
Equivariant Bootstrapping for Uncertainty Quantification in Imaging Inverse Problems

Julián Tachella

Laboratoire de Physique, CNRS &
ENS de Lyon

Marcelo Pereyra

Heriot-Watt University &
Maxwell Institute for Mathematical Sciences

Abstract

Scientific imaging problems are often severely ill-posed and hence have significant intrinsic uncertainty. Accurately quantifying the uncertainty in the solutions to such problems is therefore critical for the rigorous interpretation of experimental results as well as for reliably using the reconstructed images as scientific evidence. Unfortunately, existing imaging methods are unable to quantify the uncertainty in the reconstructed images in a way that is robust to experiment replications. This paper presents a new uncertainty quantification methodology based on an equivariant formulation of the parametric bootstrap algorithm that leverages symmetries and invariance properties commonly encountered in imaging problems. Additionally, the proposed methodology is general and can be easily applied with any image reconstruction technique, including unsupervised training strategies that can be trained from observed data alone, thus enabling uncertainty quantification in situations where there is no ground truth data available. We demonstrate the proposed approach with a series of experiments and comparisons with alternative state-of-the-art uncertainty quantification strategies. In all our experiments, the proposed equivariant bootstrap delivers remarkably accurate high-dimensional confidence regions and outperforms the competing approaches in terms of estimation accuracy, uncertainty quantification accuracy, and computing time. These empirical findings are supported by a detailed theoretical analysis of equivariant bootstrap for linear estimators.

1 Introduction

Digital images inform decisions that have a major impact on the economy, society, and the environment (e.g., in medicine, agriculture, forestry, astronomy, defense). Early uses of images as evidence in decision-making or science were predominantly qualitative. However, modern strategies increasingly adopt a so-called quantitative imaging approach which recognises images as high-dimensional physical measurements and seeks to use the images as quantitative evidence. The approach is particularly prominent in scientific applications.

Performing inference in imaging problems usually requires solving a high-dimensional inverse problem that is severely ill-posed. In particular, many commonly encountered tasks related to image restoration and reconstruction involve performing inference on an unknown image x_* taking values in a signal set $\mathcal{X} \subset \mathbb{R}^n$, from a single measurement $y \in \mathbb{R}^m$ modelled as a realization of

$$Y \sim P(Ax_*) \quad (1)$$

where P is a statistical model describing stochastic aspects of the data acquisition process (e.g., Gaussian or Poisson measurement noise) and $A \in \mathbb{R}^{m \times n}$ describes deterministic instrumental aspects of the observation process. For example, such models are frequently encountered in problems related to image deblurring, inpainting, super-resolution, compressive sensing reconstruction, and tomographic reconstruction (see, e.g., (Kaipio and Somersalo, 2005; Ongie et al., 2020)), as well as the examples provided in Appendix A). Unfortunately, the models that accurately describe the data acquisition process are usually either weakly or not identifiable, and hence there is significant uncertainty about x_* after observing y .

State-of-the-art imaging methods address this difficulty by intimately combining computational imaging and machine learning techniques, which learn to recover x_* from y by leveraging information from large training datasets (Ongie et al., 2020; Mukherjee et al., 2023). Early learning-based imaging methods required abun-

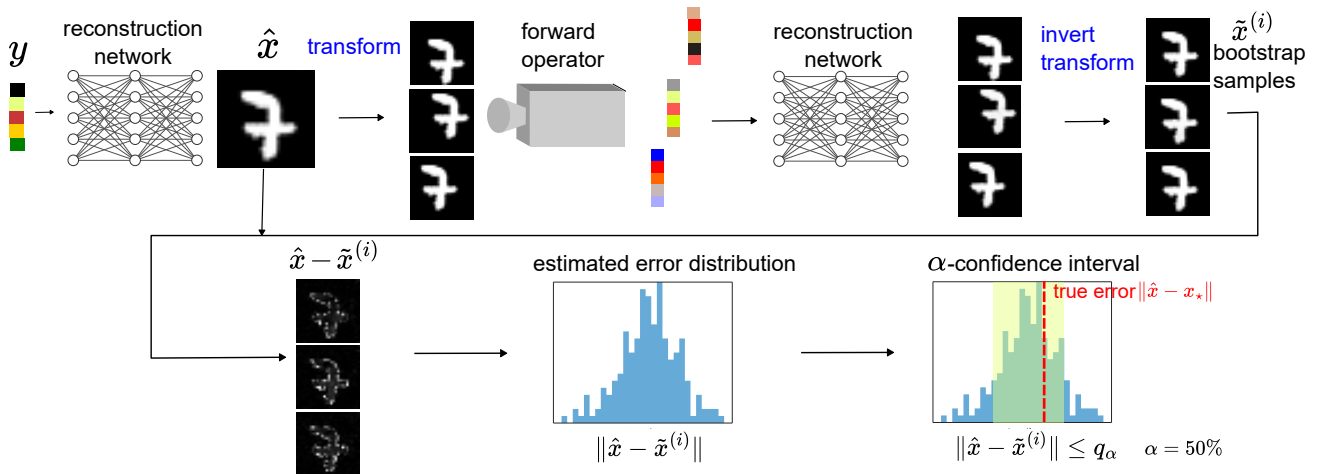


Figure 1: **Illustration of the equivariant bootstrap method.** Additional measurement vectors are bootstrapped by applying random transformations such as translations and/or rotations, which are then used to quantify the error associated with the reconstruction network.

dant ground truth data for training. However, new unsupervised strategies exploit symmetries or equivariance properties in the problem to learn estimators of x_* from the observed data alone (Chen et al., 2023).

Despite sustained progress in image estimation accuracy, important aspects of quantitative imaging sciences are still in their infancy. In particular, most imaging methods cannot reliably quantify the uncertainty in the delivered solutions, which is essential for the rigorous interpretation of experiments and robust interfacing of imaging pipelines with decision-making processes. Furthermore, as we illustrate through numerical experiments in Section 5, even the most powerful uncertainty quantification (UQ) methods currently available cannot meaningfully quantify the uncertainty in the reconstructed images in a manner that is robust to frequentist validation and experiment replications. This critical limitation hinders the value of images as quantitative evidence for decision-making and science.

This paper presents the following main contributions:

1. We propose the equivariant bootstrap algorithm, a new parametric bootstrapping technique that exploits symmetries in the problem in order to construct accurate confidence regions for x_* , even in situations where the model is not identifiable.
2. We present a theoretical analysis of the equivariant bootstrap for the case of a linear estimator, which provides clear insights into the key factors driving the accuracy of the method.
3. We demonstrate the effectiveness of the proposed equivariant bootstrap through a series of experiments involving three different imaging inverse

problems, where the method consistently outperforms alternative strategies from the state-of-the-art in terms of estimation accuracy, uncertainty quantification accuracy, and computing time.

2 Related Work

Bayesian Methods Modern imaging methods rely strongly on the Bayesian statistical framework to perform UQ tasks. In particular, Bayesian imaging strategies with data-driven priors encoded by neural networks have received a lot of attention lately. Within this context, state-of-the-art methods involve the prior through its score function, which can be related to an image denoising network via Tweedie’s formula (see, e.g., plug-and-play ULA (Laumont et al., 2022), and the denoising diffusion methods DPS (Chung et al., 2022), DDRM (Kawar et al., 2022), and DiffDIPR (Zhu et al., 2023)). Alternatively, many modern methods rely on deep generative architectures; e.g., variational autoencoders (Holden et al., 2022; Zhang et al., 2021), generative adversarial networks (Arridge et al., 2019), and normalizing flows (Altekrüger and Hertrich, 2023). These can be used to encode the prior within a Bayesian model or to approximate the posterior distribution directly. Otherwise, some methods perform UQ via Laplace approximations (e.g., (Antorán et al., 2022)).

Stein’s Unbiased Risk Estimator (SURE) SURE is a highly effective approach for estimating the mean squared error (MSE) in denoising problems (Stein, 1981). The method can be extended to other inverse problems (Eldar, 2008), but it can only reliably quantify the error in the range space of the forward operator. As a result, it is not directly useful for UQ in inverse

problems with a non-trivial nullspace.

Conformal Prediction Conformal prediction is a model-free UQ approach that leverages a calibration dataset and exchangeability to build marginally exact confidence intervals (i.e., probabilities are computed w.r.t. the joint distribution of the unknown image and the measurement data, without conditioning of the observed measurement data). Several recent papers propose conformal methods for imaging problems (see, e.g., Angelopoulos et al. (2022)), focusing predominantly on UQ tasks at the scale of a single pixel. Scaling conformal prediction to large image structures requires careful construction of the conformity function that determines the shape of the confidence region, a difficulty that is usually addressed by using the regions provided by a different UQ technique and applying a conformal calibration step (see, e.g., (Narnhofer et al., 2022)). The method proposed in this paper could be easily combined with a conformal correction step if a calibration set is available.

Bootstrap Methods Bootstrap resampling is a powerful general strategy for assessing the variability of estimators and other statistics w.r.t. the data sampling distribution (Efron and Tibshirani, 1994). Although widely used for UQ in other disciplines, the bootstrap is difficult to apply directly to imaging problems, as these are often severely ill-posed because of identifiability issues in the likelihood function. The method proposed in this paper builds on the parametric formulation of the bootstrap, which we adapt in order to tackle models that are weakly or not identifiable (i.e., with a forward operator A that is not invertible).

Equivariant Imaging Equivariance plays an important role in imaging inverse problems, enabling unsupervised learning (Chen et al., 2021, 2022) and improving the generalization of the estimators (Mohan et al., 2019; Celledoni et al., 2021) (e.g., see the recent tutorial Chen et al. (2023)). The proposed bootstrap endows the equivariant imaging approach (Chen et al., 2021, 2022) with powerful UQ capabilities.

3 Parametric Bootstrap

The standard parametric bootstrap method probes the accuracy of $\hat{x}(y)$ by constructing an i.i.d sample $\{\tilde{x}^{(1)}, \dots, \tilde{x}^{(n)}\}$ which approximates the sampling distribution of $\hat{x}(Y)$ when $Y \sim P(Ax_*)$. The sample is constructed as follows: for any $i = \{1, \dots, N\}$

1. Given $\hat{x}(y)$, draw a bootstrap measurement $\tilde{y}^{(i)}$ from the sampling distribution

$$\tilde{Y}^{(i)} \sim P(A\hat{x}(y)).$$

2. Given the bootstrap measurement $\tilde{y}^{(i)}$, compute

$$\tilde{x}^{(i)} = \hat{x}(\tilde{y}^{(i)}).$$

The realized sample $\{\tilde{x}^{(1)}, \dots, \tilde{x}^{(n)}\}$ can then be used, for example, to construct confidence regions for x_* . In our experiments, we consider the confidence region

$$\mathcal{C}_\alpha = \{x \in \mathcal{X} : \|x - \hat{x}(y)\|_2^2 < q_\alpha\}, \quad (2)$$

with q_α set to the top α -quantile of the sample $\{\|\tilde{x}^{(i)} - \hat{x}(y)\|_2^2\}_{i=1}^n$. One can also design more compact and geometrically interesting regions by using the sample covariance of $\{\tilde{x}^{(1)}, \dots, \tilde{x}^{(n)}\}$. The sample is also useful for exploring the accuracy of $\hat{x}(Y)$ (e.g., bias, variance).

Although widely used in other settings, the parametric bootstrap often performs poorly in imaging problems. In particular, it severely underestimates the uncertainty in the solution because of the bias arising from using an estimate $\hat{x}(y)$ of x_* in the sampling distribution of Y , as opposed to the true value x_* which generated y .

For illustration, consider a severely under-determined problem of the form $y = Ax_* + e$ where $m \ll n$ and e is a realization of Gaussian noise. In such problems, there is significant uncertainty about the solution on the nullspace of A . For simplicity, consider the linear estimator $\hat{x}(y) = My$ for all $y \in \mathbb{R}^m$, that seeks to invert A on \mathcal{X} . Suppose that we use the conventional parametric bootstrap to estimate the MSE $\|\hat{x}(y) - x_*\|_2^2$. For the considered setup, we have that the true MSE is given by

$$\|\hat{x}(y) - x_*\|_2^2 = \|MAx_* + Me - x_*\|_2^2, \quad (3)$$

which is approximately

$$\|\hat{x}(y) - x_*\|_2^2 \approx \|(MA - I)x_*\|_2^2, \quad (4)$$

when the error stemming from the noise e is negligible relative to the error from incorrectly inverting A on \mathcal{X} . Conversely, in that same scenario, the parametric bootstrap produces the estimate

$$E_{\tilde{Y}}\{\|\hat{x}(\tilde{Y}) - \hat{x}(y)\|_2^2\} \approx \|(MA - I_n)MAx_*\|_2^2, \quad (5)$$

which can be potentially much smaller than the true error $\|(MA - I)x_*\|_2^2$. For example, if MA is a projection matrix, the bootstrap error $\|(MA - I_n)MAx_*\|_2^2 = 0$ regardless of the true error. This drawback of the parametric bootstrap is illustrated through a series of experiments in Section 4, where we consistently observe an underestimation of the uncertainty in the solution.

4 Equivariant Bootstrap

4.1 Proposed method

The proposed equivariant bootstrap method is a parametric bootstrapping technique designed for situations in which the set of signals \mathcal{X} is known to be invariant to a certain group of transformations. The aim is to leverage these symmetries to mitigate the bias arising from using an estimate $\hat{x}(y)$ of x_* in the approximation of the sampling distribution of Y . The method is particularly useful in situations where access to ground truth data is difficult or not possible, and it is necessary to quantify the uncertainty in the delivered solutions from the observed measurement data alone.

Let \mathcal{G} be a compact group acting on \mathcal{X} , whose action is represented by the invertible linear mappings T_g . We assume that \mathcal{X} is \mathcal{G} -invariant; i.e., for all $x \in \mathcal{X}$ and all $g \in \mathcal{G}$, we have that

$$T_g x \in \mathcal{X}.$$

In a manner akin to the conventional parametric bootstrap, the proposed equivariant bootstrap method probes the accuracy of $\hat{x}(y)$ by constructing an i.i.d sample $\{\tilde{x}^{(1)}, \dots, \tilde{x}^{(n)}\}$ which approximates the sampling distribution of $\hat{x}(Y)$ when $Y \sim P(Ax_*)$. The sample is constructed as follows: for any $i = \{1, \dots, N\}$

1. Draw a random transform g_i uniformly from \mathcal{G} .
2. Given g_i and $\hat{x}(y)$, draw a bootstrap measurement $\tilde{y}^{(i)}$ from the sampling distribution

$$\tilde{Y}^{(i)} \sim P(AT_{g_i}\hat{x}(y)).$$

3. Given the generated bootstrap measurement $\tilde{y}^{(i)}$ and g_i , compute

$$\tilde{x}^{(i)} = T_{g_i}^{-1}\hat{x}(\tilde{y}^{(i)}).$$

where $T_{g_i}^{-1}$ inverts the action g_i .

It is useful to view the equivariant bootstrap procedure as a data augmentation strategy that constructs an augmented sampling distribution

$$G \sim \mathcal{U}(\mathcal{G}), \quad (Y|G = g, x) \sim P(AT_g x),$$

where G is an auxiliary variable that takes values uniformly in \mathcal{G} . Without loss of generality, we view the observed data y is a realization of this augmented model with $g = 1_{\mathcal{G}}$, the identity element of \mathcal{G} .

As mentioned previously, in many imaging inverse problems the uncertainty about the solution stems predominantly from the fact that A has a large nullspace, which

leads to significant non-identifiability issues in the likelihood function. The equivariant bootstrap exploits the fact that when A is not \mathcal{G} -equivariant, the composition AT_g can span a different subspace than A , such that the operator resulting from averaging AT_g over \mathcal{G} can be full rank. The equivariant bootstrap leverages this property to probe the variability of the estimator $\hat{x}(Y)$ and characterise the uncertainty in x_* .

4.2 Analysis With a Linear Estimator

In order to develop an intuition for the equivariant bootstrap and the effect of introducing the actions \mathcal{G} , we consider a noise-free problem $y = Ax_*$ with $m \ll n$, where there is significant uncertainty about the solution because A has non-trivial nullspace. For simplicity, we consider again a linear estimator $\hat{x}(y) = My$ for all $y \in \mathbb{R}^m$ that seeks to invert A on \mathcal{X} , and use the proposed equivariant bootstrap to estimate the MSE $\|\hat{x}(y) - x_*\|_2^2$. We assume that \mathcal{X} is a low-dimensional subspace that is \mathcal{G} -invariant for some compact group \mathcal{G} . For example, \mathcal{G} could be the group of cyclic shifts associated with circulant matrices T_g which shift the image by g pixels, with $g \in \{0, \dots, n-1\}$ (see (Tachella et al., 2023a, Section 4.1) for other commonly encountered examples). In practice, \mathcal{X} is unknown so the estimator M will fail to perfectly invert A . We let $B = MA$, and use the decomposition $B = B_* + R$, where B_* is a \mathcal{G} -equivariant orthogonal projection onto \mathcal{X} , and $R = B - B_*$ is a residual term representing the estimator error. We will see that the equivariant bootstrap is effective when R is small but also far from being \mathcal{G} -equivariant. In that case, the actions of \mathcal{G} will “average out” the bias introduced by R and improve the quality of the bootstrap sample as a result.

We now introduce some elements of linear representation theory that are essential to our analysis (Serre, 1977). For any compact group \mathcal{G} , T_g with $g \in \mathcal{G}$ admits a linear representation $T_g = F^{-1}\Lambda_g F$ where Λ_g is a block-diagonal matrix and F is an orthonormal basis on \mathbb{C}^n associated with \mathcal{G} , but independent of g . For example, when \mathcal{G} represents the group of cyclic shifts, we have that F is the discrete Fourier transform and Λ_g is a diagonal matrix containing the Fourier transform of the discrete shift operation (see (Tachella et al., 2023a, Section 4.1) for more details and other examples). More generally, any \mathcal{G} -equivariant matrix $C \in \mathbb{R}^{n \times n}$ also admits a linear representation $C = F^{-1}\Lambda_C F$, where Λ_C is a block-diagonal with the same structure or support as Λ_g . It follows that any \mathcal{G} -equivariant matrix C commutes with T_g for any $g \in \mathcal{G}$. Conversely, R is not \mathcal{G} -equivariant, and therefore FRF^{-1} is usually a dense matrix that does not commute with T_g for any $g \in \mathcal{G}$.

We analyze the estimation by equivariant bootstrapping of the MSE between $\hat{x}(y) - x_*$, which also underpins

\mathcal{C}_α . For the considered setup,

$$\|\hat{x}(y) - x_\star\|^2 = \|Bx_\star - x_\star\|^2 = \|Rx_\star\|^2. \quad (6)$$

The equivariant bootstrap estimate of (6) is given by¹

$$\begin{aligned} \mathbb{E}_{\tilde{Y}} \|\tilde{x}(\tilde{Y}) - Bx_\star\|_2^2 &= \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} \|T_g^{-1} B T_g B x_\star - B x_\star\|_2^2 \\ &= \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} x_\star^\top B^\top (I_n + T_g^{-1} B^\top B T_g - 2T_g^{-1} B T_g) B x_\star \\ &= x_\star^\top B^\top (I_n + \Pi_{\mathcal{G}}(B^\top B) - 2\Pi_{\mathcal{G}}(B)) B x_\star, \end{aligned}$$

where for any matrix $C \in \mathbb{R}^{n \times n}$, $\Pi_{\mathcal{G}}(C) \triangleq \sum_g T_g^{-1} C T_g / |\mathcal{G}|$ denotes the so-called Reynolds averaging operator associated with \mathcal{G} (Serre, 1977). $\Pi_{\mathcal{G}}$ is a projection operator onto the linear subspace of matrices that are diagonal on the basis F associated with \mathcal{G} . Consequently, $\Pi_{\mathcal{G}}(C) = C$ for any matrix $C \in \mathbb{R}^{n \times n}$ that is \mathcal{G} -equivariant. Note that if C is not \mathcal{G} -equivariant, $\Pi_{\mathcal{G}}$ will potentially significantly shrink the Frobenius norm of C by setting all its off-block-diagonal elements (in the basis F) to zero.

By using the decomposition $B = B_\star + R$ introduced previously and that $\Pi_{\mathcal{G}}$ is a linear operator, we obtain

$$\begin{aligned} \mathbb{E}_{\tilde{Y}} \|\tilde{x}(\tilde{Y}) - Bx_\star\|_2^2 &= \underbrace{\|Rx_\star\|^2}_{\text{true error}} - \underbrace{\|Bx_\star\|_{B_\star}^2}_{\text{bias term 1}} \\ &\quad + \underbrace{x_\star^\top B^\top \Pi_{\mathcal{G}}(R^\top R + R^\top B_\star + B_\star R - 2R) B x_\star}_{\text{bias term 2}} \end{aligned}$$

where we have also used the that B_\star is the orthogonal projection onto \mathcal{X} and therefore $B_\star x_\star = x_\star$, $B_\star^\top = B_\star$ and $\Pi_{\mathcal{G}}(B_\star) = B_\star$.

We now analyze the two bias terms which can lead to a biased estimation of the true error. If the estimator is measurement consistent, i.e., if it verifies $A\hat{x}(y) = y$, then the first bias term is zero, see appendix B for a detailed derivation. This is a simple property that most estimators verify in practice.

The second bias term depends on whether the matrix R is \mathcal{G} -equivariant or not. Due to the averaging operator, we can rewrite this second term as

$$x_\star^\top B^\top F^{-1} \Lambda F B x_\star,$$

where F is an orthonormal basis and Λ is a block-diagonal matrix with

$$0 \leq \|\Lambda\|_{\text{Frob}} \leq \|R^\top R + R^\top B_\star + B_\star R - 2R\|_{\text{Frob}}.$$

¹We use a finite group here for the sake of simplicity, however, the results presented here apply also to any infinite compact group by replacing sums for integrals, see (Serre, 1977, Chapter 4).

If the error term R is not \mathcal{G} -equivariant (and thus dense on the basis F), then the norm of Λ will be much smaller than the norm of R , and consequently, the second bias term will be small. Conversely, if R is \mathcal{G} -equivariant, Λ will have a norm similar to R and the averaging will not mitigate the second bias term.

We conclude that the equivariant bootstrap can potentially significantly reduce the bias inherent in parametric bootstrapping, especially when the estimator is measurement consistent, the signal \mathcal{X} is \mathcal{G} -invariant and low-dimensional and either the forward operator and/or the estimator are not \mathcal{G} equivariant. Natural signal sets often exhibit a range of symmetries (e.g., invariance to rotation, translations, permutations, etc. (Tachella et al., 2023a)). The key to leveraging these symmetries to reduce the bootstrap bias is to identify a symmetry group \mathcal{G} that ‘‘averages out’’ the estimation error while leaving the signal set \mathcal{X} unchanged. This arises when \mathcal{X} is \mathcal{G} -invariant but the estimation error is far from \mathcal{G} -equivariance.

5 Experimental Results

We now present a series of numerical experiments and comparisons with alternative statistical UQ imaging approaches from the state of the art. We use each method to compute a confidence region for x_\star , derived from the pivotal statistic $\|x_\star - \hat{x}(Y)\|_2^2$ related to the estimation MSE. We evaluate the accuracy of these confidence regions by calculating the empirical coverage probabilities on a test set, as measured by the proportion of test images that lie within the confidence regions for a range of specified confidence levels between 0% and 100%. We perform three types of experiments: compressed sensing image reconstruction, sparse-angle tomography, and image inpainting. All our experiments were performed using the `deepinv` open-source library (Tachella et al., 2023b) on a local cluster with 4 NVIDIA RTX 3090 GPUs.

Compressed Sensing We use the MNIST dataset, which consists of images of 28×28 pixels. We use 6×10^4 images for training, and 384 images for testing. The forward operator is defined using $m = 256$ measurements with entries sampled from a Gaussian distribution with zero mean and variance $1/m$. Measurements are corrupted with Gaussian noise of standard deviation 0.05.

Inpainting The DIV2K dataset (Agustsson and Timofte, 2017) contains 1000 high-resolution RGB images. We generate a dataset of 2×10^4 crops of 256×256 pixels for training and 200 for testing. The measurement data is obtained by applying an inpainting mask with binary entries sampled from a Bernoulli distribution with a probability of 0.5 and adding white Gaussian

Table 1: Average test PSNR in dB for the evaluated methods.

	Diffusion (DDRM)	Diffusion (DiffPIR)	ULA	Proposed bstrap (unsup. model)	Proposed bstrap (sup. model)
C. Sensing (MNIST)	-	-	28.54 ± 2.25	34.11 ± 2.09	33.9 ± 2.32
Inpainting (DIV2K)	32.27 ± 3.95	30.51 ± 3.74	30.52 ± 3.35	31.56 ± 4.12	32.47 ± 3.87
Tomography (LIDC)	-	37.02 ± 0.79	35.85 ± 0.54	37.38 ± 0.65	41.03 ± 0.91

noise with a standard deviation of 0.05.

Sparse-Angle Tomography The LIDC-IDRI dataset (Armato III et al., 2011) consists of computed tomography scans from 1010 patients. We re-scale the 2D slices to a size of 256×256 pixels and use the central slices of 100 patients for testing and 5000 slices from the remaining patients for training. The measurement data are obtained by using 40 projections, taken at equally-distanced angles, corrupted with Gaussian noise of standard deviation 0.1.

Moreover, we report comparisons between our method and the following techniques from the state of the art:

- Diffusion methods:** We evaluate two recently proposed score-based diffusion models, namely DDRM (Kawar et al., 2022) and DiffPIR Zhu et al. (2023). These are Bayesian posterior sampling strategies that rely on pre-trained denoising score-matching networks that encode the prior information available. We use the denoiser network in (Choi et al., 2021) for RGB images and the one in (Zhang et al., 2017b) for single-channel images. By running the diffusion model multiple times we obtain a set of Monte Carlo samples from the posterior distribution associated with each method, which we then use to approximate the posterior mean $\hat{x}(y) = \mathbb{E}\{x|y\}$ as well as a posterior credible region based on an ℓ_2 -ball around $\hat{x}(y)$, derived from the posterior distribution of the error $\|x - \hat{x}(y)\|_2^2$.
- Unadjusted Langevin Algorithm:** We consider the plug-and-play Unadjusted Langevin Algorithm (ULA) introduced by Laumont et al. (2022). This Bayesian posterior sampling strategy implements a Langevin MCMC scheme with a pre-trained MMSE denoiser in lieu of the gradient of the log-prior density. For stability, the denoiser is trained with a controlled Lipschitz constant. We use the stable DnCNN denoiser introduced in (Terris et al., 2020). As with the diffusion methods, we use the Monte Carlo samples generated by ULA to approximate $\hat{x}(y) = \mathbb{E}\{x|y\}$ and to compute a posterior

credible region based on an ℓ_2 -ball around $\hat{x}(y)$.

- Equivariant bootstrap:** We evaluate the proposed method on estimator networks $\hat{x}(y)$ trained to minimize the reconstruction MSE. We consider both standard supervised training using a dataset pairs $\{(x_i, y_i)\}_i$, as well as a unsupervised training using the method in Chen et al. (2022) which only requires a dataset of measurement vectors $\{y_i\}_i$. We consider two forms of group actions: rotations and two-dimensional shifts. Rotations are sampled from a Gaussian distribution with zero mean and standard deviation of σ_θ , while horizontal and vertical shifts are sampled from a uniform distribution on $[-\Delta t, \Delta t]$ pixels. In practice, we compute the MSE samples as $\|T_{g_i} \hat{x}(y) - \hat{x}(\tilde{y}^{(i)})\|^2$, avoiding² the use of the inverse transform $T_{g_i}^{-1}$.
- Naive bootstrap:** For completeness, we also evaluate the conventional parametric bootstrap introduced in Section 3, based on the sampling distribution $\tilde{Y} \sim P(A\hat{x}(y))$, without using a group action. We use the bootstrap sample to determine the sampling distribution of the error $\|\hat{x}(\tilde{Y}) - \hat{x}(y)\|_2^2$. The only source of randomness in this case is the stochasticity from P (e.g., measurement noise).

All methods are used to generate 100 independent Monte Carlo samples, from which we compute the confidence regions for various confidence levels (in the case of PnP-ULA we achieve quasi-independent samples by applying thinning). Using a larger number of Monte Carlo samples was not possible because of the large amount of repetitions and deep neural function evaluations (NFEs) required by the diffusion models. Table 2 compares the computational load of the evaluated methods in terms of the number of NFEs. Diffusion-based methods require 100 NFEs per sample, and they are the most computationally demanding method across the evaluated algorithms. ULA requires a single NFE per MCMC iteration, however, we use a 1-in-30 thinning

²This bypasses interpolation or border artifacts associated with rotations that are not multiple of 90 degrees.

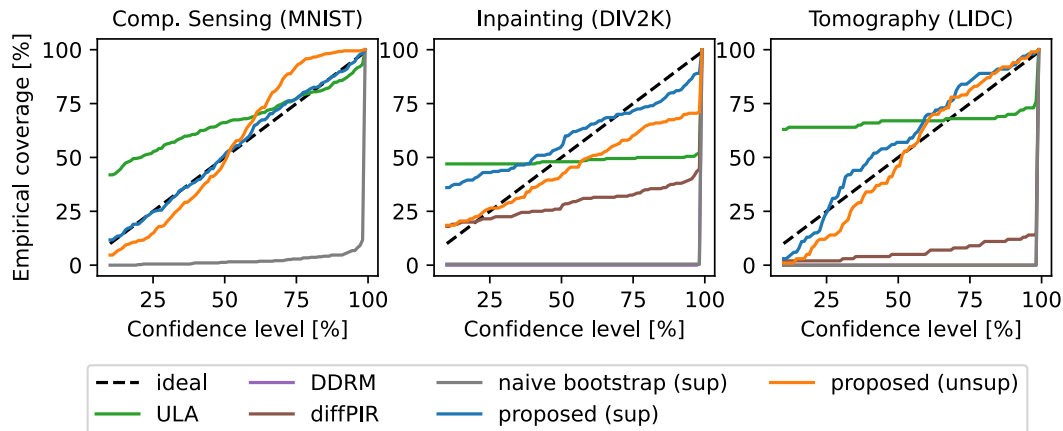


Figure 2: Coverage plots for all the evaluated methods on 3 inverse problems. Methods with a coverage closer to the dotted line provide a more precise quantification of the uncertainty of the estimates.

factor to reduce the correlation in the chain, so the cost is approximately 30 NFEs per sample. Remarkably, the bootstrap method requires only a single NFE per sample and thus can obtain uncertainty estimates in the order of seconds even for large images.

Table 2: Neural function evaluations (NFEs) per Monte Carlo (MC) sample.

Method	Diffusion	ULA	Bootstrap
NFEs/MC sample	100	30	1

Table 1 shows the peak signal-to-noise ratio (PSNR) for each evaluated method and each imaging problem considered, while Figure 2 shows the coverage probabilities of their delivered confidence regions. Examples of measurements and reconstructed images are included in Appendix D. For these experiments, the hyper-parameters of all the methods (bootstrap parameters in Table 3, regularisation parameter in ULA, etc.), were calibrated on a small evaluation set of 16 images per problem.

Notice that while most competing methods produce relatively accurate UQ results for the moderately low-dimensional images of the MNIST dataset, they perform poorly in the high-dimensional setting of DIV2K and LIDC. Conversely, the proposed equivariant bootstrap algorithm delivers remarkably accurate UQ results even in large-scale problems. ULA tends to provide confidence regions that are only accurate at a specific confidence level, and fail at other levels, whereas diffusion-based methods provide over-confident confidence regions that severely underestimate the uncertainty in the solution. The naive bootstrap also severely underestimates uncertainty, as discussed in Section 3.

Remarkably, the proposed method, applied with a supervised reconstruction network, simultaneously

achieves the most accurate uncertainty quantification results across all tasks as well as the highest image estimation accuracy. The proposed approach also delivers accurate uncertainty estimates when used with an unsupervised network that has been trained from observed data alone, thus removing the need for extensive ground-truth data (we only require a small calibration dataset to set hyperparameters σ_θ and Δt).

Diffusion-based methods provide the most accurate point estimation results in the inpainting problem. This is related to the fact that the denoisers underpinning the diffusion method were trained on large datasets of natural images which are similar to those in the DIV2K dataset. However, their performance deteriorates in domain-specific images such as tomography scans, even after fine-tuning. For the tomography task, we fine-tuned the DRUNet denoiser Zhang et al. (2017b) used by diffPIR on the LIDC training set, resulting in an improvement of 0.3 dB in test PSNR (note that Table 1 shows the performance after fine-tuning).

Table 3: Bootstrap parameters.

	Unsup. model	Sup. model
C. Sensing	$\sigma_\theta = 4$	$\sigma_\theta = 0$
	$\Delta t = 3$	$\Delta t = 3$
Inpainting	$\sigma_\theta = 5$	$\sigma_\theta = 5$
	$\Delta t = 10$	$\Delta t = 10$
Tomography	$\sigma_\theta = 10$	$\sigma_\theta = 8$
	$\Delta t = 5$	$\Delta t = 0$

Lastly, Figure 3 show the reconstructions, the estimation of the marginal per-pixel error (as calculated by the standard deviation), and the true absolute error for each of the evaluated methods. In order to avoid boundary issues in the estimation of the per-pixel errors, we restricted the rotations in the equivariant

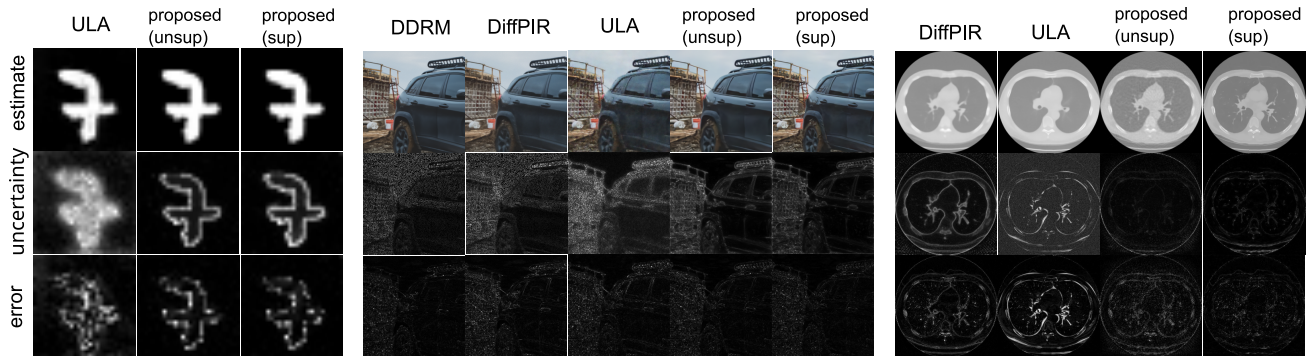


Figure 3: Per-pixel estimation of the error by the evaluated methods. The last row shows the true estimation error. The reader can zoom in to observe the details of the uncertainty maps.

bootstrap method to multiples of 90 degrees, which have an exact inverse. We observe that the per-pixel error estimates³ produced by the proposed approach are in good agreement with the true per-pixel errors, whereas the competing methods deliver per-pixel error estimates that are less accurate.

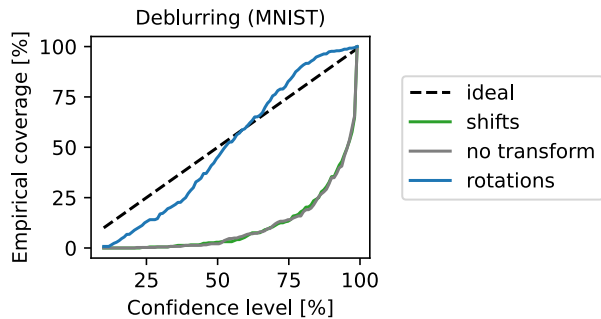


Figure 4: Coverage plots for the anisotropic deblurring problem obtained by bootstrapping with shifts only, rotations only, and no transformations (noise only).

Role of Equivariance We now present a final experiment designed to demonstrate the role of the group action in the uncertainty estimates. We consider a deblurring inverse problem with a highly anisotropic blur kernel, using the MNIST dataset (see appendix C for more details). We use the equivariant bootstrap with an estimator \hat{x} that is implemented by using a fully convolutional network, which is by construction shift-equivariant and is trained in a supervised fashion. Figure 4 shows the coverage obtained by the proposed equivariant bootstrap technique when \mathcal{G} includes: i) shifts only ($\Delta t = 5$ and $\sigma_\theta = 0$), and ii) rotations only ($\Delta t = 0$ and $\sigma_\theta = 5$). As suggested by the analysis in Section 4.2, the bootstrap estimates obtained with

³This uncertainty visualization analysis can be easily conducted in a multi-resolution manner (see, e.g., Laumont et al. (2022)).

the shifts obtain the same performance as the naive bootstrap approach, thus under-estimating the true error, as the composition of the forward operator and the reconstruction network is \mathcal{G} -equivariant. Conversely, using rotations provides excellent UQ results.

6 Limitations and future extensions

The performance of the proposed method is sensitive to the choice of group action. This involves selecting the type of group action (e.g., rotation, shifts), but also choosing the specific actions that will be included in \mathcal{G} (e.g., angles, magnitude of shifts). This calibration can be done on a small evaluation set. However, in some applications, obtaining a reliable evaluation set might be difficult. We leave for future work the study of automatic calibration techniques that can completely bypass the need for ground-truth data.

The proposed method also requires knowledge of the forward operator and noise distribution, which might be mildly misspecified or not fully known in some applications. We leave the analysis of the impact of model misspecification for future work. Future work could also explore blind and semi-blind variants of the proposed equivariant bootstrap.

Beyond imaging problems, we envisage applying the proposed method to inverse problems involving other data structures, such as time series by relying on translations as group actions, or graphs, where permutations are often used as group symmetries.

7 Conclusion

Uncertainty quantification is a critical missing component in the modern computational imaging toolbox. Accurately quantifying the uncertainty in restored images is a highly challenging task because of the dimensionality involved and because uncertainty estimates depend

strongly on tail probabilities and higher-order moments that are difficult to estimate precisely. The proposed equivariant bootstrap method offers a conceptually simple, yet highly effective strategy for obtaining accurate uncertainty quantification results in imaging inverse problems. The method is also very computationally efficient and scales seamlessly to large settings.

In addition, the method does not require ground-truth data, or a very small amount for calibration tests. When combined with modern unsupervised learning techniques that require measurement data only for training, this enables, for the first time to the best of our knowledge, to perform image point estimation and uncertainty quantification reliably in a fully unsupervised manner. This potent combination is of paramount importance to a wide range of impactful application domains. In particular, applications in science and medicine, where uncertainty quantification is critically important and where obtaining reliable ground-truth data can be extremely expensive or even impossible.

Acknowledgements

This work was supported by the French National Research Agency (ANR) through grant JCJC UNLIP and by the UK Research and Innovation (UKRI) Engineering and Physical Sciences Research Council (EPSRC) through grants BLOOM (EP/V006134/1) and LEXCI (EP/W007681/1).

References

- Agustsson, E. and Timofte, R. (2017). Ntire 2017 challenge on single image super-resolution: Dataset and study. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- Altekrüger, F. and Hertrich, J. (2023). Wppnets and wppflows: The power of wasserstein patch priors for superresolution. *SIAM Journal on Imaging Sciences*, 16(3):1033–1067.
- Angelopoulos, A. N., Kohli, A. P., Bates, S., Jordan, M., Malik, J., Alshaabi, T., Upadhyayula, S., and Romano, Y. (2022). Image-to-image regression with distribution-free uncertainty quantification and applications in imaging. In *Proc. of the 39th International Conference on Machine Learning*, volume 162, pages 717–730. PMLR.
- Antorán, J., Barbano, R., Leuschner, J., Hernández-Lobato, J. M., and Jin, B. (2022). Uncertainty estimation for computed tomography with a linearised deep image prior.
- Armato III, S. G., McLennan, G., Bidaut, L., McNitt-Gray, M. F., Meyer, C. R., Reeves, A. P., Zhao, B., Aberle, D. R., Henschke, C. I., Hoffman, E. A., et al. (2011). The lung image database consortium (lidc) and image database resource initiative (idri): a completed reference database of lung nodules on ct scans. *Medical physics*, 38(2):915–931.
- Arridge, S., Maass, P., Öktem, O., and Schönlieb, C.-B. (2019). Solving inverse problems using data-driven models. *Acta Numerica*, 28:1–174.
- Celledoni, E., Ehrhardt, M. J., Etmann, C., Owren, B., Schönlieb, C.-B., and Sherry, F. (2021). Equivariant neural networks for inverse problems. *Inverse Problems*, 37(8):085006.
- Chambolle, A. and Pock, T. (2016). An introduction to continuous optimization for imaging. *Acta Numerica*, 25:161–319.
- Chen, D., Davies, M., Ehrhardt, M. J., Schönlieb, C.-B., Sherry, F., and Tachella, J. (2023). Imaging with equivariant deep learning: From unrolled network design to fully unsupervised learning. *IEEE Signal Processing Magazine*, 40(1):134–147.
- Chen, D., Tachella, J., and Davies, M. (2021). Equivariant imaging: Learning beyond the range space. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4379–4388.
- Chen, D., Tachella, J., and Davies, M. (2022). Robust equivariant imaging: a fully unsupervised framework for learning to image from noisy and partial measurements. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Choi, J., Kim, S., Jeong, Y., Gwon, Y., and Yoon, S. (2021). ILVR: Conditioning method for denoising diffusion probabilistic models. *arXiv preprint arXiv:2108.02938*.
- Chung, H., Kim, J., McCann, M. T., Klasky, M. L., and Ye, J. C. (2022). Diffusion posterior sampling for general noisy inverse problems. *arXiv preprint arXiv:2209.14687*.
- Efron, B. and Tibshirani, R. J. (1994). *An introduction to the bootstrap*. CRC press.
- Eldar, Y. C. (2008). Generalized sure for exponential families: Applications to regularization. *IEEE Transactions on Signal Processing*, 57(2):471–481.
- Holden, M., Pereyra, M., and Zygalakis, K. C. (2022). Bayesian imaging with data-driven priors encoded by neural networks. *SIAM Journal on Imaging Sciences*, 15(2):892–924.
- Jin, K. H., McCann, M. T., Froustey, E., and Unser, M. (2017). Deep convolutional neural network for inverse problems in imaging. *IEEE Transactions on Image Processing*, 26(9):4509–4522.
- Kaipio, J. P. and Somersalo, E. (2005). *Statistical and Computational Inverse Problems*. Springer New York.

- Kawar, B., Elad, M., Ermon, S., and Song, J. (2022). Denoising diffusion restoration models. *Advances in Neural Information Processing Systems*, 35:23593–23606.
- Laumont, R., Bortoli, V. D., Almansa, A., Delon, J., Durmus, A., and Pereyra, M. (2022). Bayesian imaging using plug & play priors: when langevin meets tweedie. *SIAM Journal on Imaging Sciences*, 15(2):701–737.
- Mohan, S., Kadkhodaie, Z., Simoncelli, E. P., and Fernandez-Granda, C. (2019). Robust and interpretable blind image denoising via bias-free convolutional neural networks. *arXiv preprint arXiv:1906.05478*.
- Mukherjee, S., Hauptmann, A., Öktem, O., Pereyra, M., and Schönlieb, C.-B. (2023). Learned reconstruction methods with convergence guarantees: A survey of concepts and applications. *IEEE Signal Processing Magazine*, 40(1):164–182.
- Narnhofer, D., Habring, A., Holler, M., and Pock, T. (2022). Posterior-variance-based error quantification for inverse problems in imaging.
- Ongie, G., Jalal, A., Metzler, C. A., Baraniuk, R. G., Dimakis, A. G., and Willett, R. (2020). Deep learning techniques for inverse problems in imaging. *IEEE Journal on Selected Areas in Information Theory*, 1(1):39–56.
- Serre, J.-P. (1977). *Linear representations of finite groups*, volume 42. Springer.
- Stein, C. M. (1981). Estimation of the mean of a multivariate normal distribution. *The annals of Statistics*, pages 1135–1151.
- Tachella, J., Chen, D., and Davies, M. (2023a). Sensing theorems for learning from incomplete measurements. *Journal of Machine Learning Research*, 24(39):1–45.
- Tachella, J., Chen, D., Hurault, S., and Terris, M. (2023b). DeepInverse: A deep learning framework for inverse problems in imaging.
- Terris, M., Repetti, A., Pesquet, J.-C., and Wiaux, Y. (2020). Building firmly nonexpansive convolutional neural networks. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8658–8662. IEEE.
- Zhang, C., Barbano, R., and Jin, B. (2021). Conditional variational autoencoder for learned image reconstruction. *Computation*, 9(11).
- Zhang, K., Zuo, W., Chen, Y., Meng, D., and Zhang, L. (2017a). Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE transactions on image processing*, 26(7):3142–3155.
- Zhang, K., Zuo, W., Gu, S., and Zhang, L. (2017b). Learning deep cnn denoiser prior for image restoration. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3929–3938.
- Zhu, Y., Zhang, K., Liang, J., Cao, J., Wen, B., Timofte, R., and Van Gool, L. (2023). Denoising diffusion models for plug-and-play image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1219–1229.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes] We detail the inverse problem setting in Equation (1) and provide a detailed description of the standard bootstrap in Section 3 and the proposed equivariant bootstrap in Section 4.1.
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes] The mathematical properties of the proposed method (in the setting of a linear estimator) are presented in Section 4.2. The complexity of all evaluated methods is summarized in Table 2.
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes] We have included the Pytorch source code in our submission.
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. [Not Applicable] We do not have proofs, however we provide all the assumptions in the analysis of the linear predictor in Section 4.2.
 - (b) Complete proofs of all theoretical results. [Not Applicable] We do not have formal proofs in this work.
 - (c) Clear explanations of any assumptions. [Yes] The assumptions used in the analysis of the linear estimator are clearly stated in Section 4.2.
3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL).

- [Yes] The source code included in the submission can reproduce all the results stated in Section 5.
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes] The training details are included in the supplementary materials.
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes] We provide the standard deviation of the test PSNR of all evaluated algorithms in Table 1.
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes] We detail the local cluster of 4 GPUs used for the experiments in Section 5.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
- (a) Citations of the creator If your work uses existing assets. [Yes] We have cited the public datasets (Armato III et al., 2011; Agustsson and Timofte, 2017) and open-source libraries Tachella et al. (2023b) used in the paper.
 - (b) The license information of the assets, if applicable. [Not Applicable] The assets do not require providing license information.
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Yes] We have provided all the codes used in this paper in the official submission.
 - (d) Information about consent from data providers/curators. [Not Applicable] We only use public datasets in this work, which do not require providing such information.
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable] We have not used new datasets that could involve sensible content.
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
- (a) The full text of instructions given to participants and screenshots. [Not Applicable] We have not conducted research with human subjects or used crowdsourcing. All our experiments were performed on public datasets.
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

A Examples of Inverse Problems in Imaging Sciences

Inverse problems commonly encountered in image restoration are often addressed by assuming an additive white Gaussian noise degradation model. In that case, (1) takes the form $Y \sim \mathcal{N}(Ax_*, \sigma^2 \mathbb{I})$ where σ is related to the intensity of the observation noise. Four canonical examples include image denoising (A is the identity matrix \mathbb{I}); image inpainting (i.e., the problem of recovering a full image from a subset of observed pixels - A is a diagonal binary matrix encoding the pattern of observed and unobserved pixels); image deblurring (i.e., the problem of restoring fine detail in an image that has been degraded by blur stemming from limited resolution or motion - A is often a Toeplitz matrix chosen to model the blurring operation); and tomographic image (i.e., the problem of recovering an image from a set of finite projections - A is often the discrete Radon transform). We refer the reader to (Kaipio and Somersalo, 2005; Chambolle and Pock, 2016) for an excellent introduction to the topic.

B Detailed Proofs (Section 4.2)

We begin by providing a detailed derivation of the following equality shown in Section 4.2 of the main paper:

$$\mathbb{E}_{\tilde{Y}} \|\tilde{x}(\tilde{Y}) - Bx_*\|_2^2 = \|Rx_*\|^2 - \|Rx_*\|_{B_*}^2 + x_*^\top B^\top \Pi_{\mathcal{G}} (R^\top R + R^\top B_* + B_* R - 2R) Bx_*$$

Starting with the definition of the expectation of the equivariant bootstrap error, we have

$$\mathbb{E}_{\tilde{Y}} \|\tilde{x}(\tilde{Y}) - Bx_*\|_2^2 = \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} \|T_g^{-1} B T_g Bx_* - Bx_*\|_2^2 \quad (7)$$

$$= \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} [x_*^\top B^\top (I_n + T_g^{-1} B^\top B T_g - 2T_g^{-1} B T_g) Bx_*] \quad (8)$$

$$= x_*^\top B^\top (I_n + \Pi_{\mathcal{G}}(B^\top B - 2B)) Bx_* \quad (9)$$

Using the decomposition $B = B_* + R$, we have that

$$\begin{aligned} \Pi_{\mathcal{G}}(B^\top B - 2B) &= \Pi_{\mathcal{G}}(B_*^\top B_* + B_*^\top R + R^\top B_* + R^\top R - 2B_* - 2R) \\ &= \Pi_{\mathcal{G}}(B_*^\top R + R^\top B_* + R^\top R - B_* - 2R) \\ &= -\Pi_{\mathcal{G}}(B_*) + \Pi_{\mathcal{G}}(B_*^\top R + R^\top B_* + R^\top R - 2R) \\ &= -B_* + \Pi_{\mathcal{G}}(B_*^\top R + R^\top B_* + R^\top R - 2R) \end{aligned}$$

where we used that $B_*^\top B_* = B_*$ and $\Pi_{\mathcal{G}}(B_*) = B_*$. Plugging-in this result into Equation (9), we obtain

$$\mathbb{E}_{\tilde{Y}} \|\tilde{x}(\tilde{Y}) - Bx_*\|_2^2 = x_*^\top B(I - B_*)Bx_* + x_*^\top B^\top \Pi_{\mathcal{G}} (B_*^\top R + R^\top B_* + R^\top R - 2R) Bx_*,$$

We can conclude our derivation by showing that the first term can be written as

$$x_*^\top B(I - B_*)Bx_* = (Rx_* + x_*)^\top (I - B_*)(Rx_* + x_*) \quad (10)$$

$$= (Rx_*)^\top Rx_* - (Rx_*)^\top B_* Rx_* + \underbrace{\|x_*\|^2 - \|B_* x_*\|^2}_{=0} + \underbrace{2x_*^\top R^\top x_* - 2x_*^\top R^\top B_* x_*}_{=0} \quad (11)$$

$$= \|Rx_*\|^2 - \|Rx_*\|_{B_*}^2 \quad (12)$$

We now show that an estimator $\hat{x}(\cdot)$ verifying $A\hat{x}(y) = y$ implies that $\|Rx_*\|_{B_*}^2 = 0$, which is used in the analysis of the equivariant bootstrap with a linear estimator in Section 4.2 of the main paper.

Using that $\hat{x}(y) = Bx_*$ and $y = Ax_*$, we have that

$$A\hat{x}(y) = y \quad (13)$$

$$ABx_* = Ax_* \quad (14)$$

$$A(Bx_* - x_*) = 0 \quad (15)$$

Applying the decomposition $B = B_\star + R$ where $B_\star x_\star = x_\star$, we obtain

$$A((B_\star + R)x_\star - x_\star) = 0 \tag{16}$$

$$ARx_\star = 0 \tag{17}$$

Thus, writing $B_\star = M_\star A$ where $M_\star \in \mathbb{R}^{n \times n}$ is the oracle linear estimator, we have

$$M_\star ARx_\star = 0 \tag{18}$$

$$B_\star Rx_\star = 0 \tag{19}$$

and consequently $\|Rx_\star\|_{B_\star}^2 = 0$ since B_\star is positive semi-definite.

C Architecture and Training Details

1. ULA: We use the DnCNN denoiser introduced by Terris et al. (2020) which is trained on the dataset in (Zhang et al., 2017b), with a dedicated loss that forces the network to be firmly non-expansive.
2. DiffPIR: For the tomography problem (single channel images), we finetune the (pretrained) DRUNet denoiser (Zhang et al., 2017b) on the LIDC train set. For the inpainting task (RGB images), we use the attention-based denoiser introduced by Choi et al. (2021) trained on the FFHQ dataset.
3. DDRM: We use the same denoiser as diffPIR (Choi et al., 2021) for the image inpainting task. We didn't evaluate the DDRM method on the compressed sensing and tomography problems since the algorithm can only be applied to problems with a simple singular value decomposition of the forward operator.
4. Reconstruction networks:
 - (a) Compressed Sensing (MNIST): For this task, we use an architecture based on an unrolled proximal gradient algorithm with 4 iterations, where the proximal operator was replaced by a trainable U-Net architecture with 4 scales. The same proximal operator network is applied at each iteration (i.e., weight-tied). The resulting unrolled network has 2065033 trainable parameters. The unsupervised network is trained using the robust EI loss (Chen et al., 2022) with random shifts as transformations.
 - (b) Inpainting (DIV2K): We use an architecture based on an unrolled proximal half-quadratic splitting algorithm with 3 iterations, where the proximal operator was replaced by a trainable U-Net architecture with 4 scales. The same proximal operator network is applied at each iteration (i.e., weight-tied). The resulting unrolled network has 8554375 trainable parameters. The unsupervised network is trained using the robust EI loss (Chen et al., 2022) with random shifts as transformations.
 - (c) Tomography (LIDC-IDRI): We use a U-Net architecture applied to the filtered-back projected images, similarly to Jin et al. (2017). The unsupervised network is trained using the robust EI loss (Chen et al., 2022) with random rotations as transformations. The network has 8553088 trainable parameters.
 - (d) Anisotropic deblurring (MNIST): We choose a blur forward operator with a vertical 7×1 kernel with entries $\frac{1}{7}$ (i.e., summing to one). We use an architecture based on an unrolled proximal gradient algorithm with 3 iterations, where the proximal operator was replaced by a DnCNN architecture (Zhang et al., 2017a) with circular padding. The same proximal operator network is applied at each iteration (i.e., weight-tied). The resulting unrolled network has 664711 trainable parameters and is fully shift-equivariant (see (Celledoni et al., 2021) for more details on unrolled equivariant networks).

The unrolled architectures are generated using the DeepInverse library (Tachella et al., 2023b). We use the same architecture for supervised and unsupervised learning. All the U-Net architectures used for the supervised and unsupervised reconstruction networks are based on the backbone in (Chen et al., 2021). All networks are trained using the Adam optimizer with standard hyper-parameters ($\beta_1 = 0.9$ and $\beta_2 = 0.999$) and a step size of 10^{-4} .

D Additional Results

Figures 5 to 7 show test reconstructions obtained by the different evaluated methods for the problems of compressed sensing, image inpainting, and sparse-angle tomography.

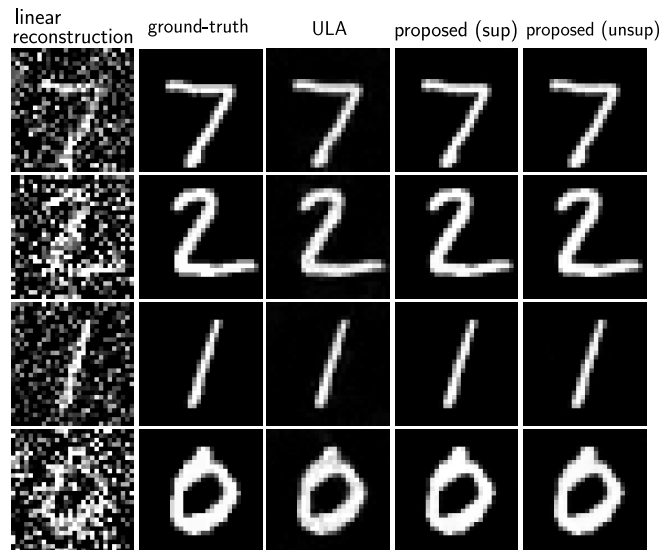


Figure 5: Reconstructed test images by the different evaluated methods for the compressed sensing task using MNIST. The first column shows a simple linear reconstruction of the measurements, i.e., $\hat{x}(y) = A^T y$.

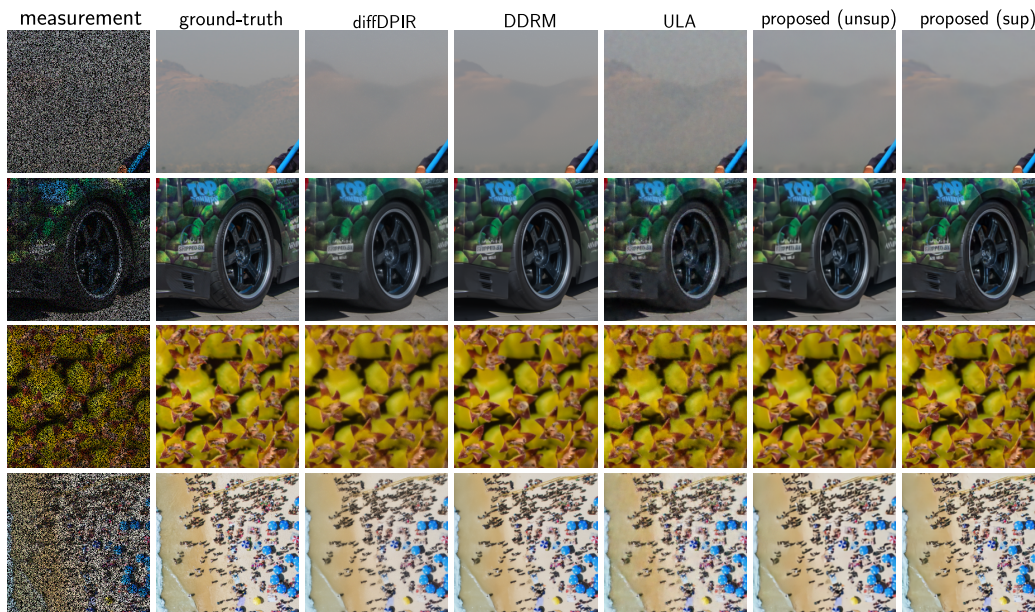


Figure 6: Reconstructed test images by the different evaluated methods for the image inpainting task using the DIV2K dataset.

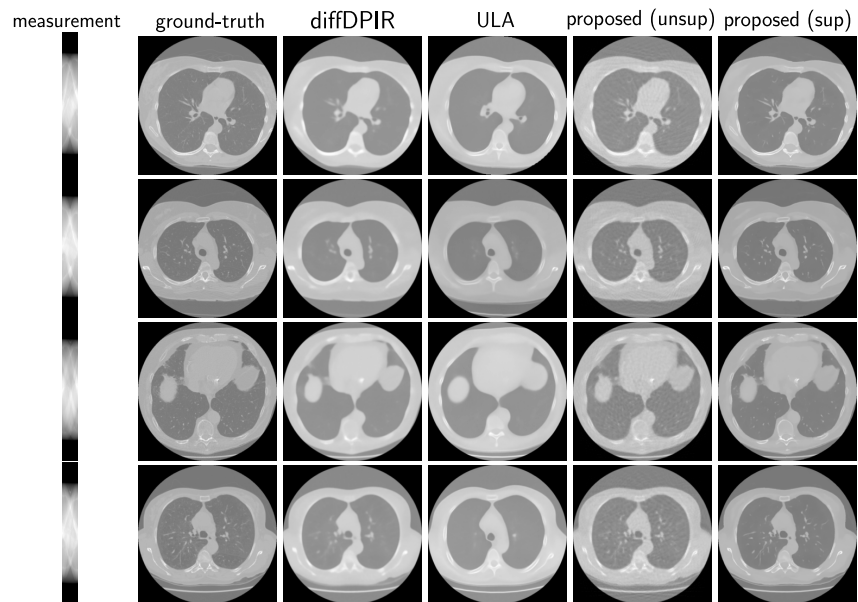


Figure 7: Reconstructed test images by the different evaluated methods for the sparse angle tomography task using the LIDC-IDRI dataset. The first column shows the sinograms y .