

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 EQUIVARIANT SPLITTING: SELF-SUPERVISED LEARNING FROM INCOMPLETE DATA

Anonymous authors
Paper under double-blind review

ABSTRACT

Self-supervised learning for inverse problems allows to train a reconstruction network from noise and/or incomplete data alone. These methods have the potential of enabling learning-based solutions when obtaining ground-truth references for training is expensive or even impossible. In this paper, we propose a new self-supervised learning strategy devised for the challenging setting where measurements are observed via a single incomplete observation model. We introduce a new definition of equivariance in the context of reconstruction networks, and show that the combination of self-supervised splitting losses and equivariant reconstruction networks results in unbiased estimates of the supervised loss. Through a series of experiments on image inpainting, accelerated magnetic resonance imaging, sparse-view computed tomography, and compressive sensing, we demonstrate that the proposed loss achieves state-of-the-art performance in settings with highly rank-deficient forward models.

1 INTRODUCTION

Inverse problems are ubiquitous in many sensing and imaging applications. They are written as

$$y = Ax + \varepsilon \quad (1)$$

where $A \in \mathbb{R}^{m \times n}$ is the known forward matrix, $x \in \mathbb{R}^n$ is the ground truth image to be estimated, $y \in \mathbb{R}^m$ is the observed measurement vector, and $\varepsilon \in \mathbb{R}^m$ is the unknown noise, generally assumed to follow a Gaussian distribution. This model is suitable for many imaging modalities, including magnetic resonance imaging (MRI) for medical imaging (Zbontar et al., 2019), computed tomography (CT) (Withers et al., 2021), microscopy (Ragone et al., 2023), remote sensing (Fassnacht et al., 2024), and astronomical imaging (Vojtekova et al., 2021).

The number of linearly independent measurements is often smaller than the number of pixels in the target images due to physical and practical constraints. In this case, the forward matrix is rank-deficient, the forward process discards some of the information present in the target image. Further information about the target images is needed to solve the problem and different methods make different assumptions about the signal distribution.

Modern learning-based solvers generally obtain state-of-the-art performance by training on supervised pairs of ground-truth images and measurements. For certain applications, it is expensive or even impossible to obtain enough ground-truth data for supervised training (Belthangady & Royer, 2019). This is notably the case in astronomical imaging, microscopy and medical imaging.

Recent self-supervised methods overcome this limitation by learning to reconstruct without ground-truth data, relying only on a dataset of measurements, and they generally differ in the assumption they make on the forward model. For denoising problems where the forward matrix is the identity mapping, certain methods rely on knowing the exact noise distribution (Eldar, 2009; Pang et al., 2021; Monroy et al., 2025), others only assume it is entry-wise independent (Krull et al., 2019), while some make intermediate assumptions (Tachella et al., 2025a).

In settings where measurements are observed via multiple incomplete forward operators, such as accelerated MRI with masks varying across acquisitions or inpainting problems with missing pixels varying across images, the main approaches are splitting (Yaman et al., 2020; Millard & Chiew,

054 2023) and consistency across operators (Tachella et al., 2022). Splitting losses divide the measurements
 055 into input and target components, and are unbiased estimators of the supervised loss if there
 056 is enough diversity of operators in the dataset (Daras et al., 2023; Millard & Chiew, 2023).
 057

058 In the more challenging case of measurement data obtained via a single incomplete forward operator,
 059 the main self-supervised approach is equivariant imaging (Chen et al., 2021; 2022) which makes
 060 the assumption that the target distribution is invariant to a certain group of transformations including
 061 geometric transformations (Wang & Davies, 2024; Scanvic et al., 2025) and range transformations
 062 such as intensity scalings (Sechaud et al., 2024). Experimental results show that it can obtain com-
 063 petitive performances to supervised methods even though it does not require ground-truth references
 064 for training (Chen et al., 2022). However, training with EI is typically slower than supervised learn-
 065 ing, as it requires two to three evaluations of the model for every iteration, and can obtain subpar
 066 performances if the operator is highly incomplete.
 067

068 In this work, we propose equivariant splitting (ES), a new self-supervised method for learning from
 069 measurements obtained via a single forward operator that combines the invariance to transforma-
 070 tions assumption of equivariant imaging and the simplicity and computational efficiency of splitting
 071 methods. The key idea behind our method is the use of recent developments in equivariant architec-
 072 tures (Chaman & Dokmanić, 2021a; Puny et al., 2022) to design a training loss that performs implicit
 073 ground truth data augmentation without any transformation overhead. Our theoretical results show
 074 that our method yields (in expectation) the minimum mean squared error (MMSE) estimator as
 075 long as the model is expressive enough, which is not guaranteed for equivariant imaging and pre-
 076 vious splitting methods. **Additionally, our experiments demonstrate state-of-the-art performance on**
 077 **a wide range of self-supervised imaging problems including compressed sensing, image inpainting,**
 078 **accelerated MRI and sparse-view CT.** This work is additional evidence that architectural constraints
 079 built upon equivariance are a powerful tool to solve ill-posed imaging inverse problems.
 080

081 Our contributions are the following:
 082

- 083 1. We propose a new definition for equivariance in inverse problems, and propose architec-
 084 tures that satisfy this property, including unrolled architectures.
- 085 2. We propose a new self-supervised loss that leverages equivariant networks (according to
 086 our definition above) whose global minimizer is the gold standard MMSE estimator under
 087 the assumption of an invariant signal distribution.
- 088 3. **We demonstrate the performance of our method on a wide range of inverse problems in-
 089 cluding compressed sensing, inpainting, accelerated MRI and sparse-view computed to-
 090 mography.**

091 2 RELATED WORK

092 **Measurement splitting** Various self-supervised losses consist of dividing measurement vectors into
 093 two components, one used as input and the other as target. It has been used to solve oversam-
 094 pled single-operator inverse problems including full-view CT (Hendriksen et al., 2020), and also
 095 undersampled multi-operator inverse problems with theoretical guarantees of producing estimates
 096 equivalent to supervised estimates in expectation, notably accelerated MRI and image inpainting
 097 with varying masks (Daras et al., 2023; Millard & Chiew, 2023). To the best of our knowledge,
 098 this work is the first to extend these methods to the more challenging single-operator undersampled
 099 setting, which notably includes sparse-view CT and accelerated MRI (with fixed mask).
 100

101 **Equivariant imaging** It is possible to learn from incomplete data associated with a single degra-
 102 dation operator, as long as the underlying signal distribution is invariant to a group of transforma-
 103 tions (Tachella et al., 2023). Equivariant imaging (Chen et al., 2021; 2022) assumes that the dis-
 104 tribution of clean images remains unchanged under certain transformations, including transla-
 105 tions, rotations and flips, and introduces a training loss that enforces the equivariance to these trans-
 106 formations of the entire measurement-reconstruction process, thereby constraining the set of learnable
 107 models. It has been used effectively to solve various inverse problems using adequate groups of
 108 transformations (Wang & Davies, 2024; Sechaud et al., 2024), but it is computationally expensive
 109 due to the two to three network evaluations. Moreover, the EI loss is not necessarily an unbiased
 110 estimator of the supervised loss, and it is unclear whether this approach recovers the optimal MMSE
 111 estimator in expectation.
 112

Equivariant neural networks The design of equivariant networks is an active research topic and many different approaches exist. Some rely on data augmentation to make neural networks more equivariant to rotations and flips, while others rely on averaging on the group of transformations at test time (Rivera et al., 2021; Puny et al., 2022; Kaba et al., 2023; Sannai et al., 2024). A third line of work focuses on architectures that are equivariant by design: Cohen & Welling (2016; 2017) propose convolutional layers that are equivariant to rotations and flips, and other works also rely on the design of more equivariant layers to improve translation-equivariance. Zhang (2019); Chaman & Dokmanić (2021b) propose equivariant pooling and downsampling/upsampling layers, and Karras et al. (2021); Michaeli et al. (2023) equivariant non-linearities and activation layers. Closer to our work, in the specific setting of inverse problems, Celledoni et al. (2021) propose the use of equivariant denoising blocks within unrolled architectures, and Terris et al. (2024) similarly propose to render plug-and-play denoisers more equivariant by averaging over transformations at test time. We take these analyses further by providing a clear definition of equivariance in inverse problems and showing which popular posterior estimators and related architectures verify these properties.

3 BACKGROUND

Solving the inverse problem in eq. (1) amounts to designing a reconstruction function $f(\mathbf{y}, \mathbf{A}) \approx \mathbf{x}$ estimating a ground truth signal $\mathbf{x} \in \mathbb{R}^n$ from its measurement vector $\mathbf{y} \in \mathbb{R}^m$ and forward matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$. In practice, it is often implemented as a neural network parametrized by a set of weights. Supervised methods assume the existence of a finite dataset containing pairs of ground truth signals and measurements $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i \in \mathcal{I}}$ that are used to learn the reconstruction function $f(\mathbf{y}, \mathbf{A})$. The main approach is to minimize a training loss equal to the mean squared error (Ongie et al., 2019)

$$\min_f \left\{ \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} \mathcal{L}_{\text{SUP}}(\mathbf{x}_i, \mathbf{y}_i, \mathbf{A}, f) \right\}, \quad \mathcal{L}_{\text{SUP}}(\mathbf{x}, \mathbf{y}, \mathbf{A}, f) = \|f(\mathbf{y}, \mathbf{A}) - \mathbf{x}\|^2. \quad (2)$$

While this approach obtains state-of-the-art performance, it cannot be used in the absence of ground-truth data. Self-supervised methods overcome this limitation using a finite dataset containing only measurements $\{\mathbf{y}_i\}_{i \in \mathcal{I}}$, and a training loss $\mathcal{L}(\mathbf{y}, \mathbf{A}, f)$ that need no ground truth data to be evaluated, and which is designed to approximate well the supervised objective in eq. (2).

In denoising problems ($\mathbf{A} = \mathbf{I}$) with Gaussian noise of known variance, Recorrupted2Recorrupted (R2R) (Pang et al., 2021) and SURE (Metzler et al., 2020) provide unbiased estimators of the supervised loss. The R2R loss is computed by adding synthetic Gaussian noise $\omega \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$ to the measurements, creating input-target pairs as $(\mathbf{y} + \alpha\omega, \mathbf{y} - \frac{\omega}{\alpha})$ for some $\alpha \in (0, +\infty)$. However, if measurements are observed by an incomplete operator \mathbf{A} with a non-trivial nullspace, these losses fail to learn in the nullspace, approximating $f(\mathbf{y}, \mathbf{A}) = \mathbf{A}^\dagger \mathbf{A} \mathbb{E}_{\mathbf{x}|\mathbf{y}, \mathbf{A}} \{\mathbf{x}\} + v(\mathbf{y}, \mathbf{A})$, with v being any function taking values in the nullspace of \mathbf{A} (Chen et al., 2021).

In the rest of this section, we present the two main self-supervised losses that can learn beyond the nullspace, measurement splitting in Section 3.1 and equivariant imaging in Section 3.2. By combining their main ideas, we obtain our new self-supervised loss introduced in Section 4.

3.1 MEASUREMENT SPLITTING

One strategy to address the limitations imposed by the nullspace of a single operator is to employ multiple operators (Millard & Chiew, 2023). The key idea is that different operators generally do not share the same nullspace; thus, observing measurements through the image spaces of multiple operators allows access to the whole space \mathbb{R}^n . Formally, they assume that measurements are obtained according to $\mathbf{y} \sim p(\mathbf{y}|\mathbf{A}\mathbf{x})$ where the measurement operator \mathbf{A} is itself drawn from a distribution $p(\mathbf{A})$, and differ for each acquisition.

Splitting losses divide the measurements into two components $\mathbf{y} = [\mathbf{y}_1^\top, \mathbf{y}_2^\top]^\top$ with corresponding operators $\mathbf{A} = [\mathbf{A}_1^\top, \mathbf{A}_2^\top]^\top$, where $\mathbf{A}_1 \in \mathbb{R}^{m_1 \times n}$ and $\mathbf{A}_2 \in \mathbb{R}^{m_2 \times n}$ with $m_1 + m_2 = m$. The network is then trained to predict the entire measurements vector \mathbf{y} from only one of its two components \mathbf{y}_1 , using the training loss¹

$$\mathcal{L}_{\text{SPLIT}}(\mathbf{y}, \mathbf{A}, f) = \mathbb{E}_{\mathbf{y}_1, \mathbf{A}_1|\mathbf{y}, \mathbf{A}} \{\|\mathbf{A}f(\mathbf{y}_1, \mathbf{A}_1) - \mathbf{y}\|^2\}, \quad (3)$$

¹We use the notation $\mathbb{E}_{a|b}\{g(a)\}$ for $\int_a g(a)p(a|b)da$.

162 where $p(\mathbf{y}_1, \mathbf{A}_1 \mid \mathbf{y}, \mathbf{A})$ is a random splitting distribution, which is chosen on a per-problem basis.
 163 The loss encourages the model to learn in the nullspace of each operator by predicting the
 164 unobserved part. In practice, the expectation in eq. (3) is estimated using a single split \mathbf{y}_1 for each
 165 training batch.
 166

167 3.2 EQUIVARIANT IMAGING

169 Equivariant imaging (EI) relies on the assumption that the distribution of images is invariant under
 170 a group of transformations $\mathbf{T}_g \in \mathbb{R}^{n \times n}$, $g \in \mathcal{G}$, to learn beyond the nullspace of measurements ob-
 171 tained via a *single* operator (Chen et al., 2021; 2022). In this setting, the reconstruction model is
 172 expected to be able to estimate ground truth images \mathbf{x} as well as their transformations $\mathbf{T}_g \mathbf{x}$ in a co-
 173 herent manner, i.e., such that the entire measurement-reconstruction pipeline $f(\mathbf{Ax})$ is equivariant
 174 with respect to the transformations $f(\mathbf{AT}_g \mathbf{x}, \mathbf{A}) = \mathbf{T}_g f(\mathbf{Ax}, \mathbf{A})$. In order to achieve this, they pro-
 175 pose a self-supervised loss which consists in a traditional measurement consistency term (replaced
 176 by SURE in the presence of noise), along with an equivariance-promoting term

$$177 \mathcal{L}_{\text{EI}}(\mathbf{y}, \mathbf{A}, f) = \|\mathbf{Af}(\mathbf{y}, \mathbf{A}) - \mathbf{y}\|^2 + \lambda \mathbb{E}_g \{ \|\mathbf{T}_g f(\mathbf{y}, \mathbf{A}) - f(\mathbf{AT}_g \mathbf{y}, \mathbf{A})\|^2 \}, \quad (4)$$

178 where $\lambda > 0$ is a trade-off coefficient. Even though it has been shown to be particularly effective on
 179 a wide variety of problems (Wang & Davies, 2024; Sechaud et al., 2024), it typically requires from
 180 two to three evaluations of the neural network which makes it very computationally expensive (Xu
 181 et al., 2025). Moreover, the equivariant loss is only effective at enforcing equivariance when the
 182 learned estimator achieves almost perfect reconstructions, $f(\mathbf{y}, \mathbf{A}) \approx \mathbf{x}$ (Chen et al., 2021), which
 183 is not the case for very ill-conditioned problems and which leads to the method having multiple
 184 possible solutions in general.
 185

186 4 METHOD

188 In this section, we present our method that combines measurement splitting and EI introduced in
 189 Section 3. In order to learn from incomplete measurements obtained via a single operator, we rely
 190 on the same assumption of EI:

191 **Assumption 1.** *The distribution of ground truth images $p(\mathbf{x})$ is invariant to the transformations*
 192 $\{\mathbf{T}_g\}_{g \in \mathcal{G}}$

$$193 p(\mathbf{T}_g \mathbf{x}) = p(\mathbf{x}), \forall g \in \mathcal{G}, \forall \mathbf{x} \in \mathbb{R}^n. \quad (5)$$

195 The set of transformations is a design choice of the method to be chosen on a per-problem basis.
 196 Corollary 1 helps to choose them for specific problems, and we use it in the experiments. **This**
 197 **assumption applies in many different settings, natural image distributions are generally invariant to**
 198 **rigid transformations (translations, rotations, flips, scalings), especially microscopic, aerial or**
 199 **remote sensing images which have no privileged orientation.** Moreover, our experiments show that
 200 our method performs well even in the presence of approximate invariance, e.g., for medical im-
 201 ages. Refer to Appendix A.2 for more details about how to choose the transformations for a given
 202 application.
 203

204 In Section 4.1, we present our proposed loss and state its optimality under mild assumptions in The-
 205 orem 1 and Proposition 1. In Section 4.2, we present a new definition of equivariance for reconstruc-
 206 tion functions and state sufficient conditions for common architectures to be equivariant in The-
 207 orem 2. In Section 4.3, we finally present a computational synergy between our loss and these
 208 equivariant architectures stated in Theorem 3. **See Appendix A.3 for a description of the end-to-end**
 209 **algorithm, and Appendix C for the detailed proofs.**
 210

211 4.1 PROPOSED LOSS

212 Under Assumption 1, the measurements can be understood in a different way than they are tradi-
 213 tionally. Indeed, measurements \mathbf{y} are generally thought of as being associated to the ground truth
 214 image \mathbf{x} and the forward matrix \mathbf{A} , but they can equally be understood as being associated to the
 215 virtual ground truth image $\mathbf{x}_g = \mathbf{T}_g^{-1} \mathbf{x}$ and the virtual forward matrix $\mathbf{A}_g = \mathbf{AT}_g$, with

$$216 \mathbf{y} = \mathbf{Ax} + \varepsilon = \mathbf{AT}_g \mathbf{T}_g^{-1} \mathbf{x} + \varepsilon = \mathbf{A}_g \mathbf{x}_g + \varepsilon. \quad (6)$$

216 The implicit multi-operator structure (with operators $\{\mathbf{A}\mathbf{T}_g\}_{g \in \mathcal{G}}$) of the problem hints that we should
 217 be able to leverage the splitting approaches presented in Section 3. Moreover, since we use the same
 218 invariance assumption as EI, we can combine the two approaches to obtain equivariant splitting
 219 (ES), a new self-supervised loss \mathcal{L}_{ES} that has the advantages of both methods.
 220

221 **Noiseless measurements** The ES self-supervised loss is expressed as
 222

$$\mathcal{L}_{\text{ES}}(\mathbf{y}, \mathbf{A}, f) \triangleq \mathbb{E}_g \{\mathcal{L}_{\text{SPLIT}}(\mathbf{y}, \mathbf{A}\mathbf{T}_g, f)\} \quad (7)$$

$$= \mathbb{E}_g \{\mathbb{E}_{\mathbf{y}_1, \mathbf{A}_1 | \mathbf{y}, \mathbf{A}\mathbf{T}_g} \{\|\mathbf{A}\mathbf{T}_g f(\mathbf{y}_1, \mathbf{A}_1) - \mathbf{y}\|^2\}\} \quad (8)$$

223 where $\mathbf{A}_1 \sim p(\mathbf{A}_1 | \mathbf{A}\mathbf{T}_g) \triangleq p(\mathbf{A}_1 | g)$ is a random splitting of $\mathbf{A}\mathbf{T}_g$.
 224

225 **Theorem 1.** *In the case of noiseless measurements with $p(\mathbf{x})$ \mathcal{G} -invariant (Assumption 1), if the
 226 matrix $\mathbf{Q}_{\mathbf{A}_1} \triangleq \mathbb{E}_{g | \mathbf{A}_1} \{(\mathbf{A}\mathbf{T}_g)^\top \mathbf{A}\mathbf{T}_g\} \in \mathbb{R}^{n \times n}$ has full rank for some split \mathbf{A}_1 , then the splitting
 227 method yields the same MMSE-optimal reconstructions as the supervised method, i.e.,*

$$f^*(\mathbf{y}_1, \mathbf{A}_1) = \mathbb{E}_{\mathbf{x} | \mathbf{y}_1, \mathbf{A}_1} \{\mathbf{x}\}. \quad (9)$$

228 While it is sufficient for the matrix $\mathbf{Q}_{\mathbf{A}_1}$ to be invertible for the reconstructions to be almost or
 229 exactly optimal, in practice the spectrum (number of its non-negligible eigenvalues) of the matrix
 230 determines how close to optimal they are. In the experiments, we use a single (random) split per
 231 batch element and we average the reconstructions corresponding to 10 splits at inference.
 232

233 **Proposition 1.** *If the matrix $\bar{\mathbf{Q}}_{\mathbf{A}} \triangleq \mathbb{E}_{\mathbf{A}_1 | \mathbf{A}} \{\mathbf{Q}_{\mathbf{A}_1}\} \in \mathbb{R}^{n \times n}$ is invertible and f minimizes
 234 $\mathbb{E}_{\mathbf{y}} \{\mathcal{L}_{\text{ES}}(\mathbf{y}, \mathbf{A}, f)\}$. Then the reconstruction function*

$$\bar{f}(\mathbf{y}, \mathbf{A}) \triangleq \mathbb{E}_{\mathbf{y}_1, \mathbf{A}_1 | \mathbf{y}, \mathbf{A}} \{\bar{\mathbf{Q}}_{\mathbf{A}}^{-1} \mathbf{Q}_{\mathbf{A}_1} f(\mathbf{y}_1, \mathbf{A}_1)\} \quad (10)$$

235 satisfies

$$\bar{f}(\mathbf{y}, \mathbf{A}) = \mathbb{E}_{\mathbf{y}_1, \mathbf{A}_1 | \mathbf{y}, \mathbf{A}} \{\bar{\mathbf{Q}}_{\mathbf{A}}^{-1} \mathbf{Q}_{\mathbf{A}_1} \mathbb{E}_{\mathbf{x} | \mathbf{y}_1, \mathbf{A}_1} \{\mathbf{x}\}\}, \quad (11)$$

236 where eq. (11) is a convex combination of MMSE estimators for different splittings.
 237

238 In practice, often neither $\bar{\mathbf{Q}}_{\mathbf{A}}$ nor $\mathbf{Q}_{\mathbf{A}_1}$ can be computed in closed-form, and we use a non-weighted
 239 average over random splittings

$$\bar{f}(\mathbf{y}, \mathbf{A}) := \frac{1}{J} \sum_{j=1}^J f(\mathbf{y}_1^{(j)}, \mathbf{A}_1^{(j)}) \text{ with } (\mathbf{y}_1^{(j)} \mathbf{A}_1^{(j)}) \sim p(\mathbf{y}_1 \mathbf{A}_1 | \mathbf{y}, \mathbf{A}\mathbf{T}_g) \quad (12)$$

240 where g is chosen randomly over the group of transformations for each split.
 241

242 As with EI, the forward operator should not be equivariant with respect to the choice of transformations,
 243 in order to learn beyond the nullspace of the operator:
 244

245 **Corollary 1.** *In order for the matrices $\mathbf{Q}_{\mathbf{A}_1}$ or $\bar{\mathbf{Q}}_{\mathbf{A}}$ to have full rank, it is necessary that \mathbf{A} is not
 246 equivariant:*

$$\exists g \in \mathcal{G}, \mathbf{A}\mathbf{T}_g \neq \mathbf{T}_g \mathbf{A}. \quad (13)$$

247 **Noisy measurements** The ES loss can be split into two separate terms, one enforcing measure-
 248 ment consistency, and the other prediction accuracy:
 249

$$\mathcal{L}_{\text{ES}}(\mathbf{y}, \mathbf{A}, f) = \mathbb{E}_g \{\mathbb{E}_{\mathbf{y}_1, \mathbf{A}_1 | \mathbf{y}, \mathbf{A}\mathbf{T}_g} \{\|\mathbf{A}_1 f(\mathbf{y}_1, \mathbf{A}_1) - \mathbf{y}_1\|^2 + \|\mathbf{A}_2 f(\mathbf{y}_1, \mathbf{A}_1) - \mathbf{y}_2\|^2\}\},$$

250 where \mathbf{A}_1 and \mathbf{A}_2 are a splitting of $\mathbf{A}\mathbf{T}_g$. If the measurements are noisy, the first term can be re-
 251 placed by a self-supervised denoising loss. In particular, if measurements are corrupted by Gaussian
 252 noise of standard deviation σ , we replace the first term by the R2R loss, yielding:
 253

$$\mathcal{L}_{\text{G-ES}}(\mathbf{y}, \mathbf{A}, f) = \mathbb{E}_{g, \mathbf{y}_1, \mathbf{A}_1, \omega | \mathbf{y}, \mathbf{A}\mathbf{T}_g} \left\{ \left\| \mathbf{A}_1 f(\mathbf{y}_1 + \alpha\omega, \mathbf{A}_1) - \left(\mathbf{y}_1 - \frac{\omega}{\alpha}\right) \right\|^2 \right. \\ \left. + \|\mathbf{A}_2 f(\mathbf{y}_1 + \alpha\omega, \mathbf{A}_1) - \mathbf{y}_2\|^2 \right\}$$

270 with $\omega \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$ and a hyper-parameter $\alpha \in (0, +\infty)$. Since R2R provides an unbiased
 271 estimate of the clean measurement consistency term (Pang et al., 2021), we can apply Theorem 1 to
 272 show that minimizing this loss (in expectation) also results in MMSE estimators (if the conditions
 273 on Q_{A_1} or Q_A are verified). In the case of non-Gaussian noise, the R2R loss can be replaced by
 274 its non-Gaussian extension (Monroy et al., 2025). As with random splits, the expectation over ω is
 275 computed using a random realization per batch. At test time, we modify eq. (12) to average over
 276 both splits and synthetic noise additions.

277 4.2 EQUIVARIANT RECONSTRUCTORS

278 The ES loss requires a model evaluation for every mask and transformation. We show that, instead of
 279 sampling a random transformation each evaluation, imposing architectural equivariance constraints
 280 removes the need to explicitly compute the transforms.

281 Image-to-image functions $\phi(\mathbf{x})$ are equivariant if they satisfy (Cohen & Welling, 2016)

$$282 \phi(\mathbf{T}_g \mathbf{x}) = \mathbf{T}_g \phi(\mathbf{x}), \forall \mathbf{x} \in \mathbb{R}^n, \forall g \in \mathcal{G}. \quad (14)$$

283 In this work, we introduce an extension of this definition to reconstruction functions $f(\mathbf{y}, \mathbf{A})$. To
 284 the best of our knowledge, this is the first work that introduces this definition.

285 **Definition 1.** *We say that the reconstruction function $f(\mathbf{y}, \mathbf{A})$ is an equivariant reconstructor if*

$$286 f(\mathbf{y}, \mathbf{A} \mathbf{T}_g) = \mathbf{T}_g^{-1} f(\mathbf{y}, \mathbf{A}), \forall \mathbf{y} \in \mathbb{R}^m, \forall g \in \mathcal{G}, \forall \mathbf{A} \in \mathbb{R}^{m \times n}. \quad (15)$$

287 This property is very general and the class of classical reconstruction functions that satisfy it is large.

288 **Theorem 2.** *The reconstruction functions defined in points below are all equivariant as in eq. (15).*

289 1. **Artifact removal network.** For a denoiser $\phi(\mathbf{x})$ equivariant in the sense of eq. (14),

$$290 f(\mathbf{y}, \mathbf{A}) = \phi(\mathbf{A}^\top \mathbf{y}), \text{ or } f(\mathbf{y}, \mathbf{A}) = \phi(\mathbf{A}^\dagger \mathbf{y}). \quad (16)$$

291 2. **Unrolled network.** For $\phi(\mathbf{x})$ equivariant, any $\gamma \in \mathbb{R}$ and data fidelity $d(\mathbf{A}\mathbf{x}, \mathbf{y})$, with

$$292 \mathbf{x}_0 = \mathbf{0}, \quad \mathbf{x}_{k+1} = \phi(\mathbf{x}_k - \gamma \nabla_{\mathbf{x}_k} d(\mathbf{A}\mathbf{x}_k, \mathbf{y})) \quad (17)$$

293 for $k = 0, \dots, L-1$ and $f(\mathbf{y}, \mathbf{A}) = \mathbf{x}_L$.

294 3. **Reynolds averaging.** For a possibly non-equivariant reconstructor $r(\mathbf{y}, \mathbf{A})$, with

$$295 f(\mathbf{y}, \mathbf{A}) = \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} \mathbf{T}_g r(\mathbf{y}, \mathbf{A} \mathbf{T}_g). \quad (18)$$

296 4. **Maximum a posteriori (MAP).** For a distribution $p(\mathbf{x})$ invariant as in eq. (5), with

$$297 f(\mathbf{y}, \mathbf{A}) = \operatorname{argmax}_{\mathbf{x} \in \mathbb{R}^n} \left\{ p(\mathbf{x} \mid \mathbf{y}, \mathbf{A}) \right\}. \quad (19)$$

298 5. **Minimum mean squared error (MMSE).** For a distribution $p(\mathbf{x})$ invariant as in eq. (5),

$$299 f(\mathbf{y}, \mathbf{A}) = \mathbb{E}_{\mathbf{x} \mid \mathbf{y}, \mathbf{A}} \{ \mathbf{x} \}. \quad (20)$$

300 For additional motivation and details about these reconstructor architectures, see Appendix A.

301 4.3 EFFICIENT LOSS EVALUATION WITH EQUIVARIANT RECONSTRUCTORS

302 For equivariant reconstructors, the ES loss in eq. (7) reduces to the splitting loss in eq. (3).

303 **Theorem 3.** *If $f(\mathbf{A}, \mathbf{x})$ is an equivariant reconstructor, then ES is equivalent to the splitting loss*

$$304 \mathcal{L}_{\text{ES}}(\mathbf{y}, \mathbf{A}, f) = \mathcal{L}_{\text{SPLIT}}(\mathbf{y}, \mathbf{A}, f). \quad (21)$$

305 We emphasize that the condition for a reconstructor to be equivariant is different from the condition
 306 enforced by the EI loss, i.e., $f(\mathbf{A} \mathbf{T}_g \mathbf{x}, \mathbf{A}) = \mathbf{T}_g f(\mathbf{A}\mathbf{x}, \mathbf{A})$. They are only equivalent if the
 307 reconstruction function is a perfect one-to-one mapping over all possible images.

308 In our experiments, we build equivariant reconstructors using i) artifact removal networks with a
 309 translation equivariant UNet denoiser Chaman & Dokmanić (2021b) and ii) unrolled networks with
 310 a denoiser architecture equivariant to rotations and flips via averaging (Sannai et al., 2021).

324

5 EXPERIMENTS

326 We assess the effectiveness of the proposed self-supervised loss using experiments conducted on
 327 different inverse problems. For each experiment, we train a model corresponding to our method as
 328 well as baseline methods and compare their performance. The inverse problems we consider are 1)
 329 **inpainting**, 2) **compressive sensing**, 3) **accelerated MRI** and 4) **sparse-view CT**. We also validate our
 330 theoretical predictions by testing the effect of using an equivariant architecture. In Section 5.1 we
 331 detail our experiments on compressive sensing, in Section 5.2 on image inpainting, in Section 5.3
 332 on accelerated MRI and in Section 5.4 on sparse-view CT, and in Section 5.5 we present an abla-
 333 tion study on the effect of equivariant architectures. For additional details about the experiments,
 334 see Appendix B.

335 In each experiment, we compare against multiple baselines: a supervised baseline using the super-
 336 vised loss described in eq. (2), the EI (Chen et al., 2021) baseline described in eq. (4), a measurement
 337 consistency baseline (MC and SURE) Eldar (2009) and a learning-free baseline which are either the
 338 measurements directly, or their image under the adjoint or the pseudo-inverse of the forward oper-
 339 ator. We use the same architecture and the same training procedure for every method to ensure a
 340 fair comparison. We report the peak signal-to-noise ratio (PSNR) and the structural similarity index
 341 measure (SSIM) (Wang et al., 2004) of the final reconstructions for the different methods. They are
 342 distortion metrics indicating how close the reconstructions are to the ground truth images. We do
 343 not include perception metrics which are known to be at odds with them (Blau & Michaeli, 2018).
 344 In each case, we also compute equivariance metrics (EQUIV) for translations or rotations and flips.

345 We design a model architecture from the principles introduced in Section 4, with a variant equivari-
 346 ant to shifts, one equivariant to rotations and flips, and one without equivariance. It uses an existing
 347 unrolled architecture (Aggarwal et al., 2019) with a prior step implemented as a standard UNet (Ron-
 348 neberger et al., 2015), or that of Chaman & Dokmanić (2021b) to enforce the equivariance to shifts.
 349 It also optionally uses Reynolds' averaging to enforce the equivariance to rotations and flips. We
 350 use a single equivariant variant dictated by Corollary 1 for each problem, the shift-equivariant one
 351 for compressive sensing and inpainting, and that equivariant to rotations and flips for MRI and CT.

352

5.1 COMPRESSIVE SENSING

354 The 28×28 ground truth images are obtained from the MNIST dataset and are measured without
 355 additional noise through compressive matrices $\mathbf{A} \in \mathbb{R}^{m \times n}$, with $m < n$ varying from training
 356 to training to assess the impact of the compression rate. These matrices are obtained by sampling
 357 $A_{i,j} \sim \mathcal{N}(0, 1/m)$ for $i = 1, \dots, m$ and $j = 1, \dots, m$, where $n = 28 \times 28$. Figure 1 shows
 358 the performance of the different methods for the different compression rates. Our method performs
 359 almost as well as the supervised baseline, while the equivariant imaging baseline performs close to
 360 the supervised baseline only for higher compression rates.

361

5.2 IMAGE INPAINTING

363 The dataset consists of 128×128 images from DIV2K (Agustsson & Timofte, 2017) measured
 364 without additional noise through a single subsampling matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ selecting about 30% of
 365 the pixels. This matrix is obtained by sampling $a_1, \dots, a_n \sim \mathcal{B}(0.3)$ where $n = 3 \times 128 \times 128$
 366 and letting $A_{i,j} = \delta_{j_i,j}$ for $i = 1, \dots, m$ and $j = 1, \dots, n$, where $m \approx 0.3n$ is the number of
 367 nonzero values in a , and j_i is the i -th index in a corresponding to a nonzero value. Among the 900
 368 images in the dataset, 800 are used for training while the remaining 100 are used for testing. For the
 369 supervised method, we use different crops at each evaluation. Table 2 and Figure 2 show that ES
 370 performs almost as well as the supervised baseline, and better than EI.

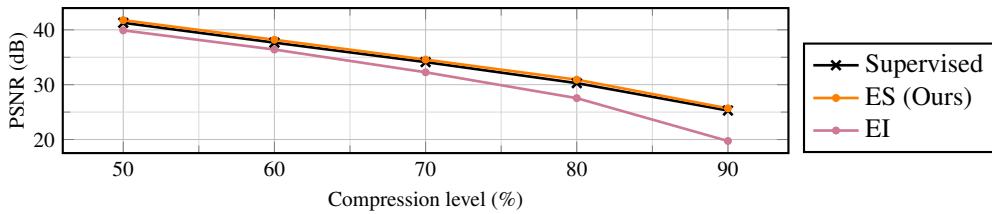
371

5.3 MAGNETIC RESONANCE IMAGING

372 The dataset contains 320×320 images from FastMRI (Zbontar et al., 2019) subsampled in the
 373 Fourier domain by a single binary mask corresponding to an acceleration of 8, as well as Gaus-
 374 sian noise with a standard deviation of 0.005 corresponding to a signal-to-noise ratio (SNR) of 40
 375 dB. Mathematically, the forward operator $\mathbf{A} \in \mathbb{R}^{m \times n}$ is expressed as $\mathbf{A} = \mathbf{MF}$ where $\mathbf{F} \in \mathbb{R}^{n \times n}$
 376 denotes the $n \times n$ discrete Fourier transform matrix and where $\mathbf{M} \in \mathbb{R}^{m \times n}$ is the subsampling mask

378
 379 **Table 1: Medical imaging results.** ES (ours) performs better than EI, SURE and MC (baselines),
 380 while performing almost as well as the supervised baseline in reconstruction quality (PSNR, SSIM)
 381 and measured equivariance (EQUIV). In **bold**, the best self-supervised metrics (avg \pm st.d.).

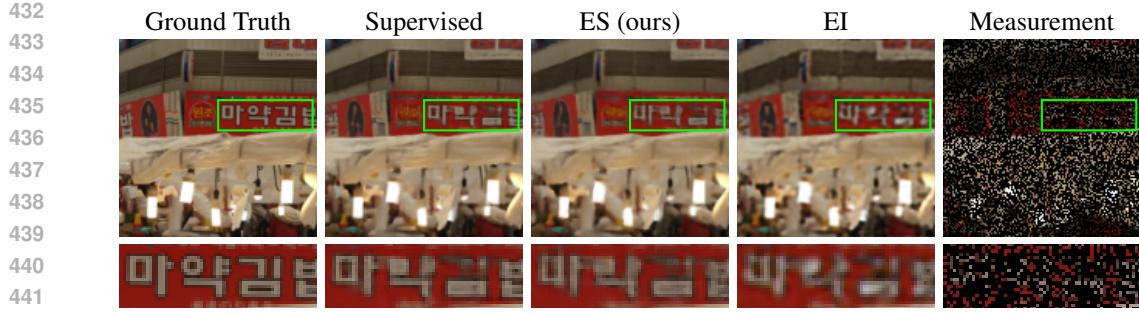
MRI ($\times 8$ Accel., 40 dB SNR)			
Method	PSNR \uparrow	SSIM \uparrow	EQUIV \uparrow
Supervised	28.74 ± 2.81	0.6445 ± 0.1094	31.71 ± 2.83
ES (Ours)	28.54 ± 2.75	0.6195 ± 0.1188	31.53 ± 2.74
EI	27.88 ± 2.64	0.5731 ± 0.1299	30.79 ± 2.64
SURE	24.45 ± 1.86	0.5479 ± 0.0740	27.35 ± 1.90
IDFT	23.62 ± 1.90	0.5052 ± 0.0900	25.99 ± 1.94
Real MRI measurements ($\times 8$ Accel.)			
Method	PSNR \uparrow	SSIM \uparrow	EQUIV \uparrow
Supervised	28.81 ± 2.85	0.6480 ± 0.1103	31.81 ± 2.84
ES (Ours)	28.30 ± 2.64	0.6151 ± 0.1179	31.29 ± 2.62
EI	27.88 ± 2.61	0.5740 ± 0.1290	30.80 ± 2.61
MC	23.63 ± 1.90	0.5061 ± 0.0904	26.00 ± 1.94
IDFT	23.63 ± 1.90	0.5060 ± 0.0904	26.00 ± 1.94
CT (50 views, 50 dB SNR)			
Method	PSNR \uparrow	SSIM \uparrow	EQUIV \uparrow
Supervised	33.99 ± 2.48	0.8819 ± 0.0585	34.00 ± 2.49
ES (Ours)	32.62 ± 2.16	0.8570 ± 0.0596	32.60 ± 2.17
EI	28.61 ± 1.28	0.7400 ± 0.0466	28.61 ± 1.29
FBP	25.59 ± 0.69	0.4805 ± 0.0363	25.59 ± 0.70



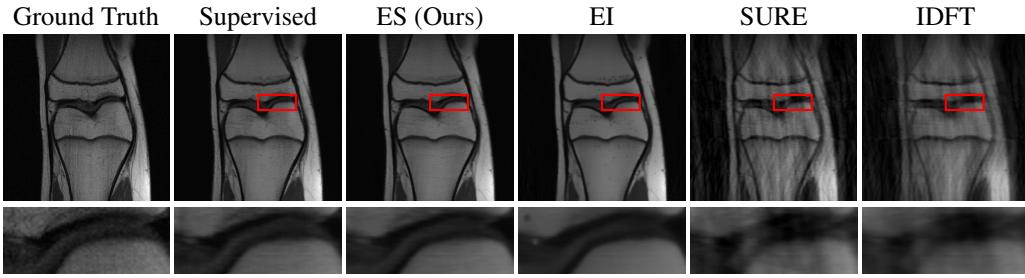
416 **Figure 1: Compressive sensing results.** ES (ours) performs similarly as the supervised baseline,
 417 unlike EI (baseline) whose performance gap widens with higher compression levels.
 418
 419
 420

421 **Table 2: Inpainting results.** ES (ours) performs better than EI (baseline), both in terms of recon-
 422 struction quality (PSNR, SSIM) and measured equivariance (EQUIV), while performing competi-
 423 tively against the supervised baseline. In **bold**, the best self-supervised metrics (avg \pm st.d.).

Method	PSNR \uparrow	SSIM \uparrow	EQUIV \uparrow
Supervised	28.46 ± 2.97	0.8982 ± 0.0411	28.46 ± 2.97
ES (Ours)	27.45 ± 2.86	0.8737 ± 0.0461	27.46 ± 2.85
EI	25.89 ± 2.65	0.8332 ± 0.0521	25.89 ± 2.65
MC	8.22 ± 2.47	0.0983 ± 0.0551	8.22 ± 2.47
Incomplete image	8.22 ± 2.47	0.0973 ± 0.0542	N/A



443 Figure 2: **Sample reconstructions for image inpainting.** ES (ours) produces images perceptually
444 closer to the supervised baseline than EI (baseline) which appears blurry.



456 Figure 3: **Sample reconstructions for MRI ($\times 8$ Accel., 40 dB SNR).** Unlike EI (baseline) which
457 suffers from dot-shaped artifacts, ES (ours) is perceptually closer to the supervised baseline. In line
458 with the theoretical predictions, SURE and IDFT (baselines) fail to recover information beyond the
459 observed frequencies.

460
461
462
463 defined as $M_{i,j} = \delta_{j_i,j}$ for $i = 1, \dots, m$ and $j = 1, \dots, n$, where j_i denotes the i -th component in
464 \mathbb{R}^n corresponding to a pixel in one of the subsampled vertical lines in a random Gaussian mask. Out
465 of the 973 images in the full dataset, 900 are selected for training and the remaining 73 are used for
466 testing. We also test our method on a noise dominated setting using a different mask corresponding
467 to an acceleration of 6, with a higher noise level of 0.1 corresponding to an SNR of only 10 dB,
468 see these results in Appendix B.4. The variant of our proposed loss we use in this experiment is
469 the R2R one introduced in Section 4 with $\alpha = 0.5$. Table 1 shows the performance of the different
470 methods on the test set. Table 3 shows the synergy of our method with the equivariant architecture.
471 Figure 3 shows sample reconstructions from the trained models.

472 We also evaluate our method on real MRI measurements from FastMRI instead of synthetic
473 measurements. We normalize the k-space data that are sampled on different grids for different scans by
474 resampling them using aliasing-free sinc interpolation on the 320×320 sampling grid of the ground
475 truth scans. The resulting k-space data are subsampled using the same $\times 8$ acceleration mask used
476 in the setting with synthetic measurements. Table 1 shows that our method performs competitively.

478 5.4 COMPUTED TOMOGRAPHY

479
480 The dataset consists in pairs of ground truth CT scans and corresponding sinograms. In order to
481 obtain the CT scans, we resize to 256×256 pixels and clip between -1,000 and 1,000 HUs the
482 ground truth scans from the LIDC-IDRI dataset (Armato III et al., 2011). For each scan, we compute
483 the corresponding sinogram using a discrete Radon transform with 50 views and additive white
484 Gaussian noise with a standard deviation of 0.001 corresponding to a SNR of about 50 dB. The
485 resulting 1,010 pairs are further split into 900 training pairs and 110 test pairs. Table 1 shows that
486 ES performs almost as well as the supervised baseline, and better than EI.

486 5.5 ABLATION STUDY
487

488 Table 3 shows that architectures designed to be equivariant are measurably more equivariant across
489 imaging modalities and training losses. Moreover, it shows that networks trained using the splitting
490 loss perform better for equivariant architectures than for non-equivariant architectures, with an even
491 greater gap than for supervised inpainting baselines, which confirms the theoretical analysis made
492 in Section 4. Finally, we observe that non-equivariant architectures still lead to fairly measurably
493 equivariant models. While surprising, this phenomenon has already been witnessed and is usually
494 referred to as learned equivariance, whereby the training data and inductive biases lead to fairly
495 equivariant learned models (Gruver et al., 2024). We believe that this learned equivariance is responsible
496 for the high performance of splitting methods even when using non-equivariant architectures.
497 For additional results on the impact of equivariant architectures, see Appendix B.4.

498
499 **Table 3: Impact of using equivariant architectures.** In accordance with the theoretical results
500 described in Section 4, there is a synergy between the splitting loss and equivariant architectures
501 resulting in higher performance. Non-equivariant models have surprisingly high equivariance mea-
502 sures (EQUIV) which might explain their high performance when using the splitting loss. Eq. arch.
503 denotes whether the architecture is equivariant. In **bold**, the best self-supervised metrics (avg \pm
504 st.d.).

Image inpainting				
Training loss	Eq. arch.	PSNR \uparrow	SSIM \uparrow	EQUIV \uparrow
Supervised	✓	28.46 \pm 2.97	0.8982 \pm 0.0411	28.46 \pm 2.97
	✗	28.62 \pm 3.03	0.9002 \pm 0.0414	27.85 \pm 2.71
Splitting (Ours)	✓	27.45 \pm 2.86	0.8737 \pm 0.0461	27.46 \pm 2.85
	✗	27.20 \pm 2.83	0.8652 \pm 0.0461	26.52 \pm 2.60
MRI ($\times 8$ Accel., 40 dB SNR)				
Training loss	Eq. arch.	PSNR \uparrow	SSIM \uparrow	EQUIV \uparrow
Supervised	✓	28.74 \pm 2.81	0.6445 \pm 0.1094	31.71 \pm 2.83
	✗	28.48 \pm 2.68	0.6381 \pm 0.1082	28.78 \pm 1.95
Splitting (Ours)	✓	28.54 \pm 2.75	0.6195 \pm 0.1188	31.53 \pm 2.74
	✗	28.18 \pm 2.58	0.6104 \pm 0.1176	27.28 \pm 2.10

519
520 6 CONCLUSION
521

522 In this work, we propose a new self-supervised loss for solving inverse problems which bridges
523 the gap between existing equivariance and splitting-based self-supervised losses. We motivate the
524 design of our loss by showing that minimizing the expected loss results in MMSE estimators. **We**
525 **further validate our method using numerical simulations on different image distributions and imaging**
526 **modalities, including inpainting of natural images, MRI and CT.** These results suggest that the
527 proposed method compares favorably to the equivariant imaging baseline and is close to supervised
528 methods. To the best of our knowledge, this work is the first to leverage equivariant networks to
529 learn from incomplete data alone, going beyond the usual goal of improving the generalization of
530 the networks to unseen transformations at test time. Our method provides a new way to evaluate
531 the benefits of using different equivariant architectures, and can benefit from future advances made
532 in this field. More broadly, our work is further evidence that invariance is a promising prior for
533 learning from incomplete data.

534
535 **REPRODUCIBILITY STATEMENT**
536

537 We share the implementation of our method and experiments² to make our work easier to reproduce.
538

539
540 ²<https://anonymous.4open.science/r/Equivariant-Splitting-ICLR2026>

540 REFERENCES
541

542 Hemant Kumar Aggarwal, Merry P. Mani, and Mathews Jacob. MoDL: Model Based Deep Learning
543 Architecture for Inverse Problems. *IEEE Trans. Med. Imaging*, 38(2):394–405, February 2019.

544 Eirikur Agustsson and Radu Timofte. NTIRE 2017 Challenge on Single Image Super-Resolution:
545 Dataset and Study. In *CVPRW*, pp. 1122–1131, July 2017.

546 Samuel G. Armato III, Geoffrey McLennan, Luc Bidaut, Michael F. McNitt-Gray, Charles R. Meyer,
547 Anthony P. Reeves, Binsheng Zhao, Denise R. Aberle, Claudia I. Henschke, Eric A. Hoffman,
548 Ella A. Kazerooni, Heber MacMahon, Edwin J. R. van Beek, David Yankelevitz, Alberto M.
549 Biancardi, Peyton H. Bland, Matthew S. Brown, Roger M. Engelmann, Gary E. Laderach, Daniel
550 Max, Richard C. Pais, David P.-Y. Qing, Rachael Y. Roberts, Amanda R. Smith, Adam Starkey,
551 Poonam Batra, Philip Caligiuri, Ali Farooqi, Gregory W. Gladish, C. Matilda Jude, Reginald F.
552 Munden, Iva Petkovska, Leslie E. Quint, Lawrence H. Schwartz, Baskaran Sundaram, Lori E.
553 Dodd, Charles Fenimore, David Gur, Nicholas Petrick, John Freymann, Justin Kirby, Brian
554 Hughes, Alessi Vande Casteele, Sangeeta Gupte, Maha Sallam, Michael D. Heath, Michael H.
555 Kuhn, Ekta Dharaiya, Richard Burns, David S. Fryd, Marcos Salganicoff, Vikram Anand, Uri
556 Shreter, Stephen Vastagh, Barbara Y. Croft, and Laurence P. Clarke. The Lung Image Database
557 Consortium (LIDC) and Image Database Resource Initiative (IDRI): A Completed Reference
558 Database of Lung Nodules on CT Scans. *Medical Physics*, 38(2):915–931, 2011. ISSN 2473-
559 4209. doi: 10.1111/1.3528204.

560 Chinmay Belthangady and Loic A Royer. Applications, promises, and pitfalls of deep learning for
561 fluorescence image reconstruction. *Nature methods*, 16(12):1215–1225, 2019.

562 Yochai Blau and Tomer Michaeli. The Perception-Distortion Tradeoff. In *CVPR*, pp. 6228–6237,
563 2018.

564 Elena Celledoni, Matthias J. Ehrhardt, Christian Etmann, Brynjulf Owren, Carola-Bibiane
565 Schönlieb, and Ferdia Sherry. Equivariant neural networks for inverse problems. *Inverse Prob-
566 lems*, 37(8):085006, July 2021.

567 Anadi Chaman and Ivan Dokmanić. Truly shift-invariant convolutional neural networks. In *CVPR*,
568 pp. 3773–3783, 2021a.

569 Anadi Chaman and Ivan Dokmanić. Truly shift-equivariant convolutional neural networks with
570 adaptive polyphase upsampling. In *Asilomar Conf Signals Syst Comput*, pp. 1113–1120, 2021b.

571 Dongdong Chen, Julián Tachella, and Mike E. Davies. Equivariant Imaging: Learning Beyond the
572 Range Space. In *ICCV*, pp. 4379–4388, 2021.

573 Dongdong Chen, Julián Tachella, and Mike E. Davies. Robust Equivariant Imaging: A fully unsu-
574 pervised framework for learning to image from noisy and partial measurements. In *CVPR*, pp.
575 5647–5656, 2022.

576 Taco S. Cohen and Max Welling. Group Equivariant Convolutional Networks. In *ICML*, pp. 2990–
577 2999, 2016.

578 Taco S. Cohen and Max Welling. Steerable CNNs. In *ICLR*, 2017.

579 Giannis Daras, Kulin Shah, Yuval Dagan, Aravind Gollakota, Alexandros G. Dimakis, and Adam
580 Klivans. Ambient Diffusion: Learning Clean Distributions from Corrupted Data. In *NeurIPS*, pp.
581 288–313, 2023.

582 Leo Davy, Luis M. Briceno-Arias, and N. Pustelnik. Restarted contractive operators to learn at
583 equilibrium, June 2025.

584 Yonina C. Eldar. Generalized SURE for Exponential Families: Applications to Regularization. *IEEE
585 Trans Signal Process*, 57(2):471–481, 2009.

586 Fabian Ewald Fassnacht, Joanne C White, Michael A Wulder, and Erik Næsset. Remote sensing in
587 forestry: current challenges, considerations and directions. *Forestry: An International Journal of
588 Forest Research*, 97(1):11–37, 2024.

594 Nate Gruver, Marc Finzi, Micah Goldblum, and Andrew Gordon Wilson. The Lie Derivative for
 595 Measuring Learned Equivariance, 2024.

596

597 Allard A. Hendriksen, Daniel M. Pelt, and K. Joost Batenburg. Noise2Inverse: Self-supervised deep
 598 convolutional denoising for tomography. *IEEE Trans. Comput. Imaging*, 6:1320–1335, 2020.

599

600 Kyong Hwan Jin, Michael T McCann, Emmanuel Froustey, and Michael Unser. Deep convolutional
 601 neural network for inverse problems in imaging. *IEEE Trans Image Process*, 26(9):4509–4522,
 602 2017.

603

604 Sékou-Oumar Kaba, Arnab Kumar Mondal, Yan Zhang, Yoshua Bengio, and Siamak Ravanbakhsh.
 605 Equivariance with Learned Canonicalization Functions. In *ICML*, pp. 15546–15566, 2023.

606

607 Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and
 608 Timo Aila. Alias-free generative adversarial networks. In *NeurIPS*, volume 34, pp. 852–863,
 609 2021.

610

611 Achim Klenke. *Probability theory: a comprehensive course*. Springer, 2008.

612

613 Alexander Krull, Tim-Oliver Buchholz, and Florian Jug. Noise2Void - Learning Denoising From
 614 Single Noisy Images. In *CVPR*, pp. 2129–2137, 2019.

615

616 Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization, 2019.

617

618 Christopher A. Metzler, Ali Mousavi, Reinhard Heckel, and Richard G. Baraniuk. Unsupervised
 619 Learning with Stein’s Unbiased Risk Estimator, July 2020.

620

621 Hagay Michaeli, Tomer Michaeli, and Daniel Soudry. Alias-Free Convnets: Fractional Shift Invari-
 622 ance via Polynomial Activations. In *CVPR*, pp. 16333–16342, 2023.

623

624 Charles Millard and Mark Chiew. A theoretical framework for self-supervised MR image recon-
 625 struction using sub-sampling via variable density Noisier2Noise. *IEEE Trans Comput Imaging*,
 626 9:707–720, 2023.

627

628 Brayan Monroy, Jorge Bacca, and Julián Tachella. Generalized Recorrupted-to-Recorrupted: Self-
 629 Supervised Learning Beyond Gaussian Noise. In *CVPR*, pp. 28155–28164, 2025.

630

631 Greg Ongie, Rebecca Willett, Daniel Soudry, and Nathan Srebro. A Function Space View of
 632 Bounded Norm Infinite Width ReLU Nets: The Multivariate Case, October 2019.

633

634 Tongyao Pang, Huan Zheng, Yuhui Quan, and Hui Ji. Recorrupted-to-Recorrupted: Unsupervised
 635 Deep Learning for Image Denoising. In *CVPR*, pp. 2043–2052, June 2021.

636

637 Javier Portilla and Eero P Simoncelli. A parametric texture model based on joint statistics of com-
 638 plex wavelet coefficients. *International journal of computer vision*, 40(1):49–70, 2000.

639

640 Omri Puny, Matan Atzmon, Heli Ben-Hamu, Ishan Misra, Aditya Grover, Edward J. Smith, and
 641 Yaron Lipman. Frame Averaging for Invariant and Equivariant Network Design, March 2022.

642

643 Marco Ragone, Reza Shahabazian-Yassar, Farzad Mashayek, and Vitaliy Yurkiv. Deep learning
 644 modeling in microscopy imaging: A review of materials science applications. *Progress in Mate-
 645 rials Science*, 138:101165, 2023.

646

647 Sebastián Rivera, Iván Ortíz, Tatiana Gelvez, Fernando Rojas, and Henry Arguello. Rotation invari-
 648 ant deep learning approach for image inpainting. In *STSIIVA*, pp. 1–5, 2021.

649

650 Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomed-
 651 ical Image Segmentation. In *MICCAI*, pp. 234–241, 2015.

652

653 Leonid I. Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal
 654 algorithms. *Physica D: Nonlinear Phenomena*, 60(1):259–268, November 1992.

655

656 Akiyoshi Sannai, Makoto Kawano, and Wataru Kumagai. Equivariant and Invariant Reynolds Net-
 657 works, October 2021.

648 Akiyoshi Sannai, Makoto Kawano, and Wataru Kumagai. Invariant and equivariant reynolds net-
 649 works. *Journal of Machine Learning Research*, 25(42):1–36, 2024.
 650

651 Jérémie Scanvic, Mike Davies, Patrice Abry, and Julián Tachella. Scale-Equivariant Imaging: Self-
 652 Supervised Learning for Image Super-Resolution and Deblurring, March 2025.
 653

654 Victor Sechaud, Laurent Jacques, Patrice Abry, and Julián Tachella. Equivariance-based self-
 655 supervised learning for audio signal recovery from clipped measurements. In *EUSIPCO*, August
 656 2024.
 657

658 Julián Tachella, Dongdong Chen, and Mike Davies. Unsupervised learning from incomplete mea-
 659 surements for inverse problems. *NeurIPS*, 35:4983–4995, 2022.
 660

661 Julián Tachella, Dongdong Chen, and Mike Davies. Sensing Theorems for Unsupervised Learning
 662 in Linear Inverse Problems. *Journal of Machine Learning Research*, 24(39):1–45, 2023. ISSN
 663 1533-7928.
 664

665 Julián Tachella, Mike Davies, and Laurent Jacques. UNSURE: Self-supervised learning with Un-
 666 known Noise level and Stein’s Unbiased Risk Estimate. In *ICLR*, 2025a.
 667

668 Julián Tachella, Matthieu Terris, Samuel Hurault, Andrew Wang, Dongdong Chen, Minh-Hai
 669 Nguyen, Maxime Song, Thomas Davies, Leo Davy, Jonathan Dong, Paul Escande, Johannes
 670 Hertwich, Zhiyuan Hu, Tobías I. Liaudat, Nils Laurent, Brett Levac, Mathurin Massias, Thomas
 671 Moreau, Thibaut Modrzyk, Brayan Monroy, Sebastian Neumayer, Jérémie Scanvic, Florian Sar-
 672 rron, Victor Sechaud, Georg Schramm, Chao Tang, Romain Vo, and Pierre Weiss. DeepInverse: A
 673 Python package for solving imaging inverse problems with deep learning, May 2025b.
 674

675 Matthieu Terris, Thomas Moreau, Nelly Pustelnik, and Julian Tachella. Equivariant plug-and-play
 676 image reconstruction. In *CVPR*, pp. 25255–25264, 2024.
 677

678 Singanallur V Venkatakrishnan, Charles A Bouman, and Brendt Wohlberg. Plug-and-play priors for
 679 model based reconstruction. In *IEEE Glob Conf Signal Inf Process*, pp. 945–948, 2013.
 680

681 Antonia Vojtekova, Maggie Lieu, Ivan Valtchanov, Bruno Altieri, Lyndsay Old, Qifeng Chen, and
 682 Filip Hroch. Learning to denoise astronomical images with u-nets. *Monthly Notices of the Royal
 683 Astronomical Society*, 503(3):3204–3215, 2021.
 684

685 Andrew Wang and Mike Davies. Perspective-equivariance for unsupervised imaging with camera
 686 geometry. In Alessio Del Bue, Cristian Canton, Jordi Pont-Tuset, and Tatiana Tommasi (eds.),
 687 *ECCV Workshops*, pp. 116–134, 2024.
 688

689 Andrew Wang, Steven McDonagh, and Mike Davies. Benchmarking Self-Supervised Learning
 690 Methods for Accelerated MRI Reconstruction, May 2025.
 691

692 Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: From error
 693 visibility to structural similarity. *IEEE Trans Image Process*, 13(4):600–612, April 2004.
 694

695 Philip J Withers, Charles Bouman, Simone Carmignato, Veerle Cnudde, David Grimaldi, Char-
 696 lotte K Hagen, Eric Maire, Marena Manley, Anton Du Plessis, and Stuart R Stock. X-ray com-
 697 puted tomography. *Nature Reviews Methods Primers*, 1(1):18, 2021.
 698

699 Guixian Xu, Jinglai Li, and Junqi Tang. Fast Equivariant Imaging: Acceleration for Unsupervised
 700 Learning via Augmented Lagrangian and Auxiliary PnP Denoisers, July 2025.
 701

702 Xuemiao Xu, Huaidong Zhang, Guoqiang Han, Kin Chung Kwan, Wai-Man Pang, Jiaming Fang,
 703 and Gansen Zhao. A two-phase space resection model for accurate topographic reconstruction
 704 from lunar imagery with pushbroomscanners. *Sensors*, 16(4):507, 2016.
 705

706 Burhaneddin Yaman, Seyed Amir Hosseini, Steen Moeller, Jutta Ellermann, Kâmil Uğurbil,
 707 and Mehmet Akçakaya. Self-Supervised Learning of Physics-Guided Reconstruction Neural Net-
 708 works without Fully-Sampled Reference Data. *Magnetic Resonance in Med*, 84(6):3172–3191,
 709 2020.

Jure Zbontar, Florian Knoll, Anuroop Sriram, Tullie Murrell, Zhengnan Huang, Matthew J. Muckley, Aaron Defazio, Ruben Stern, Patricia Johnson, Mary Bruno, Marc Parente, Krzysztof J. Geras, Joe Katsnelson, Hersh Chandarana, Zizhao Zhang, Michal Drozdzal, Adriana Romero, Michael Rabbat, Pascal Vincent, Nafissa Yakubova, James Pinkerton, Duo Wang, Erich Owens, C. Lawrence Zitnick, Michael P. Recht, Daniel K. Sodickson, and Yvonne W. Lui. fastMRI: An Open Dataset and Benchmarks for Accelerated MRI, December 2019.

Richard Zhang. Making Convolutional Networks Shift-Invariant Again. In *ICML*, pp. 7324–7334, 2019.

A DETAILS ABOUT THE METHOD

A.1 EQUIVARIANT RECONSTRUCTOR

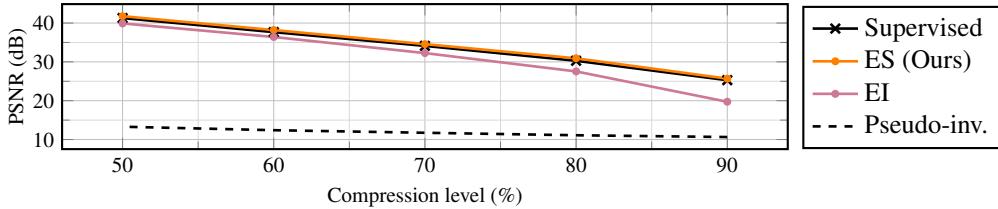


Figure 4: **Compressive sensing results.** Adds the pseudo-inverse reconstruction (pseudo-inv.) as a baseline to Figure 1

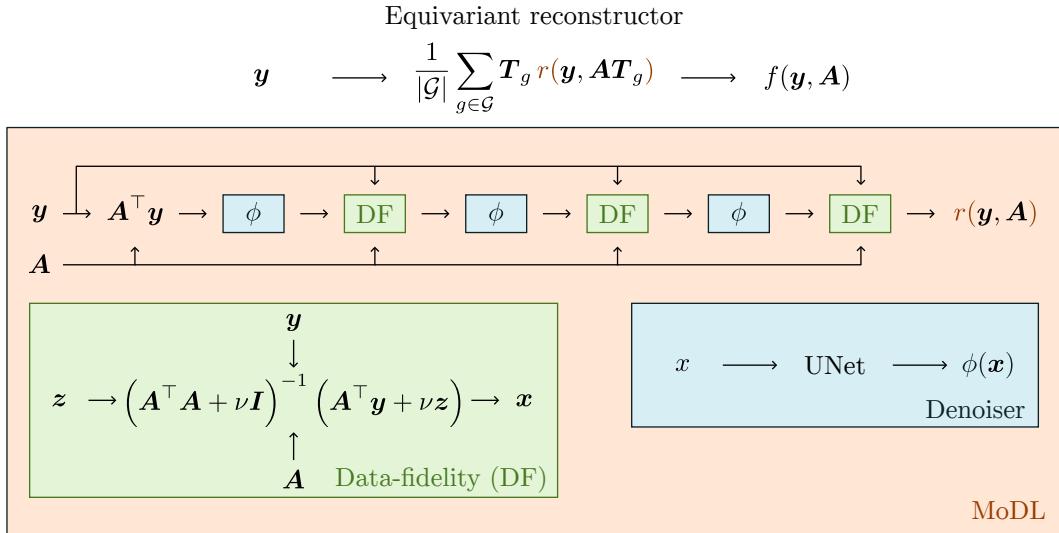


Figure 5: **Reconstructor equivariant to rotations and flips.** It is the Reynolds averaging of 90° rotations and horizontal and vertical flips as in eq. (18) of a non-equivariant reconstructor of the MAP type, as in eq. (19), implemented using a MoDL unrolled algorithm (Aggarwal et al., 2019) with 3 iterations, shared weights, and using a non-equivariant residual UNet (Ronneberger et al., 2015) as the denoiser architecture.

In Theorem 2, we consider reconstruction functions of 4 different structures. In this section, we give additional details about them.

The artifact removal reconstructor architecture (Jin et al., 2017) in eq. (16) consists in a projection step that maps the measurements back into the image space using the adjoint or the pseudo-inverse

756 of the forward matrix, which is immediately followed by a very general trainable network, often
 757 a UNet or another encoder-decoder type of network, or sometimes simply a fully convolutional
 758 network. Theorem 2 states that this reconstructor architecture produces equivariant reconstructors
 759 as long as the denoiser architecture is, itself, equivariant in the sense of eq. (14).

760 Reynolds averaging for possibly non-equivariant reconstruction functions $r(\mathbf{y}, \mathbf{A})$ in eq. (18)
 761

$$762 \quad f(\mathbf{y}, \mathbf{A}) = \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} \mathbf{T}_g r(\mathbf{y}, \mathbf{A} \mathbf{T}_g). \quad (18)$$

763 has, as far as we know, not been defined in previous works. It is a natural extension of Reynolds
 764 averaging for possibly non-equivariant image-to-image functions or denoisers $\psi(\mathbf{x})$ (Sannai et al.,
 765 2024; Terris et al., 2024),

$$766 \quad \phi(\mathbf{x}) = \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} \mathbf{T}_g^{-1} \psi(\mathbf{T}_g \mathbf{x}) \quad (22)$$

767 which makes them equivariant in the sense of eq. (14), to reconstruction functions in order to make
 768 them equivariant in the sense of eq. (15). It is a fairly simple way to make a reconstructor equivariant
 769 and it is relatively inexpensive for small groups of transformations such as the group of 90° rotations
 770 and horizontal and vertical flips (Cohen & Welling, 2016). It is however too expensive to be used in
 771 practice when the group is relatively large, like the group of shifts, since it would require to evaluate
 772 the neural network for as many times as there are pixels in the input image, for each image.
 773

774 The maximum a posteriori (MAP) reconstructor architecture defined in eq. (19)
 775

$$776 \quad f(\mathbf{y}, \mathbf{A}) = \underset{\mathbf{x} \in \mathbb{R}^n}{\operatorname{argmax}} \left\{ p(\mathbf{x} \mid \mathbf{y}, \mathbf{A}) \right\}. \quad (19)$$

777 is a classical reconstructor architecture that has been used in different ways, including iterative
 778 algorithms with a hand-crafted prior (Rudin et al., 1992; Davy et al., 2025), plug-and-play archi-
 779 tectures using an iterative approach but with the hand-crafted prior replaced with a pre-trained de-
 780 noiser (Venkatakrishnan et al., 2013), and unrolled architectures where the optimization problem is
 781 done with a fixed number of steps and where parts of the algorithm are replaced with trainable mod-
 782 ules (Aggarwal et al., 2019). It is often interpreted as a Bayesian maximum a posteriori estimator,
 783 but also commonly outside of a Bayesian framework as a variational approach with a data-fidelity
 784 term and a regularization term. It is one of the reconstructor architectures that we use for our ex-
 785 periments in Section 5. Theorem 2 states that these reconstructors are equivariant as long as the
 786 prior distribution is invariant in the sense of eq. (5), or equivalently, if the associated regularization
 787 function or negative log-prior is invariant.

788 The minimum mean squared error (MMSE) estimator in eq. (20)
 789

$$790 \quad f(\mathbf{y}, \mathbf{A}) = \mathbb{E}_{\mathbf{x} \mid \mathbf{y}, \mathbf{A}} \{ \mathbf{x} \} \quad (20)$$

791 is generally the target theoretical reconstructor as it achieves the highest theoretical PSNR (Tachella
 792 et al., 2023). It is generally estimated by reconstructors trained with a mean squared error loss (Chen
 793 et al., 2021), which is notably how we train the supervised baselines in the experiments in Section 5.
 794 Theorem 2 states that as long as the prior distribution is invariant in the sense of eq. (5), the MMSE
 795 reconstructor is equivariant in the sense of eq. (15).
 796

800 A.2 HOW TO CHOOSE TRANSFORMATIONS FOR A GIVEN INVERSE PROBLEM

801 There are two major criteria for choosing the transformations for a given application. First, the im-
 802 age distribution of interest should be invariant to the chosen transformations. Aerial, remote sensing
 803 and microscopic images are invariant to translations and rotations as the scenes and subjects they
 804 measure exhibit no privileged position and orientation with respect to the image plane. Natural im-
 805 age distributions and texture distributions (Portilla & Simoncelli, 2000) are also generally invariant
 806 to translations but they are less invariant to rotations as natural images are typically oriented upward
 807 and texture distributions might be anisotropic.

808 Second, the transformations should also be chosen in accordance with the measurement operator.
 809 Corollary 1 shows that transformations for which the operator is equivariant do not improve the

810
 811 **Table 4: MRI results in another setting.** Supplementary results for a different MRI problem ($\times 6$
 812 Accel., 10 dB SNR) than in Table 1 In **bold**, the best self-supervised metrics. Values: avg \pm st.d.

813 814 815 816 817 818 819 820 821 822 823 824 825 826 827 828 829	MRI ($\times 6$ Accel., 10 dB SNR)		
	Method	PSNR \uparrow	SSIM \uparrow
Supervised	27.39 \pm 2.44	0.5243 \pm 0.1373	30.38 \pm 2.43
ES (Ours)	27.33 \pm 2.45	0.5126 \pm 0.1444	30.32 \pm 2.44
EI	27.23 \pm 2.41	0.5110 \pm 0.1421	30.21 \pm 2.40
SURE	27.08 \pm 2.29	0.5097 \pm 0.1372	30.06 \pm 2.28
IDFT	23.85 \pm 1.05	0.3878 \pm 0.0272	25.14 \pm 0.79

830
 831 **Table 5: Decision table for the transformations.** Corollary 1 shows that not all transformations are
 832 well-suited for all problems. Namely, transformations for which the operator is equivariant introduce
 833 no additional information and should not be used. This table specifies which transformations are
 834 well-suited for which operator.

834 835 836 837 838 839	Operator	Translation	Rotation	Permutation	Amplitude
Isotropic blur	\times	\times	\checkmark	\times	
Image inpainting	\checkmark	\checkmark	\checkmark	\times	
Sparse-view CT	\times	\checkmark	\checkmark	\times	
Accelerated MRI	\times	\checkmark	\checkmark	\times	
Compressive sensing	\checkmark	\checkmark	\checkmark	\times	

840
 841
 842
 843
 844
 845
 846
 847
 848 reconstruction process. It is a well-known criterion introduced in the original work on equivariant
 849 imaging (Chen et al., 2021) which remains correct in our setting where measurement splitting
 850 is added to the theoretical analysis. Table 5 lists correct choices of transformations for common
 851 measurement operators.

852 853 A.3 END-TO-END ALGORITHM

854
 855
 856
 857
 858 In this section, we present the end-to-end ES algorithm. It consists in a training step where an equivariant
 859 reconstructor is trained using backpropagation against a training dataset of measurements
 860 only, after which it is applied to obtain the reconstructions associated to the test measurements. The
 861 ES loss used in the training step is computed using the expression in Equation (3). The expectations
 862 are estimated using Monte Carlo sampling where a single sample is used at training time and $T = 10$
 863 samples are used at inference. In the experiments, we use the splitting ratio $s = 0.8$ corresponding
 to $m_1 = 0.8m$ and $m_2 = 0.2m$. Algorithms 1 and 2 show the detailed algorithms in pseudo-code.

864

Algorithm 1: Equivariant Splitting (Training procedure)

865

Input: Dataset $D = \{y_i\}_{i \in I}$, forward operator $A \in \mathbb{R}^{m \times n}$, split ratio s , learning rate η , number of epochs E , equivariant reconstructor f_θ

866

Output: Trained model f_θ

867

for $epoch = 1$ **to** E **do**

868

 Shuffle dataset D ;

869

foreach $mini\text{-}batch B \subset D$ **do**

870

init: $L = 0$

871

foreach y **in** B **do**

872

init: $M \in \mathbb{R}^{m_1 \times m}$ a random split s.t $m_1 = s \times m$.

873

 Split measurement and operator: $y_1 = My, A_1 = MA$;

874

 Compute predictions: $\hat{x} = f_\theta(y_1, A_1)$;

875

 Compute sample loss $\ell = \|A\hat{x} - y\|_2^2$;

876

 Aggregate loss: $L \leftarrow L + \ell$;

877

end

878

 Mean: $L \leftarrow \frac{1}{|B|}L$;

879

 Backpropagation: compute gradients $\nabla_\theta L$;

880

 Update parameters: $\theta \leftarrow \text{ADAM}(\theta, \nabla_\theta L, \eta)$;

881

end

882

end

883

884

Algorithm 2: Equivariant Splitting (Inference)

885

Input: Measurement $y \in \mathbb{R}^{m \times n}$, forward operator $A \in \mathbb{R}^{m \times n}$, split ratio s , number of samples T , trained reconstructor f_θ

886

Output: Reconstructed images $\hat{x} = \bar{f}_\theta(y, A)$

887

init: $\hat{x} = 0$

888

for $i = 1$ **to** T **do**

889

init: $M \in \mathbb{R}^{m_1 \times m}$ a random split s.t $m_1 = s \times m$.

890

 Split measurement and operator: $y_1 = My, A_1 = MA$;

891

 Compute one predictions: $\hat{x}_i = f_\theta(y_1, A_1)$;

892

 Aggregate: $\hat{x} \leftarrow \hat{x} + \hat{x}_i$;

893

end

894

 Mean: $\hat{x} \leftarrow \frac{1}{T}\hat{x}$;

895

896

B DETAILS ABOUT THE EXPERIMENTS

897

We use the optimizer AdamW (Loshchilov & Hutter, 2019) for every training with different learning rates for the different inverse problems, a weight decay of 10^{-8} , beta coefficients equal to 0.9 and 0.999 and without the AMSGrad option. For longer trainings, we use step schedulers that divide the learning rate by a factor ranging from 2 to 10 at specific epochs, up to 3 or 4 times.

898

In our experiments, we make extensive use of the DeepInverse library (Tachella et al., 2025b) that provides an implementation of the various forward operators and training losses that we use. Every model is trained for up to 50 hours on a single GPU, either an NVIDIA H100, GH200 or RTX 4090. See Appendix B.4 for more details about the training durations.

899

B.1 DATASETS

900

Image inpainting The ground truth images are images obtained from the dataset DIV2K (Agustsson & Timofte, 2017) which contains pictures of natural scenes (landscapes, animals) by first resizing them to a resolution of 256×256 pixels before extracting a central 128×128 pixels crop. We synthesize the measurements by corrupting the ground truth images using a single binary mask sampled from a pixel-wise Bernoulli distribution with a 30% chance of keeping each pixel value. In this setting, we do not corrupt the measurements further with additional noise.

911

MRI The dataset consists in 973 pairs of ground truth images from the FastMRI dataset (Zbontar et al., 2019) and associated k-space measurements synthesized using a single coil sensitivity map.

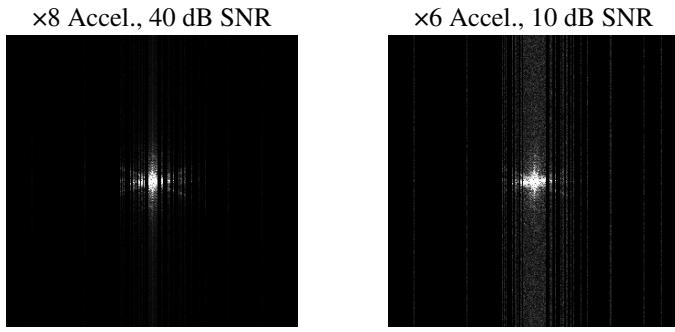


Figure 6: **Examples of k-space measurements in MRI.** (left) a less noisy problem with less measurements, (right) a noisier problem with more measurements.

The ground truth images have a size of 320×320 pixels and each corresponds to the middle slice of a different 3d knee acquisition. The k-spaces are synthesized as discrete Fourier transforms subsampled on a single non-regular grid and further corrupted with additive white Gaussian noise with a standard deviation of 0.005 corresponding to an average SNR of about 40 dB. The subsampling grid models the coil sensitivity map, is sampled from a Gaussian distribution and corresponds to an acceleration of 8. The entire dataset is finally split into a train/val split containing 900 images and a test split containing the remaining 73 images. Our implementation uses the work from Wang et al. (2025).

We also consider an additional dataset obtained in the same way except using a different mask corresponding to an acceleration of 6 and with a higher noise level corresponding to a standard deviation of 0.1 or an average SNR of 10 dB in the k-space domain.

B.2 NETWORK ARCHITECTURES

The unrolled architecture uses 3 iterations and the weights are shared across different iterations. Every UNet is residual, has 4 scales and has no normalization layer as we find them to be detrimental to the performance.

The transforms we use in the model architecture and in the metrics are grid-preserving: grid-aligned shifts, 90° rotations, and vertical and horizontal flips. The transforms we use in the EI are grid-aligned shifts and 1° rotations following the prior art Chen et al. (2021). Reynolds' averaging is implemented using an unbiased Monte-Carlo estimator whereby a single random transform is sampled at every evaluation to save on computational cost.

B.3 METRICS

In the experiments, we use three different performance metrics including two standard distortion metrics (PSNR, SSIM) and a new equivariance metric for reconstructors (EQUIV). The peak signal-to-noise ratio is defined, for images with a dynamic range normalized to $[0, 1]$, as the mean squared error expressed in decibels (dB)

$$\text{PSNR} = \mathbb{E}_{\mathbf{x}, \mathbf{y}} \left\{ -10 \log_{10} (\|f(\mathbf{y}, \mathbf{A}) - \mathbf{x}\|^2) \right\}. \quad (23)$$

The structural similarity index measure (Wang et al., 2004) is a more perceptual metric which is a combination of the empirical means, standard deviations and correlation coefficient associated to the reference and compared images. The complete definition of the metric is too long to be included in this work and we refer the reader to the original publication for more details. In addition to these two standard distortion metrics, we use a new equivariance metric similar to that used by Chaman & Dokmanić (2021b) adapted for our proposed definition of equivariant reconstructors. It is the average mean squared error associated with eq. (15) and expressed in dB for readability

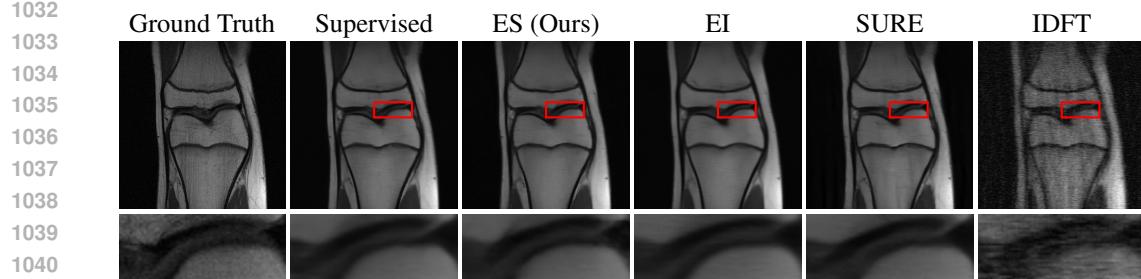
$$\text{EQUIV} = -10 \log_{10} \left(\mathbb{E}_{\mathbf{y}, g} \left\{ \|f(\mathbf{y}, \mathbf{A}T_g) - T_g^{-1}f(\mathbf{y}, \mathbf{A})\|^2 \right\} \right). \quad (24)$$

It can be roughly understood as a PSNR for equivariance. In every experiment, we use the same group of transformations for EQUIV as we use for the equivariant reconstructor architectures.

972 B.4 RESULTS
973974 In this section, we provide more details about the main experiments and present additional experi-
975 ments.976 For each training done in the main experiments, we report the average epoch duration and the number
977 of epochs until the model has finished training. For the sake of comparability, we made sure to
978 conduct the different trainings for the same imaging modality on the same GPU. Namely, we used
979 a single NVIDIA RTX 3090 Ti GPU for every inpainting experiment and a single NVIDIA H100
980 GPU for every MRI experiment. Table 10 shows that self-supervised methods including ours have
981 generally longer epochs and require more epochs than the fully supervised gold standard. We believe
982 that this is due to a fundamental trade-off whereby the use of ground truth data accelerates the
983 learning procedure while also enabling the use of simpler training algorithms. Moreover, unlike
984 standard implementations of EI and SURE, ES requires only a single network pass per iteration
985 resulting in significantly faster epochs. Overall, ES is computationally more efficient than EI, in
986 addition to being more performant in terms of reconstruction quality.987 Figure 4 shows additional compressive sensing results, Table 4 and Figure 7 show results on a MRI
988 experiment with settings different from the main one, Table 6 shows extended results for the ablation
989 study on equivariant architectures, and Figure 6 shows sample k-spaces from the MRI experiments.
990 Figure 5 shows the equivariant reconstructor architecture that we adopt in the MRI experiments.
991992 Table 6: **Extended results on the impact of equivariant architectures.** Adds to Table 3 the results
993 for EI with a non-equivariant architecture for the inpainting task, results for the noise-dominated
994 MRI task. In **bold**, the best self-supervised metrics. Values: avg \pm st.d.

Image inpainting				
Training loss	Eq. arch.	PSNR \uparrow	SSIM \uparrow	EQUIV \uparrow
Supervised	✓	28.46 \pm 2.97	0.8982 \pm 0.0411	28.46 \pm 2.97
	✗	28.62 \pm 3.03	0.9001 \pm 0.0415	27.85 \pm 2.71
Splitting (Ours)	✓	27.45 \pm 2.86	0.8737 \pm 0.0461	27.46 \pm 2.85
	✗	27.20 \pm 2.83	0.8651 \pm 0.0463	26.52 \pm 2.60
EI loss	✓	25.89 \pm 2.65	0.8332 \pm 0.0521	25.89 \pm 2.65
	✗	26.33 \pm 2.81	0.8451 \pm 0.0536	25.58 \pm 2.52
MC loss	✓	8.22 \pm 2.47	0.098 \pm 0.055	8.22 \pm 2.47
	✗	8.24 \pm 2.48	0.100 \pm 0.056	8.24 \pm 2.48
MRI (x8 Accel., 40 dB SNR)				
Training loss	Eq. arch.	PSNR \uparrow	SSIM \uparrow	EQUIV \uparrow
Supervised	✓	28.74 \pm 2.81	0.6445 \pm 0.1094	31.71 \pm 2.83
	✗	28.48 \pm 2.68	0.6381 \pm 0.1082	28.78 \pm 1.95
Splitting (Ours)	✓	28.54 \pm 2.75	0.6195 \pm 0.1188	31.53 \pm 2.74
	✗	28.18 \pm 2.58	0.6104 \pm 0.1176	27.28 \pm 2.10
MRI (x6 Accel., 10 dB SNR)				
Training loss	Eq. arch.	PSNR \uparrow	SSIM \uparrow	EQUIV \uparrow
Supervised	✓	27.39 \pm 2.44	0.5243 \pm 0.1373	30.38 \pm 2.43
	✗	27.33 \pm 2.42	0.5174 \pm 0.1410	29.73 \pm 2.20
Splitting (Ours)	✓	27.33 \pm 2.45	0.5126 \pm 0.1444	30.32 \pm 2.44
	✗	27.20 \pm 2.38	0.5095 \pm 0.1430	28.66 \pm 1.76

1021 We present an additional inpainting experiment, designed to simulate a more realistic acquisition
1022 scenario and study the consistency of our method. Specifically, we consider a noisy inpainting setting
1023 in which entire image columns are randomly removed. Such sampling patterns naturally arise in
1024 satellite imaging systems based on push-broom scanners (PBS) (Xu et al., 2016). As reported in Ta-
1025 ble 7, the proposed method achieves performance comparable to that of the EI baseline, confirming
its consistency under this more practical acquisition model.

1026
1027
1028
1029
1030
1031

1042 **Figure 7: Sample reconstructions for noise dominated MRI (x6 Accel., 10 dB SNR).** In the noise
 1043 dominated setting, the different models perform more similarly than in the less noisy setting shown
 1044 in Figure 3.

1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058

1059 **Table 7: Inpainting with push-broom masks.** We consider an unevenly-spaced mask sampled
 1060 randomly to test our method in a realistic inpainting setting, and an evenly-spaced mask to verify
 1061 the claim in Corollary 1 empirically in a case of almost-equivariance.

1062	1063 Image inpainting (Randomly-spaced push-broom mask)		
	1064 Method	1065 PSNR \uparrow	1066 SSIM \uparrow
1067	Supervised	23.72 ± 2.10	0.743 ± 0.051
	ES (Ours)	23.04 ± 1.81	0.734 ± 0.050
	EI	23.05 ± 2.13	0.707 ± 0.066
	Incomplete image	9.44 ± 2.31	0.141 ± 0.048
1069	1070 Image inpainting (Evenly-spaced push-broom mask)		
	1071 Method	1072 PSNR \uparrow	1073 SSIM \uparrow
1074	Supervised	28.37 ± 2.15	0.873 ± 0.035
	ES (Ours)	21.94 ± 2.12	0.617 ± 0.091
	Incomplete image	9.61 ± 2.42	0.152 ± 0.061

1075
1076
1077
1078
1079

To clarify the behavior of ES when the forward operator A is equivariant or nearly equivariant, we conducted the following experiment: we performed an PBS inpainting operation in which every other column of the image is removed. In this setting, the operator A becomes “almost” equivariant in the sense that for any even horizontal shift g , we have $AT_g = T_g A$, and the same property holds for all vertical shifts. Table 7 shows the reconstruction performance deteriorates significantly.

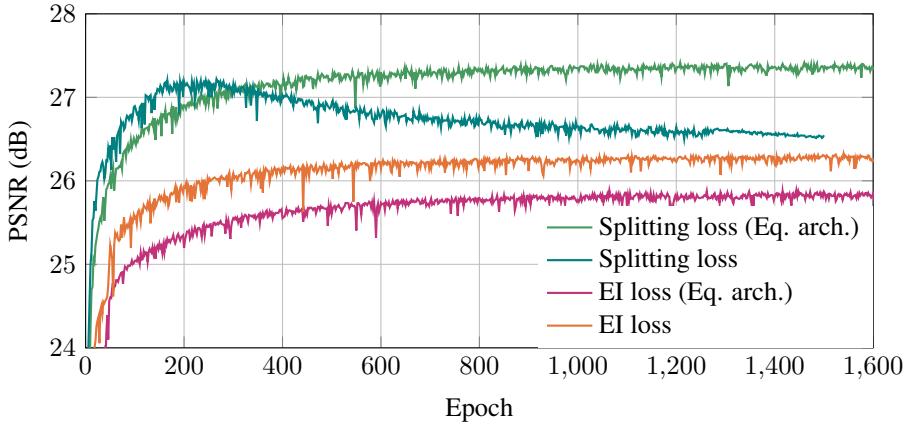


Figure 8: **Performance evolution during training for inpainting.** Splitting methods perform better than EI independent of the network architecture.

B.4.1 TESTING THE EFFECT OF UNKNOWN NOISE DISTRIBUTIONS

In addition to the main experiments where we use the knowledge of the noise distribution in the reconstruction algorithm, we test a more realistic scenario where the noise distribution is unknown and is estimated to be zero for a lack of a better estimate. To do so, we train models using variants of ES and EI where the measurements are assumed to be noiseless even though the training and testing data are nonetheless corrupted with noise. Table 8 shows that ES also performs better than EI when the noise distribution is unknown. Moreover, it shows that the performance of ES tends to be lower when the noise distribution is unknown, but it does not drop exceedingly which demonstrates the stability of ES to unexpected noise.

B.4.2 EMPIRICAL VERIFICATION OF THE EQUIVARIANCE OF MAP RECONSTRUCTORS

We verify empirically that MAP reconstructors are equivariant as long as the prior is itself equivariant, i.e., the claim made in Theorem 2. Since they cannot be computed exactly in general, we consider a specific scenario where they can. We assume that 1) $x \sim \mathcal{N}(0, \tau^2 I_n)$, 2) $A \in \mathbb{R}^{m \times n}$ is the two-dimensional decimation operator with decimation rate 2, 3) $y \mid Ax \sim \mathcal{N}(0, \sigma^2 I_m)$, and 4) that T_g denotes the rotation by angle $g \in \{0^\circ, 90^\circ, 180^\circ, 270^\circ\}$. Under these assumptions, the prior distribution is equivariant and the MAP estimator in eq. (19) can be expressed in closed-form as

$$f(\mathbf{y}, AT_g) = \frac{\tau^2}{\tau^2 + \sigma^2} T_g^{-1} A^\top \mathbf{y}, \quad (25)$$

with $f(\mathbf{y}, A)$ being the special case where $T_g = I_n$. In the experiment, we set $n = 128 \times 128$ for a grayscale image with 128 rows and 128 columns and we compute the equivariance metric in eq. (24) (EQUIV) for the MAP reconstructor using 256 i.i.d. samples from the joint distribution. Table 9 shows the results for every angle and for the average over all angles. As predicted theoretically, perfect equivariance is achieved.

C PROOFS

For the sake of clarity, we state the propositions and theorems a second time before their proofs. We also state and prove the additional Lemma 1 which helps prove Theorem 1.

1134
 1135
 1136
 1137
 1138
 1139
 1140
 1141
 1142
 1143
 1144
 1145

1146 **Table 8: Performance when the noise distribution is unknown.** In the unknown noise scenarios,
 1147 we use the variants of ES and EI corresponding to assuming that the measurements are noiseless.

		Image inpainting	
		Method	Known
			PSNR \uparrow
1151	Supervised		23.72 ± 2.10
1152	ES (Ours)	✓	23.04 ± 1.81
1153	ES	✗	22.05 ± 1.36
1154	EI	✓	23.05 ± 2.13
1155	EI	✗	22.01 ± 1.71
1156	Incomplete image		9.44 ± 2.31
1157			MRI ($\times 8$ Accel., 40 dB SNR)
		Method	Known
			PSNR \uparrow
1160	Supervised		28.74 ± 2.81
1161	ES (Ours)	✓	28.54 ± 2.75
1162	ES	✗	28.52 ± 2.75
1163	EI	✓	27.88 ± 2.64
1164	EI	✗	27.89 ± 2.59
1165	IDFT		23.62 ± 1.90
1166			MRI ($\times 6$ Accel., 10 dB SNR)
		Method	Known
			PSNR \uparrow
1170	Supervised		27.39 ± 2.44
1171	ES (Ours)	✓	27.33 ± 2.45
1172	ES	✗	25.73 ± 1.49
1173	EI	✓	27.23 ± 2.41
1174	EI	✗	26.02 ± 1.65
1175	IDFT		23.85 ± 1.05

1177
 1178
 1179
 1180
 1181
 1182
 1183
 1184
 1185
 1186
 1187

1188

1189

Table 9: Empirical validation of the equivariance of MAP estimators.

1190

1191

	0°	90°	180°	270°	Average
EQUIV	∞	∞	∞	∞	∞

1192

1193

1194

Table 10: Training durations. For each training, we report the average epoch duration and the number of epochs until the model is trained. Inpainting trainings are conducted on a single NVIDIA RTX 3090 Ti GPU and MRI trainings on a NVIDIA H100 GPU.

1195

1196

1197

Image inpainting		
Method	Epoch duration (s)	Epochs
Supervised	12	200
ES (Ours)	12	1000
EI	14	1000
MRI (x8 Accel., 40 dB SNR)		
Method	Epoch duration (s)	Epochs
Supervised	29	200
ES (Ours)	24	13800
EI	53	7800
SURE	36	5700
MRI (x6 Accel., 10 dB SNR)		
Method	Epoch duration (s)	Epochs
Supervised	19	70
ES (Ours)	19	3100
EI	75	2400
SURE	35	1200

1198

1199

Lemma 1. *The minimization problem*

1200

1201

$$\min_f \mathbb{E}_{\mathbf{x}, \mathbf{y}} \{ \| \mathbf{A}f(\mathbf{y}) - \mathbf{A}\mathbf{x} \|^2 \} \quad (26)$$

1202

1203

admits as solutions the functions of the form

1204

1205

$$f(\mathbf{y}) = \mathbf{A}^\dagger \mathbf{A} \mathbb{E}_{\mathbf{x}|\mathbf{y}} \{ \mathbf{x} \} + (\mathbf{I} - \mathbf{A}^\dagger \mathbf{A}) v(\mathbf{y}) \quad (27)$$

1206

1207

where $v(\mathbf{y})$ is any function.

1208

1209

Proof. Let's start by stating that $f(\mathbf{y}) = \mathbb{E}_{\mathbf{x}|\mathbf{y}} \{ \mathbf{x} \}$ is the only solution (Klenke, 2008)

1210

1211

$$\min_f \mathbb{E}_{\mathbf{x}, \mathbf{y}} \{ \| f(\mathbf{y}) - \mathbf{x} \|^2 \}. \quad (28)$$

1212

1213

For f any solution of eq. (26), applying it to $\tilde{f}(\mathbf{y}) = \mathbf{A}f(\mathbf{y})$ and $\tilde{x} = \mathbf{A}\mathbf{x}$ gives

1214

1215

$$\mathbf{A}f(\mathbf{y}) = \mathbf{A} \mathbb{E}_{\mathbf{x}|\mathbf{y}} \{ \mathbf{x} \}, \quad (29)$$

1216

1217

and applying $f(\mathbf{y}) = \mathbf{A}^\dagger \mathbf{A}f(\mathbf{y}) + (\mathbf{I} - \mathbf{A}^\dagger \mathbf{A})f(\mathbf{y})$ with $v(\mathbf{y}) := f(\mathbf{y})$ yields eq. (27). Conversely, the objective in eq. (26) has the same value no matter the f satisfying eq. (27) and since at least one of them is solution of eq. (26), they all are. \square

1218

1219

Theorem 1. *In the case of noiseless measurements with $p(\mathbf{x})$ \mathcal{G} -invariant (Assumption 1), if the matrix $\mathbf{Q}_{\mathbf{A}_1} \triangleq \mathbb{E}_{g|\mathbf{A}_1} \{ (\mathbf{A}\mathbf{T}_g)^\top \mathbf{A}\mathbf{T}_g \} \in \mathbb{R}^{n \times n}$ has full rank for some split \mathbf{A}_1 , then the splitting method yields the same MMSE-optimal reconstructions as the supervised method, i.e.,*

1220

1221

$$f^*(\mathbf{y}_1, \mathbf{A}_1) = \mathbb{E}_{\mathbf{x}|\mathbf{y}_1, \mathbf{A}_1} \{ \mathbf{x} \}. \quad (9)$$

1222

1242 *Proof.*

$$\begin{aligned}
& \mathbb{E}_{\mathbf{y}} \{ \mathcal{L}_{\text{ES}}(\mathbf{y}, \mathbf{A}, f) \} \\
&= \mathbb{E}_{\mathbf{y}} \{ \mathbb{E}_g \{ \mathbb{E}_{\mathbf{y}_1, \mathbf{A}_1 | \mathbf{y}, \mathbf{AT}_g} \{ \| \mathbf{AT}_g f(\mathbf{y}_1, \mathbf{A}_1) - \mathbf{y} \|^2 \} \} \} \\
&= \mathbb{E}_{\mathbf{y}, g} \{ \mathbb{E}_{\mathbf{y}_1, \mathbf{A}_1 | \mathbf{y}, g} \{ \| \mathbf{AT}_g f(\mathbf{y}_1, \mathbf{A}_1) - \mathbf{y} \|^2 \} \} \\
&= \mathbb{E}_{\mathbf{y}, \mathbf{y}_1, \mathbf{A}_1, g} \{ \| \mathbf{AT}_g f(\mathbf{y}_1, \mathbf{A}_1) - \mathbf{y} \|^2 \} \\
&= \mathbb{E}_{\mathbf{x}, \mathbf{y}_1, \mathbf{A}_1, g} \{ \| \mathbf{AT}_g f(\mathbf{y}_1, \mathbf{A}_1) - \mathbf{AT}_g \mathbf{x} \|^2 \} \\
&= \mathbb{E}_{\mathbf{y}_1, \mathbf{A}_1, \mathbf{x}} \mathbb{E}_{g | \mathbf{y}_1, \mathbf{A}_1} \{ \| \mathbf{AT}_g (f(\mathbf{y}_1, \mathbf{A}_1) - \mathbf{x}) \|^2 \} \\
&= \mathbb{E}_{\mathbf{y}_1, \mathbf{A}_1, \mathbf{x}} \{ (f(\mathbf{y}_1, \mathbf{A}_1) - \mathbf{x})^\top \mathbf{Q}_{\mathbf{A}_1} (f(\mathbf{y}_1, \mathbf{A}_1) - \mathbf{x}) \},
\end{aligned}$$

where the third line use that $p(\mathbf{y}_1, \mathbf{A}_1 | \mathbf{y}, \mathbf{AT}_g) = p(\mathbf{y}_1, \mathbf{A}_1 | \mathbf{y}, g)$ as \mathbf{A} is fixed. The fifth line use the noiseless measurements assumption and the invariance of the distribution $p(\mathbf{x})$. The last line uses definition of $\mathbf{Q}_{\mathbf{A}_1}$. By applying Lemma 1, the global minimizer of the expected loss is given by:

$$f^*(\mathbf{y}_1, \mathbf{A}_1) = \mathbf{Q}_{\mathbf{A}_1}^\dagger \mathbf{Q}_{\mathbf{A}_1} \mathbb{E}_{\mathbf{x} | \mathbf{y}_1, \mathbf{A}_1} \{ \mathbf{x} \} + (\mathbf{I} - \mathbf{Q}_{\mathbf{A}_1}^\dagger \mathbf{Q}_{\mathbf{A}_1}) v(\mathbf{y}_1) \quad (30)$$

where $v : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is any function. Moreover, since \mathbf{Q} has full-rank, $\mathbf{Q}_{\mathbf{A}_1}^\dagger = \mathbf{Q}_{\mathbf{A}_1}^{-1}$ and then

$$f^*(\mathbf{y}_1, \mathbf{A}_1) = \mathbb{E}_{\mathbf{x} | \mathbf{y}_1, \mathbf{A}_1} \{ \mathbf{x} \}. \quad (31)$$

□

Proposition 1. *If the matrix $\bar{\mathbf{Q}}_{\mathbf{A}} \triangleq \mathbb{E}_{\mathbf{A}_1 | \mathbf{A}} \{ \mathbf{Q}_{\mathbf{A}_1} \} \in \mathbb{R}^{n \times n}$ is invertible and f minimizes $\mathbb{E}_{\mathbf{y}} \{ \mathcal{L}_{\text{ES}}(\mathbf{y}, \mathbf{A}, f) \}$. Then the reconstruction function*

$$\bar{f}(\mathbf{y}, \mathbf{A}) \triangleq \mathbb{E}_{\mathbf{y}_1, \mathbf{A}_1 | \mathbf{y}, \mathbf{A}} \{ \bar{\mathbf{Q}}_{\mathbf{A}}^{-1} \mathbf{Q}_{\mathbf{A}_1} f(\mathbf{y}_1, \mathbf{A}_1) \} \quad (10)$$

satisfies

$$\bar{f}(\mathbf{y}, \mathbf{A}) = \mathbb{E}_{\mathbf{y}_1, \mathbf{A}_1 | \mathbf{y}, \mathbf{A}} \{ \bar{\mathbf{Q}}_{\mathbf{A}}^{-1} \mathbf{Q}_{\mathbf{A}_1} \mathbb{E}_{\mathbf{x} | \mathbf{y}_1, \mathbf{A}_1} \{ \mathbf{x} \} \}. \quad (11)$$

where eq. (11) is a convex combination of MMSE estimators for different splittings.

Proof. By applying eq. (30) to f in the definition of \bar{f} we obtain:

$$\begin{aligned}
\bar{f}(\mathbf{y}, \mathbf{A}) &= \mathbb{E}_{\mathbf{y}_1, \mathbf{A}_1 | \mathbf{y}, \mathbf{A}} \left\{ \bar{\mathbf{Q}}_{\mathbf{A}}^{-1} \mathbf{Q}_{\mathbf{A}_1} (\mathbf{Q}_{\mathbf{A}_1}^\dagger \mathbf{Q}_{\mathbf{A}_1} \mathbb{E}_{\mathbf{x} | \mathbf{y}_1, \mathbf{A}_1} \{ \mathbf{x} \} + (\mathbf{I} - \mathbf{Q}_{\mathbf{A}_1}^\dagger \mathbf{Q}_{\mathbf{A}_1}) v(\mathbf{y}_1)) \right\} \\
&= \mathbb{E}_{\mathbf{y}_1, \mathbf{A}_1 | \mathbf{y}, \mathbf{A}} \{ \bar{\mathbf{Q}}_{\mathbf{A}}^{-1} \mathbf{Q}_{\mathbf{A}_1} \mathbb{E}_{\mathbf{x} | \mathbf{y}_1, \mathbf{A}_1} \{ \mathbf{x} \} \} + \mathbb{E}_{\mathbf{y}_1, \mathbf{A}_1 | \mathbf{y}, \mathbf{A}} \left\{ \bar{\mathbf{Q}}_{\mathbf{A}}^{-1} \mathbf{Q}_{\mathbf{A}_1} ((\mathbf{I} - \mathbf{Q}_{\mathbf{A}_1}^\dagger \mathbf{Q}_{\mathbf{A}_1}) v(\mathbf{y}_1)) \right\} \\
&= \mathbb{E}_{\mathbf{y}_1, \mathbf{A}_1 | \mathbf{y}, \mathbf{A}} \{ \bar{\mathbf{Q}}_{\mathbf{A}}^{-1} \mathbf{Q}_{\mathbf{A}_1} \mathbb{E}_{\mathbf{x} | \mathbf{y}_1, \mathbf{A}_1} \{ \mathbf{x} \} \}.
\end{aligned}$$

□

Corollary 1. *In order for the matrices $\mathbf{Q}_{\mathbf{A}_1}$ or $\bar{\mathbf{Q}}_{\mathbf{A}}$ to have full rank, it is necessary that \mathbf{A} is not equivariant:*

$$\exists g \in \mathcal{G}, \mathbf{AT}_g \neq \mathbf{T}_g \mathbf{A}. \quad (13)$$

Proof. Let's assume by contradiction that \mathbf{A} is equivariant with respect to \mathbf{T}_g

$$\mathbf{AT}_g = \mathbf{T}_g \mathbf{A}, \quad (32)$$

and let $\mathbf{x} \in \ker(\mathbf{A})$.

$$\begin{aligned}
\mathbf{Q}_{\mathbf{A}_1} \mathbf{x} &= \left(\mathbb{E}_{g | \mathbf{A}_1} \{ (\mathbf{AT}_g)^\top \mathbf{AT}_g \} \right) \mathbf{x} \\
&= \mathbb{E}_{g | \mathbf{A}_1} \{ (\mathbf{AT}_g)^\top \mathbf{AT}_g \mathbf{x} \} \\
&= \mathbb{E}_{g | \mathbf{A}_1} \{ (\mathbf{AT}_g)^\top \mathbf{T}_g \mathbf{A} \mathbf{x} \} \\
&= \mathbb{E}_{g | \mathbf{A}_1} \{ (\mathbf{AT}_g)^\top \mathbf{T}_g \mathbf{0} \} \\
&= \mathbf{0}
\end{aligned}$$

Therefore,

$$\ker(\mathbf{Q}_{\mathbf{A}_1}) \supseteq \ker(\mathbf{A}) \supsetneq \{ \mathbf{0} \}. \quad (33)$$

The matrix $\mathbf{Q}_{\mathbf{A}_1}$ has a non-trivial nullspace and thus cannot have full rank. Moreover, since this non-trivial nullspace is the same for all virtual operators \mathbf{AT}_g , then $\bar{\mathbf{Q}}_{\mathbf{A}}$ shares the same non-trivial nullspace.

□

1296 **Theorem 2.** *The reconstruction functions defined in points below are all equivariant as in eq. (15).*

1297 **1. Artifact removal network.** For a denoiser $\phi(\mathbf{x})$ equivariant in the sense of eq. (14),

$$1298 \quad f(\mathbf{y}, \mathbf{A}) = \phi(\mathbf{A}^\top \mathbf{y}), \text{ or } f(\mathbf{y}, \mathbf{A}) = \phi(\mathbf{A}^\dagger \mathbf{y}). \quad (16)$$

1300 **2. Unrolled network.** For $\phi(\mathbf{x})$ equivariant, any $\gamma \in \mathbb{R}$ and data fidelity $d(\mathbf{A}\mathbf{x}, \mathbf{y})$, with

$$1301 \quad \mathbf{x}_0 = \mathbf{0}, \quad \mathbf{x}_{k+1} = \phi(\mathbf{x}_k - \gamma \nabla_{\mathbf{x}_k} d(\mathbf{A}\mathbf{x}_k, \mathbf{y})) \quad (17)$$

1304 for $k = 0, \dots, L-1$ and $f(\mathbf{y}, \mathbf{A}) = \mathbf{x}_L$.

1305 **3. Reynolds averaging.** For a possibly non-equivariant reconstructor $r(\mathbf{y}, \mathbf{A})$, with

$$1307 \quad f(\mathbf{y}, \mathbf{A}) = \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} \mathbf{T}_g r(\mathbf{y}, \mathbf{A} \mathbf{T}_g). \quad (18)$$

1310 **4. Maximum a posteriori (MAP).** For a distribution $p(\mathbf{x})$ invariant as in eq. (5), with

$$1311 \quad f(\mathbf{y}, \mathbf{A}) = \operatorname{argmax}_{\mathbf{x} \in \mathbb{R}^n} \{p(\mathbf{x} \mid \mathbf{y}, \mathbf{A})\}. \quad (19)$$

1314 **5. Minimum mean squared error (MMSE).** For a distribution $p(\mathbf{x})$ invariant as in eq. (5),

$$1315 \quad f(\mathbf{y}, \mathbf{A}) = \mathbb{E}_{\mathbf{x} \mid \mathbf{y}, \mathbf{A}} \{\mathbf{x}\}. \quad (20)$$

1317 *Proof.* We prove each case separately.

1319 1. Denoting $\mathbf{A}^\times := \mathbf{A}^\top$ or $\mathbf{A}^\times := \mathbf{A}^\dagger$, eq. (16) gives, as $(\mathbf{A} \mathbf{T}_g)^\times = \mathbf{T}_g^{-1} \mathbf{A}^\times$,

$$1320 \quad f(\mathbf{y}, \mathbf{A} \mathbf{T}_g) = \phi(\mathbf{T}_g^{-1} \mathbf{A}^\times \mathbf{y}), \quad (34)$$

1322 and since $\phi(\mathbf{x})$ is equivariant, i.e., eq. (14) holds, it simplifies to eq. (15).

1323 2. We start by making the notation show the explicit dependency on \mathbf{y} and \mathbf{A} :

$$1325 \quad \mathbf{x}_0(\mathbf{y}, \mathbf{A}) = \mathbf{0}, \quad \mathbf{x}_{k+1}(\mathbf{y}, \mathbf{A}) = \phi\left(\mathbf{x}_k(\mathbf{y}, \mathbf{A}) - \gamma \nabla_{\mathbf{x}_k(\mathbf{y}, \mathbf{A})} d(\mathbf{A}\mathbf{x}_k(\mathbf{y}, \mathbf{A}), \mathbf{y})\right). \quad (35)$$

1327 and proceed to show that for $k = 0, \dots, L$ it holds that

$$1328 \quad \mathbf{x}_k(\mathbf{y}, \mathbf{A} \mathbf{T}_g) = \mathbf{T}_g^{-1} \mathbf{x}_k(\mathbf{y}, \mathbf{A}). \quad (36)$$

1329 For $k = 0$, it holds as

$$1330 \quad \mathbf{x}_0(\mathbf{y}, \mathbf{A} \mathbf{T}_g) = \mathbf{0} = \mathbf{T}_g^{-1} \mathbf{x}_0(\mathbf{y}, \mathbf{A}). \quad (37)$$

1331 Let's assume that eq. (36) holds for $k < L$. Applying it and the chain rule in eq. (35) yields

$$1333 \quad \mathbf{x}_{k+1}(\mathbf{y}, \mathbf{A} \mathbf{T}_g) = \phi\left(\mathbf{T}_g^{-1} (\mathbf{x}_k(\mathbf{y}, \mathbf{A}) - \gamma \nabla_{\mathbf{x}_k(\mathbf{y}, \mathbf{A})} d(\mathbf{A}\mathbf{x}_k(\mathbf{y}, \mathbf{A}), \mathbf{y}))\right). \quad (38)$$

1334 Finally, applying eq. (14) in this equation gives

$$1335 \quad \mathbf{x}_{k+1}(\mathbf{y}, \mathbf{A} \mathbf{T}_g) = \mathbf{T}_g^{-1} \mathbf{x}_{k+1}(\mathbf{y}, \mathbf{A}), \quad (39)$$

1336 and by induction, as $f(\mathbf{y}, \mathbf{A}) = \mathbf{x}_L(\mathbf{y}, \mathbf{A})$, eq. (15) holds.

1338 3. From eq. (18), it holds that

$$1340 \quad f(\mathbf{y}, \mathbf{A} \mathbf{T}_g) = \frac{1}{|\mathcal{G}|} \sum_{h \in \mathcal{G}} \mathbf{T}_h r(\mathbf{y}, \mathbf{A} \mathbf{T}_g \mathbf{T}_h), \quad (40)$$

1342 which, as the group action property holds $\mathbf{T}_g \mathbf{T}_h = \mathbf{T}_{gh}$, rewrites as

$$1343 \quad f(\mathbf{y}, \mathbf{A} \mathbf{T}_g) = \frac{1}{|\mathcal{G}|} \sum_{h \in \mathcal{G}} \mathbf{T}_h r(\mathbf{y}, \mathbf{A} \mathbf{T}_{gh}), \quad (41)$$

1346 Applying the change of variable $h' = gh$ in this equation gives

$$1347 \quad f(\mathbf{y}, \mathbf{A} \mathbf{T}_g) = \frac{1}{|\mathcal{G}|} \sum_{h \in \mathcal{G}} \mathbf{T}_{g^{-1}h} r(\mathbf{y}, \mathbf{A} \mathbf{T}_h), \quad (42)$$

1348 which finally, using group action property again $\mathbf{T}_{g^{-1}h} = \mathbf{T}_g^{-1} \mathbf{T}_h$, gives eq. (15).

1349

1350 4. Taking the negative natural logarithm in eq. (19) and using Bayes' theorem gives
 1351

$$1352 \quad f(\mathbf{y}, \mathbf{A}) = \underset{\mathbf{x} \in \mathbb{R}^n}{\operatorname{argmin}} \left\{ d(\mathbf{Ax}, \mathbf{y}) + \rho(\mathbf{x}) \right\}, \quad (43)$$

1354 where $d(\mathbf{Ax}, \mathbf{y}) = -\log p(\mathbf{y} \mid \mathbf{Ax})$ and $\rho(\mathbf{x}) = -\log p(\mathbf{x})$. Applying $\mathbf{x}' = \mathbf{T}_g \mathbf{x}$,
 1355

$$1356 \quad f(\mathbf{y}, \mathbf{AT}_g) = \mathbf{T}_g^{-1} \underset{\mathbf{x} \in \mathbb{R}^n}{\operatorname{argmin}} \left\{ d(\mathbf{Ax}, \mathbf{y}) + \rho(\mathbf{T}_g^{-1} \mathbf{x}) \right\}, \quad (44)$$

1358 and eq. (5) makes $\rho(\mathbf{x})$ invariant as well $\rho(\mathbf{T}_g^{-1} \mathbf{x}) = \rho(\mathbf{x})$. Therefore, eq. (15) holds.
 1359

1360 5. Let's assume that $p(\mathbf{x})$ and $p(\mathbf{A})$ are invariant in the sense of eq. (5). We first prove that
 1361

$$1361 \quad p(\mathbf{y}, \mathbf{AT}_g) = p(\mathbf{y}, \mathbf{A}). \quad (45)$$

1362 Using eq. (5), the invariance of $p(\mathbf{A})$ and the independence of \mathbf{x} and \mathbf{A} , we compute
 1363

$$\begin{aligned} 1364 \quad p(\mathbf{y}, \mathbf{AT}_g) &= \mathbb{E}_{\mathbf{x}} \{p(\mathbf{y}, \mathbf{AT}_g \mid \mathbf{x})\} = \mathbb{E}_{\mathbf{x}} \{p(\mathbf{y} \mid \mathbf{AT}_g, \mathbf{x}) p(\mathbf{AT}_g \mid \mathbf{x})\} \\ 1365 &= \mathbb{E}_{\mathbf{x}} \{p(\mathbf{y} \mid \mathbf{AT}_g, \mathbf{x}) p(\mathbf{AT}_g)\} = \mathbb{E}_{\mathbf{x}} \{p(\mathbf{y} \mid \mathbf{AT}_g \mathbf{x}) p(\mathbf{A})\} \\ 1366 &= \mathbb{E}_{\mathbf{x}} \{p(\mathbf{y} \mid \mathbf{Ax}) p(\mathbf{A})\} = \mathbb{E}_{\mathbf{x}} \{p(\mathbf{y} \mid \mathbf{A}, \mathbf{x}) p(\mathbf{A} \mid \mathbf{x})\} \\ 1367 &= \mathbb{E}_{\mathbf{x}} \{p(\mathbf{y}, \mathbf{A} \mid \mathbf{x})\} = p(\mathbf{y}, \mathbf{A}). \end{aligned}$$

1368 Next we start from
 1369

$$1370 \quad f(\mathbf{y}, \mathbf{AT}_g) = \mathbb{E}_{\mathbf{x} \mid \mathbf{y}, \mathbf{AT}_g} \{\mathbf{x}\}. \quad (46)$$

1371 and applying the integral formula for expectations gives
 1372

$$1373 \quad f(\mathbf{y}, \mathbf{AT}_g) = \int \mathbf{x} p(\mathbf{x} \mid \mathbf{y}, \mathbf{AT}_g) d\mathbf{x}. \quad (47)$$

1374 Using Bayes' formula and $p(\mathbf{y} \mid \mathbf{A}, \mathbf{x}) = p(\mathbf{y} \mid \mathbf{Ax})$, it becomes
 1375

$$1376 \quad f(\mathbf{y}, \mathbf{AT}_g) = \int \mathbf{x} \frac{p(\mathbf{y} \mid \mathbf{AT}_g \mathbf{x})}{p(\mathbf{AT}_g, \mathbf{y})} p(\mathbf{AT}_g, \mathbf{x}) d\mathbf{x}. \quad (48)$$

1378 By using eq. (45), the invariance of $p(\mathbf{A})$ and the independence of \mathbf{A} with \mathbf{x} , we obtain
 1379

$$1380 \quad f(\mathbf{y}, \mathbf{AT}_g) = \int \mathbf{x} \frac{p(\mathbf{y} \mid \mathbf{AT}_g \mathbf{x})}{p(\mathbf{A}, \mathbf{y})} p(\mathbf{A}) p(\mathbf{x}) d\mathbf{x}. \quad (49)$$

1382 With the change of variable $\mathbf{x}' = \mathbf{T}_g \mathbf{x}$, and since \mathbf{T}_g is unitary, we arrive at
 1383

$$1384 \quad f(\mathbf{y}, \mathbf{AT}_g) = \int \mathbf{T}_g^{-1} \mathbf{x} \frac{p(\mathbf{y} \mid \mathbf{Ax})}{p(\mathbf{A}, \mathbf{y})} p(\mathbf{A}) p(\mathbf{T}_g^{-1} \mathbf{x}) d\mathbf{x}. \quad (50)$$

1386 Finally, applying eq. (5) in this equation yields eq. (15).
 1387

□

1389 **Theorem 3.** If $f(\mathbf{A}, \mathbf{x})$ is an equivariant reconstructor, then ES is equivalent to the splitting loss
 1390

$$1391 \quad \mathcal{L}_{\text{ES}}(\mathbf{y}, \mathbf{A}, f) = \mathcal{L}_{\text{SPLIT}}(\mathbf{y}, \mathbf{A}, f). \quad (21)$$

1392 *Proof.* We start from eq. (7)

$$1394 \quad \mathcal{L}_{\text{ES}}(\mathbf{y}, \mathbf{A}, f) \triangleq \mathbb{E}_g \{\mathcal{L}_{\text{SPLIT}}(\mathbf{y}, \mathbf{AT}_g, f)\} \quad (51)$$

$$1395 \quad = \mathbb{E}_g \left\{ \mathbb{E}_{\mathbf{y}_1, \mathbf{A}_1 \mid \mathbf{y}, \mathbf{AT}_g} \left\{ \|\mathbf{AT}_g f(\mathbf{y}_1, \mathbf{A}_1) - \mathbf{y}\|^2 \right\} \right\} \quad (52)$$

1396 As \mathbf{A}_1 is a splitting of \mathbf{AT}_g , we can write $\mathbf{A}_1 = \mathbf{M} \mathbf{AT}_g$ for \mathbf{M} a splitting matrix. We obtain
 1397

$$1398 \quad \mathcal{L}_{\text{ES}}(\mathbf{y}, \mathbf{A}, f) = \mathbb{E}_g \left\{ \mathbb{E}_{\mathbf{M} \mid \mathbf{y}, g} \left\{ \|\mathbf{AT}_g f(\mathbf{My}, \mathbf{MAT}_g) - \mathbf{y}\|^2 \right\} \right\} \quad (53)$$

1400 Applying eq. (15) and cancelling out \mathbf{T}_g with \mathbf{T}_g^{-1} yields
 1401

$$1402 \quad \mathcal{L}_{\text{ES}}(\mathbf{y}, \mathbf{A}, f) = \mathbb{E}_g \left\{ \mathbb{E}_{\mathbf{M} \mid \mathbf{y}, g} \left\{ \|\mathbf{Af}(\mathbf{My}, \mathbf{MA}) - \mathbf{y}\|^2 \right\} \right\}. \quad (54)$$

1403 By dropping the expectation in g and rewriting $(\mathbf{My}, \mathbf{MA})$ as $(\mathbf{y}_1, \mathbf{A}_1)$, this yields in eq. (21). □