

Equivariant plug-and-play image reconstruction

Matthieu Terris, Thomas Moreau,
 Université Paris-Saclay, Inria, CEA
 Palaiseau, 91120, France

Nelly Pustelnik, Julian Tachella
 ENSL, CNRS UMR 5672
 Lyon, 69342, France

Abstract

Plug-and-play algorithms constitute a popular framework for solving inverse imaging problems that rely on the implicit definition of an image prior via a denoiser. These algorithms can leverage powerful pre-trained denoisers to solve a wide range of imaging tasks, circumventing the necessity to train models on a per-task basis. Unfortunately, plug-and-play methods often show unstable behaviors, hampering their promise of versatility and leading to suboptimal quality of reconstructed images. In this work, we show that enforcing equivariance to certain groups of transformations (rotations, reflections, and/or translations) on the denoiser strongly improves the stability of the algorithm as well as its reconstruction quality. We provide a theoretical analysis that illustrates the role of equivariance on better performance and stability. We present a simple algorithm that enforces equivariance on any existing denoiser by simply applying a random transformation to the input of the denoiser and the inverse transformation to the output at each iteration of the algorithm. Experiments on multiple imaging modalities and denoising networks show that the equivariant plug-and-play algorithm improves both the reconstruction performance and the stability compared to their non-equivariant counterparts.

1. Introduction

Linear inverse imaging problems are ubiquitous in imaging sciences, famous instances of which include image restoration, magnetic resonance imaging (MRI), computed tomography, and astronomical imaging to name a few. In this setting, the aim is to recover an image $x \in \mathbb{R}^n$ from measurements $y \in \mathbb{R}^m$ acquired through

$$y = Ax + \epsilon, \quad (1)$$

where $A: \mathbb{R}^n \rightarrow \mathbb{R}^m$ is a linear operator and ϵ the realisation of some random noise. A myriad of methods for solving problems of the likes of (1) have been proposed in the literature, ranging from variational methods [6, 12] (solution of a cost function minimisation problem) to end-to-end

reconstruction with deep neural networks [19, 22, 52] and, more recently, diffusion algorithms [9, 14, 21, 54].

In this work, we focus on approaches relying on implicit denoising priors. For instance, plug-and-play (PnP) algorithms [20, 35, 45, 50] propose to replace the proximity operator involved in variational methods (solved by proximal algorithms) by a denoiser modeling an implicit image prior. Similarly, regularization-by-denoising (RED) approaches propose to replace gradient steps on the prior term by a denoiser [33]. While the denoiser is typically trained as a Gaussian denoiser using grayscale or color natural images, it can be plugged into algorithms designed to tackle a wide range of image-related problems, without being constrained by the nature of the input images (e.g. MRI images, CT scans, etc) [1, 26, 39, 43].

Yet, despite successful performance in various applications, these algorithms tend to suffer from instability issues, making it difficult to transfer across imaging tasks or to derive statistical estimates from the reconstruction. In particular, RED and PnP algorithms often require careful fine-tuning [27, 32], or departing from the original optimization algorithm for efficient application [10, 50]. Thus, several works have been proposed to restore the convergence of PnP and RED algorithms while establishing a clear connection with an associated cost function [10, 17, 30, 35, 42]. However, these efforts introduce notable constraints on the denoiser, leading to a trade-off between stability and reconstruction performance. Typically, convergent PnP and RED algorithms tend to transfer easily to new problems, but perform less well than their fine-tuned and early-stopped non-convergent counterparts.

In this work, we propose to investigate theoretically and empirically the effect of equivariance on algorithms relying on implicit denoising priors. More precisely, we prove that enforcing equivariance on the denoiser improves the stability of the resulting PnP algorithm and plays a symmetrizing role on the Jacobian of the implicit prior. Our experiments show that the proposed equivariant approach can significantly improve the quality of images reconstructed with PnP algorithms as well as with popular algorithmic frameworks such as RED or Langevin algorithms. We give an

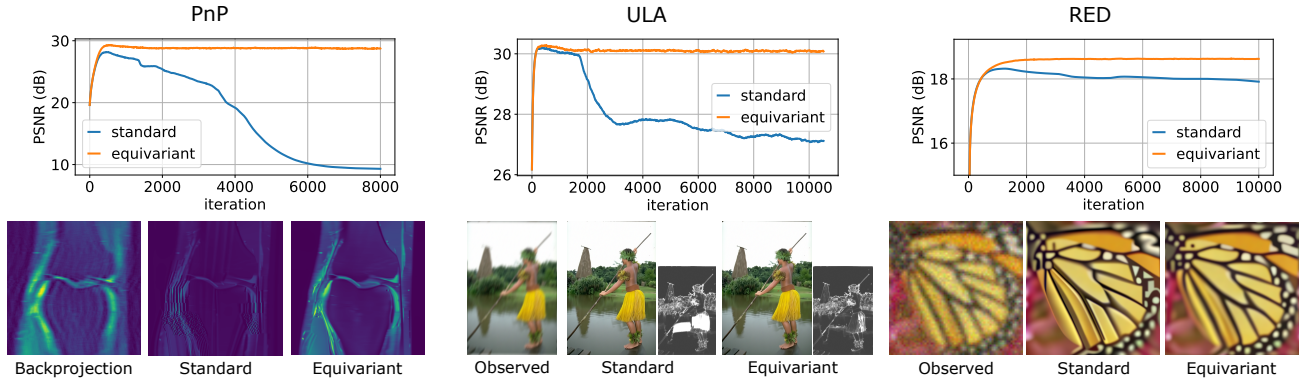


Figure 1. **Instability of algorithms relying on implicit denoising priors can be solved by incorporating equivariance.** Enforcing approximate equivariance of the denoiser at test time allows to both stabilize the algorithm and to improve the reconstruction quality without needing to retrain the implicit prior. **Left:** PnP algorithm applied to an accelerated MRI problem. **Middle:** Unadjusted Langevin sampling algorithm for a motion blur problem; estimated mean and variance of the associated Markov chain are displayed. **Right:** RED algorithm on a $4\times$ super-resolution problem.

overview of the possibilities offered by the proposed approach on popular algorithms relying on implicit denoising priors in Figure 1.

2. Related Work

Stable plug-and-play algorithms A significant amount of work has recently been proposed to provide stable PnP algorithms. A popular line of research in that direction consists in constraining the denoiser’s architecture to ensure its stability when plugged into a PnP scheme. Some approaches propose to regularize the denoiser’s training loss with a term penalizing the Lipschitz constant of the denoiser [30, 35], which can be combined with architectural constraints [17, 42]. Other works propose to alter the optimization algorithm itself in order to ensure its stability [10, 18, 27, 50]. In turn, ensuring convergence of the algorithm allows to ensure better transferability to new imaging tasks [43], but also to perform iteration-intensive tasks, such as sampling from the posterior distribution [23]. However, all the aforementioned methods come at the cost of either algorithmic modifications or strong constraints on the design of the denoiser, as opposed to the original, denoiser agnostic approaches [33, 45].

Equivariance in imaging inverse problems Equivariance to certain transformations, such as rotations or translations, has often been a desired property when designing handcrafted variational priors [13, 36, 37]. Geometric ensembling techniques have been known in the computer vision literature, where it has been shown to improve reconstruction quality for image super-resolution [44]. In the case of linear denoisers, symmetrization of the Jacobian has been shown to improve performance of linear denoisers [28]. Yet, to the notable exception of [43, 50], this strategy went unnoticed in the PnP literature, and to the best of our knowl-

edge, the role of equivariance has not been explored in the context of algorithms relying on implicit denoising priors. In a different line of work, but still in the context of imaging inverse problems, recent works have exploited equivariance for the design of unrolled network architectures [5], or the construction of self-supervised learning losses [7] (e.g., see the recent review [8]).

Equivariant neural networks More generally, invariance of the underlying prior appears as a natural assumption in a large number of applications, and there exists a vast literature on building equivariant neural network architectures [4, 11, 47]. Typical applications involve segmentation on spherical manifolds, robotics, point cloud analysis, data augmentation to name a few [7, 15, 29, 34, 46, 53]. However, these networks often fail to perform as well as other state-of-the-art architectures. Moreover, it can be challenging to incorporate complex layers (e.g., upsampling/downsampling, attention-like layers, etc.) without breaking the equivariance of the resulting network. In this work, we provide a simple method for rendering any denoiser equivariant, without any architectural constraints.

3. PnP algorithms

Traditional variational approaches for solving (1) consist in reformulating it as a minimization problem. Following a maximum-a-posteriori approach, one can derive an estimate \hat{x} as

$$\hat{x} = \underset{x}{\operatorname{argmin}} f(x) + \lambda r(x) \quad (2)$$

where f is a data-fidelity enforcing term, r is a regularization enforcing prior knowledge about the solution, and $\lambda > 0$ is a regularization parameter.

PnP approaches propose to replace the proximity operator of r [2] (implicit gradient step) arising in algorithms

for solving (2) by a denoiser D [45]. For the standard case of a quadratic data-fidelity term $f(x) = \frac{1}{2}\|Ax - y\|_2^2$, the classical PnP algorithm reads

$$x_{k+1} = D(x_k - \gamma A^\top(Ax_k - y)) \quad (\text{PnP})$$

where $\gamma > 0$ is a stepsize.

Similarly, regularization by denoising (RED) algorithms approximate the gradient of r as $\nabla r(x) \propto x - D(x)$ using Tweedie’s formula [33]. A simple explicit gradient descent optimization based on this definition yields

$$x_{k+1} = x_k - \gamma A^\top(Ax_k - y) - \gamma \lambda(x_k - D(x_k)). \quad (\text{RED})$$

We also consider the Unadjusted Langevin algorithm (ULA) [23], which aims to obtain samples associated with the negative log posterior density $-\log p(x|y) \propto f(x) + \lambda r(x)$, and requires adding noise to the iterates in (RED), i.e.,

$$x_{k+1} = x_k - \gamma A^\top(Ax_k - y) - \gamma \lambda(x_k - D(x_k)) + \sqrt{2\gamma}\epsilon_k, \quad (\text{ULA})$$

where $\epsilon_k \sim \mathcal{N}(0, I)$ is a standard Gaussian vector.

While there exist many different variants of (PnP), (RED), and (ULA), we here focus on their most standard formulations. Interestingly, these algorithms have shown impressive performance on a wide variety of imaging tasks while relying on Gaussian denoisers agnostic to the imaging modality of interest. Yet, these algorithms suffer from a lack of stability and potential divergence effects, hurting their versatility.

3.1. Proposed equivariant approach

Intuitively, imaging priors should have some invariance properties with respect to certain groups of transformations, such as rotations, translations, and reflections. We denote transformations associated with a group \mathcal{G} as $\{T_g\}_{g \in \mathcal{G}}$ where $T_g \in \mathbb{R}^{n \times n}$ is a unitary matrix¹. We say that D is equivariant to the group action $\{T_g\}_{g \in \mathcal{G}}$ if $D(T_g x) = T_g D(x)$ for all x and $g \in \mathcal{G}$. At the algorithmic level, this requirement translates into the equivariance of the denoiser with respect to the transforms of interest: if $r(x)$ is a \mathcal{G} -invariant function, its proximal operator and gradient (if they exist) are necessarily \mathcal{G} -equivariant functions [5].

A simple way of rendering any function \mathcal{G} -equivariant is by averaging over the group². The associated averaged denoiser

$$D_{\mathcal{G}}(x) \stackrel{\text{def}}{=} \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} T_g^{-1} D(T_g x). \quad (3)$$

¹While it is possible to define group actions with non-unitary matrices [38], here we focus on the unitary matrices, which is the case of translations, rotations and reflections of images.

²This construction is known as Reynolds averaging, see e.g. [38].

is equivariant by construction. For large groups, or in the case of large denoising architectures D , computing the averaged denoiser might be too computationally demanding. However, in this work, we propose to use a simple Monte Carlo approximation by sampling a single transformation at each step of the algorithm, i.e.,

$$g \sim \mathcal{G} \\ \tilde{D}_{\mathcal{G}}(x) = T_g^{-1} D(T_g x). \quad (4)$$

The equivariant counterpart of the (PnP), (RED) and (ULA) algorithms is thus simply obtained by replacing the denoiser D by a sample of the Monte Carlo estimate (4) at each iterate of the algorithm. Explicit versions of these algorithms can be found in the Supplementary Material (SM).

4. Theoretical analysis

In this section, we provide a theoretical analysis of the advantages in terms of performance and stability of equivariant denoisers compared to their non-equivariant counterparts. In this section, we denote the Jacobian of the denoiser as $J_x \stackrel{\text{def}}{=} \frac{\delta D}{\delta x}(x)$.

Optimality of equivariant denoisers We first show that if the signal distribution is \mathcal{G} -invariant, then for any denoiser D , its averaged version (3) obtains an equal or better denoising performance. This can be shown by computing the expected ℓ_2 error with respect to the signal and noise distributions, i.e.,

$$\begin{aligned} & \mathbb{E} \left\| \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} T_g^{-1} D(T_g(x + \varepsilon)) - x \right\| \\ & \leq \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} \mathbb{E} \|T_g^{-1} D(T_g(x + \varepsilon)) - x\| \\ & \leq \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} \mathbb{E} \|T_g^{-1} D(T_g T_g^{-1}(x + \varepsilon)) - T_g^{-1} x\| \\ & \leq \mathbb{E} \|D(x + \varepsilon) - x\| \end{aligned}$$

where the expectation is taken with respect to x and ε , and where the second line uses the triangle inequality and the third line uses that (i) $\mathbb{E}h(x) = \mathbb{E}h(T_g^{-1}x)$ for any function $h: \mathbb{R}^n \mapsto \mathbb{R}$ and T_g in the group action if the distribution of x and ε is \mathcal{G} -invariant, and that (ii) the transformations T_g are isometries.

Existence of an explicit prior A necessary condition for a denoiser to be associated with an explicit (PnP or RED) prior is to have a symmetric Jacobian, i.e., $J_x = J_x^\top$, see [31, Theorem 1]. Unfortunately, most state-of-the-art denoisers do not exhibit this property. Averaging a denoiser

over a sufficiently large group can lead to symmetric Jacobians. In particular, if the denoiser is linear and \mathcal{G} includes translations and reflections, then the denoiser is assured to have a symmetric Jacobian:

Proposition 1. *Any linear denoiser D that is equivariant to the action of 2-dimensional shifts, and vertical and horizontal reflections, has a symmetric Jacobian.*

Proof. Let $D_{\mathcal{G}}(x) = Mx$ with Jacobian $M \in \mathbb{R}^{n \times n}$. A matrix M that is equivariant to the action of 2-dimensional has a circulant form, i.e., $M = \text{circ}(d)$ where $d \in \mathbb{R}^n$ is a filter. Thus the transposed Jacobian is also a circulant matrix $M^{\top} = \text{circ}(d')$ where $d' \in \mathbb{R}^n$ is the transposed filter. Since D is also equivariant to vertical and horizontal reflections, we have that d is even and real, and thus $M = M^{\top}$. \square

While this result applies only to linear denoisers, the symmetry of non-linear denoisers also improves when incorporating equivariance: Table 1 shows the relative symmetry error $\|J_x - J_x^{\top}\|_F^2 / \|J_x\|_F^2$ of popular non-linear denoisers averaged over 10 different patches of 64×64 pixels. The \mathcal{G} -equivariant denoisers $D_{\mathcal{G}}$ have significantly smaller errors than their non-equivariant counterparts.

	Lipschitz DnCNN	DnCNN	SCUNet	SwinIR	DRUNet ($\sigma_d=0.01$)
Standard D	0.014	0.022	0.954	0.604	0.030
Equivariant $D_{\mathcal{G}}$	0.003	0.005	0.710	0.291	0.008

Table 1. Mean Jacobian Symmetry error $\|J_x - J_x^{\top}\|_F^2 / \|J_x\|_F^2$. Equivariant denoisers are obtained by averaging over the group of 90-degree rotations and reflections.

Lipschitz constant of the denoiser The stability of PnP algorithms depends crucially on the Lipschitz constant of the denoiser [20, 35]. For example, if the Lipschitz constant of the denoiser is lower than 1, both the (PnP) and (RED) iterates converge under a good choice of step size. Since the Lipschitz constant of the sum of two mappings is smaller or equal than the sum of their Lipschitz constants, we have that the Lipschitz constant of the averaged equivariant denoiser is necessarily equal or lower than the non-equivariant one. If we restrict to linear denoisers, we can show that the equivariant denoiser will have a strictly smaller constant, as long as the dominant singular vector is not equivariant:

Proposition 2. *Let $D(x) = Mx$ be a linear denoiser with singular value decomposition $M = \sum_{i=1}^n \lambda_i u_i v_i^{\top}$ and $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n \geq 0$. If the principal component $u_1 v_1^{\top}$ is not \mathcal{G} -equivariant, the averaged denoiser $D_{\mathcal{G}}$ has a strictly smaller Lipschitz constant than D .*

Proof. For any p matrices A_1, \dots, A_p , we have that $\|\frac{1}{p} \sum_{g=1}^p A_g\| = \frac{1}{p} \sum_{g=1}^p \|A_g\|$ if and only if all matrices share the same leading left and right singular vectors. The \mathcal{G} -averaged denoiser can be written as $D_{\mathcal{G}} =$

$\frac{1}{|\mathcal{G}|} \sum_{g=1}^{|\mathcal{G}|} A_g$ where $A_g := T_g M T_g^{-1}$. We have that A_g has the same singular values as M since singular vectors are defined as $u'_i = T_g u_i$ and $v'_i = T_g v_i$ for $i = 1, \dots, n$. Since $u_1 v_1^{\top}$ is not equivariant, we have that $T_g u_1 v_1^{\top} T_g^{-1} \neq u_1 v_1^{\top}$ for some $g \in \mathcal{G}$. Thus, there exist at least 2 terms in the sum $\sum_g T_g M T_g^{-1}$ which do not share the same leading singular vectors, and consequently $\|M\| > \|\frac{1}{|\mathcal{G}|} \sum_g T_g M T_g^{-1}\|$. \square

In practice, we observe a significantly smaller constant for most popular non-linear denoisers. Table 2 shows the local Lipschitz constant (i.e., the spectral norm of the Jacobian) of various denoisers averaged over 10 different patches of 64×64 pixels. The averaged denoiser can have a significantly smaller constant.

	Lipschitz DnCNN	DnCNN	SCUNet	SwinIR	DRUNet ($\sigma_d=0.01$)
Standard D	1.06	1.44	5.78	6.28	1.57
Equivariant $D_{\mathcal{G}}$	0.92	1.18	4.19	4.05	1.26

Table 2. Local Lipschitz constant of the denoiser averaged over 16 image patches. Equivariant denoisers are obtained by averaging over the group of 90 degree rotations and reflections.

Interplay between the group action and forward operator So far we have focused on the properties associated with an equivariant denoiser, however, the (lack of) equivariance of A also plays an important role in the convergence of PnP algorithms. The iterates in (PnP) can converge even for denoisers with Lipschitz constant larger than 1, as long as the Lipschitz constant of the composition

$$D \circ (I - \gamma A^{\top} A) \quad (5)$$

is smaller than 1. If the spectra of the matrices $A^{\top} A$ and D are incoherent, i.e., $A^{\top} A$ and the Jacobian of the denoiser are diagonalized in different bases, the Lipschitz constant of (5) is likely to be smaller than that of D . A similar stability phenomenon happens for the (RED) and (ULA) iterates due to the incoherence between the spectrum of the forward operator and the one of the equivariant denoiser. The following proposition formalizes this intuition in the case of linear denoisers:

Proposition 3. *Let x be a grayscale image and $\{T_g\}_{g \in \mathcal{G}}$ be a group of transformations that includes 2-dimensional shifts and $D_{\mathcal{G}}$ be a linear \mathcal{G} -equivariant denoiser. If $A^{\top} A$ is not \mathcal{G} -equivariant, it does not share the same singular vectors as the Jacobian of D .*

Proof. Let $D(x) = Mx$ with Jacobian $M \in \mathbb{R}^{n \times n}$. A matrix $B \in \mathbb{R}^{n \times n}$ is \mathcal{G} -equivariant to the action of 2-dimensional shifts if and only if it admits a diagonal decomposition as $B = F \text{diag}(d) F^*$ where F is the 2-dimensional Fourier transform, see e.g. [40, Section 4.1]. Thus, if D is

\mathcal{G} -equivariant, then M is diagonal in the Fourier domain, whereas the non-equivariant $A^\top A$ matrix does not admit such a diagonalization. \square

Table 3 shows the local Lipschitz constant of the mapping in (5) for various popular non-linear denoisers, where A is a random inpainting operator, which is not equivariant to shifts, rotations nor reflections. The constants are smaller than those shown in Table 2 and are below 1, ensuring contraction of the PnP iterates for that specific operator. The equivariant denoisers show smaller constants than the non-equivariant counterparts.

	Lipschitz DnCNN	DnCNN	SCUNet	SwinIR	DRUNet ($\sigma_d=0.01$)
Standard	0.91	0.91	0.62	0.69	0.83
Equivariant	0.79	0.78	0.52	0.67	0.70

Table 3. Local Lipschitz constant of PnP iteration $\|J_x(I - A^\top A)\|$, with A a random inpainting operator. Equivariant denoisers are obtained by averaging over the group of 90 rotations and reflections.

4.1. Non-linear example

We demonstrate some of the properties analysed in the previous subsection on a non-linear setting of a neural network denoiser with a single hidden layer. More precisely, we consider the case where D is a slight perturbation of a proximity operator, i.e.

$$D(x) = B_2 \operatorname{prox}_{\gamma\lambda\|\cdot\|_1}(B_1x) \quad (6)$$

where $B_1 \in \mathbb{R}^{n \times n}$ satisfies $B_1 B_1^\top = B_1^\top B_1 = I$ and where $B_2 = (B_1^\top + P)$ where $P \in \mathbb{R}^{n \times n}$ is a small random perturbation. In particular, if $P = 0$, then the denoiser is a well-defined proximal operator, i.e. $D(x) = \operatorname{prox}_{\gamma\lambda\|B_1\cdot\|_1}(x)$, with a well-defined prior $r(x) = \|B_1x\|_1$. We stress that, for $P \neq 0$, there exists a priori no loss function g associated with the (PnP) iterates with the denoiser in (6). Furthermore, we assume that B_1 is \mathcal{G} -equivariant.

Using the fact that B_1 and $\operatorname{prox}_{\gamma\lambda\|\cdot\|_1}$ are \mathcal{G} -equivariant functions, we can write the \mathcal{G} -equivariant denoiser as (see a detailed derivation in the SM)

$$D_{\mathcal{G}}(x) = (B_1^\top + P_{\mathcal{G}}) \operatorname{prox}_{\gamma\lambda\|\cdot\|_1}(B_1x)$$

where $P_{\mathcal{G}} = \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} T_g^{-1} P T_g$ is the \mathcal{G} -averaged perturbation. If the original perturbation is not equivariant we have $0 \leq \|P_{\mathcal{G}}\|_F^2 < \|P\|_F^2$, and the equivariant denoiser will be closer to the proximal operator $\operatorname{prox}_{\gamma\lambda\|B_1\cdot\|_1}(x)$.

Figure 2 illustrates the behaviour of the (PnP) sequence with and without the group averaging in (3) for two specific choices of A in a 2D toy example, where \mathcal{G} is the group of flips (see the SM for more details). In both cases, the algorithm involving the equivariant denoiser converges to a point close to the global minima associated with the




Figure 2. Behaviour of the (PnP) algorithm with an approximated proximity operator (blue curve) and its equivariant counterpart (red curve). Contour lines show the loss in (2) with $r(x) = \|B_1x\|_1$. Stars denote the limit point of each sequence (when it exists) and green dots show the initialization points.

prior $\lambda r(x) = \lambda \|B_1x\|_1$, whereas sequence generated by the non-equivariant algorithm diverges.

5. Experimental results

In this section, we evaluate the influence of the proposed equivariant approach for different algorithms and linear inverse imaging problems. Our implementation³ relies on the DeepInverse library [41].

5.1. Problems considered

Image deblurring and image super-resolution In this setting, we set $y = h * x + \epsilon$ in (1) where h is a convolutional kernel and $*$ the circular convolution. We consider either Gaussian deblurring, in which case h is a Gaussian kernel of standard deviation 1, and motion deblurring, in which case h is the first kernel from [24]. Unless specified otherwise, ϵ is a Gaussian noise with standard deviation 0.01. In the image super-resolution (SR) setting, (1) writes as $y = (h * x)_{\downarrow S} + \epsilon$ where h is a Gaussian kernel of standard deviation 1, and S denotes the undersampling factor. When $S = 2$ (resp. $S = 4$), ϵ is a Gaussian noise with standard deviation 0.01 (resp. 0.05). We test the proposed method on the Set3C dataset as well as the BSD10 dataset, a subset of 10 images from the BSD68 dataset.

MRI In this setting, we consider $y = MFx$ in (1) where M is a binary mask and F the 2D Fourier transform. Following [48], we consider the $\times 4$ and $\times 8$ acceleration factors. In contrast to the previous problem, no noise is added to the measurements in this case. We test the method on a subset of 10 images from the validation set of the fully acquired k-space data of [48].

Algorithms and backbone denoisers We consider several backbone pretrained denoiser, namely DRUNet [50], SCUNet [51], SwinIR [25], DiffUNet [16], DnCNN [49] as well as its 1-Lipschitz version (LipDnCNN) from [30] and the gradient-step denoiser GSNet [17]. These architectures are representative of state-of-the-art image reconstruction

³available at <https://github.com/matthieutrs/EquivariantPnP>.

	Deblur (motion)		Deblur (Gaussian)		MRI	
	Set3C	BSD10	Set3C	BSD10	$\times 4$	$\times 8$
LipDnCNN [30]	31.7 ± 2.0	30.9 ± 0.7	32.1 ± 2.5	32.6 ± 0.9	30.6 ± 2.6	26.4 ± 2.3
Eq. LipDnCNN	31.8 ± 2.0	31.1 ± 0.6	32.1 ± 2.5	32.7 ± 0.9	30.7 ± 2.7	26.6 ± 2.2
DnCNN	29.9 ± 1.9	30.4 ± 0.3	22.0 ± 6.0	29.6 ± 4.6	28.6 ± 4.3	div.
Eq. DnCNN	30.7 ± 1.6	30.9 ± 0.1	29.8 ± 3.1	33.0 ± 0.6	30.1 ± 3.9	24.6 ± 2.9
DRUNet	10.1 ± 1.5	16.5 ± 9.8	14.6 ± 8.8	17.5 ± 10.5	27.7 ± 4.0	16.9 ± 8.6
Eq. DRUNet	18.0 ± 6.0	28.4 ± 2.3	25.3 ± 10.1	31.9 ± 1.1	30.7 ± 3.0	22.0 ± 6.0
wavelets	29.6 ± 2.0	27.2 ± 1.0	31.2 ± 2.5	30.9 ± 0.8	28.6 ± 2.1	25.4 ± 1.8
TGV	29.2 ± 2.1	26.7 ± 1.3	30.7 ± 2.5	30.5 ± 0.9	28.5 ± 2.3	24.8 ± 2.1
GSPnP [17]	34.6 ± 0.2	33.5 ± 1.9	35.1 ± 0.7	31.4 ± 3.3	n/a	n/a
DPIR [50]	33.9 ± 1.9	35.0 ± 0.5	33.0 ± 2.5	34.6 ± 0.7	28.4 ± 2.2	25.1 ± 1.9

Table 4. Mean reconstruction PSNR for various image restoration problems using the (PnP) algorithm with different backbone denoisers. The first, third, and fifth rows present results for non-equivariant denoisers (LipDnCNN, DnCNN, and DRUNet, respectively), while their equivariant counterparts (Eq. LipDnCNN, Eq. DnCNN, and Eq. DRUNet) are shown in the second, fourth, and sixth rows. The bottom four rows offer benchmarks with standard reconstruction methods. The notation “div.” indicates cases where the method diverged.

architectures, involving both convolutional and attention layers. We stress that none of these networks is equivariant to translations or rotations, to the exception of DnCNN that show approximate translation equivariance. The choice of the backbone denoiser may influence the chosen algorithm. For instance, the DnCNN and SwinIR denoisers on which this article relies are trained for a fixed level of noise, limiting the ability to finetune the algorithm. Moreover, DiffUNet and SCUNet architectures can only be applied to color images, preventing their usage on the MRI problem. Unless specified, each algorithm is run for 10^4 iterations.

Baselines In this paper, we mainly investigate the influence of equivariant priors on the stability of PnP algorithms. We thus compare the proposed approach with variational approaches which can be seen as a class of convergent PnP algorithms. In particular, we use wavelet denoisers (i.e. $D(x) = \text{prox}_{\lambda \|\Psi \cdot\|_1}(x)$ for Ψ a redundant wavelet dictionary), as well as total generalized variation (TGV) denoisers [3]. The (PnP) algorithm with Lipschitz denoiser corresponds to the method from [30] which ensures convergence of the (PnP) algorithm. We also compare our approach with the state-of-the-art DPIR algorithm [50] which runs a small number of steps of a half-quadratic splitting algorithm with fine-tuned decaying stepsizes. As a consequence, DPIR can be seen as representative of non-convergent, fine-tuned PnP algorithm. We also compare our results with the gradient-step PnP algorithm [17], leveraging a nonconvex implicit prior and backtracking and that is representative of state-of-the-art convergent PnP algorithms.

5.2. Stability of PnP and RED algorithms

The instability of PnP algorithms often translates into unrealistic artifacts in the reconstructed image. Figure 3 illustrates this phenomenon and shows that the proposed approach allows us to circumvent this drawback.

Figure 4 shows the PSNR and convergence criterion $\|x_{k+1} - x_k\| / \|x_k\|$ on different problems and backbone ar-

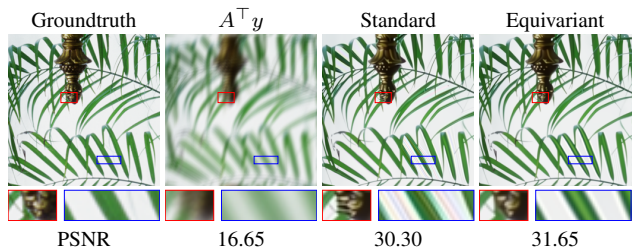


Figure 3. Motion deblurring on a Set3C sample with (PnP) relying on a DRUNet backbone denoiser.

chitecture. For each of these problems, the denoising level σ in (PnP) is set to $\sigma = 0.01$ (resp. $\sigma = 0.015$) for DnCNN (resp. DRUNet) backbone architectures. We notice that the (PnP) algorithm with equivariant denoiser shows a more stable PSNR along iterations than its non-equivariant counterpart. In particular, we observe in the right panel of Figure 4 that the convergence rate of the (PnP) algorithm with \mathcal{G} -equivariant DnCNN matches the behaviour of the algorithm with Lipschitz denoisers. Lastly, we notice that the proposed equivariant approach also benefits convergent (PnP) algorithms relying on Lipschitz backbone denoisers. We further stress that reconstructions obtained with the equivariant Monte-Carlo estimates are consistent with those obtained with deterministic Reynolds averaging (see Tab. 5).

In the case of the (RED) algorithm, we observe similar behaviours. Reconstructions on a $\times 2$ SR problem are shown in Figure 6 and associated convergence plots can be found in Figure 7. In the case of a non-equivariant DRUNet backbone, the reconstructed image shows important geometric artifacts that disappear in the equivariant case.

We however stress that the proposed approach may fail to stabilize the algorithm in certain settings. For example, in the case of a SCUNet backbone denoiser, strong artifacts are visible from 50 iterations only in the reconstruction in both the classical and equivariant version of the (PnP) algorithm as seen in Figure 5. Similarly, we did not observe convergence of the (RED) algorithm with the DiffUNet backbone,

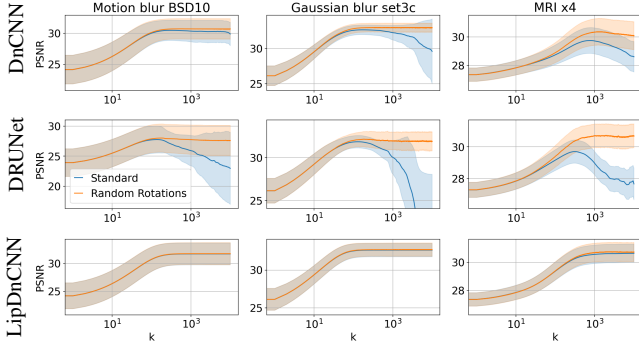


Figure 4. Average PSNR (left) and convergence criterion $\|x_{k+1} - x_k\| / \|x_k\|$ (right) for 3 different imaging problems as a function of the iteration number with different backbone denoisers plugged in the (PnP) algorithm. Top row: DnCNN, middle row: DRUNet, bottom row: 1-Lipschitz DnCNN.

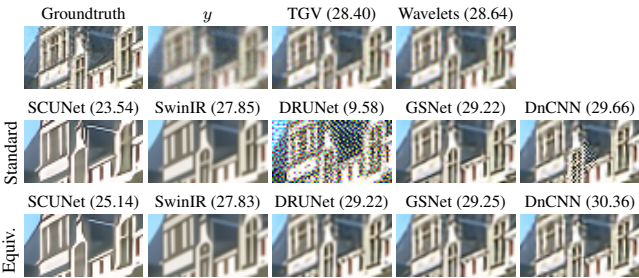


Figure 5. Gaussian deblurring with standard deviation $\sigma = 0.02$ on a BSD10 sample (detail) for different denoising backbone plugged in the (PnP) algorithm.

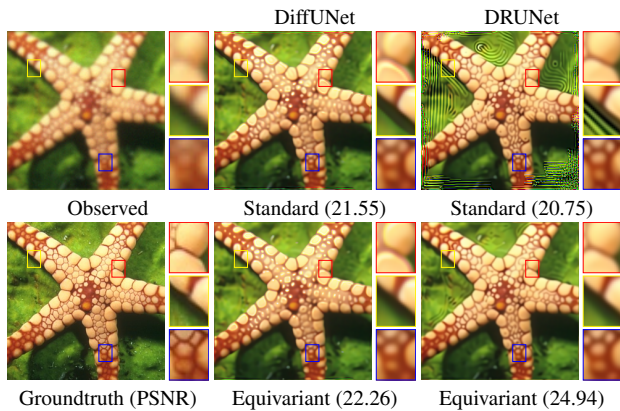


Figure 6. Results of the (RED) algorithm on a $\times 2$ SR problem for different backbone denoisers. Middle column: DiffUNet; right column: DRUNet. Top row: standard algorithm; Bottom row: equivariant algorithm.

thus requiring early stopping of the algorithm in order to reach good reconstruction results.

Lastly, several studies have shown that the noise level σ plays an important role in the stability of the PnP algorithm. Increasing σ may help to solve the instability issue at the cost of over-smoothed reconstructions.

5.3. Interplay with the kernel of A

In the previous section, we have seen that equivariant denoisers can prevent the emergence of artifacts along their-

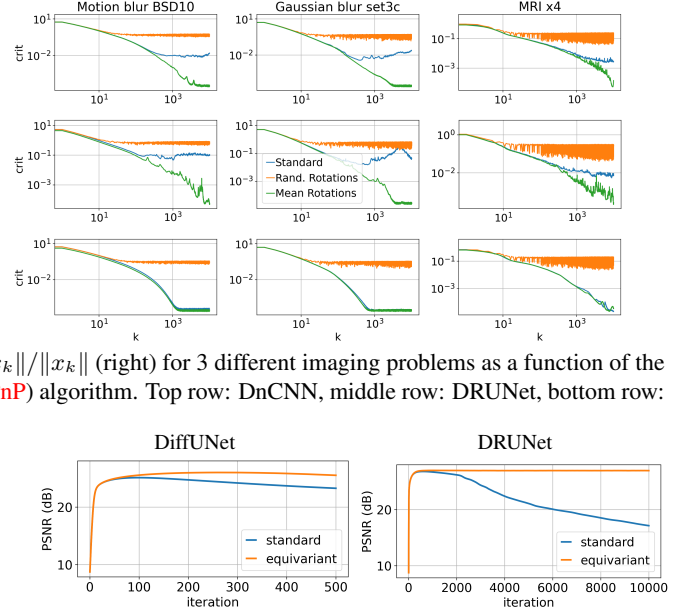


Figure 7. Evolution of the PSNR along the (RED) algorithm for different backbone denoisers associated to the reconstructions shown in Figure 6.

ations of (PnP). Despite their unnatural aspect, these artifacts are not incompatible with good data-fidelity measures. In fact, PnP algorithms offer no control over $\ker(A)$, which is nontrivial by nature of the ill-posed inverse problem (1); artifacts appearing during the reconstruction are therefore likely to belong to $\ker(A)$.

This phenomenon can be illustrated in the case of MRI, where $\ker(A)$ corresponds to the non-sampled subspace of the Fourier domain. In this case, A is not rotation equivariant and Proposition 3 suggests that enforcing equivariance of the prior improves the stability. This is illustrated in Figure 8 showing the reconstruction with and without the proposed equivariant algorithmic update. In the standard (non-equivariant) setting, mild artifacts appear between iteration $i = 10^3$ and $j = 10^4$. These may appear unnoticed in x_j ; however, they clearly appear when plotting the difference $x_j - x_i$. Interestingly, the Fourier spectrum of these artifacts shows significantly more energy at frequencies that were not sampled. The same experiment with the equivariant algorithm shows no such artifacts and a more uniform Fourier spectrum.

5.4. Sampling from RED prior

The gain in stability provided by the proposed approach opens the door to more robust sampling algorithms relying on implicit denoising priors, such as (ULA), where a sufficiently large number of iterations is required in order to obtain good estimators. We show in Figure 9 the estimated mean and variance obtained with the (ULA) algorithm, with both equivariant and non-equivariant DRUNet

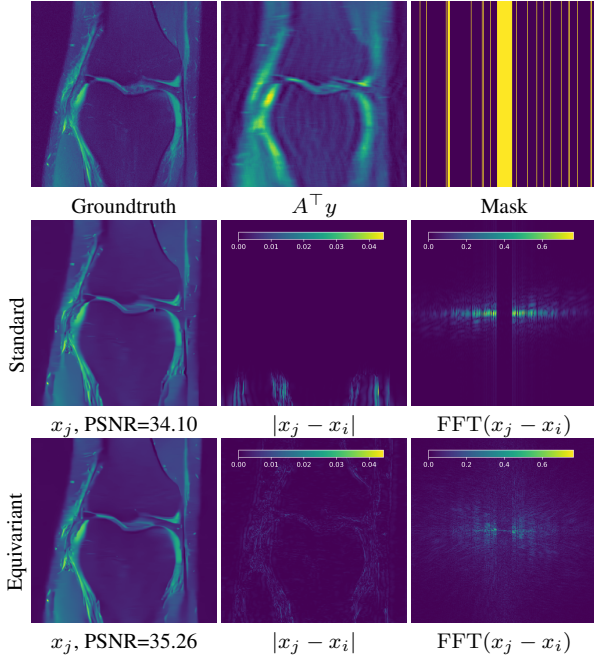


Figure 8. Evolution of the reconstruction along the iterates of the (PnP) with a (non 1-Lipschitz) DnCNN backbone. The top row shows the ground truth, back-projected data, and Fourier mask. The second (resp. third) row shows, from left to right: reconstruction at iteration $j = 10^4$; difference $|x_j - x_i|$ between reconstructions at iteration $j = 10^4$ and at iteration $i = 10^3$; the difference in the Fourier domain (displayed in logarithmic scale).

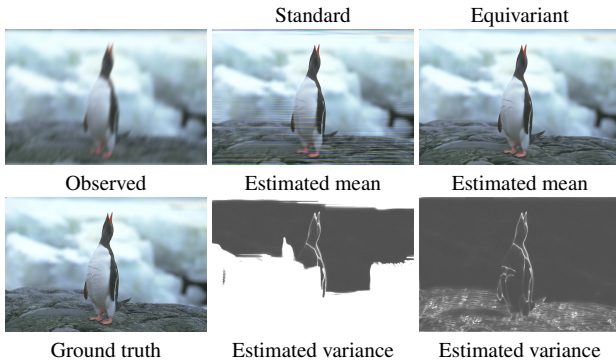


Figure 9. Sampling posterior mean and variance of a motion deblurring problem on a BSD10 sample for a DRUNet backbone plugged in the (ULA) algorithm.

backbone denoisers. In the non-equivariant setting, we observe similar artifacts to those obtained in the deterministic case on both the estimated mean and variance; these artifacts vanish in the equivariant case.

The leftmost plot of Figure 10 further illustrates this phenomenon. We notice that after a few hundred iterations, the reconstruction quality with the non-equivariant (ULA) algorithm collapses, yielding an irrelevant MCMC chain. While enforcing equivariance improves the situation, increasing the noise level in the denoiser in (PnP) can significantly improve the situation, see the rightmost plot of Figure 10.




Figure 10. Sampling with the (ULA) algorithm on a motion deblurring and the DRUNet prior, in the standard and equivariant cases. Left: PSNR of a sample from the MCMC chain for a sample of the Set3C dataset. Right: average PSNR over the Set3C dataset as a function of σ in the denoising prior.

6. Limitations

While our experiments show that the proposed method improves both the stability of the algorithmic scheme and the reconstructed image quality over its non-equivariant counterpart, we stress that it remains prone to divergence and/or hallucinating artifacts. This can for example be seen in Figure 4 in the case of the DRUNet backbone, which shows decreasing mean PSNR after a certain number of iterations, or in Figure 5 for the SCUNet backbone, which shows important reconstruction artifacts. In particular, the proposed equivariant approach does not clearly improve over its non-equivariant counterpart when using a SCUNet backbone.

7. Conclusion

In this work, we proposed a simple yet efficient method for ensuring approximate equivariance of implicit denoising priors. In essence, the method amounts to sampling and applying a group action at random at each step of the algorithm. In spite of its simplicity, this method shows interesting theoretical properties. In the case of linear denoisers, for example, it allows us to enforce the symmetry of the Jacobian, which is a cornerstone property when aiming to derive explicit priors from implicit denoising priors. Furthermore, it can be shown that the Lipschitz constant of the equivariant linear denoiser can only be lower than that of its non-equivariant counterpart, thus improving the stability of the resulting PnP algorithm. We showcase the symmetrization effect of equivariance on the Jacobian of several architectures, and illustrate its stabilization effect for three families of imaging algorithms, namely PnP, RED, and ULA. Importantly, this stabilization procedure is not detrimental to the reconstruction quality, as often observed in the literature. However, despite the significant improvements brought by equivariance in terms of stability and image reconstruction quality, the proposed method remains prone to divergence and artifact contamination in reconstructions.

Acknowledgements

This work was supported by the BrAIN grant (ANR-20-CHIA-0016) and was granted access to the HPC resources of IDRIS under the allocation 2023-AD011014344 made by GENCI. J. T. is supported by the ANR's JCJC grant UNLIP.

References

- [1] Rizwan Ahmad, Charles A Bouman, Gregory T Buzzard, Stanley Chan, Sizhuo Liu, Edward T Reehorst, and Philip Schniter. Plug-and-play methods for magnetic resonance imaging: Using denoisers for image recovery. *IEEE signal processing magazine*, 37(1):105–116, 2020. [1](#)
- [2] Heinz H Bauschke, Patrick L Combettes, Heinz H Bauschke, and Patrick L Combettes. *Correction to: Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer, 2017. [2](#)
- [3] Kristian Bredies, Karl Kunisch, and Thomas Pock. Total generalized variation. *SIAM Journal on Imaging Sciences*, 3(3):492–526, 2010. [6](#)
- [4] Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021. [2](#)
- [5] Elena Celledoni, Matthias J Ehrhardt, Christian Etmann, Brynjulf Owren, Carola-Bibiane Schönlieb, and Ferdia Sherry. Equivariant neural networks for inverse problems. *Inverse Problems*, 37(8):085006, 2021. [2](#), [3](#)
- [6] Antonin Chambolle and Thomas Pock. An introduction to continuous optimization for imaging. *Acta Numerica*, 25: 161–319, 2016. [1](#)
- [7] Dongdong Chen, Julián Tachella, and Mike E Davies. Equivariant imaging: Learning beyond the range space. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4379–4388, 2021. [2](#)
- [8] Dongdong Chen, Mike Davies, Matthias J Ehrhardt, Carola-Bibiane Schönlieb, Ferdia Sherry, and Julián Tachella. Imaging with equivariant deep learning: From unrolled network design to fully unsupervised learning. *IEEE Signal Processing Magazine*, 40(1):134–147, 2023. [2](#)
- [9] Hyungjin Chung, Jeongsol Kim, Michael T Mccann, Marc L Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. *arXiv preprint arXiv:2209.14687*, 2022. [1](#)
- [10] Regev Cohen, Michael Elad, and Peyman Milanfar. Regularization by denoising via fixed-point projection (red-pro). *SIAM Journal on Imaging Sciences*, 14(3):1374–1406, 2021. [1](#), [2](#)
- [11] Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International conference on machine learning*, pages 2990–2999. PMLR, 2016. [2](#)
- [12] Patrick L Combettes, Laurent Condat, J-C Pesquet, and BC Vũ. A forward-backward view of some primal-dual optimization methods in image recovery. In *2014 IEEE International Conference on Image Processing (ICIP)*, pages 4141–4145. IEEE, 2014. [1](#)
- [13] Laurent Condat. Discrete total variation: New definition and minimization. *SIAM Journal on Imaging Sciences*, 10(3): 1258–1290, 2017. [2](#)
- [14] Valentin De Bortoli, James Thornton, Jeremy Heng, and Arnaud Doucet. Diffusion schrödinger bridge with applications to score-based generative modeling. *Advances in Neural Information Processing Systems*, 34:17695–17709, 2021. [1](#)
- [15] Fabian Fuchs, Daniel Worrall, Volker Fischer, and Max Welling. Se (3)-transformers: 3d roto-translation equivariant attention networks. *Advances in neural information processing systems*, 33:1970–1981, 2020. [2](#)
- [16] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. [5](#)
- [17] Samuel Hurault, Arthur Leclaire, and Nicolas Papadakis. Gradient step denoiser for convergent plug-and-play. In *International Conference on Learning Representations*, 2021. [1](#), [2](#), [5](#), [6](#)
- [18] Samuel Hurault, Antonin Chambolle, Arthur Leclaire, and Nicolas Papadakis. Convergent plug-and-play with proximal denoiser and unconstrained regularization parameter. *arXiv preprint arXiv:2311.01216*, 2023. [2](#)
- [19] 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*, 26(9):4509–4522, 2017. [1](#)
- [20] Ulugbek S Kamilov, Charles A Bouman, Gregory T Buzzard, and Brendt Wohlberg. Plug-and-play methods for integrating physical and learned models in computational imaging: Theory, algorithms, and applications. *IEEE Signal Processing Magazine*, 40(1):85–97, 2023. [1](#), [4](#)
- [21] Bahjat Kawar, Michael Elad, Stefano Ermon, and Jiaming Song. Denoising diffusion restoration models. *Advances in Neural Information Processing Systems*, 35:23593–23606, 2022. [1](#)
- [22] Florian Knoll, Jure Zbontar, Anuroop Sriram, Matthew J Muckley, Mary Bruno, Aaron Defazio, Marc Parente, Krzysztof J Geras, Joe Katsnelson, Hersh Chandarana, et al. fastmri: A publicly available raw k-space and dicom dataset of knee images for accelerated mr image reconstruction using machine learning. *Radiology: Artificial Intelligence*, 2(1):e190007, 2020. [1](#)
- [23] Rémi Laumont, Valentin De Bortoli, Andrés Almansa, Julie Delon, Alain Durmus, and Marcelo Pereyra. Bayesian imaging using plug & play priors: when langevin meets tweedie. *SIAM Journal on Imaging Sciences*, 15(2):701–737, 2022. [2](#), [3](#)
- [24] Anat Levin, Yair Weiss, Fredo Durand, and William T Freeman. Understanding and evaluating blind deconvolution algorithms. In *2009 IEEE conference on computer vision and pattern recognition*, pages 1964–1971. IEEE, 2009. [5](#)
- [25] Jingyun Liang, Jie Zhang Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. Swinir: Image restoration using swin transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1833–1844, 2021. [5](#)
- [26] Renhao Liu, Yu Sun, Jiabei Zhu, Lei Tian, and Ulugbek S Kamilov. Recovery of continuous 3d refractive index maps from discrete intensity-only measurements using neural fields. *Nature Machine Intelligence*, 4(9):781–791, 2022. [1](#)
- [27] Tim Meinhardt, Michael Moller, Caner Hazirbas, and Daniel Cremers. Learning proximal operators: Using denoising networks for regularizing inverse imaging problems. In *Pro-*

- ceedings of the IEEE International Conference on Computer Vision*, pages 1781–1790, 2017. 1, 2
- [28] Peyman Milanfar. Symmetrizing smoothing filters. *SIAM Journal on Imaging Sciences*, 6(1):263–284, 2013. 2
- [29] Philip Müller, Vladimir Golkov, Valentina Tomassini, and Daniel Cremers. Rotation-equivariant deep learning for diffusion mri. *arXiv preprint arXiv:2102.06942*, 2021. 2
- [30] Jean-Christophe Pesquet, Audrey Repetti, Matthieu Terris, and Yves Wiaux. Learning maximally monotone operators for image recovery. *SIAM Journal on Imaging Sciences*, 14(3):1206–1237, 2021. 1, 2, 5, 6
- [31] Edward T Reehorst and Philip Schniter. Regularization by denoising: Clarifications and new interpretations. *IEEE transactions on computational imaging*, 5(1):52–67, 2018. 3
- [32] JH Rick Chang, Chun-Liang Li, Barnabas Póczos, BVK Vijaya Kumar, and Aswin C Sankaranarayanan. One network to solve them all—solving linear inverse problems using deep projection models. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5888–5897, 2017. 1
- [33] Yaniv Romano, Michael Elad, and Peyman Milanfar. The little engine that could: Regularization by denoising (red). *SIAM Journal on Imaging Sciences*, 10(4):1804–1844, 2017. 1, 2, 3
- [34] Cédric Rommel, Thomas Moreau, and Alexandre Gramfort. Deep invariant networks with differentiable augmentation layers. *Advances in Neural Information Processing Systems*, 35:35672–35683, 2022. 2
- [35] Ernest Ryu, Jialin Liu, Sicheng Wang, Xiaohan Chen, Zhangyang Wang, and Wotao Yin. Plug-and-play methods provably converge with properly trained denoisers. In *International Conference on Machine Learning*, pages 5546–5557. PMLR, 2019. 1, 2, 4
- [36] Andrew K Saydjari and Douglas P Finkbeiner. Equivariant wavelets: Fast rotation and translation invariant wavelet scattering transforms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2):1716–1731, 2022. 2
- [37] Ivan W Selesnick, Richard G Baraniuk, and Nick C Kingsbury. The dual-tree complex wavelet transform. *IEEE signal processing magazine*, 22(6):123–151, 2005. 2
- [38] Jean-Pierre Serre et al. *Linear representations of finite groups*. Springer, 1977. 3
- [39] Vasiliki Stergiopoulou, Subhadip Mukherjee, Luca Calatroni, and Laure Blanc-Féraud. Fluctuation-based deconvolution in fluorescence microscopy using plug-and-play denoisers. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pages 498–510. Springer, 2023. 1
- [40] Julián Tachella, Dongdong Chen, and Mike Davies. Sensing theorems for unsupervised learning in linear inverse problems. *Journal of Machine Learning Research*, 24(39):1–45, 2023. 4
- [41] Julian Tachella, Dongdong Chen, Samuel Hurault, and Matthieu Terris. DeepInverse: A deep learning framework for inverse problems in imaging, 2023. 5
- [42] Matthieu Terris, Audrey Repetti, Jean-Christophe Pesquet, and Yves Wiaux. Building firmly nonexpansive convolutional neural networks. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8658–8662. IEEE, 2020. 1, 2
- [43] Matthieu Terris, Arwa Dabbech, Chao Tang, and Yves Wiaux. Image reconstruction algorithms in radio interferometry: From handcrafted to learned regularization denoisers. *Monthly Notices of the Royal Astronomical Society*, 518(1):604–622, 2023. 1, 2
- [44] Radu Timofte, Rasmus Rothe, and Luc Van Gool. Seven ways to improve example-based single image super resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1865–1873, 2016. 2
- [45] Singanallur V Venkatakrishnan, Charles A Bouman, and Brendt Wohlberg. Plug-and-play priors for model based reconstruction. In *2013 IEEE global conference on signal and information processing*, pages 945–948. IEEE, 2013. 1, 2, 3
- [46] Dian Wang, Robin Walters, Xupeng Zhu, and Robert Platt. Equivariant q learning in spatial action spaces. In *Conference on Robot Learning*, pages 1713–1723. PMLR, 2022. 2
- [47] Maurice Weiler and Gabriele Cesa. General $e(2)$ -equivariant steerable cnns. *Advances in neural information processing systems*, 32, 2019. 2
- [48] 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. 5
- [49] 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*, 26(7):3142–3155, 2017. 5
- [50] 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*, 44(10):6360–6376, 2021. 1, 2, 5, 6
- [51] Kai Zhang, Yawei Li, Jingyun Liang, Jiezhong Cao, Yulun Zhang, Hao Tang, Radu Timofte, and Luc Van Gool. Practical blind denoising via swin-conv-unet and data synthesis. *arXiv preprint arXiv:2203.13278*, 2022. 5
- [52] Yulun Zhang, Kai Zhang, Zheng Chen, Yawei Li, Radu Timofte, Junpei Zhang, Kexin Zhang, Rui Peng, Yanbiao Ma, Licheng Jia, et al. Ntire 2023 challenge on image super-resolution (x4): Methods and results. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1864–1883, 2023. 1
- [53] Minghan Zhu, Shizhong Han, Hong Cai, Shubhankar Borse, Maani Ghaffari, and Fatih Porikli. 4d panoptic segmentation as invariant and equivariant field prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22488–22498, 2023. 2
- [54] Yuanzhi Zhu, Kai Zhang, Jingyun Liang, Jiezhong Cao, Bihan Wen, Radu Timofte, and Luc Van Gool. Denoising diffusion models for plug-and-play image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1219–1229, 2023. 1