

PARAMETER-EFFICIENT ADAPTATION FOR COMPUTATIONAL IMAGING

Nebiyou Yismaw¹, Ulugbek S. Kamilov², M. Salman Asif^{1,*}

¹University of California Riverside

²Washington University in St. Louis

ABSTRACT

Deep learning-based methods provide remarkable performance in a number of computational imaging problems. Examples include end-to-end trained networks that map measurements to unknown signals, plug-and-play (PnP) methods that use pretrained denoisers as image prior, and model-based unrolled networks that train artifact removal blocks. Many of these methods lack robustness and fail to generalize with distribution shifts in data, measurements, and noise. In this paper, we present a simple framework to perform domain adaptation as data and measurement distribution shifts. Our method learns a small number of factors to add in a pretrained model to bridge the gap in performance. We present a number of experiments on accelerated magnetic resonance imaging (MRI) reconstruction and image deblurring to demonstrate that our method requires a small amount of memory and parameter overhead to adapt to new domains.

1. INTRODUCTION

Computational imaging (CI) is an imaging paradigm that integrates optics and signal processing, wherein measurements and reconstruction algorithms are jointly designed to address numerous applications [1]. These applications include Magnetic Resonance Imaging (MRI) [2, 3], Computed Tomography (CT) [4, 5], and Phase Imaging [6], particularly in cases where conventional imaging is infeasible. In instances where conventional imaging remains possible but not optimal, computational imaging methods, such as super-resolution [7], deblurring [8], and denoising [9], are used to enhance reconstruction quality and reduce measurement costs. Our proposed method solves such imaging problems while demonstrating robustness to shifts in data and distribution.

The measurement process in CI problems is formulated as

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

where $\mathbf{y}$ represents measurements, $\mathbf{A}(\cdot)$ represents the forward operator, $\mathbf{x}$ represents the image to be reconstructed, and η represents additive measurement noise. The problem of reconstructing $\mathbf{x}$ from a set of imperfect and noisy measurements $\mathbf{y}$ falls under a class of problems known as *inverse problems*. Many inverse problems are ill-posed by nature and

*This paper is partially based on work supported by the NSF CAREER awards under grants CCF-2043134 and CCF-2046293.




Fig. 1: A reconstruction network trained on brain scans (a) shows poor performance on knee MRI images. Our proposed modulated network (b) learns domain specific rank-1 factors to perform modulated convolutions. This enables the pretrained network to successfully reconstruct knee scans.

require solving a regularized optimization problem to recover the original image. For instance, in the case of compressive sensing (CS), we seek to recover a high-dimensional signal from noisy and under-sampled measurements [10]. The optimization algorithm for such problems typically uses the forward model to enforce data consistency with the measurements and a regularizer to enforce a certain prior belief on the image [11]. Priors that promote sparsity in some representation space are commonly used as regularizers [12, 13].

Recently, deep learning-based end-to-end trained networks have emerged as effective solutions for various inverse problems. These methods are entirely data-driven and trained to directly map measurements to ground-truth images through paired training data. In this paper, we focus on end-to-end networks denoted as f , with parameters θ , that take an initial estimate $\bar{\mathbf{x}}$ and output final reconstruction $\hat{\mathbf{x}} = f_{\theta}(\bar{\mathbf{x}})$. The training objective for these networks aims to minimize the distance between the ground truth and the reconstructed images, as formulated by the equation

$$\min_{\theta} \frac{1}{N} \sum_{i=1}^N \|\mathbf{x}_i - \hat{\mathbf{x}}_i\|_2^2, \quad (2)$$

where $\hat{\mathbf{x}}_i$ represents reconstructed images and $\mathbf{x}_i$ represent ground truth images. These methods are proven effective in various reconstruction and restoration tasks, including CT [14], MRI [15], denoising [16], and super-resolution [17].

Despite their impressive performance, end-to-end methods are not robust to distribution shifts [18, 19, 20]. The study in [18] concluded that deep learning methods yield unstable results for image reconstruction. A small perturbation in the input can result in severe artifacts in the reconstructed output. Robustness in accelerated MR image reconstruction was studied in [19]. The study showed that end-to-end methods lack robustness when faced with distribution shifts. The findings from the paper also indicated that this lack of robustness is not unique to deep learning-based end-to-end methods; traditional sparsity-based methods are also affected by these shifts. Additionally, the work in [20] showed that training a single denoiser for all noise levels leads to suboptimal performance.

In this paper, we propose a simple and efficient framework that enables end-to-end reconstruction networks to adapt to new domains. Our method factorizes these networks into shared and domain-specific modules. Once the shared modules are trained on the source domain, they are fixed. Our method then learns target domain-specific modulations. We apply these modulations to the shared modules using element-wise products during inference. We demonstrate our methods' effectiveness on a number of accelerated MRI reconstruction and image restoration tasks.

In Figure 1 (a), we present a network $\mathbf{f}$ with parameters θ initially trained to reconstruct brain MRI scans. While this network performs well on brain scans, its performance is poor when tested on knee scans. Our proposed network, shown in Figure 1 (b), introduces domain-specific modulated convolution blocks. These blocks apply modulations specific to the knee domain. The number of parameters needed for the modulations is less than 0.5% of the parameters in the pretrained network. As illustrated in the figure, the modulated network has an improved performance on knee scans.

2. BACKGROUND AND RELATED WORK

2.1. Accelerated MRI reconstruction

MR scanners acquire measurements in the frequency domain known as k -space. These measurements are related to the scanned image via a Fourier transformation, $\mathbf{F}$. Most modern scanners employ acceleration and multiple receiver coils to speed up the acquisition process. The measurement from the i -th coil of a multi-coil MR scanner can be written as

$$\mathbf{y}_i = \mathbf{MFS}_i \mathbf{x} + \eta, \quad i = 1, \dots, K, \quad (3)$$

where $\mathbf{M}$ represents the under-sampling matrix, $\mathbf{F}$ is the 2D discrete Fourier transform, $\mathbf{S}_i$ represents the sensitivity map

of the i -th coil, η represents a complex measurement noise and $\mathbf{x}$ represents the underlying image to be scanned.

Convolutional Neural Network (CNN) based deep networks have recently been introduced for MR image reconstruction [21]. In this paper, we follow a U-Net [22]-like architecture to solve (2), which takes an initial estimate $\bar{\mathbf{x}}$ computed as:

$$\bar{\mathbf{x}} = \sqrt{\left(\sum_{c=1}^C |\mathbf{F}^{-1}(\mathbf{y}_i)|^2 \right)}, \quad (4)$$

where C represents the number of coils, $\mathbf{F}^{-1}$ represents the inverse Fourier transform, and $\mathbf{y}_i$ represents the under-sampled measurements.

2.2. Image restoration

Image restoration is the task of recovering clean images from noisy and corrupted measurements. Denoising, deblurring, and super-resolution are common examples of restoration problems. For denoising, the forward model can be expressed as $\mathbf{y} = \mathbf{x} + \eta$, where the forward operator is the identity, and η represents the measurement noise. In deblurring, the forward model can be modeled as convolution with a blur kernel. For a detailed review of deep learning methods for inverse problems, we refer the reader to [23].

3. METHODS

Our method adapts a pretrained network denoted as $\mathbf{f}$ with parameters θ to a target domain by applying a unique set of modulations denoted as M . These modulations are specific to each new domain and are applied to the convolution layers of the network. At the l -th convolution layer, the pretrained network has kernels represented as $\theta_l \in \mathbb{R}^{k \times k \times C_{in} \times C_{out}}$, where k represents the kernel size, and C_{in} and C_{out} represent the input and output channels, respectively. To adapt to new target domains, we learn a set of modulations for each convolutional layer and apply them to the pretrained weights using element-wise product, denoted as $\theta_l \odot M_l$.

Inspired by [24], we aim to find a parameter-efficient approximation for our modulation, M_l . We decompose M_l into four rank-1 factors: $M_l^1 \in \mathbb{R}^k, M_l^2 \in \mathbb{R}^k, M_l^3 \in \mathbb{R}^{C_{in}}, M_l^4 \in \mathbb{R}^{C_{out}}$. Our proposed modulation technique is described as follows. Prior to performing convolution, we modulate the i -th channel of the input I with C_{in} features as

$$\tilde{I}(:, :, i) = I(:, :, i) \times M_l^3(i), \quad (5)$$

where $\times$ represents ordinary multiplication. We then perform convolution on the i -th modulated input using a modulated kernel $\tilde{\theta}_l$, where $\tilde{\theta}_l(:, :, i, j) = \theta_l(:, :, i, j) \odot (M_l^1 \otimes M_l^2)$, and $\otimes$ represents the outer product. Finally, we apply post-convolution modulation using M_l^4 , similar to (5). In summary, our modulated convolution operation to obtain the j -th

channel of the output feature map O can be expressed as

$$O(:, :, j) = M_l^A(j) \left[\sum_{i=1}^{C_{in}} \tilde{\theta}_l(:, :, i, j) * \tilde{I}(:, :, i) \right], \quad (6)$$

where $*$ represents 2D convolution, $\tilde{\theta}$ represents the modulated kernel, $\tilde{I}$ represents modulated inputs, and M_l^A represent the output modulation.

Once the base network parameters θ are trained on the source domain, they are kept frozen. To adapt to target domains, we only optimize over the rank-1 modulations as

$$\min_M \frac{1}{N} \sum_i \|\mathbf{x}_i - \mathbf{f}(\bar{\mathbf{x}}; \theta, M)\|_2^2, \quad (7)$$

where $\bar{\mathbf{x}}$ represents the initial estimate, and $\mathbf{x}_i$ represents the ground truth image. This approach significantly reduces the number of additional trainable parameters. The number of parameters needed to modulate the l -th convolution layer is $2K + C_{in} + C_{out}$, which is notably smaller than the base network weights at that layer, which amount to $K^2 C_{in} C_{out}$. Moreover, such parameterization restricts the optimization search space, which has been shown to be effective in [25].

4. EXPERIMENTS AND RESULTS

We conducted a series of image reconstruction and restoration experiments to demonstrate the effectiveness of our method. Our first set of experiments involve domain adaptation for MR image reconstruction, where we adapt networks trained on brain MR scans to knee scans and vice versa. We then investigate adaptation related to forward model shifts, focusing on adapting to changes in sampling patterns and sampling ratios. Additionally, we conducted experiments on image blurring problems, where we adapt to shifts in the degradation model. In the following sub-sections, we provide a detailed experimental procedure as well as results.

4.1. Domain shift

We utilized the NYU fastMRI Brain and Knee datasets [21] for our domain adaptation experiments. To simulate the under-sampling operator $\mathbf{M}$, we use the sub-sampling library from the fastMRI package [21]. We use Cartesian sampling that selects 5% vertical lines in the center, followed by $4\times$ accelerated sampling at random. Our reconstruction network is an 8-layer U-Net [22] model with 64 channels. We independently trained two models on knee and brain scans. The performance of these networks, when tested on both knee and brain scans, is reported in the second and third columns of Table 1. We observed a performance drop when using the network trained on knee images to reconstruct brain scans and vice versa. To bridge this performance gap, we applied our modulation technique and reported performance of our adapted networks in the last two columns of Table 1. We successfully adapted the network trained on knee to brain scans, resulting in the network with the best average performance. Similarly, in the last column, we report the results of the

brain network after adapting it to knee scans. These networks maintain their performance when tested on the source domain and apply the learned modulation when performing inference on the target domains.

Table 1: Networks trained on **Knee** and **Brain** domains did not maintain their performance when tested on unseen domains, as shown in columns two and three. Our proposed method effectively adapted these networks to each target domain, as shown in the last two columns. (Colors indicate source domains, and bold text highlights the top-performing network).

Train \ Test	Knee	Brain	Knee→Brain	Brain→Knee
	PSNR, SSIM	PSNR, SSIM	PSNR, SSIM	PSNR, SSIM
Knee	32.47, 0.83	28.25, 0.73	32.47, 0.83	31.57, 0.80
Brain	29.93, 0.87	32.45, 0.90	32.11, 0.90	32.45, 0.90
Avg	31.20, 0.85	30.35, 0.82	32.29, 0.87	32.01, 0.85

4.2. Forward model shifts

We conduct a set of experiments to demonstrate the effectiveness of our proposed adaptation method under shifts in the forward operator, $\mathbf{A}$.

4.2.1. MRI Reconstruction

Sampling ratio shift. We now focus on the effects of sampling ratio shifts in accelerated MRI reconstruction. Similar to the previous experiments, we independently train networks on $4\times$ and $2\times$ sampling ratios and test these networks on both ratios. These results are reported in the second and third columns of Table 2. We observe a performance drop when there is a mismatch between the source and target sampling ratios, especially when testing a network trained on $4\times$ and testing it on $2\times$ accelerated measurements. To close this performance gap, we applied our modulation technique to both networks, allowing them to adapt to the new target domains. We report these results in the last two columns, where we demonstrate that our adapted networks achieve the best average performance.

Table 2: We show the performance of MRI networks trained on $2\times$ and $4\times$ acceleration factors during sampling ratio shift. Our adapted network shows an overall superior performance across all sampling ratios.

Train \ Test	2x	4x	2x→4x	4x→2x
	PSNR, SSIM	PSNR, SSIM	PSNR, SSIM	PSNR, SSIM
2x	35.36, 0.90	31.33, 0.86	35.36, 0.90	34.92, 0.89
4x	29.68, 0.78	32.47, 0.83	32.35, 0.82	32.47, 0.83
Avg	32.52, 0.84	31.90, 0.85	33.86, 0.86	33.70, 0.86

Sampling pattern shift. MR image scanners can employ various types of sampling patterns in the Fourier domain, such as Cartesian and radial sampling. In this experiment, we apply our adaptation technique to address the performance gap caused by shifts in the sampling pattern. We observe a performance drop when a network trained with Cartesian sampling is tested on radially sampled Fourier measurements, and vice versa. These results are reported in Table 3 under the second and third columns. In both cases, we observed a performance drop of approximately 5 dB. Using our proposed method, we successfully bridged the performance gap and achieved superior performance using our adapted networks.

Table 3: We trained networks on **Cartesian** and **Radial** patterns separately and tested them on both patterns. Our adapted networks, shown in the last two columns, mitigated the performance drop caused by the pattern shift.

Test \ Train	Cartesian		Radial		Cart→Rad		Rad→Cart	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
Cartesian	32.47	0.83	27.00	0.72	32.47	0.83	31.20	0.80
Radial	27.88	0.78	35.83	0.86	35.48	0.86	35.83	0.86
Avg	30.18	0.81	31.42	0.79	33.98	0.85	33.52	0.83

4.2.2. Image restoration

Degradation model shifts. We now consider image restoration problems, specifically, *image deblurring* problems, where we adapt to shifts in degradation models. We model the degradation as convolution with blur kernels and analyze the effects of shifts in these kernels.

Image blurs can be caused by several factors and can be modeled using different blur kernels. We consider motion, disk, and Gaussian blur kernels adapted from [26]. We set the kernel sizes to 35×35 for motion and Gaussian blurs and a radius of 5 for the disk blur. Our objective is to adapt a deblurring network trained on a specific kernel to unseen kernels. Our training dataset consists of images from the DIV2K [27], Flickr2K [28], and Waterloo exploration [29] datasets. We cropped and resized our training images to the size 256×256 . For testing, we used the widely-used BSD68 [30] dataset.

We used the network architecture proposed in [31], that combines U-Net [22] and ResNet [32]. We trained three networks on motion, disk, and Gaussian blur kernels. These networks learn a direct mapping from blurred to sharp images. As presented in columns 2 to 4 of Table 4 and Figure 2, these networks perform poorly when tested on images blurred with unseen kernels during training. To overcome this performance degradation, we applied our modulation technique on the network trained on motion blur to adapt to Gaussian and disk blurred images. We report the performance of our proposed method in the final column of Table 4 and showcase sample visual results in the last column of Figure 2. Our adapted network significantly outperforms the blur-specific network when tested on unseen kernels.

Table 4: The performance of deblurring networks trained on a specific blur type degrades severely when tested on unseen blur types. Our Adapted Network, which is pretrained on motion blur and applies modulations for the other blurs, achieves the best average performance.

Test \ Train	Motion		Gaussian		Disk		Adapted	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
Motion	43.99	0.99	8.41	0.02	14.03	0.21	43.99	0.99
Gaussian	19.93	0.45	43.73	0.99	15.66	0.23	39.77	0.99
Disk	21.54	0.48	8.06	0.01	35.20	0.95	29.37	0.87
Avg	28.39	0.64	20.07	0.34	21.63	0.46	37.71	0.95

4.3. Comparison experiments

We compared our proposed method with LoRA and reported the results in Table 5. We tested knee and brain scans for domain shifts, and Cartesian and radial sampling for forward model shifts. Our method outperformed LoRA in all exper-




Fig. 2: Image deblurring from three different blur kernels using blur-specific (columns 2 to 4) and adapted (last column) networks.

iments with a PSNR gains ranging from 0.3 to 1 dB (except for radial to Cartesian adaptation, where LoRA’s PSNR was 0.03 dB better).

Table 5: Distribution shift Comparison experiments. Our method outperforms LoRA in all adaptations, except for the Radial to Cartesian adaptation.

Adaptation \ Method	LoRA		Our method	
	PSNR	SSIM	PSNR	SSIM
Knee → Brain	31.28	0.89	32.11	0.90
Brain → Knee	31.31	0.79	31.57	0.80
Radial → Cartesian	31.23	0.80	31.20	0.80
Cartesian → Radial	34.63	0.85	35.48	0.86

4.4. Number of parameters

Our proposed method is parameter-efficient as it requires significantly fewer parameters compared to the base network for adaptation to new domains. The U-Net model used in the MRI reconstruction experiments has 31 million parameters. To adapt it to D new domains using full model tuning, we would need $D \times 31\text{M}$ parameters. In contrast, our method only requires 11.9 thousand modulations to adapt to new domains, totaling $D \times 11.9\text{K}$ parameters for adaptation to D new domains. Similarly, the network used for image restoration has 1.3 million parameters, while our modulations only require 4.7 thousand parameters per domain. These results show that our method is remarkably parameter efficient.

5. CONCLUSION

We explored the effects of domain shifts on deep learning-based end-to-end reconstruction networks for solving inverse problems. We observed that these shifts often lead to a severe performance degradation. To overcome this, we proposed a flexible and parameter-efficient domain adaptation method. We demonstrated the effectiveness of our method on a variety of image reconstruction and restoration problems, encompassing different network architectures. Our adapted networks achieve competitive performance compared to single-domain fully trained networks while requiring only a fraction of the number of parameters.

6. REFERENCES

- [1] Joseph N Mait, Gary W Euliss, and Ravindra A Athale, "Computational imaging," *Advances in Optics and Photonics*, vol. 10, no. 2, pp. 409–483, 2018.
- [2] Mark A Griswold, Martin Blaimer, Felix Breuer, Robin M Heidemann, Matthias Mueller, and Peter M Jakob, "Parallel magnetic resonance imaging using the grappa operator formalism," *Magnetic resonance in medicine*, vol. 54, no. 6, pp. 1553–1556, 2005.
- [3] Michael Lustig, David Donoho, and John M Pauly, "Sparse mri: The application of compressed sensing for rapid mr imaging," *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007.
- [4] Frank Natterer, "The mathematics of computerized tomography (classics in applied mathematics, vol. 32)," *Inverse Problems*, vol. 18, pp. 283–284, 2001.
- [5] Emil Y Sidky and Xiaochuan Pan, "Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization," *Physics in Medicine & Biology*, vol. 53, no. 17, pp. 4777, 2008.
- [6] Mustafa Mir, Basanta Bhaduri, Ru Wang, Ruoyu Zhu, and Gabriel Popescu, "Quantitative phase imaging," *Progress in optics*, vol. 57, no. 133-37, pp. 217, 2012.
- [7] Jia-Bin Huang, Abhishek Singh, and Narendra Ahuja, "Single image super-resolution from transformed self-exemplars," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 5197–5206.
- [8] Amir Beck and Marc Teboulle, "Fast gradient-based algorithms for constrained total variation image denoising and deblurring problems," *IEEE transactions on image processing*, vol. 18, no. 11, pp. 2419–2434, 2009.
- [9] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang, "Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising," *IEEE Transactions on Image Processing*, vol. 26, no. 7, pp. 3142–3155, 2017.
- [10] David L Donoho, "Compressed sensing," *IEEE Transactions on information theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [11] Jennifer L Mueller and Samuli Siltanen, *Linear and nonlinear inverse problems with practical applications*, SIAM, 2012.
- [12] Emmanuel J Candes, Justin K Romberg, and Terence Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, vol. 59, no. 8, pp. 1207–1223, 2006.
- [13] Michael Elad and Michal Aharon, "Image denoising via sparse and redundant representations over learned dictionaries," *IEEE Transactions on Image processing*, vol. 15, no. 12, pp. 3736–3745, 2006.
- [14] Kyong Hwan Jin, Michael T McCann, Emmanuel Froustey, and Michael Unser, "Deep convolutional neural network for inverse problems in imaging," *IEEE transactions on image processing*, vol. 26, no. 9, pp. 4509–4522, 2017.
- [15] Anuroop Sriram, Jure Zbontar, Tullie Murrell, Aaron Defazio, C Lawrence Zitnick, Nafissa Yakubova, Florian Knoll, and Patricia Johnson, "End-to-end variational networks for accelerated mri reconstruction," in *Medical Image Computing and Computer Assisted Intervention—MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part II* 23. Springer, 2020, pp. 64–73.
- [16] Bo Zhu, Jeremiah Z Liu, Stephen F Cauley, Bruce R Rosen, and Matthew S Rosen, "Image reconstruction by domain-transform manifold learning," *Nature*, vol. 555, no. 7697, pp. 487–492, 2018.
- [17] 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 conference on computer vision and pattern recognition*, 2020, pp. 2437–2445.
- [18] Vegard Antun, Francesco Renna, Clarice Poon, Ben Adcock, and Anders C Hansen, "On instabilities of deep learning in image reconstruction and the potential costs of ai," *Proceedings of the National Academy of Sciences*, vol. 117, no. 48, pp. 30088–30095, 2020.
- [19] Mohammad Zalbagi Darestani, Akshay S Chaudhari, and Reinhard Heckel, "Measuring robustness in deep learning based compressive sensing," in *International Conference on Machine Learning*. PMLR, 2021, pp. 2433–2444.
- [20] Abhiram Gnanasambandam and Stanley Chan, "One size fits all: Can we train one denoiser for all noise levels?," in *International Conference on Machine Learning*. PMLR, 2020, pp. 3576–3586.
- [21] Jure Zbontar, Florian Knoll, Anuroop Sriram, Tullie Murrell, Zhengnan Huang, Matthew J Muckley, Aaron Defazio, Ruben Stern, Patricia Johnson, Mary Bruno, et al., "fastmri: An open dataset and benchmarks for accelerated mri," *arXiv preprint arXiv:1811.08839*, 2018.
- [22] Olaf Ronneberger, Philipp Fischer, and Thomas Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III* 18. Springer, 2015, pp. 234–241.
- [23] Gregory Ongie, Ajil Jalal, Christopher A Metzler, Richard G Baraniuk, Alexandros G Dimakis, and Rebecca Willett, "Deep learning techniques for inverse problems in imaging," *IEEE Journal on Selected Areas in Information Theory*, vol. 1, no. 1, pp. 39–56, 2020.
- [24] Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al., "Lora: Low-rank adaptation of large language models," in *International Conference on Learning Representations*, 2021.
- [25] Chunyuan Li, Heerad Farkhoor, Rosanne Liu, and Jason Yosinski, "Measuring the intrinsic dimension of objective landscapes," in *International Conference on Learning Representations*, 2018.
- [26] Kai Zhang, Wangmeng Zuo, and Lei Zhang, "Deep plug-and-play super-resolution for arbitrary blur kernels," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 1671–1681.
- [27] Eirikur Agustsson and Radu Timofte, "Ntire 2017 challenge on single image super-resolution: Dataset and study," in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2017, pp. 126–135.
- [28] Radu Timofte, Eirikur Agustsson, Luc Van Gool, Ming-Hsuan Yang, and Lei Zhang, "Ntire 2017 challenge on single image super-resolution: Methods and results," in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2017, pp. 114–125.
- [29] Kede Ma, Zhengfang Duanmu, Qingbo Wu, Zhou Wang, Hongwei Yong, Hongliang Li, and Lei Zhang, "Waterloo exploration database: New challenges for image quality assessment models," *IEEE Transactions on Image Processing*, vol. 26, no. 2, pp. 1004–1016, 2016.
- [30] David Martin, Charles Fowlkes, Doron Tal, and Jitendra Malik, "A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics," in *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*. IEEE, 2001, vol. 2, pp. 416–423.
- [31] Kai Zhang, Yawei Li, Wangmeng Zuo, Lei Zhang, Luc Van Gool, and Radu Timofte, "Plug-and-play image restoration with deep denoiser prior," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 10, pp. 6360–6376, 2021.
- [32] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.