



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Robust Equivariant Imaging: a fully unsupervised framework for learning to image from noisy and partial measurements

Citation for published version:

Chen, D, Tachella, J & Davies, ME 2022, Robust Equivariant Imaging: a fully unsupervised framework for learning to image from noisy and partial measurements. in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Conference on Computer Vision and Pattern Recognition (CVPR) , Institute of Electrical and Electronics Engineers, pp. 5637-5646, IEEE/CVF Conference on Computer Vision and Pattern Recognition 2022 , New Orleans, Louisiana, United States, 19/06/22. <https://doi.org/10.1109/CVPR52688.2022.00556>

Digital Object Identifier (DOI):

[10.1109/CVPR52688.2022.00556](https://doi.org/10.1109/CVPR52688.2022.00556)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Robust Equivariant Imaging: a fully unsupervised framework for learning to image from noisy and partial measurements

Dongdong Chen*
School of Engineering
University of Edinburgh

Julián Tachella*
School of Engineering
University of Edinburgh

Mike E. Davies
School of Engineering
University of Edinburgh

Abstract

Deep networks provide state-of-the-art performance in multiple imaging inverse problems ranging from medical imaging to computational photography. However, most existing networks are trained with clean signals which are often hard or impossible to obtain. Equivariant imaging (EI) is a recent self-supervised learning framework that exploits the group invariance present in signal distributions to learn a reconstruction function from partial measurement data alone. While EI results are impressive, its performance degrades with increasing noise. In this paper, we propose a Robust Equivariant Imaging (REI) framework which can learn to image from noisy partial measurements alone. The proposed method uses Stein’s Unbiased Risk Estimator (SURE) to obtain a fully unsupervised training loss that is robust to noise. We show that REI leads to considerable performance gains on linear and nonlinear inverse problems, thereby paving the way for robust unsupervised imaging with deep networks. Code is available at <https://github.com/edongdongchen/REI>.

1. Introduction

Inverse problems play a fundamental role in computer vision and signal processing. Applications such as computed tomography (CT), magnetic resonance imaging (MRI), super-resolution and image inpainting have been extensively explored. The goal in an inverse problem is to reconstruct a signal $x \in \mathbb{R}^n$ from measurements $y \in \mathbb{R}^m$, that is inverting the forward process

$$y = A(x) + \epsilon, \quad (1)$$

which is generally a challenging task due to the noise ϵ and the ill-conditioned or rank deficient forward operator A (e.g., $m < n$). Traditional approaches exploit prior knowledge about x (sparsity [10], total variation [36] or more re-




Figure 1. **REI vs EI**: EI is not robust to noise and the performance degrades with increasing noise. From top to bottom: reconstructions of EI, supervised (Sup) baseline, and the proposed REI on $4\times$ accelerated MRI with Gaussian noise level $\sigma = 0.01, 0.1, 0.2$. PSNR values are shown in the top right corner of the images.

cently the deep image prior [39]) to regularize the reconstructions. A different approach is taken by more recent end-to-end learning solutions [17] which aim to learn the inverse mapping directly from (x, y) pairs. However, in many inverse problems, it is expensive or impossible to obtain clean signals x , and only noisy measurement data y are available for training.

Without any assumptions on the signal distribution, it is fundamentally impossible to learn the reconstruction mapping from partial measurements alone, as there are infinitely many inverses for each observed measurement y [4, 10]. Some unsupervised methods alleviate this limitation by us-

*The first two authors have contributed equally to this paper.

ing multiple forward operators [1, 9, 32], however this setting is not often met in practice. Recently, the EI framework [4] showed that learning from partial measurements y is possible if the underlying signal set is invariant to transformations, such as shifts or rotations. The EI strategy offers an elegant way to learn the model in a self-supervised fashion, as most natural signals present certain invariances (e.g., invariance to translations and rotations in images). EI was shown to achieve similar performance to supervised methods on a number of noiseless image reconstruction tasks such as sparse-view CT and image inpainting. However, as shown in Figure 1, EI’s performance degrades rapidly in the presence of measurement noise.

In this paper, we propose a novel robust EI (REI) framework which can learn from only noisy measurements data. In particular, we use Stein’s Unbiased Risk Estimator (SURE) [38] to derive an unsupervised loss that takes into account the noise present in the measurements. The proposed learning method is architecture agnostic and can handle a wide range of different noise models, including Gaussian noise, Poisson noise and mixed Poisson-Gaussian noise. We experimentally show that our REI method outperforms the vanilla EI for all noise levels in 3 inverse problems: sparse-view CT, accelerated MRI and image inpainting. In summary, our contributions are as follows:

1. We present a new end-to-end unsupervised learning framework for training deep networks to solve imaging problems from noisy partial measurements alone.
2. Our REI learning framework can handle a wide range of the most frequently encountered noise models in practice, including: Gaussian noise, Poisson noise, and mixed Gaussian-Poisson noise.
3. We evaluate the proposed method on accelerated MRI, image inpainting, and a nonlinear imaging model for sparse-view CT, showing substantial improvements over vanilla EI in the presence of noise and approaching the quality of reconstructions trained with ground truth data using supervised learning.

2. Background

We begin with some basic definitions. Let $A : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a (possibly nonlinear) smooth measurement operator with constant rank $m < n$. For a given measurement vector $y \in \mathbb{R}^m$, the set of feasible solutions $A^{-1}(y) = \{x \in \mathbb{R}^n | A(x) = y\}$ is a submanifold of \mathbb{R}^n of dimension $n - m$ [23, Chapter 4].

2.1. Learning to Solve Inverse Problems

Modern learning-based approaches learn a reconstruction function $f_\theta : \mathbb{R}^m \rightarrow \mathbb{R}^n$ which maps the observed measurements y to the reconstructed signal x . The function




Figure 2. **Equivariant imaging systems.** If the set of signals is invariant to a certain set of transformations, the composition of imaging operator A with the reconstruction function f_θ should be equivariant to these transformations.

is generally parameterized by a convolutional network and trained using N pairs of measurements and associated reconstructions $\{(y_i, x_i)\}_{i=1, \dots, N}$. Learning only with partial measurements y_i is impossible without any further assumptions [4, 10] as $f_\theta(y)$ can output any of the $x \in A^{-1}(y)$ and still perfectly fit the training set.

Equivariant Imaging (EI) Recently, the EI framework [4] showed that learning with only measurement data y_i is possible if the signal set is assumed to be group invariant. That is, the signal set $\mathcal{X} \subseteq \mathbb{R}^n$ is invariant to a group of transformations (e.g., shifts, rotations, etc.) $\{T_g\}_{g=1, \dots, |\mathcal{G}|}$ if for all $x \in \mathcal{X}$, $T_g x$ also belongs to \mathcal{X} for all transformations $g = 1, \dots, |\mathcal{G}|$. Assuming that the transformation matrices are unitary¹ and there is no noise affecting the measurements, for all group elements $g = 1, \dots, |\mathcal{G}|$, we have that

$$y = A(x) = A(T_g T_g^\top x) = A_g(\tilde{x}) \quad (2)$$

where $A_g = A \circ T_g$ and both x and $\tilde{x} = T_g^\top x$ belong to the signal set \mathcal{X} due to the invariance property. Thus, the invariance property gives access to multiple *virtual* operators $\{A_g\}_{g=1, \dots, |\mathcal{G}|}$. For a given measurement y , each operator A_g might have a different submanifold of possible inverses $A_g^{-1}(y)$. Thus, it is feasible to learn f_θ from only measurements if the intersection of these submanifolds $\bigcap_g A_g^{-1}(y)$ contains a single signal x .

The invariance property also tells us that the composition $f_\theta \circ A$ should be equivariant to the transformations (see Figure 2):

$$T_g f_\theta(A(x)) = f_\theta(A(T_g x)) \quad (3)$$

for $g = 1, \dots, |\mathcal{G}|$. To exploit this idea the EI training strategy enforces both measurement consistency $A(f_\theta(y)) = y$ and the equivariance condition in (3) using a dataset of measurements $\{y_i\}_{i=1, \dots, N}$. The reconstruction function f_θ is then learned by minimizing the following loss:

$$\mathcal{L}_{EI}(\theta) = \mathcal{L}_{MC}(\theta) + \alpha \mathcal{L}_{EQ}(\theta), \quad (4)$$

¹This can always be obtained after a change of basis if the group is compact, see for example [20].

with

$$\mathcal{L}_{\text{MC}}(\theta) = \sum_{i=1}^N \frac{1}{m} \|y_i - A(f_\theta(y_i))\|^2, \quad (5)$$

$$\mathcal{L}_{\text{EQ}}(\theta) = \sum_{i=1}^N \sum_{g=1}^{|\mathcal{G}|} \frac{1}{|\mathcal{G}|n} \|T_g f_\theta(y_i) - f_\theta(A(T_g f_\theta(y_i)))\|^2 \quad (6)$$

where the first term (5) enforces measurement consistency (MC) and the second term (6) enforces system equivariance, and α is a trade-off parameter.

While EI proved to be an efficient self-supervised learning method in noiseless settings [4], its performance degrades rapidly in the presence of noise, as shown in Figure 1, hindering its use in noisy inverse problems.

2.2. Self-supervised denoising via SURE

In most real-world settings, observed measurements are corrupted by noise. For example, measurements can contain Gaussian noise, that is

$$y|u \sim \mathcal{N}(u, I\sigma^2), \quad (7)$$

where u denotes the clean measurements, σ is the standard deviation of the noise and I denotes the identity matrix. A denoiser network h_θ that maps $y \mapsto u$ is generally trained with a dataset of noisy and clean pairs (y, u) via the supervised loss

$$\mathcal{L}_{\text{MSE}}(\theta) = \frac{1}{m} \sum_{i=1}^N \|y_i - h_\theta(y_i)\|_2^2. \quad (8)$$

However, in practice sometimes only the noisy observations y_i are available. In settings where the noise distribution is Gaussian, it is possible to bypass the need of clean training data with a SURE-based loss [38]

$$\mathcal{L}_{\text{SURE}}(\theta) = \sum_{i=1}^N \frac{1}{m} \|y_i - h_\theta(y_i)\|^2 - \sigma^2 + \frac{2\sigma^2}{m} \nabla \cdot h_\theta(y_i) \quad (9)$$

where $\nabla \cdot h_\theta(y) = \sum_{j=1}^m \frac{\partial h_{\theta,j}(y)}{\partial y_j}$ is the divergence of the function h_θ with respect to the noisy input y . The first two terms in (9) penalize the error between the noisy observations and denoised ones (bias of the denoiser), whereas the last term penalizes the variance of the denoiser. Crucially, the SURE loss is an unbiased estimator of the supervised loss:

Theorem 1 [38] *Let $h_\theta : \mathbb{R}^m \rightarrow \mathbb{R}^m$ be a (weakly) differentiable real-valued function such that $\forall j = 1, 2, \dots, m$, $\mathbb{E}_\epsilon \{|\partial h_j(y)/\partial y_j|\} < -\infty$. Then, the SURE loss is an unbiased estimator of the supervised mean squared loss, that is*

$$\mathbb{E}_y \{\mathcal{L}_{\text{SURE}}(\theta)\} = \mathbb{E}_{y,u} \{\mathcal{L}_{\text{MSE}}(\theta)\} \quad (10)$$

The SURE framework has been generalized by a number of authors, [8, 25, 34] beyond simple additive Gaussian noise, including Poisson and mixed Poisson Gaussian noise, resulting in different loss expressions depending on the noise type. In each case, however, there is always the need to calculate the divergence of the denoiser function (c.f. (9)) which may not have a simple analytic form, e.g., for deep networks. To alleviate this limitation, Ramani et al. [33] proposed a Monte Carlo method to estimate the divergence which only requires one additional evaluation of h_θ :

Theorem 2 [33] *Let b be a zero-mean i.i.d. random vector with unit variance and bounded higher order moments. Then*

$$\nabla \cdot h_\theta(y) = \lim_{\tau \rightarrow 0} \mathbb{E}_b \left\{ \frac{1}{\tau} b^\top (h_\theta(y + \tau b) - h_\theta(y)) \right\} \quad (11)$$

provided that $h_\theta(y)$ admits a well-defined second-order Taylor expansion. If not, this is still valid in the weak sense provided that $h_\theta(y)$ is tempered.

Thus, the divergence term can be approximated by sampling a single i.i.d. vector, e.g., $b \sim \mathcal{N}(0, 1)$, and fixing a small positive value τ :

$$\nabla \cdot h_\theta(y) \approx \frac{1}{\tau} b^\top (h_\theta(y + \tau b) - h_\theta(y)). \quad (12)$$

Incorporating this expression into the associated SURE loss has been shown to yield accurate unbiased estimates may not have a simple analytic form MSE for many various denoising methods $h_\theta(y)$ [33]. Furthermore, combining the SURE loss (9) with a parameterized denoiser enables unsupervised learning via minimization of this proxy for MSE without access to ground truth signals.

Recently, it has become popular to use the SURE loss for unsupervised training of denoising networks [37, 42], and to learn solutions to compressed sensing problems where reconstruction can be cast as an iterative sequence of denoising steps, e.g., [12, 28–30]. However, as explained in Sec. 2.1, learning from measurement alone is not possible without further assumptions. We therefore consider next how to exploit SURE within the EI framework.

3. Robust Equivariant Imaging

Our aim is to develop a *Robust Equivariant Imaging* (REI) framework to learn the reconstruction function $f_\theta : y \rightarrow x$ in a fully unsupervised manner from only noisy partial measurement data $\{y\}_{i=1, \dots, N}$ in inverse problems whose acquisition models are rank deficient (e.g., with a non-trivial nullspace \mathcal{N}_A when A is linear). We assume the following generative model:

$$y|u \sim q_u(y), \quad u = A(x) \quad (13)$$




Figure 3. **REI training strategy.** $x^{(1)}$ represents the estimated image, while $x^{(2)}$ and $x^{(3)}$ represent $T_g x^{(1)}$ and the estimate of $x^{(2)}$ from the (noisy) measurements $\tilde{y} = A(x^{(2)})$ respectively.

where $q_u(y)$ is the noise distribution and $x \sim p(x)$, the unknown signal distribution, is supported on a signal set \mathcal{X} which is assumed to be invariant to the transformations $\{T_g\}_{g=1, \dots, |\mathcal{G}|}$.

The REI model replaces the two loss terms in (4) with noise-robust counterparts:

$$\mathcal{L}_{\text{REI}}(\theta) = \mathcal{L}_{\text{SURE}}(\theta) + \alpha \mathcal{L}_{\text{REQ}}(\theta) \quad (14)$$

The right hand term, presented in Section 3.1, aims to enforce equivariance of the imaging system $f_\theta \circ A$ taking into account that the input to f_θ should be noisy. The left hand term is a SURE-based approximation of the unknown clean measurement consistency loss in (4), which depends on the noise model and it is presented in Section 3.2. The network can be trained in an end-to-end fashion using stochastic optimization [19], see Figure 3.

The proposed unsupervised training framework is agnostic of the network architecture used for f_θ . Thus, in principle, REI can be used to train any existing model in an unsupervised way. In our experiments, we use a trainable neural network $f_\theta = G_\theta \circ A^\dagger$, where $G_\theta : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a neural network and $A^\dagger : \mathbb{R}^m \rightarrow \mathbb{R}^n$ is any (fixed) suitable backprojection.

3.1. Enforcing equivariance with noisy data

Following [4], we wish to enforce that $f_\theta \circ A$ is equivariant to transformations. We adapt the equivariant loss in (6) to handle noisy measurements:

$$\mathcal{L}_{\text{REQ}}(\theta) = \sum_{i=1}^N \sum_{g=1}^{|\mathcal{G}|} \frac{1}{|\mathcal{G}|n} \|T_g f_\theta(y_i) - f_\theta(\tilde{y}_i)\|^2 \quad (15)$$

where \tilde{y}_i is sampled from the distribution $\tilde{y}_i \sim q_{\tilde{y}_i}$ with underlying measurement $\tilde{u}_i = A(T_g f_\theta(y_i))$. The sampled measurements \tilde{y}_i can be interpreted as noisy measurements obtained via the virtual operator $A \circ T_g$. In the absence of noise, $\mathcal{L}_{\text{REQ}}(\theta)$ is equivalent to the equivariant loss in (6).

3.2. SURE-based measurement consistency

In this work, while the methodology is more general, we consider 3 specific noise models $q_u(y)$ which appear

often in practice: Gaussian, Poisson and mixed Poisson-Gaussian. In each case, we develop a noise-specific measurement consistency SURE term $\mathcal{L}_{\text{SURE}}(\theta)$ which is an unbiased estimator of the unknown clean measurement consistency loss. That is we develop a SURE estimate for the denoising function $A \circ f_\theta$ that denoises the measurements:

$$\mathbb{E}_y \{\mathcal{L}_{\text{SURE}}(\theta)\} = \mathbb{E}_{y,u} \left\{ \sum_{i=1}^N \frac{1}{m} \|u_i - A(f_\theta(y_i))\|_2^2 \right\}. \quad (16)$$

A detailed derivation of all the formulas is included in the supplementary material (SM).

Gaussian noise We first consider the standard Gaussian noise model, that is

$$y \sim \mathcal{N}(u, \sigma^2 I) \quad \text{with} \quad u = A(x). \quad (17)$$

Following Sec. 2.2, in this case we have

$$\begin{aligned} \mathcal{L}_{\text{SURE}}(\theta) &= \sum_{i=1}^N \frac{1}{m} \|y_i - A(f_\theta(y_i))\|_2^2 - \sigma^2 \\ &+ \frac{2\sigma^2}{m\tau} b_i^\top (A(f_\theta(y_i + \tau b_i)) - A(f_\theta(y_i))) \end{aligned} \quad (18)$$

where $b_i \sim \mathcal{N}(0, I)$ and τ is a small positive number.

Poisson noise In various imaging settings such as low-photon imaging, the observed measurements follow a Poisson distribution. This noise is modeled as:

$$y = \gamma z \quad \text{with} \quad z \sim \text{Poisson}\left(\frac{u}{\gamma}\right), \quad u = A(x). \quad (19)$$

where $\gamma > 0$ is the gain of the acquisition process which controls the noise level. In this case, we have

$$\begin{aligned} \mathcal{L}_{\text{SURE}}(\theta) &= \frac{1}{m} \sum_{i=1}^N \|y_i - A(f_\theta(y_i))\|_2^2 - \gamma \mathbf{1}^\top y_i \\ &+ \frac{2\gamma}{\tau} (b_i \odot y_i)^\top (A(f_\theta(y_i + \tau b_i)) - A(f_\theta(y_i))) \end{aligned} \quad (20)$$

where here we select b_i to be a Bernoulli random variable [22] taking values of -1 and 1 each with a probability of 0.5 , τ is a small positive number, and \odot is an element-wise multiplication.

Mixed Gaussian-Poisson (MPG) model The Gaussian and Poisson noise models usually do not individually account for the various phenomena involved with real image acquisition processes in inverse problems such as fluorescence microscopy and X-ray computed tomography

(CT) [22]. The noise in these inverse problems is better captured by a mixed Poisson-Gaussian noise model [16, 22]:

$$y = \gamma z + \epsilon \quad \text{with} \quad \begin{cases} u = A(x) \\ z \sim \text{Poisson}\left(\frac{u}{\gamma}\right) \\ \epsilon \sim \mathcal{N}(0, \sigma^2 I) \end{cases}. \quad (21)$$

The Gaussian model is recovered as $\gamma \rightarrow 0$ whereas the Poisson model is obtained by setting $\sigma = 0$.

Extending the pioneering work in [22, 25], we derive the unbiased risk estimator of the MSE of clean measurement consistency for the MPG noise model:

$$\begin{aligned} \mathcal{L}_{\text{SURE}}(\theta) &= \sum_{i=1}^N \frac{1}{m} \|y_i - A(f_\theta(y_i))\|_2^2 - \gamma \mathbf{1}^\top y_i - \sigma^2 \\ &+ \frac{2}{\tau} (b_i \odot (\gamma y_i + \sigma^2 I))^\top (A(f_\theta(y_i + \tau b_i)) - A(f_\theta(y_i))) \\ &+ \frac{2\gamma\sigma^2}{\tau} c_i^\top (A(f_\theta(y_i + \tau c_i)) + A(f_\theta(y_i - \tau c_i))) \dots \\ &- 2A(f_\theta(y_i)) \end{aligned} \quad (22)$$

where $b_i \sim \mathcal{N}(0, I)$, c_i are i.i.d. random variables that follow a Bernoulli distribution taking values of -1 and 1 each with a probability of 0.5 , τ is a small positive number.

4. Experiments

We evaluated the REI framework in 3 inverse problems: accelerated MRI, image inpainting and a (nonlinear) sparse view CT problem. More details are presented in the SM.

4.1. Setup and Implementation

The forward operator A in all the evaluated tasks is known and is dimension reducing ($m < n$). We considered different noise models for each task. Throughout the experiments, we follow [4] and use a residual U-Net [35] to build G_θ . We compare our new learning method (REI) with three different learning strategies: vanilla EI given by the loss in (4), an *oracle* REI trained with noiseless measurement consistency (REI_{oracle}) that replaces the SURE loss by the true clean measurement consistency loss, and supervised learning (Sup) [17], which minimizes $\sum_{i=1}^N \frac{1}{n} \|x_i - f_\theta(y_i)\|_2^2$ using clean signal and noisy measurements pairs. The oracle REI provides us with a way to evaluate the separate contributions of the REQ and SURE losses. We use the same neural network architecture for all learning methods. For a fair comparison with EI and REI, no data augmentation of is used for supervised learning. We *emphasize* that our goal here is not to achieve state-of-the-art performance, but to show that unsupervised learning is possible on different inverse problems with noise. The choice of architecture is somewhat orthogonal to this goal.

All methods are implemented in PyTorch and optimized by Adam [19]. A full ablation study, based on the MRI experiment below, and across all the possible loss terms: EI, REI, REI_{oracle}, MC, SURE and combinations thereof is presented in the SM. This shows that replacement of REI or SURE with other losses has a negative impact on performance and fails to track the performance of supervised learning.

4.2. Accelerated MRI

MRI produces images of biological tissues by sampling the Fourier transform (k -space) of the image. In accelerated MRI, the goal is to recover high quality MR images from sparsely sampled k -space data to reduce the acquisition time. The forward operator is $A := \mathcal{S}_\omega \circ \mathcal{F}$ where \mathcal{F} is the 2D Fourier transform and \mathcal{S}_ω is a binary matrix, whose ones correspond to frequencies in ω that are measured. The pseudo-inverse $A^\dagger = \mathcal{F}^H \circ \mathcal{S}_\omega$ is the masked inverse Fourier transform. This problem has been extensively studied in the context of compressed sensing [26] and supervised deep learning [13, 14, 17].

The primary sources of noise in MR are electronic, and dielectric and inductive coupling to the conducting solution inside the body [2], and is usually modelled as Gaussian [11]. We therefore train the REI models using (18) to implement the SURE loss in (14). In this task, we exploit the invariance of MRI images to rotations, and use rotations of integer degree ($|\mathcal{G}|=360$), while α is set to one (see SM for the effect of varying α).

The MRI data used were obtained from NYU fastMRI Initiative [40]. We trained and tested on a subset of the single-coil measurements at $4\times$ acceleration. There are 973 samples in the dataset, the first 900 were used for training and the last 73 were used for testing. Each image is of size 320×320 pixels. In all tests, we use complex-valued data treating the real and imaginary parts of the images as separate channels. For the purpose of visualization, we display only the magnitude images.

The reconstruction results are reported in Table 1 and Figures 4-5. Note that EI only performs well for small noise levels, e.g., $\sigma \leq 0.05$, and degrades rapidly when the noise level is increased becoming almost as bad as the linear pseudo-inverse reconstruction $A^\dagger(y)$ for high noise levels. In comparison, REI enjoys a 3.6 dB gain at $\sigma = 0.1$ which rises to a 7 dB improvement with respect to EI when $\sigma = 0.2$. While, unsurprisingly, supervised learning always enjoys the slowest performance degradation as σ increases due to its access to ground truth training data, the performance of REI stays within 1.5 to 2.5 dB of that of supervised learning over the full range of noise levels. The oracle REI performance lies close to that for REI, suggesting that the SURE loss (18) is doing a reasonable job of approximating the (oracle) clean measurement consistency loss (see SM for the full ablation study). Overall, these results sug-




Figure 4. MRI image reconstruction on the test data with Gaussian noise (from top to bottom: $\sigma = 0.01, 0.1, 0.2$) and $4\times$ acceleration. PSNR values are shown in the top right corner of the images.




Figure 5. Reconstruction performance (PSNR) as a function of noise level σ for linear reconstruction $A^\dagger(y)$, supervised training [17], vanilla EI [4], and our method REI on MRI image reconstruction with $4\times$ compressed and noisy k -space measurements.

σ	$A^\dagger(y)$	EI	REI	REI _{oracle}	Sup
0	25.8 ± 2.7	31.0 ± 2.0	31.0 ± 2.0	31.0 ± 2.0	32.5 ± 2.1
0.01	25.7 ± 2.7	30.6 ± 1.8	30.8 ± 2.0	30.8 ± 2.0	32.2 ± 2.0
0.05	24.7 ± 2.1	27.6 ± 1.3	29.0 ± 1.8	29.6 ± 1.9	30.8 ± 1.8
0.1	22.6 ± 1.4	24.1 ± 0.9	27.7 ± 2.0	28.9 ± 1.7	29.8 ± 1.8
0.2	18.7 ± 0.7	19.7 ± 0.7	26.9 ± 2.0	27.7 ± 1.7	28.8 ± 1.7

Table 1. PSNR of $4\times$ accelerated MRI image reconstruction for different methods using noisy test measurements.

gest that REI provides an unsupervised learning framework that is robust to noise with nearly as good performance as supervised learning.

4.3. Inpainting

In image inpainting problems, the goal is to recover a complete image from a subset of noisy pixels. Inpainting can be used for example in imaging systems with hot or dead pixels. We consider the Poisson noise model in (19) to generate the noisy measurements. We therefore use (20) to implement (14) and to train REI models.

In this task, the forward model is described as $y = a \odot x$ where a is a binary mask, \odot is the Hadamard product, and the associated operator and pseudo-inverse are $A = \text{diag}(a)$ and $A = A^\dagger$. We consider random masks which drop 30%

of pixel measurements. We then train our REI models by applying random (horizontal and vertical) shift transformations with $\alpha = 1$. We evaluate the reconstruction performance on the Urban100 [15] natural image dataset. For each image, we cropped a 512×512 pixel area at the center and then resized it to 256×256 for the ground truth image. The first 90 images are used for training while the last 10 measurements are used for testing.

The results in Table 2 and Figure 6 show that even for small values of γ , EI struggles to compete with REI and the performance rapidly decreases when the noise level γ is increased. In contrast, REI performs well across all noise levels, beating EI by more than 8 dB at high noise levels, and within 1 dB of the fully supervised model, demonstrating the effectiveness of REI for the Poisson noise model.




Figure 6. Inpainting reconstructions on test images with Poisson noise (from top to bottom: $\gamma = 0.01, 0.05, 0.1$) and 30% mask rate. PSNR values are shown in the top right corner of the images.

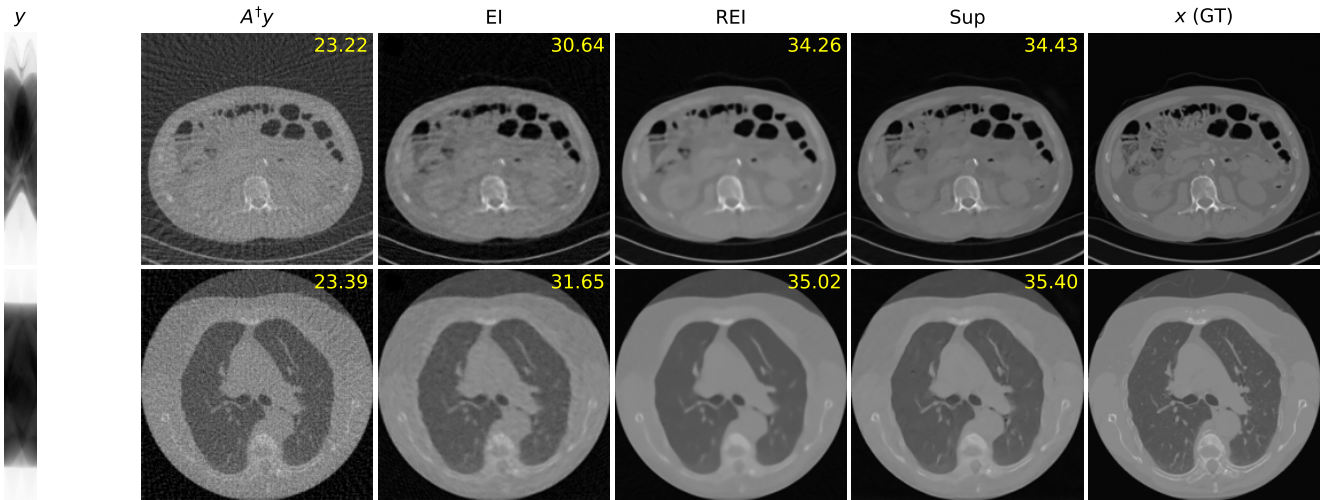


Figure 7. CT image reconstruction (50 views) on the test observations with mixed Poisson-Gaussian noise, $I_0 = 10^5$, $\sigma = 30$, $\gamma = 1$. PSNR values are shown in the top right corner of the images.

4.4. Sparse-view CT

Our final experiment provides an initial demonstration of robust equivariant imaging for a nonlinear inverse problem. Sparse-view CT can be described by a nonlinear forward model $A(x) = I_0 e^{-\text{radon}(x)}$, where radon denotes a sparse-view Radon transformation and I_0 denotes the X-ray source intensity. Quantum and electronic noise are two major noise sources in X-ray CT scanners [7] and for nor-

mal clinical exposures, measurements y are often modeled as the sum of a Poisson distribution representing photon-counting statistics and an independent Gaussian distribution representing additive electronic noise, i.e., $y = z + \epsilon$ where $z \sim \text{Poisson}(A(x))$ and $\epsilon \sim \mathcal{N}(0, \sigma^2 I)$.

Here, we exploit the invariance of CT images to rotations and consider rotations of multiples of 1 degree ($|\mathcal{G}|=360$). For A^\dagger we use the filtered backprojection (FBP), iradon ,

γ	$A^\dagger(y)$	EI	REI	REI _{oracle}	Sup
0.01	5.7 ± 1.5	18.9 ± 1.0	21.7 ± 1.3	22.0 ± 1.3	22.3 ± 1.1
0.05	5.1 ± 1.4	11.8 ± 3.0	18.2 ± 1.2	18.3 ± 0.9	19.6 ± 1.2
0.1	4.4 ± 1.3	9.8 ± 0.8	16.6 ± 1.3	17.3 ± 1.3	18.2 ± 1.2

Table 2. PSNR of pixelwise image inpainting reconstruction for different methods using noisy test measurements.

of the logarithm of the measurements, $\log(I_0/y)$. The reconstruction function is thus defined as $f_\theta(y) = G_\theta \circ \text{iradon} \circ \log(I_0/y)$. We therefore use the MPG model, (22) to implement (14) and train the REI models.

We use the CT100 clinical dataset [5], which comprises 100 real in-vivo CT images collected from the cancer imaging archive. The CT images were resized to 256×256 pixels. The `radon` function consists of 50 uniformly sampled views (angles). For the imaging noise we follow [7] and generate MPG noisy measurements with $I_0 = 10^5$, $\sigma = 30$ and $\gamma = 1$. We used the sinograms of the first 90 images for training and the remaining 10 for testing. In this task, we use the U-Net which is the same architecture as FBPCnnvnet in [17]. We trained EI, REI and REI_{oracle} with $\alpha = 10^3$ (this is much larger than in the MRI and inpainting experiments due to the scaled log, $\log(I_0/y)$)² using the sinograms y alone while the supervised training is performed with the noisy measurements and ground truth pairs (x, y) .

A quantitative comparison is presented in Table 3. While all learning-based methods outperform the standard FBP reconstruction, REI achieves a 5 dB gain when compared with EI, and performs almost as well as the oracle REI_{oracle} and the fully supervised baseline. Figure 7 shows examples of reconstructed images. The FBP reconstruction contains massive line artefacts and noise due to the sparse-view MPG measurements. Although EI performs better, it fails to fully remove the artefacts and noise. In contrast, REI achieves almost equivalent quality to the supervised method. A similarly performing low-dose example, with $I_0 = 10^4$, $\sigma = 50$, is given in the SM.

	FBP	EI	REI	REI _{oracle}	Sup
CT	22.8 ± 0.5	29.2 ± 0.6	34.0 ± 1.1	34.0 ± 1.2	34.4 ± 1.1

Table 3. Reconstruction performance (PSNR) of sparse-view CT task for different methods on the test noisy measurements.

5. Discussion

One might naturally ask whether one could hard wire the equivariance into the neural network. Indeed, there has been considerable efforts recently in designing such architectures, mostly in the context of classification [6, 21, 27].

²Generally, α should scale inversely proportional to the norm of the (linearized) measurement operator.

However, in equivariant imaging it is not the network that we want to be equivariant but the full acquisition-reconstruction system, $f_\theta \circ A$. Celledoni et al. [3] proposed an interesting extension to design equivariant networks for inverse problems by using an unrolled network with an equivariant network substituting the proximal component. However, as noted by the authors, this approach does not guarantee full system equivariance and is restricted to supervised learning applications. In contrast, (R)EI promotes full system equivariance through the (robust) equivariant loss, thereby enabling unsupervised learning.

Training using SURE requires knowledge of the noise characteristics.³ While we evaluated REI on the Gaussian, Poisson and MPG models, SURE can handle many other models including non-exponential ones, see [34] for a detailed list. SURE based training losses have also been used in other related works, [18, 30, 37, 41] for unsupervised imaging/denoising, however they are only applied to inverse problems with complete measurements. Noise2Noise [24] takes a slightly different route (and only considered denoising) by using training pairs consisting of two different noisy realisations of the same signal. While this may appear to be blind to the noise characteristics, the requirement to have two distinct noisy realisations implicitly provides such information. Indeed, we emphasise that some such prior information is always required in any fully unsupervised denoising technique.

Finally, some papers have proposed unsupervised methods for compressed sensing (with non-trivial forward operators) [12, 30]. However, these methods rely heavily on the randomness of A . In contrast, REI applies a SURE-based loss for measurement consistency while exploiting equivariance to encourage a unique reconstruction from partial measurements.

6. Conclusion

In this paper, we proposed *Robust Equivariant Imaging*, a fully unsupervised imaging method that can effectively learn to image from only noisy partial measurements. The method only requires knowledge about the (possibly non-linear) forward operator and the acquisition noise characteristics, which for most imaging systems are usually known. The framework is architecture agnostic, and can be used to train any existing model, and a wide range of noise models in a fully unsupervised way.

Acknowledgments

This work is supported by the ERC C-SENSE project (ERCADG-2015-694888).

³Robustness to different noise levels can be achieved by training over different levels, or possibly by incorporating ideas from [31].

References

- [1] Ashish Bora, Eric Price, and Alexandros G Dimakis. AmbientGAN: Generative models from lossy measurements. In *International Conference on Learning Representations*, 2018. 2
- [2] Arturo Cárdenas-Blanco, Cristian Tejos, Pablo Irarrazaval, and Ian Cameron. Noise in magnitude magnetic resonance images. *Concepts in Magnetic Resonance Part A: An Educational Journal*, 32(6):409–416, 2008. 5
- [3] Elena Celledoni, Matthias Joachim Ehrhardt, Christian Etmann, Brynjulf Owren, Carola-Bibiane Schönlieb, and Ferdia Sherry. Equivariant neural networks for inverse problems. *Inverse Problems*, 37(8):085006, 2021. 8
- [4] Dongdong Chen, Julián Tachella, and Mike E Davies. Equivariant imaging: Learning beyond the range space. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021. 1, 2, 3, 4, 5, 6
- [5] Kenneth Clark, Bruce Vendt, Kirk Smith, John Freymann, Justin Kirby, Paul Koppel, Stephen Moore, Stanley Phillips, David Maffitt, Michael Pringle, et al. The cancer imaging archive (TCIA): maintaining and operating a public information repository. *Journal of digital imaging*, 26(6):1045–1057, 2013. 8
- [6] Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International conference on machine learning*, pages 2990–2999. PMLR, 2016. 8
- [7] Qiaoqiao Ding, Yong Long, Xiaoqun Zhang, and Jeffrey A Fessler. Statistical image reconstruction using mixed Poisson-Gaussian noise model for X-ray CT. *arXiv preprint arXiv:1801.09533*, 2018. 7, 8
- [8] Yonina C Eldar. Generalized SURE for exponential families: Applications to regularization. *IEEE Transactions on Signal Processing*, 57(2):471–481, 2008. 3
- [9] Weijie Gan, Yu Sun, Cihat Eldeniz, Jiaming Liu, Hongyu An, and Ulugbek S Kamilov. Deep image reconstruction using unregistered measurements without groundtruth. In *ISBI*, pages 1531–1534, 2021. 2
- [10] Sivan Gleichman and Yonina C Eldar. Blind compressed sensing. *IEEE Transactions on Information Theory*, 57(10):6958–6975, 2011. 1, 2
- [11] Hákon Gudbjartsson and Samuel Patz. The rician distribution of noisy MRI data. *Magnetic resonance in medicine*, 34(6):910–914, 1995. 5
- [12] Chunli Guo and Mike E. Davies. Near optimal compressed sensing without priors: Parametric sure approximate message passing. *IEEE Transactions on Signal Processing*, 63(8):2130–2141, 2015. 3, 8
- [13] Kerstin Hammernik, Teresa Klatzer, Erich Kobler, Michael P Recht, Daniel K Sodickson, Thomas Pock, and Florian Knoll. Learning a variational network for reconstruction of accelerated mri data. *Magnetic resonance in medicine*, 79(6):3055–3071, 2018. 5
- [14] Yoseo Han, Leonard Sunwoo, and Jong Chul Ye. k-space deep learning for accelerated mri. *IEEE transactions on medical imaging*, 39(2):377–386, 2019. 5
- [15] 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*, pages 5197–5206, 2015. 6
- [16] Anna Jezierska, Caroline Chaux, Jean-Christophe Pesquet, and Hugues Talbot. An em approach for poisson-gaussian noise modeling. In *2011 19th European Signal Processing Conference*, pages 2244–2248. IEEE, 2011. 5
- [17] 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, 5, 6, 8
- [18] Kwanyoung Kim, Shakarim Soltanayev, and Se Young Chun. Unsupervised training of denoisers for low-dose CT reconstruction without full-dose ground truth. *IEEE Journal of Selected Topics in Signal Processing*, 14(6):1112–1125, 2020. 8
- [19] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 4, 5
- [20] Imre Risi Kondor. *Group theoretical methods in machine learning*, volume 2. PhD Thesis, 2008. 2
- [21] Risi Kondor and Shubhendu Trivedi. On the generalization of equivariance and convolution in neural networks to the action of compact groups. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 2747–2755. PMLR, 10–15 Jul 2018. 8
- [22] Yoann Le Montagner, Elsa D Angelini, and Jean-Christophe Olivo-Marin. An unbiased risk estimator for image denoising in the presence of mixed Poisson-Gaussian noise. *IEEE Transactions on Image processing*, 23(3):1255–1268, 2014. 4, 5
- [23] John M Lee. Introduction to smooth manifolds. In *Introduction to Smooth Manifolds*, pages 1–31. Springer, 2013. 2
- [24] Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, and Timo Aila. Noise2noise: Learning image restoration without clean data. In *International Conference on Machine Learning*, pages 2965–2974, 2018. 8
- [25] Florian Luisier, Thierry Blu, and Michael Unser. Image denoising in mixed Poisson-Gaussian noise. *IEEE Transactions on image processing*, 20(3):696–708, 2011. 3, 5
- [26] 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*, 58(6):1182–1195, 2007. 5
- [27] Stéphane Mallat. Group invariant scattering. *Communications on Pure and Applied Mathematics*, 65(10):1331–1398, 2012. 8
- [28] Christopher A Metzler, Arian Maleki, and Richard G Baraniuk. From denoising to compressed sensing. *IEEE Transactions on Information Theory*, 62(9):5117–5144, 2016. 3
- [29] Christopher A Metzler, Ali Mousavi, and Richard G Baraniuk. Learned D-AMP: principled neural network based compressive image recovery. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 1770–1781, 2017. 3

- [30] Christopher A Metzler, Ali Mousavi, Reinhard Heckel, and Richard G Baraniuk. Unsupervised learning with Stein’s unbiased risk estimator. *arXiv preprint arXiv:1805.10531*, 2018. 3, 8
- [31] Sreyas Mohan, Zahra Kadkhodaie, Eero P. Simoncelli, and Carlos Fernandez-Granda. Robust and interpretable blind image denoising via bias-free convolutional neural networks. In *ICLR*, 2020. 8
- [32] Arthur Pajot, Emmanuel de Bezenac, and Patrick Gallinari. Unsupervised adversarial image reconstruction. In *International Conference on Learning Representations*, 2019. 2
- [33] Sathish Ramani, Thierry Blu, and Michael Unser. Monte-Carlo SURE: A black-box optimization of regularization parameters for general denoising algorithms. *IEEE Transactions on image processing*, 17(9):1540–1554, 2008. 3
- [34] Martin Raphan and Eero Simoncelli. Least squares estimation without priors or supervision. *Neural computation*, 23:374–420, 02 2011. 3, 8
- [35] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 5
- [36] Leonid I Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: nonlinear phenomena*, 60(1-4):259–268, 1992. 1
- [37] Shakarim Soltanayev and Se Young Chun. Training deep learning based denoisers without ground truth data. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 3261–3271, 2018. 3, 8
- [38] Charles M Stein. Estimation of the mean of a multivariate normal distribution. *The annals of Statistics*, pages 1135–1151, 1981. 2, 3
- [39] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9446–9454, 2018. 1
- [40] 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
- [41] Magauiya Zhussip, Shakarim Soltanayev, and Se Young Chun. Extending Stein’s unbiased risk estimator to train deep denoisers with correlated pairs of noisy images. *Advances in neural information processing systems*, 32:1465–1475, 2019. 8
- [42] Magauiya Zhussip, Shakarim Soltanayev, and Se Young Chun. Training deep learning based image denoisers from undersampled measurements without ground truth and without image prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10255–10264, 2019. 3