet. We quantitatively and qualitatively demonstrate that it is crucial to simultaneously estimate the correct transformation during the projection step. Furthermore, we show that CMA, a non-parametric gradient-free optimization technique, significantly improves the robustness of the optimization and leads to better solutions. As shown in Figure 1, our method allows us to fine-tune our model to recover the missing details without losing the editing capabilities of the generative model.

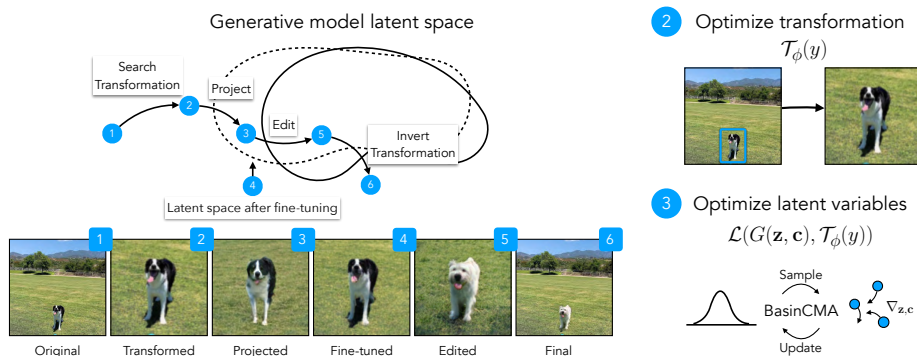


Fig. 2. **Overview:** Our method first searches for a transformation to apply to the input target image. We then solve for the latent vector that closely resembles the object in the target image, using our proposed optimization method, also referred to as “projection”. The generative model can then be further fine-tuned to reconstruct the missing details that the original model could not generate. Finally, we can edit the image by altering the latent code or the class vector (e.g., changing the border collie to a west highland white terrier), and invert and blend the edited image back into the original image.

2 Related Work

Image editing with generative models. Image editing tools allow a user to manipulate a photograph according to their goal while producing realistic visual content. Seminal work is often built on low-level visual properties, such as patch-based texture synthesis [18, 28, 17, 4], gradient-domain image blending [47], and image matting with locally affine color model [37]. Different from previous hand-crafted low-level methods, several recent works [60, 9] proposed to build editing tools based on a deep generative model, with the hope that a generative model can capture high-level information about the image manifold.

Many prior works have investigated using trained generative models as a tool to edit images [60, 9, 5, 6]. The same image prior from deep generative models has also been used in face editing, image inpainting, colorization, and deblurring prior [46, 56, 2, 25, 51]. Unlike these works that focuses on single-class and fixated image, our method presents a new ways of embedding an image into a class-conditional generative model, which allows the same GAN to be applied to many more “in-the-wild” scenarios.

Inverting networks. Our work is closely related to methods for inverting pre-trained networks. Earlier work proposes to invert CNN classifiers and intermediate features for visualizing recognition networks [41, 15, 43, 44]. More recently, researchers adopted the above methods to invert generative models. The common techniques include: (1) Optimization-based methods: they find the latent vector that can closely reconstruct the input image using gradient-based method (e.g., ADAM, LBFGS) [60, 9, 39, 56, 50, 49, 10] or MCMC [20], (2) Encoder-based methods: they learn an encoder to directly predict the latent vector given a real image [12, 16, 60, 46, 9, 13], (3) Hybrid methods [60, 5, 7]: they use the encoder to initialize the latent vector and then solve the optimiza-

tion problem. Although the optimized latent vector roughly approximates the real input image, many important visual details are missing in the reconstruction [5]. To address the issue, GANPaint [5] generates residual features to adapt to the individual image. Image2StyleGAN [1] optimizes StyleGAN’s intermediate representation rather than the input latent vector. Unfortunately, the above techniques still cannot handle images in many scenarios due to the limited model capacity [7], the lack of generalization ability [1], and their single-class assumption. As noted by prior work [1], the reconstruction quality severely degrades under simple image transformation, and translation has been found to cause most of the damage. Compared to prior work, we consider two new aspects in the reconstruction pipeline: image transformation and class vector. Together, these two aspects significantly expand the diversity of the images that we can reconstruct and edit.

3 Image projection methods

We aim to project an image into a class-conditional generative model (e.g., BigGAN [8]) for the purposes of downstream editing. We first introduce the basic objective function that we slowly build upon. Next, since BigGAN is an object-centric model for most classes, we infer an object mask from the input image and focus on fitting the pixels inside the mask.

Furthermore, to better fit our desired image into the generative model, we propose to optimize for various image transformation (scale, translation, and color) to be applied to the target image. Lastly, we explain how we optimize the aforementioned objective loss function.

3.1 Basic Loss Function

Class-conditional generative model A class-conditional generative network can synthesize an image $\hat{\mathbf{y}} \in \mathbb{R}^{H \times W \times 3}$, given a latent code $\mathbf{z} \in \mathbb{R}^Z$ that models intra-class variations and a one-hot class-conditioning vector $\tilde{\mathbf{c}} \in \Delta^C$ to choose over C classes. We focus on the 256×256 BigGAN model [8] specifically, where $Z = 128$ and $C = 1,000$ ImageNet classes.

The BigGAN architecture first maps the one-hot $\tilde{\mathbf{c}}$ into a continuous vector $\mathbf{c} \in \mathbb{R}^{128}$ with a linear layer $\mathbf{W} \in \mathbb{R}^{128 \times 1000}$, before injecting into the main network G_θ , with learned parameters θ .

$$\hat{\mathbf{y}} = G_\theta(\mathbf{z}, \mathbf{c}) = G_\theta(\mathbf{z}, \mathbf{W}\tilde{\mathbf{c}}). \quad (1)$$

Here, a choice must be made whether to optimize over the discrete $\tilde{\mathbf{c}}$ or continuous \mathbf{c} . As optimizing a discrete class vector is non-trivial, we optimize over the continuous embedding.

Optimization setup. Given a target image \mathbf{y} , we would like to find a \mathbf{z}^* and \mathbf{c}^* that generates the image.

$$\mathbf{z}^*, \mathbf{c}^* = \arg \min_{\mathbf{z}, \mathbf{c}} \mathcal{L}(G_\theta(\mathbf{z}, \mathbf{c}), \mathbf{y}) \quad \text{s.t. } C(\mathbf{z}) \leq C_{\max}. \quad (2)$$

During training, the latent code is sampled from a multivariate Gaussian $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Interestingly, recent methods [8, 34] find that restricting the distribution at *test time* produces higher-quality samples. We follow this and constrain our search space to match the sampling distribution from Brock et al. [8]. Specifically, we use $C(\mathbf{z}) = \|\mathbf{z}\|_\infty$ and $C_{\max} = 2$. During optimization, elements of \mathbf{z} that fall outside the threshold are clamped to +2, if positive, or -2, if negative. Allowing larger values of \mathbf{z} produces better fits but compromises editing ability.

Loss function. The loss function \mathcal{L} attempts to capture how close the approximate solution is to the target. A loss function that perfectly corresponds to human perceptual similarity is a longstanding open research problem [54], and evaluating the difference solely on a per-pixel basis leads to blurry results [58]. Distances in the feature space of a pre-trained CNN correspond more closely with human perception [30, 14, 21, 59]. We use the LPIPS metric [59], which calibrates a pre-trained model using human perceptual judgments. Here, we define our basic loss function, which combines per-pixel ℓ_1 and LPIPS.

$$\mathcal{L}_{\text{basic}}(\mathbf{y}, \hat{\mathbf{y}}) = \frac{1}{HW} \|\hat{\mathbf{y}} - \mathbf{y}\|_1 + \beta \mathcal{L}_{\text{LPIPS}}(\hat{\mathbf{y}}, \mathbf{y}). \quad (3)$$

In preliminary experiments, we tried various loss combinations and found $\beta = 10$ to work well. We now expand upon this loss function by leveraging object mask information.

3.2 Object Localization

Real images are often more complex than the ones generated by BigGAN. For example, objects may be off-centered and partially occluded, or multiple objects appear in an image. Moreover, it is possible that the object in the image can be approximated by GANs but not the background.

Accordingly, we focus on fitting a single foreground object in an image and develop a loss function to emphasize foreground pixels. We automatically produce a foreground rectangular mask $\mathbf{m} \in [0, 1]^{H \times W \times 1}$ using the bounding box of an object detector [27]. Here, we opt for bounding boxes for simplicity, but one could consider using segmentation mask, saliency maps, user-provided masks, etc. The foreground and background values within mask \mathbf{m} are set to 1 and 0.3, respectively. We adjust the objective function to spatially weigh the loss:

$$\mathcal{L}_{\text{mask}}(\mathbf{y}, \hat{\mathbf{y}}, \mathbf{m}) = \frac{1}{M} \|\mathbf{m} \odot (\hat{\mathbf{y}} - \mathbf{y})\|_1 + \beta \mathcal{L}_{\text{mLPIPS}}(\hat{\mathbf{y}}, \mathbf{y}, \mathbf{m}), \quad (4)$$

where normalization parameter $M = \|\mathbf{m}\|_1$ and \odot represents element-wise multiplication across the spatial dimensions. Given a mask of all foreground (all ones), the objective function is equivalent to Equation 3. We calculate the masked version of the perceptual loss $\mathcal{L}_{\text{mLPIPS}}(\hat{\mathbf{y}}, \mathbf{y}, \mathbf{m})$ by bilinearly downsampling the mask at the resolution of the intermediate spatial feature maps within the perceptual loss. The details are described in Appendix B. With the provided mask, we now explore how one can optimize for image transformation to better fit the object in the image.

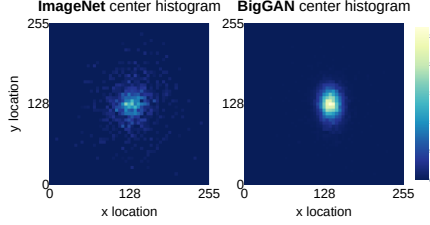


Fig. 3. **Object center comparison:** We use an object detector to compute the histogram of object locations. Note that ImageNet (**left**) is biased towards the center but exhibits a long-tail. BigGAN (**right**) is further biased towards center.

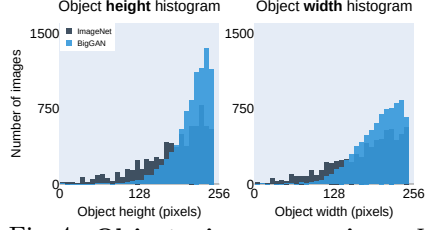


Fig. 4. **Object size comparison:** We use an object detector to compute the distribution of object widths (**left**) and heights (**right**). Note that ImageNet (black) has a long-tail, whereas the BigGAN (blue) accentuates the mode.

3.3 Transformation Model and Loss

Generative models may exhibit biases for two reasons: (a) inherited biases from the training distribution and (b) bias introduced by mode collapse [23], where the generative model only captures a portion of the distribution. We mitigate two types of biases, *spatial* and *color* during image reconstruction process.

Studying spatial biases. To study spatial bias, we first use a pre-trained object detector, MaskRCNN [27], over 10,000 real and generated images to compute the statistics of object locations. We show the statistics regarding the center locations and object sizes in Figures 3 and 4, respectively.

Figure 3 (left) demonstrates that ImageNet images exhibit clear center bias over the location of objects, albeit with a long tail. While the BigGAN learns to mimic this distribution, it further accentuates the bias [7, 29], largely forgoing the long tail to generate high-quality samples in the middle of the image. In Figure 4, we see similar trends with object height and width. Abdal et al. [1] noted that the quality of image reconstruction degrades given a simple translation in the target image. Motivated by this, we propose to incorporate spatial alignment in the inversion process.

Searching over spatial alignments. We propose to transform the generated image using $\mathcal{T}_{\psi}^{\text{spatial}}(\cdot)$, which shifts and scales the image using parameters $\psi = [s_x, s_y, t_x, t_y]$. The parameters ψ are used to generate a sampling grid which in turn is used by a grid-sampler to construct a new transformed image [42]. The corresponding inverse parameters are $\psi^{-1} = [\frac{1}{s_x}, \frac{1}{s_y}, -\frac{t_x}{s_x}, -\frac{t_y}{s_y}]$.

Transforming the generated image allows for more flexibility in the optimization. For example, if G can perfectly generate the target image, but at different scales or at off-centered locations, this framework allows it to do so.

Searching over color transformations. Furthermore, we show that the same framework allows us to search over color transformations $\mathcal{T}_{\gamma}^{\text{color}}(\cdot)$. We experimented with various color transformations such as hue, brightness, gamma, saturation, contrast, and found brightness and contrast to work the best. Specifically, we optimize for brightness, which is parameterized by scalar γ with inverse




Fig. 5. **Initialization from various methods:** We show samples drawn from different methods, before the final gradient descent optimization. In “random initialization”, seeds are drawn from the normal distribution; the results show higher variation. For the “encoder initialization”, we use a trained encoder network to predict the latent vector and apply a minor perturbation. Our method uses CMA to find a good starting distribution. For “Encoder+BasinCMA”, we initialize CMA with the output of the encoder. The results are more consistent and better reconstruct the target image.

value $\gamma^{-1} = -\gamma$. If the generator can perfectly generate the target image, but slightly darker or brighter, this allows a learned brightness transformation to compensate for the difference.

Final objective. Let transformation function $\mathcal{T}_\phi = \mathcal{T}_\psi^{\text{spatial}} \circ \mathcal{T}_\gamma^{\text{color}}$ be a composition of spatial and color transformation functions, where transformation parameters ϕ is a concatenation of spatial and color parameters ψ, γ , respectively. The inverse function is $\mathcal{T}_{\phi^{-1}}$. Our final optimization objective function, with consideration for (a) the foreground object and (b) spatial and color biases, is to minimize the following loss:

$$\arg \min_{\mathbf{z}, \mathbf{c}, \phi} \mathcal{L}_{\text{mask}}(\mathcal{T}_{\phi^{-1}}(G_\theta(\mathbf{z}, \mathbf{c})), \mathbf{y}, \mathbf{m}) \quad \text{s.t. } C(\mathbf{z}) \leq C_{\max} \quad (5)$$

Our optimization algorithm, described next, has a mix of gradient-free and gradient-based updates. Alternatively, instead of inverse transforming the *generated* image, we can transform the *target* and mask images during gradient-based updates and compute the following loss: $\mathcal{L}_{\text{mask}}(G_\theta(\mathbf{z}, \mathbf{c}), \mathcal{T}_\phi(\mathbf{y}), \mathcal{T}_\phi(\mathbf{m}))$. We will discuss when to use each variant in the next section.

3.4 Optimization Algorithms

Unfortunately, the objective function is highly non-convex. Gradient-based optimization, as used in previous inversion methods, frequently fall into poor local minima. Bau et al. [7] note that recent large-scale GAN models [31, 32] are significantly harder to invert due to a large number of layers, compared to earlier models [48]. Thus, formulating an optimizer that reliably finds good solutions is a significant challenge. We evaluate our method against various baselines and ablations in Section 4. Given the input image \mathbf{y} and foreground rectangular mask \mathbf{m} (which is automatically computed), we present the following algorithm.

Class and transform initialization We first predict the class of the image with a pre-trained ResNeXt101 classifier [55] and multiply it by \mathbf{W} to obtain our initial class vector \mathbf{c}_0 .

Algorithm 1 Transformation-aware projection algorithm

Input: Image \mathbf{y} , initial class vector \mathbf{c}_0 , mask \mathbf{m}
Output: Transformation parameter ϕ^* , latent variable \mathbf{z}^* , class vector \mathbf{c}^*

```

1: # Optimize for transformation  $\phi$ 
2: Initialize  $(\mu_\phi, \Sigma_\phi) \leftarrow (\phi_0, 0.1 \cdot \mathbf{I})$  ▷  $\phi_0$  precomputed in Section 3.3
3: for  $n$  iterations do
4:    $\phi_{1:N} \sim \text{SampleCMA}(\mu_\phi, \Sigma_\phi)$  ▷ Draw  $N$  samples of  $\phi$ 
5:    $\mathbf{z}_{1:N} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ , reset  $\mathbf{c}_{1:N} \leftarrow \mathbf{c}_0$  ▷ Reinitialize  $\mathbf{z}$  and  $\mathbf{c}$ 
6:   for  $m$  iterations do
7:     for  $i \leftarrow 1$  to  $N$  do ▷ This loop is batched
8:        $g_i \leftarrow \mathcal{L}_{\text{mask}}(G_\theta(\mathbf{z}_i, \mathbf{c}_i), \mathcal{T}_{\phi_i}(\mathbf{y}), \mathcal{T}_{\phi_i}(\mathbf{m}))$ 
9:        $(\mathbf{z}_i, \mathbf{c}_i) \leftarrow (\mathbf{z}_i, \mathbf{c}_i) - \eta \cdot \nabla_{\mathbf{z}, \mathbf{c}} g_i$  ▷ Update each sample  $\mathbf{z}, \mathbf{c}$ 
10:       $g_{1:N}^{\text{inv}} \leftarrow \mathcal{L}_{\text{mask}}(\mathcal{T}_{\phi_{1:N}}^{-1}(G_\theta(\mathbf{z}_{1:N}, \mathbf{c}_{1:N})), \mathbf{y}, \mathbf{m})$  ▷ Recompute loss with inverse
11:       $\mu_\phi, \phi_{1:N} \leftarrow \text{UpdateCMA}(\phi_{1:N}, g_{1:N}^{\text{inv}}, \mu_\phi, \Sigma_\phi)$  ▷ Section 3.4
12: Set  $\phi^* \leftarrow \mu_\phi$ 
13: # Optimize for latent variables  $\mathbf{z}, \mathbf{c}$ 
14: Initialize  $(\mu_{\mathbf{z}}, \Sigma_{\mathbf{z}}) \leftarrow (\mathbf{0}, \mathbf{I})$ 
15: for  $p$  iterations do
16:    $\mathbf{z}_{1:M} \sim \text{SampleCMA}(\mu_{\mathbf{z}}, \Sigma_{\mathbf{z}})$ , reset  $\mathbf{c}_{1:M} \leftarrow \mathbf{c}_0$  ▷ Draw  $M$  samples of  $\mathbf{z}$ 
17:   for  $q$  iterations do
18:     for  $i \leftarrow 1$  to  $M$  do ▷ This loop is batched
19:        $g_i \leftarrow \mathcal{L}_{\text{mask}}(G_\theta(\mathbf{z}_i, \mathbf{c}_i), \mathcal{T}_{\phi^*}(\mathbf{y}), \mathcal{T}_{\phi^*}(\mathbf{m}))$ 
20:        $(\mathbf{z}_i, \mathbf{c}_i) \leftarrow (\mathbf{z}_i, \mathbf{c}_i) - \nabla_{\mathbf{z}, \mathbf{c}} g_i$ 
21:       $g_{1:M}^{\text{inv}} \leftarrow \mathcal{L}_{\text{mask}}(\mathcal{T}_{\phi_{1:M}}^{-1}(G_\theta(\mathbf{z}_{1:M}, \mathbf{c}_{1:M})), \mathbf{y}, \mathbf{m})$  ▷ Recompute loss with inverse
22:       $\mu_{\mathbf{z}}, \Sigma_{\mathbf{z}} \leftarrow \text{UpdateCMA}_{\mathbf{z}}(\mathbf{z}_{1:M}, g_{1:M}^{\text{inv}}, \mu_{\mathbf{z}}, \Sigma_{\mathbf{z}})$  ▷ Section 3.4
23: Set  $\mathbf{z}^*, \mathbf{c}^* \leftarrow \arg \min_{\mathbf{z}, \mathbf{c}} (g_{1:M})$  ▷ Choose the best  $\mathbf{z}, \mathbf{c}$ 

```

Next, we initialize the spatial transformation vector $\psi_0 = [s_{x_0}, s_{y_0}, t_{y_0}, t_{x_0}]$ such that the foreground object is well-aligned with the statistics of the BigGAN model. As visualized in Figures 3 and 4, $(\bar{h}, \bar{w}) = (137, 127)$ is the center of BigGAN-generated objects and $(\bar{y}, \bar{x}) = (213, 210)$ is the mode of object sizes. We define $(h_{\mathbf{m}}, w_{\mathbf{m}})$ to be the height and width and $(y_{\mathbf{m}}, x_{\mathbf{m}})$ to be center of the masked region. We initialize scale factors as $s_{y_0} = s_{x_0} = \max(\frac{h_{\mathbf{m}}}{\bar{h}}, \frac{w_{\mathbf{m}}}{\bar{w}})$ and translations as $(t_{y_0}, t_{x_0}) = (\frac{\bar{y} - y_{\mathbf{m}}}{2}, \frac{\bar{x} - x_{\mathbf{m}}}{2})$. Finally, initial brightness transformation parameter is initialized as $\gamma_0 = 1$.

Choice of optimizer. We find the choice of optimizer critical and that BasinCMA [7] provides better results than previously used optimizers for the GAN inversion problem. Previous work [60, 1] has exclusively used gradient-based optimization, such as LBFGS [40] and ADAM [33]. However, such methods are prone to obtaining poor results due to local minima, requiring the use of multiple random initial seeds. Covariance Matrix Adaptation (CMA) [26], a *gradient-free* optimizer, finds better solutions than gradient-based methods. CMA maintains a Gaussian distribution in parameter space $\mathbf{z} \sim \mathcal{N}(\mu, \Sigma)$. At each iteration, N samples are drawn, and the Gaussian is updated using the loss. The details of this update are described in Hansen and Ostermeier [26]. A weakness of CMA is that when it nears a solution, it is slow to refine results, as it does not use gradients. To address this, we use a variant, BasinCMA [53], that alternates between CMA updates and ADAM optimization, where CMA distribution is updated after taking M gradient steps.

Next, we describe the optimization procedure between the transformation parameters ϕ and latent variables \mathbf{z}, \mathbf{c} .

Choice of Loss function In Equation 5, we described two variants of our optimization objective. Ideally, we would like to optimize the former variant $\mathcal{L}_{\text{mask}}(\mathcal{T}_{\phi^{-1}}(G_{\theta}(\mathbf{z}, \mathbf{c})), \mathbf{y}, \mathbf{m})$ such that the target image \mathbf{y} is consistent throughout optimization; and we do so for all CMA updates. However for gradient optimization, we found that back-propagating through a grid-sampler to hurt performance, especially for small objects. A potential reason is that when shrinking a generated image, the grid-sampling operation sparsely samples the image. Without low-pass filtering, this produces a noisy and aliased result [22, 45]. Therefore, for gradient-based optimization, we optimize the latter version $\mathcal{L}_{\text{mask}}(G_{\theta}(\mathbf{z}, \mathbf{c}), \mathcal{T}_{\phi}(\mathbf{y}), \mathcal{T}_{\phi}(\mathbf{m}))$.

Two-stage approach. Historically, searching over spatial transformations with reconstruction loss as guidance has proven to be a difficult task in computer vision [3]. We find this to be the case in our application as well, and that joint optimization over the transformation ϕ , and variables \mathbf{z}, \mathbf{c} is unstable. We use a two-stage approach, as shown in Algorithm 1, where we first search for ϕ^* and use ϕ^* to optimize for \mathbf{z}^* and \mathbf{c}^* . In both stages, a gradient-free CMA outer loop maintains a distribution over the variable of interest in that stage. In the inner loop, ADAM is used to quickly find the local optimum over latent variables \mathbf{z}, \mathbf{c} .

To optimize for the transformation parameter, we initialize CMA distribution for ϕ . The mean μ_{ϕ} is initialized with pre-computed statistics ϕ_0 , and Σ_{ϕ} is set to $0.1 \cdot \mathbf{I}$ (Alg. 1, line 2). A set of transformations $\phi_{1:N}$ is drawn from CMA, and latent variables $\mathbf{z}_{1:N}$ are randomly initialized (Alg. 1, line 4–5). To evaluate the sampled transformation, we take gradient updates w.r.t. $\mathbf{z}_{1:N}, \mathbf{c}_{1:N}$ for $m = 30$ iterations (Alg. 1, line 6–9). This inner loop can be interpreted as quickly assessing the viability of a given spatial transform. The final samples of $\mathbf{z}_{1:N}, \mathbf{c}_{1:N}, \phi_{1:N}$ are used to compute the loss for the CMA update (Alg. 1, line 10–11). This procedure is repeated for $n = 30$ iterations, and the final transformation ϕ^* is set to the mean of the current estimate of CMA (Alg. 1, line 12).

After solving for the transformation ϕ^* , a similar procedure is used to optimize for \mathbf{z} . We initialize CMA distribution for \mathbf{z} with $\mu_{\mathbf{z}} = \mathbf{0}$ and $\Sigma_{\mathbf{z}} = \mathbf{I}$ (Alg. 1, line 14). M samples of $\mathbf{z}_{1:M}$ are drawn from the CMA distribution and $\mathbf{c}_{1:M}$ is set to the initial predicted class vector (Alg. 1, line 16). The drawn samples are evaluated by taking $q = 30$ gradient updates w.r.t $\mathbf{z}_{1:M}$ and $\mathbf{c}_{1:M}$ (Alg. 1, line 17–20). The optimized samples are used to compute the loss for the CMA update (Alg. 1, line 21–22). This procedure is repeated for $p = 30$ iterations. On the final iteration, we take 300 gradient updates instead to obtain the final solution \mathbf{z}, \mathbf{c} (Alg. 1, line 23).

3.5 Fine-tuning

So far, we have located an approximate match within a generative model. We hypothesize that if a high-quality match is found, fine-tuning to fit the image will preserve the editability of the generative model. On the contrary, if a poor match is found, the fine-tuning will corrupt the network and result in low-quality images after editing. Next, we describe this fine-tuning process.




Fig. 6. **ImageNet comparisons:** Comparison across various methods on inverting ImageNet images without fine-tuning. A rectangular mask centered around the object of interest is provided for all methods using MaskRCNN [27]. The losses are weighted by the mask. BasinCMA+Transform is our full method.

To synthesize the missing details that the generator could not produce, we wish to fine-tune our model after solving for the latent vector \mathbf{z} , the class vector \mathbf{c} , and transformation parameters ϕ . Unlike previous work [5], which proposed to produce the residual features using a small, auxiliary network, we update the weights of the original GAN directly. This allows us to perform edits that spatially deform the image. After obtaining the values for $\phi, \mathbf{z}, \mathbf{c}$ in our projection step, we fine-tune the weights of the generative model. During fine-tuning, the full objective function is:

$$\arg \min_{\mathbf{z}, \mathbf{c}, \phi, \theta} \mathcal{L}_{\text{mask}}(\mathcal{T}_{\phi^{-1}}(G_{\theta}(\mathbf{z}, \mathbf{c})), \mathbf{y}, \mathbf{m}) + \lambda \|\theta - \theta_0\|_2 \quad \text{s.t.} \quad C(\mathbf{z}) \leq C_{\max} \quad (6)$$

We put an ℓ_2 -regularization on the weights, such that the fine-tuned weights do not deviate too much from the original weights θ_0 . In doing so, we can prevent overfitting and preserve the generative model’s ability to edit the final image. We use $\lambda = 10^3$ for our results with fine-tuning.

4 Results

We demonstrate results on images from ImageNet [11], compare against baselines and ablations, examine cases that BigGAN cannot generate, and show failure cases. We further demonstrate the validity of our method on out-of-distribution data such as COCO and conduct perceptual studies on the edited images.

The ImageNet dataset consists of 1.3 million images with 1,000 classes. We construct a test set by using PASCAL [19] classes as super-classes. There are a total of 229 classes from ImageNet that map to 16 out of 20 classes in PASCAL. We select 10 images at random from each super-class to construct a dataset of 160 images. We run off-the-shelf Mask-RCNN [27] and take the highest activating

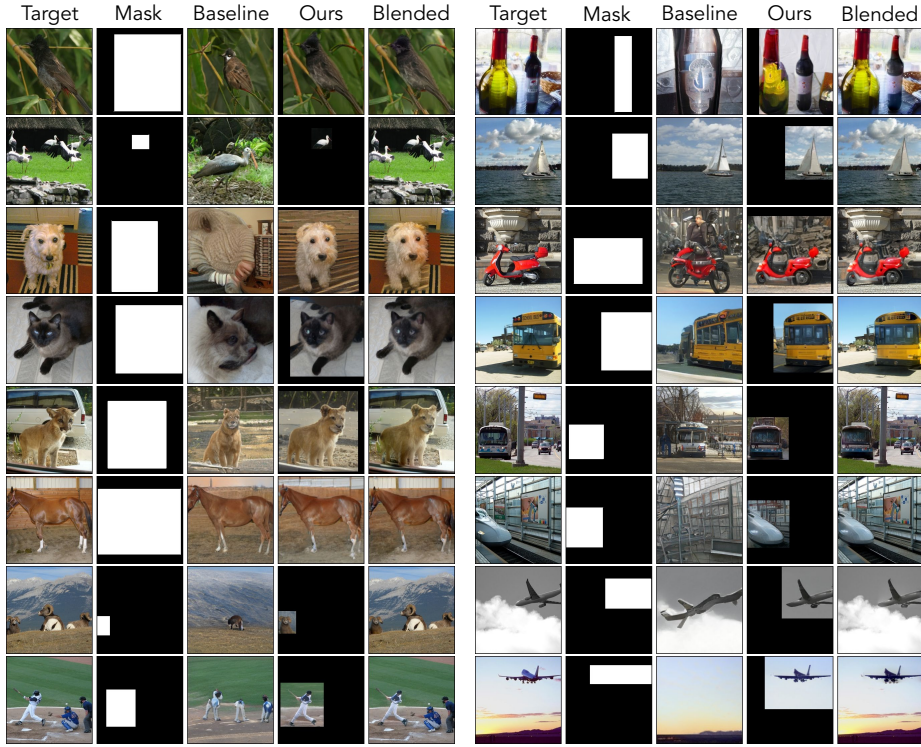


Fig. 7. **ImageNet results:** Results using our final method without fine-tuning. The final method uses BasinCMA as well as spatial and color transformation. Our generated results are inverted back for visualization. We also provide the ADAM baseline along with the blended result using Poisson blending [47].

class to generate the detection boxes. We use the same bounding box for all baselines, and the optimization hyper-parameters are tuned on a separate set of ImageNet images.

Experimental details. We use a learning rate of 0.05 for \mathbf{z} and 0.0001 for \mathbf{c} . We use AlexNet-LPIPS [36, 59] as our perceptual loss for all our methods. We did observe an improvement using VGG-LPIPS [52, 59] but found it to be 1.5 times slower. In our experiments, we use a total of 18 seeds for each method. After we project and edit the object, we blend the newly edited object with the original background using Poisson blending [47].

For all of our baselines, we optimize both the latent vector \mathbf{z} and class embedding \mathbf{c} . We use the same mask \mathbf{m} , and the same loss function throughout all of our experiments. The optimization details of our method and the baselines are in the Appendix A.

Experiments. We show qualitative comparisons of various optimization methods for ImageNet images in Figure 6. We show results of our final method with blending in Figure 7. We then quantify these results by comparing against each

Optimizer	Method			Average of 18 seeds				Best of 18 seeds			
	Spatial Transform	Color Transform	Encoder	Per-pixel		LPIPS		Per-pixel		LPIPS	
				L1	L2	Alex	VGG	L1	L2	Alex	VGG
ADAM				0.98	0.62	0.41	0.58	0.83	0.47	0.33	0.51
L-BFGS				1.04	0.68	0.45	0.61	0.85	0.49	0.35	0.53
CMA				0.96	0.61	0.39	0.55	0.91	0.54	0.37	0.54
None			✓	1.61	1.39	0.62	0.68	1.35	1.00	0.55	0.64
ADAM			✓	0.96	0.60	0.39	0.56	0.82	0.46	0.32	0.51
ADAM		✓		0.98	0.62	0.42	0.58	0.83	0.47	0.33	0.51
ADAM	✓			0.90	0.54	0.44	0.57	0.76	0.41	0.36	0.50
ADAM	✓	✓	✓	0.88	0.52	0.42	0.55	0.76	0.40	0.36	0.49
CMA+ADAM				0.93	0.57	0.37	0.55	0.83	0.47	0.32	0.51
BasinCMA				0.82	0.48	0.29	0.51	0.78	0.43	0.26	0.49
BasinCMA			✓	0.82	0.47	0.29	0.50	0.78	0.43	0.26	0.49
BasinCMA		✓		0.81	0.46	0.29	0.50	0.77	0.42	0.25	0.49
BasinCMA	✓			0.72	0.38	0.33	0.48	0.69	0.35	0.31	0.46
BasinCMA	✓	✓	✓	0.71	0.37	0.32	0.47	0.68	0.34	0.31	0.46

Table 1. **ImageNet:** We compare various methods for inverting images from ImageNet (lower is better). The last row is our full method. The model is optimized using $L1$ and AlexNet-LPIPS perceptual loss. The mask and ground-truth class vector is provided for each method. We show the error using different metrics: per-pixel and perceptual [59]. We show the average and the best score among 18 random seeds. Methods that optimized for transformation are inverted to the original location and the loss is computed on the masked region for a fair comparison. All the results here are not fine-tuned.




Fig. 8. **Failure cases:** Our method fails to invert images that are not well represented by BigGAN. The mask is overlayed on the target image in blue.

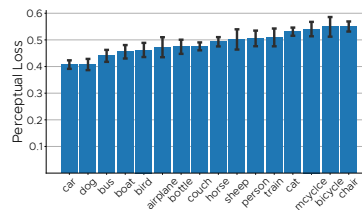


Fig. 9. **Projection error by class:** The average VGG-perceptual loss with standard error. The ImageNet images are sampled from the PASCAL super-class.

method using various metrics in Table 1. For all methods, we do *not* fine-tune our results and we only compute the loss inside the mask for a fair comparison. For methods optimized with transformation, the projected images are inverted back before computing the loss. We further evaluate on COCO dataset [38] in Table 3, and observed our findings to hold true on out-of-distribution dataset. The success of hybrid optimization over purely gradient-based optimization techniques may indicate that the generative model latent space is locally smooth but not globally.

Without transforming the object, we observed that the optimization often fails to find an approximate solution, specifically when the objects are off-centered or contain multiple objects. We observed that optimizing over color transformation does not lead to drastic improvements. Possibly because BigGAN can closely match the color gamut statistics of ImageNet images. Nonetheless, we found that optimizing for color transformation can slightly improve visual aesthetics. Out of the experimented color transformations, optimizing for brightness gave us the best result, and we use this for color transformation throughout our

Class search	Best of 18 seeds			
	Per-pixel		LPIPS	
	L1	L2	Alex	VGG
Random Gaussian	1.26	0.88	0.69	0.86
Random Class	0.88	0.51	0.40	0.59
Predicted	0.84	0.47	0.33	0.52
Ground Truth	0.83	0.47	0.33	0.51

Table 2. **Class search:** Given a fixed optimization method (ADAM), we compare different methods for initializing the class vector (lower is better). Baselines are: initialized from $\mathcal{N}(\mathbf{0}, \mathbf{I})$, a random class, and the ground truth class.

Method	Best of 18 seeds			
	Per-pixel		LPIPS	
	L1	L2	Alex	VGG
ADAM	0.96	0.57	0.32	0.56
ADAM + Transform	0.81	0.45	0.39	0.52
BasinCMA	0.93	0.18	0.81	0.53
BasinCMA + Transform	0.78	0.42	0.36	0.49

Table 3. **Out-of-distribution:** We compare different methods on the COCO dataset (lower is better). BigGAN was not trained on COCO images. The class labels are predicted using ResNext-101 and the masks are predicted using MaskRCNN.

experiments. We further experimented with composing multiple color transformations but did not observe additional improvements.

We found that using CMA/BasinCMA is robust to initialization and is a better optimization technique regardless of whether the transform was applied. Note that we did not observe any benefits of optimizing the class vectors \mathbf{c} with CMA compared to gradient-based methods, perhaps because the embedding between the continuous class vectors is not necessarily meaningful. Qualitatively, we often found the class embeddings to be meaningful when it is either in the close vicinity of original class embeddings or between the interpolation of 2 similar classes and not more. As a result, we use gradient descent to search within the local neighborhood of the initial class embedding space.

We also provide ablation study on how the number of CMA and ADAM updates for BasinCMA affects performance, and how other gradient-free optimizers compare against CMA in Appendix D. We further provide additional qualitative results for our final method in Appendix C.

Class initialization. In downstream editing application, the user may not know the exact ImageNet class the image belongs to. In Table 2, we compare different strategies for initializing the class vector. Here the classifier makes an incorrect prediction 20% of the time. We found that using the predicted class of an ImageNet classifier performs almost as well as the ground truth class. Since we optimize the class vector, we can potentially recover from a wrong initial guess if the predicted class is sufficiently close to the ground-truth.

Failure cases. Figure 8 shows some typical failure cases. We observed that our method fails to embed images that are not well modeled by BigGAN – outlier modes that may have been dropped. For example, we failed to project images that are unique, complicated, rotated, or heavily occluded. More sophisticated transformations such as rotations and perspective transformation could address many of these failure cases and are left for future work.

Which classes does BigGAN struggle to generate? Given our method, we analyze which classes BigGAN, or our method has difficulty generating. In Figure 9, we plot the mean and the standard error for each class. The plot is

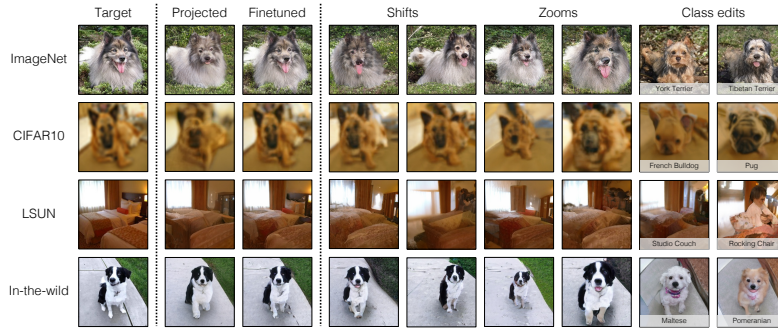


Fig. 10. **Fine-tuned edits:** Inversion results on various datasets. We use BasinCMA and transformation to optimize for the latent variables. After obtaining the projections, we fine-tune the model weights and perform edits in the latent and class vector space.

from the output of the method optimized with ADAM + CMA + Transform. We observed a general tendency for the model to struggle in generating objects with delicate structures or with large inter-class variance.

Image Edits. A good approximate solution allows us to fine-tune the generative model and recover the details easily. Good approximations require less fine-tuning and therefore preserve the original generative model editing capabilities. In Figure 10, we embed images from various datasets including CIFAR [35], LSUN [57], and images in-the-wild. We then fine-tune and edit the results by changing the latent vector or class vector. Prior works [48, 29] have found that certain latent vectors can consistently control the appearance change of GANs-generated images such as shifting an image horizontally or zooming an image in and out. We used the “shift” and “zoom” vectors [29] to modify our images. Additionally, we also varied the class vector to a similar class and observed the editability to stay consistent. Even for images like CIFAR, our method was able to find good solutions that allowed us to edit the image. In cases like LSUN, where there is no corresponding class for the scene, we observed that the edits ended up being meaningless.

5 Discussion

Projecting an image into the “space” of a generative model is a crucial step for editing applications. We have systematically explored methods for this projection. We show that using a gradient-free optimizer, CMA, produces higher quality matches. We account for biases in the generative model by enabling spatial and color transformations in the search, and the combination of these techniques finds a closer match and better serves downstream editing pipelines. Future work includes exploring more transformations, such as local geometric changes and global appearance changes, as well as modeling generation of multiple objects or foreground/background.

Acknowledgements. We thank David Bau, Phillip Isola, Lucy Chai, and Erik Härkönen for discussions, and David Bau for encoder training code.

References

1. Abdal, R., Qin, Y., Wonka, P.: Image2stylegan: How to embed images into the stylegan latent space? In: International Conference on Computer Vision (2019) [4](#), [6](#), [8](#)
2. Asim, M., Shamshad, F., Ahmed, A.: Blind image deconvolution using deep generative priors. In: British Machine Vision Conference (2018) [3](#)
3. Baker, S., Matthews, I.: Lucas-kanade 20 years on: A unifying framework. International Journal of Computer Vision **56**(3), 221–255 (2004) [9](#)
4. Barnes, C., Shechtman, E., Finkelstein, A., Goldman, D.B.: Patchmatch: A randomized correspondence algorithm for structural image editing. ACM Transactions on Graphics (TOG) (2009) [3](#)
5. Bau, D., Strobel, H., Peebles, W., Wulff, J., Zhou, B., Zhu, J.Y., Torralba, A.: Semantic photo manipulation with a generative image prior. ACM Transactions on Graphics (TOG) (2019) [1](#), [3](#), [4](#), [10](#)
6. Bau, D., Zhu, J.Y., Strobel, H., Zhou, B., Tenenbaum, J.B., Freeman, W.T., Torralba, A.: Gan dissection: Visualizing and understanding generative adversarial networks. In: International Conference on Learning Representations (2019) [1](#), [3](#)
7. Bau, D., Zhu, J.Y., Wulff, J., Peebles, W., Strobel, H., Zhou, B., Torralba, A.: Seeing what a gan cannot generate. In: International Conference on Computer Vision (2019) [1](#), [3](#), [4](#), [6](#), [7](#), [8](#)
8. Brock, A., Donahue, J., Simonyan, K.: Large scale gan training for high fidelity natural image synthesis. In: International Conference on Learning Representations (2019) [2](#), [4](#), [5](#)
9. Brock, A., Lim, T., Ritchie, J.M., Weston, N.: Neural photo editing with introspective adversarial networks. In: International Conference on Learning Representations (2017) [1](#), [3](#)
10. Creswell, A., Bharath, A.A.: Inverting the generator of a generative adversarial network. IEEE Transactions on Neural Networks and Learning Systems **30**(7), 1967–1974 (2018) [3](#)
11. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: IEEE Conference on Computer Vision and Pattern Recognition (2009) [10](#)
12. Donahue, J., Krähenbühl, P., Darrell, T.: Adversarial feature learning. In: International Conference on Learning Representations (2017) [3](#)
13. Donahue, J., Simonyan, K.: Large scale adversarial representation learning. In: Advances in Neural Information Processing Systems (2019) [3](#)
14. Dosovitskiy, A., Brox, T.: Generating images with perceptual similarity metrics based on deep networks. In: Advances in Neural Information Processing Systems (2016) [5](#)
15. Dosovitskiy, A., Brox, T.: Inverting visual representations with convolutional networks. In: IEEE Conference on Computer Vision and Pattern Recognition (2016) [3](#)
16. Dumoulin, V., Belghazi, I., Poole, B., Lamb, A., Arjovsky, M., Mastropietro, O., Courville, A.: Adversarially learned inference. In: International Conference on Learning Representations (2017) [3](#)
17. Efros, A.A., Freeman, W.T.: Image quilting for texture synthesis and transfer. In: ACM SIGGRAPH (2001) [3](#)
18. Efros, A.A., Leung, T.K.: Texture synthesis by non-parametric sampling. In: International Conference on Computer Vision (1999) [3](#)

19. Everingham, M., Eslami, S.M.A., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The pascal visual object classes challenge: A retrospective. *International Journal of Computer Vision* **111**(1), 98–136 (2015) 10
20. Fang, T., Schwing, A.: Co-generation with gans using ais based hmc. In: *Advances in Neural Information Processing Systems* (2019) 3
21. Gatys, L.A., Ecker, A.S., Bethge, M.: Image style transfer using convolutional neural networks. In: *IEEE Conference on Computer Vision and Pattern Recognition* (2016) 5
22. Gonzalez, R.C., Woods, R.E.: *Digital Image Processing*. Pearson, 2nd edn. (1992) 9
23. Goodfellow, I.: NIPS 2016 tutorial: Generative adversarial networks. arXiv preprint arXiv:1701.00160 (2016) 6
24. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: *Advances in Neural Information Processing Systems* (2014) 1
25. Gu, J., Shen, Y., Zhou, B.: Image processing using multi-code gan prior. In: *IEEE Conference on Computer Vision and Pattern Recognition* (2020) 3
26. Hansen, N., Ostermeier, A.: Completely derandomized self-adaptation in evolution strategies. *evolutionary computation*. *Evolutionary Computation* (2001) 2, 8
27. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: *International Conference on Computer Vision* (2017) 5, 6, 10
28. Hertzmann, A., Jacobs, C.E., Oliver, N., Curless, B., Salesin, D.H.: Image analogies. In: *ACM SIGGRAPH* (2001) 3
29. Jahanian, A., Chai, L., Isola, P.: On the “steerability” of generative adversarial networks. In: *International Conference on Learning Representations* (2020) 1, 6, 14
30. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: *European Conference on Computer Vision* (2016) 5
31. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of gans for improved quality, stability, and variation. In: *International Conference on Learning Representations* (2018) 7
32. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: *IEEE Conference on Computer Vision and Pattern Recognition* (2019) 1, 7
33. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: *International Conference on Learning Representations* (2015) 2, 8
34. Kingma, D.P., Dhariwal, P.: Glow: Generative flow with invertible 1x1 convolutions. In: *Advances in Neural Information Processing Systems* (2018) 5
35. Krizhevsky, A.: *Learning Multiple Layers of Features from Tiny Images*. Master’s thesis, University of Toronto (2009) 14
36. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: *Advances in Neural Information Processing Systems* (2012) 11
37. Levin, A., Lischinski, D., Weiss, Y.: A closed-form solution to natural image matting. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **30**(2), 228–242 (2007) 3
38. Lin, T., Maire, M., Belongie, S.J., Bourdev, L.D., Girshick, R.B., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: common objects in context. In: *European Conference on Computer Vision* (2014) 12
39. Lipton, Z.C., Tripathi, S.: Precise recovery of latent vectors from generative adversarial networks. *ICLR workshop* (2017) 3

40. Liu, D.C., Nocedal, J.: On the limited memory bfgs method for large scale optimization. *Mathematical programming* **45**(1-3), 503–528 (1989) [2](#), [8](#)
41. Mahendran, A., Vedaldi, A.: Understanding deep image representations by inverting them. In: *IEEE Conference on Computer Vision and Pattern Recognition* (2015) [3](#)
42. Max Jaderberg, Karen Simonyan, A.Z.K.K.: Spatial transformer networks. In: *Advances in Neural Information Processing Systems* (2015) [6](#)
43. Olah, C., Mordvintsev, A., Schubert, L.: Feature visualization. *Distill* **2**(11), e7 (2017) [3](#)
44. Olah, C., Satyanarayan, A., Johnson, I., Carter, S., Schubert, L., Ye, K., Mordvintsev, A.: The building blocks of interpretability. *Distill* **3**(3), e10 (2018) [3](#)
45. Oppenheim, A.V., Schaffer, R.W., Buck, J.R.: *Discrete-Time Signal Processing*. Pearson, 2nd edn. (1999) [9](#)
46. Perarnau, G., Van De Weijer, J., Raducanu, B., Álvarez, J.M.: Invertible conditional gans for image editing. In: *NIPS 2016 Workshop on Adversarial Training* (2016) [1](#), [3](#)
47. Pérez, P., Gangnet, M., Blake, A.: Poisson image editing. *ACM Transactions on Graphics (TOG)* **22**(3), 313–318 (2003) [3](#), [11](#)
48. Radford, A., Metz, L., Chintala, S.: Unsupervised representation learning with deep convolutional generative adversarial networks. In: *International Conference on Learning Representations* (2016) [1](#), [7](#), [14](#)
49. Raj, Ankit, Y.L., Bresler, Y.: Gan-based projector for faster recovery with convergence guarantees in linear inverse problems. In: *International Conference on Computer Vision* (2019) [3](#)
50. Shah, V., Hegde, C.: Solving linear inverse problems using gan priors: An algorithm with provable guarantees. In: *International Conference on Acoustics, Speech, and Signal Processing* (2018) [3](#)
51. Shen, Y., Gu, J., Tang, X., Zhou, B.: Interpreting the latent space of gans for semantic face editing. In: *IEEE Conference on Computer Vision and Pattern Recognition* (2020) [3](#)
52. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: *International Conference on Learning Representations* (2015) [11](#)
53. Wampler, K., Popović, Z.: Optimal gait and form for animal locomotion. *ACM Transactions on Graphics (TOG)* (2009) [8](#)
54. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P., et al.: Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* **13**(4), 600–612 (2004) [5](#)
55. Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. *arXiv preprint arXiv:1611.05431* (2016) [7](#)
56. Yeh, R.A., Chen, C., Yian Lim, T., Schwing, A.G., Hasegawa-Johnson, M., Do, M.N.: Semantic image inpainting with deep generative models. In: *IEEE Conference on Computer Vision and Pattern Recognition* (2017) [3](#)
57. Yu, F., Zhang, Y., Song, S., Seff, A., Xiao, J.: Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365* (2015) [14](#)
58. Zhang, R., Isola, P., Efros, A.A.: Colorful image colorization. In: *European Conference on Computer Vision* (2016) [5](#)
59. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep networks as a perceptual metric. In: *IEEE Conference on Computer Vision and Pattern Recognition* (2018) [5](#), [11](#), [12](#)

- 60. Zhu, J.Y., Krähenbühl, P., Shechtman, E., Efros, A.A.: Generative visual manipulation on the natural image manifold. In: European Conference on Computer Vision (2016) 1, 3, 8