

Exploring the Solution Space of Linear Inverse Problems with GAN Latent Geometry

Original

Exploring the Solution Space of Linear Inverse Problems with GAN Latent Geometry / Montanaro, Antonio; Valsesia, Diego; Magli, Enrico. - ELETTRONICO. - (2022), pp. 1381-1385. (2022 IEEE International Conference on Image Processing (ICIP) Bordeaux, France 16-19 October 2022) [10.1109/ICIP46576.2022.9897211].

Availability:

This version is available at: 11583/2977425 since: 2023-03-24T12:24:49Z

Publisher:

IEEE

Published

DOI:10.1109/ICIP46576.2022.9897211

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2022 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

EXPLORING THE SOLUTION SPACE OF LINEAR INVERSE PROBLEMS WITH GAN LATENT GEOMETRY

Antonio Montanaro, Diego Valsesia, Enrico Magli

Department of Electronics and Telecommunications – Politecnico di Torino, Italy

ABSTRACT

Inverse problems consist in reconstructing signals from incomplete sets of measurements and their performance is highly dependent on the quality of the prior knowledge encoded via regularization. While traditional approaches focus on obtaining a unique solution, an emerging trend considers exploring multiple feasible solutions. In this paper, we propose a method to generate multiple reconstructions that fit both the measurements and a data-driven prior learned by a generative adversarial network. In particular, we show that, starting from an initial solution, it is possible to find directions in the latent space of the generative model that are null to the forward operator, and thus keep consistency with the measurements, while inducing significant perceptual change. Our exploration approach allows to generate multiple solutions to the inverse problem an order of magnitude faster than existing approaches; we show results on image super-resolution and inpainting problems.

Index Terms— Inverse problems, GANs

1. INTRODUCTION

Linear inverse problems are ubiquitous in the sciences as they are tasked with reconstructing a signal of interest from a set of typically incomplete or degraded measurements. In the imaging field alone [1], numerous problems of interest such as deblurring, super-resolution, inpainting, compressed sensing, and many more fit this framework. Due to the ill-posed nature of the problem, one needs strong regularization to find reconstructions that fit the measurements and the a priori knowledge of the signal properties. Traditional approaches focused on hand-crafting regularizers to yield a unique solution by casting reconstruction as a convex optimization problem [2]. However, one must accept that the quality of this unique solution can only be as good as how well the chosen regularizer function captures the signal properties. For this reason, recently, data-driven methods based on neural networks [3, 4] started learning priors directly from the complex distributions of the signals of interest, resulting in improved reconstruction capabilities.

Nevertheless, even when using data-driven priors, we can hardly hope to capture a perfect model of our signals of interest, which in turn affects how faithful the reconstructed signal is to the true one that generated the measurements. For this reason, a novel paradigm is emerging where multiple feasible reconstructions are generated, in an effort to boost interpretability of the inversion process and expose the biases of the models.

In this paper, we use generative adversarial networks (GANs) as priors modeling the distribution of our signals of interest. We present a geometrical perspective on the latent space of such models, which allows to explore the solution space of a linear inverse problem. By exploration of the solution space, we mean finding *multiple reconstructions* that are consistent with the measurements but also consistent with the model of the data distribution. We show that it is

possible to modify an initial solution by moving towards directions in the latent space that are “null” with respect to the measurements operator (i.e., they do not significantly perturb the measurements) while inducing semantic change. Our proposed technique, called e-GLASS (exploring GAN LATent Space Solutions), is general as it can be applied to any linear inverse problem and is an order of magnitude faster than state-of-the-art methods such as PULSE [5] which generate multiple solutions by solving an optimization problem from different random initializations.

2. BACKGROUND AND RELATED WORK

Let us start from a general linear forward model of the form:

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n} \quad (1)$$

where $\mathbf{y} \in \mathbb{R}^m$ is a noisy observation from an unknown signal $\mathbf{x} \in \mathbb{R}^n$ with $m \leq n$, depending on the specific problem; \mathbf{n} is an additive noise and \mathbf{A} is a measurement matrix.

As an example, \mathbf{y} can be a degraded image, e.g., with low resolution or blurred, and we want to reconstruct the image \mathbf{x} starting from the measurements \mathbf{y} . However, this problem is ill-posed as there can be infinitely many solutions satisfying the measurements, or even none due to noise. A large body of work has been devoted to the development of priors to model \mathbf{x} as accurately as possible to regularize the problem towards admitting a unique solution. Such works frame reconstruction as a Maximum a Posteriori (MAP) estimation problem, with the unique solution obtained by solving

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \lambda R(\mathbf{x}),$$

for some handcrafted regularizer R encoding the prior. This line of works typically defines priors such that the reconstruction problem is convex with a unique global minimum. This means that a single solution to the problem can be obtained, whose properties are strictly intertwined with the ability to craft a suitable prior R .

New recent approaches involve generative models, such as GANs, to learn priors in a data-driven fashion [6], [7]. A GAN learns a function G that maps a latent vector \mathbf{z} into a sample from the data distribution. The popular approach of GAN inversion solves inverse problems by seeking the latent vector $\hat{\mathbf{z}}$ that best fits the measurements \mathbf{y} . This is done by minimizing the distance between \mathbf{y} and the degraded version of the generated data $G(\hat{\mathbf{z}})$, under the forward model \mathbf{A} :

$$\hat{\mathbf{z}} = \arg \min_{\mathbf{z}} \|\mathbf{y} - \mathbf{A}G(\mathbf{z})\|_2^2, \quad \hat{\mathbf{x}} = G(\hat{\mathbf{z}}). \quad (2)$$

Unlike convex optimization methods with handcrafted priors, GAN inversion is non-convex due to the use of neural networks, thus admitting multiple local minima.

While most works have used GAN inversion to generate a single solution to the inverse problem, there has been recent growing interest in exploring the *solution space* of inverse problems, i.e., finding multiple solutions, among the infinitely many possible, that are

consistent with the measurements and some data prior. The seminal work on this topic is PULSE [5], which uses GAN inversion to super-resolve low-resolution faces. PULSE generates multiple plausible solutions by solving Eq. (2) via gradient descent, and starting from different random guesses of \mathbf{z} . Due to the non-convex nature of the optimization, different solutions *may* be reached when starting from different initializations. The main drawback of PULSE lies in its complexity, requiring to solve an optimization problem for each solution and the lack of any guarantee that a different initialization will converge to a different minimum. Other works [8] have sought to generate multiple solutions for the super-resolution problem, but they lack generality and can only be applied to very specific neural networks devised only for the super-resolution task.

Finally, we remark that there is extensive literature on GAN edibility [9], [10], [11], [12], seeking to manipulate the latent space of GANs to induce semantically interesting transformations. However, such works are not in the framework of solutions to inverse problems and are not concerned with fidelity with measurements.

3. PROPOSED METHOD

In this paper, we propose a method to explore multiple solutions of a linear inverse problem, starting from a first solution $G(\mathbf{z}_0)$. The method exploits geometrical properties of the GAN latent space \mathcal{Z} to navigate in a neighborhood of \mathbf{z}_0 in such a way that the new generated data preserve the condition in Eq. (1) (i.e., they are solutions to the inverse problem) while manifesting novel semantic information with respect to $G(\mathbf{z}_0)$.

The latent space \mathcal{Z} can be seen as a Riemannian manifold, and a GAN G parametrizes a submanifold to the data space \mathcal{X} , and, ultimately to the measurement space \mathcal{Y} via the composition of generator and forward model $\phi = G \circ \mathbf{A}$. Wang and Ponce [13] argue that the geometry of \mathcal{Z} in a neighborhood of \mathbf{z}_0 can be approximated by a positive semi-definite quadratic form $H(\mathbf{z}_0)$:

$$d^2(\mathbf{z}_0, \mathbf{z}) \approx \delta \mathbf{z}^T \left. \frac{\partial^2 d^2(\mathbf{z}_0, \mathbf{z})}{\partial \mathbf{z}^2} \right|_{\mathbf{z}_0} \delta \mathbf{z}, \quad \mathbf{H}(\mathbf{z}_0) := \left. \frac{\partial^2 d^2(\mathbf{z}_0, \mathbf{z})}{\partial \mathbf{z}^2} \right|_{\mathbf{z}_0},$$

dependent on a distance metric d between latent vectors. While the authors in [13] define d between latents as the distance in the generated data space \mathcal{X} , we also consider it in the measurement space \mathcal{Y} , since we are interested in exploring how variations in the measurement space affect the latent geometry. In particular, we define $d_{\mathcal{Y}}(\mathbf{z}_1, \mathbf{z}_2) := \|\phi(\mathbf{z}_1) - \phi(\mathbf{z}_2)\|_2^2 = \|\mathbf{A}G(\mathbf{z}_1) - \mathbf{A}G(\mathbf{z}_2)\|_2^2$ and induce the corresponding manifold described by Riemannian metric:

$$\mathbf{H}_{\mathcal{Y}}(\mathbf{z}_0) = \frac{1}{2} \frac{\partial^2}{\partial \mathbf{z}^2} \|\phi(\mathbf{z}) - \phi(\mathbf{z}_0)\|_2^2|_{\mathbf{z}_0} = \mathbf{J}_{\phi}^T(\mathbf{z}_0) \mathbf{J}_{\phi}(\mathbf{z}_0),$$

being $\mathbf{J}_{\phi}(\mathbf{z}_0) = \partial_{\mathbf{z}} \phi(\mathbf{z})|_{\mathbf{z}_0}$ the Jacobian of $\phi = G \circ \mathbf{A}$ evaluated at point \mathbf{z}_0 . Similarly, metric $\mathbf{H}_{\mathcal{X}}(\mathbf{z}_0)$ is induced by a suitable distance in the data space. In this work, we will focus on images and, consequently, we use the LPIPS distance (a perceptual metric defined from features extracted by a pretrained network) [14] as $d_{\mathcal{X}}(\mathbf{z}_1, \mathbf{z}_2) := \text{LPIPS}(G(\mathbf{z}_1), G(\mathbf{z}_2))$. Backpropagation can be used to compute $\mathbf{H}_{\mathcal{Y}}$ and $\mathbf{H}_{\mathcal{X}}$.

Armed with this characterization of the geometry of the latent space, we seek to generate a new latent vector corresponding to a solution as $\mathbf{z}_1 = \mathbf{z}_0 + \eta \mathbf{d}$, i.e., by perturbing \mathbf{z}_0 along a direction \mathbf{d} that maximizes perceptual distance in the image space (large $d_{\mathcal{X}}(\mathbf{z}_1, \mathbf{z}_0)$) but minimizes distance in the measurement space (small $d_{\mathcal{Y}}(\mathbf{z}_1, \mathbf{z}_0)$). In other words, we seek to explore the subspace of \mathcal{Z} around \mathbf{z}_0 that is “null” with respect to the measurements operator but not so with respect to perceptual distance.

Algorithm 1 e-GLASS: exploring GAN Latent Space Solutions

Input: \mathbf{z}_0, K

Output: New solution $\hat{\mathbf{x}}$

Compute Hessian $\mathbf{H}_{\mathcal{Y}}(\mathbf{z}_0) = \frac{1}{2} \frac{\partial^2}{\partial \mathbf{z}^2} \|\mathbf{A}G(\mathbf{z}) - \mathbf{A}G(\mathbf{z}_0)\|_2^2$

Compute Hessian $\mathbf{H}_{\mathcal{X}}(\mathbf{z}_0) = \frac{1}{2} \frac{\partial^2}{\partial \mathbf{z}^2} \text{LPIPS}(G(\mathbf{z}), G(\mathbf{z}_0))$

Compute eigenvectors $\mathbf{H}_{\mathcal{Y}}(\mathbf{z}_0) = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^T$

Compute eigenvectors $\mathbf{H}_{\mathcal{X}}(\mathbf{z}_0) = \mathbf{V} \mathbf{\Omega} \mathbf{V}^T$

$\mathbf{d} \leftarrow \mathbf{v}^K$

$\mathcal{J} = \{\text{top-}k \text{ eigenvectors of } \mathbf{H}_{\mathcal{Y}}\}$

for $j \in \mathcal{J}$ **do**

$\mathbf{d} \leftarrow \mathbf{d} - (\mathbf{d}^T \mathbf{u}^j) \mathbf{u}^j$

$\mathbf{d} \leftarrow \mathbf{d} / \|\mathbf{d}\|$

end for

$\hat{\mathbf{x}} = G(\mathbf{z}_0 + \eta \mathbf{d})$

One might wonder whether this is possible at all, and, in fact, the answer is affirmative and relies on two main phenomena. The first was observed by Wang and Ponce [13] and it is the *anisotropy* of the space described by $\mathbf{H}_{\mathcal{X}}(\mathbf{z}_0)$, i.e., $\mathbf{H}_{\mathcal{X}}$ is described by a small number of principal components, meaning that there is a large number of directions that have little to no effect on perceptual quality and some significantly changing it¹. We empirically observe the same regarding the geometry induced by the measurements fidelity, i.e., $\mathbf{H}_{\mathcal{Y}}(\mathbf{z}_0)$. The second phenomenon, which is at the basis of our work, is that the directions from $\mathbf{H}_{\mathcal{X}}$ and $\mathbf{H}_{\mathcal{Y}}$ can be empirically decoupled. This means that it is indeed possible to find directions that significantly affect perceptual distance while having little to no impact on measurements, yielding novel solutions to the inverse problem.

Algorithm 1 summarizes our proposed e-GLASS scheme to find such directions. We first start by finding the latent code \mathbf{z}_0 corresponding to a single solution by means of any state-of-the-art GAN inversion technique. Then we compute the Hessians $\mathbf{H}_{\mathcal{Y}}(\mathbf{z}_0)$, $\mathbf{H}_{\mathcal{X}}(\mathbf{z}_0)$ and their eigenvectors: $\mathbf{H}_{\mathcal{Y}}(\mathbf{z}_0) = \mathbf{U} \mathbf{\Lambda} \mathbf{U}^T$, $\mathbf{H}_{\mathcal{X}}(\mathbf{z}_0) = \mathbf{V} \mathbf{\Omega} \mathbf{V}^T$. We then need to measure the coupling between the two sets of eigenvectors via the coupling matrix $\mathbf{C} = \mathbf{U}^T \mathbf{V}$. It is expected that the top eigenvectors in \mathbf{U} are coupled with the top eigenvectors in \mathbf{V} as large perceptual distances typically also correspond to large differences on the measurements. However, the bottom eigenvectors are not correlated, indicating that the corresponding null spaces do not necessarily intersect each other. The most interesting directions for our problem are the eigenvectors that are among the top in \mathbf{V} but do not correlate with the top eigenvectors in \mathbf{U} . However, directly choosing such direction is in general suboptimal, as it might still increase the distance in the measurement space more than desired.

To solve this problem, we propose a geometrical method that removes the most relevant correlations with the top eigenvectors of $\mathbf{H}_{\mathcal{Y}}$. This allows to obtain a new direction that is hopefully still creating perceptually significant differences but projected as much as possible onto the null space of $\mathbf{H}_{\mathcal{Y}}$ to minimally change the measurements. To do this, we first choose $\mathbf{d} = \mathbf{v}^K$ as the K -th top eigenvector to discard the very top eigenvectors that are coupled with the top ones in \mathbf{U} . Then, we project \mathbf{v}^K onto the hyperplane orthogonal to the top eigenvectors \mathbf{u}^j with correlation larger than a threshold:

$$\mathbf{d} \leftarrow \mathbf{d} - (\mathbf{d}^T \mathbf{u}^j) \mathbf{u}^j, \quad \mathbf{d} \leftarrow \mathbf{d} / \|\mathbf{d}\|.$$

This procedure is iterated until the resulting direction has no significant correlation to the top eigenvectors of $\mathbf{H}_{\mathcal{Y}}$. This leads to a pro-

¹[13] also note that the space is *homogeneous*, meaning that this property is valid everywhere, regardless of the specific \mathbf{z}_0 .

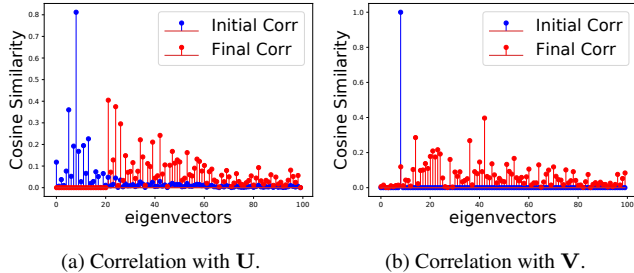


Fig. 1: Correlation of initial direction \mathbf{v}^K and final direction \mathbf{d} with eigenvectors of latent space metrics \mathbf{H}_Y and \mathbf{H}_X . The final direction is orthogonal to directions inducing large change in measurements, but correlates with directions inducing significant perceptual change.

jection of \mathbf{v}^K onto the null space of \mathbf{H}_Y . Multiple solutions to the inverse problem can be explored by either changing the step η along the direction, or trying a new direction by computing \mathbf{d} starting from \mathbf{v}^{K-1} , \mathbf{v}^{K-2} , ...

4. EXPERIMENTAL RESULTS

In this section, we experimentally evaluate the proposed method against state-of-the-art techniques to explore multiple solutions. While the proposed method is general and holds for different generative models and different inverse problems, we focus on two notable inverse problems, i.e., image super-resolution (SR) and inpainting (IP), presenting results for two different generative models, i.e. BigGAN [15] and PGGAN [16]. For super-resolution, we down-scale the 256×256 image to a 32×32 image, while for inpainting, we delete a semantically interesting area of a face from a full image of size 1024×1024 .

We first present empirical evidence about our claims in Sec.3 that directions that induce little change on measurements and significant perceptual change on the reconstructions do exist. To show this we want to see that the chosen direction correlates with the top eigenvectors of \mathbf{V} while being as orthogonal as possible to the top eigenvectors of \mathbf{U} . Fig. 1 shows in blue the correlation coefficient between the starting direction \mathbf{v}^K and the eigenvectors in \mathbf{U} and \mathbf{V} and in red the same correlations but with respect to the final direction \mathbf{d} provided by our method. It can be noticed that the final direction has been successfully orthogonalized with respect to the directions inducing significant variation on the measurements, while it retains good correlation with directions inducing perceptual change.

We now qualitatively and quantitatively examine the performance of the proposed method for the chosen inverse problems in comparison with PULSE [5]. PULSE generates multiple solutions by solving the GAN inversion problem multiple times from different random initialization, in the hope of converging to a different local minimum. For our proposed method, we generate the initial solution by means of the state-of-the-art GAN inversion technique proposed by Abu Hussein et al. [11], where the inversion problem in Eq. (2) also optimizes with respect to the GAN parameters to finetune them. Once the initial solution is computed, we apply Algorithm 1 to find a new solution to the problem.

Fig. 2 shows a few results on the SR problem. The middle row shows what reconstructions would be obtained if direction \mathbf{v}^K were used without our proposed algorithm. It can be noticed that there is significant perceptual change but the ℓ_2 norm with respect to the measurements is poorly constrained, so that these reconstructions can be hardly called feasible solutions. The last row shows the images generated by the direction found by our method. We suc-

Table 1: Computational time required to find 10 solutions, with our method and using different initializations (PULSE).

	Inverse Problem	Model	Time (s) to 10 solutions
PULSE	SR	BigGAN	3.5×10^4
e-GLASS	SR	BigGAN	3.7×10^3
PULSE	IP	PGGAN	7.9×10^3
e-GLASS	IP	PGGAN	1.3×10^3

cessfully constrain the ℓ_2 norm with respect to the measurements below the 10^{-2} threshold we consider acceptable for feasibility. At the same time, perceptual variations are still present in those regions where the highly downsampled nature of the measurements leaves more freedom to fill in information, such as the color and shape of the dog’s coat (from pale yellow to white, and the shape of the ears). Finally, the top row shows some good solutions found by PULSE. Those solutions are feasible according to our ℓ_2 criterion but are less perceptually convincing. Indeed while some solutions show different dogs, these present some artifacts around the dog’s mouth or some blurring over the whole dog’s coat. Instead, the two central solutions in Fig 2 (top row), show unnatural dogs found by the PULSE algorithm that still satisfy the ℓ_2 criterion.

Fig. 3 shows the results for the IP problem. Even in this case, PULSE (top row) found solutions with different perceptual changes, but still some of them seem to introduce unnatural variations, like the one on the nose or some distortion on the mouth. The second row shows images along two starting eigenvectors of the Hessians without applying our method. Although notable perceptual changes are visible, the ℓ_2 distance on the measure is outside of our constraint and indeed some differences on the eyes are created with respect to the observations. Finally, the last row, exploits our optimized direction showing how the ℓ_2 distance on the measurements is decreased while retaining good semantic changes in the masked area, such variations in the lip thickness or the amount of beard.

We remark that PULSE may be able to find good solutions but it has two main drawbacks that are solved by the proposed technique. First, it lacks any explicit control on the measurements distance. Constraining the ℓ_2 distance between the original measurements and the measurements of the reconstruction to a feasibility threshold can only be done by enforcing a stopping criterion on the inversion optimization problem. However, due to the non-convex nature of the problem this often results in degenerate solutions that no longer belong to the manifold of realistic images like some of the ones previously shown.

Another advantage of the proposed method with respect to PULSE is the computational complexity due to PULSE requiring to solve a full optimization problem to generate a new solution. For our proposed method, this needs to be only done once, coupled with the estimation of the Hessians, but then multiple solutions can be generated almost instantaneously. Table 1 reports the time required for the two methods to generate ten solutions. This time does not account for bad solutions: indeed, discarding bad minima found by PULSE would further increase its computational requirements.

5. CONCLUSIONS

We proposed a novel framework to explore the solution space of linear inverse problems. By exploiting the power of the geometry of the latent space around an initial solution, we showed it is possible to search for novel solutions that have semantic differences between each other. We confirmed the quality and variety of the multiple solutions as well as the reduced complexity when compared with state-of-the-art iterative optimization methods on the SR and IP problems.




Fig. 2: Top row: solutions found by PULSE (LR ℓ_2 range: $[1.8 \times 10^{-3}, 3 \times 10^{-3}]$). Mid row: solutions found by using \mathbf{v}^8 and \mathbf{v}^{12} as directions (LR ℓ_2 range: $[2.4 \times 10^{-2}, 4.5 \times 10^{-2}]$). Bottom row: solutions found by optimized \mathbf{d} as direction (LR ℓ_2 range: $[2.9 \times 10^{-3}, 4.8 \times 10^{-3}]$).



Fig. 3: Top row: solutions found by PULSE (ℓ_2 range: $[1.3 \times 10^{-4}, 1.5 \times 10^{-4}]$). Mid row: solutions found by using \mathbf{v}^8 and \mathbf{v}^{26} as directions (ℓ_2 range: $[3.9 \times 10^{-3}, 4.7 \times 10^{-3}]$). Bottom row: solutions found by optimized \mathbf{d} as direction (ℓ_2 range: $[1.2 \times 10^{-3}, 2 \times 10^{-3}]$).

6. REFERENCES

- [1] Mario Bertero, Patrizia Boccacci, and Christine De Mol, *Introduction to inverse problems in imaging*, CRC press, 2021.
- [2] Julien Mairal, Francis Bach, Jean Ponce, and Guillermo Sapiro, “Online dictionary learning for sparse coding,” in *Proceedings of the 26th annual international conference on machine learning*, 2009, pp. 689–696.
- [3] Alice Lucas, Michael Iliadis, Rafael Molina, and Aggelos K Katsaggelos, “Using deep neural networks for inverse problems in imaging: beyond analytical methods,” *IEEE Signal Processing Magazine*, vol. 35, no. 1, pp. 20–36, 2018.
- [4] Vishal Monga, Yuelong Li, and Yonina C Eldar, “Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing,” *IEEE Signal Processing Magazine*, vol. 38, no. 2, pp. 18–44, 2021.
- [5] Sachit Menon, Alexandru Damian, Shijia Hu, Nikhil Ravi, and Cynthia Rudin, “Pulse: Self-supervised photo upsampling via latent space exploration of generative models,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 2437–2445.
- [6] Antonia Creswell and Anil Anthony Bharath, “Inverting the generator of a generative adversarial network,” *IEEE transactions on neural networks and learning systems*, vol. 30, no. 7, pp. 1967–1974, 2018.
- [7] Weihao Xia, Yulun Zhang, Yujiu Yang, Jing-Hao Xue, Bolei Zhou, and Ming-Hsuan Yang, “Gan inversion: A survey,” *arXiv preprint arXiv:2101.05278*, 2021.
- [8] Andreas Lugmayr, Martin Danelljan, and Radu Timofte, “Ntire 2021 learning the super-resolution space challenge,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 596–612.
- [9] Yujun Shen, Ceyuan Yang, Xiaoou Tang, and Bolei Zhou, “Interfacegan: Interpreting the disentangled face representation learned by gans,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, 2020.
- [10] Xingang Pan, Xiaohang Zhan, Bo Dai, Dahua Lin, Chen Change Loy, and Ping Luo, “Exploiting deep generative prior for versatile image restoration and manipulation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [11] Shady Abu Hussein, Tom Tirer, and Raja Giryes, “Image-adaptive gan based reconstruction,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020, vol. 34, pp. 3121–3129.
- [12] Yujun Shen, Jinjin Gu, Xiaoou 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*, 2020, pp. 9243–9252.
- [13] Binxu Wang and Carlos R Ponce, “A geometric analysis of deep generative image models and its applications,” in *International Conference on Learning Representations*, 2021.
- [14] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang, “The unreasonable effectiveness of deep features as a perceptual metric,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 586–595.
- [15] Andrew Brock, Jeff Donahue, and Karen Simonyan, “Large scale GAN training for high fidelity natural image synthesis,” in *International Conference on Learning Representations*, 2019.
- [16] 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.