

Imaging with Confidence: Uncertainty Quantification for High-dimensional Undersampled MR Images

Frederik Hoppe^{1*}, Claudio Mayrink Verdun^{2*}, Hannah Laus^{3*}, Sebastian Endt⁴, Marion I. Menzel⁴, Felix Kraemer³, and Holger Rauhut⁵

¹ RWTH Aachen University, Germany hoppe@mathc.rwth-aachen.de

² Harvard University, USA claudioverdun@seas.harvard.edu

³ TU Munich & MCML, Germany {hannah.laus,felix.kraemer}@tum.de

⁴ TH Ingolstadt, TU Munich & GE HealthCare, Germany
{sebastian.endt,Marion.Menzel}@thi.de

⁵ LMU Munich & MCML, Germany rauhut@math.lmu.de

Abstract. Establishing certified uncertainty quantification (UQ) in imaging processing applications continues to pose a significant challenge. In particular, such a goal is crucial for accurate and reliable medical imaging if one aims for precise diagnostics and appropriate intervention. In the case of magnetic resonance imaging, one of the essential tools of modern medicine, enormous advancements in fast image acquisition were possible after the introduction of compressive sensing and, more recently, deep learning methods. Still, as of now, there is no UQ method that is both fully rigorous and scalable. This work takes a step towards closing this gap by proposing a total variation minimization-based method for pixel-wise sharp confidence intervals for undersampled MRI. We demonstrate that our method empirically achieves the predicted confidence levels. We expect that our approach will also have implications for other imaging modalities as well as deep learning applications in computer vision. Our code is available on GitHub https://github.com/HannahLaus/Project_UQ_TV.git.

Keywords: Uncertainty Quantification · Imaging Processing · Magnetic Resonance Imaging

1 Introduction

In recent years, the field of magnetic resonance imaging, like many other imaging modalities, has undergone a significant paradigm shift towards a unified design of the acquisition-reconstruction pipeline. That is, the data to be acquired is selected based on the computational methods used for reconstruction. A prominent example of this interplay is compressive imaging, where sampling trajectories chosen to be suitable for imaging reconstruction via sparse recovery schemes

* These authors contributed equally to this work.

allowed for significant undersampling [45]. Later, the same type of subsampling techniques have also proven successful in the context of deep learning reconstructions [33]. Thanks to these approaches, the retrieval of fast, undersampled MRI in clinical settings is now a reality. Indeed, in 2017, the FDA approved the first scanners employing acceleration techniques with compressive sensing [2, 3], and since then, several products and algorithms for retrieving highly undersampled images have emerged. In these advancements, mathematical theory has played a driving role, so this development has sometimes been described as *from blackboard to bedside* [18]. The main objective driving the design of the reconstruction algorithms is to use nonlinear optimization to find solutions adhering to structural constraints describing real images. These constraints could explicitly arise from a unified model, such as wavelet or gradient sparsity, or be implicitly learned from training data.

When the measurements are affected by noise, these methods search for solutions adhering to these constraints within the feasibility regions. That is, the imaging recovery method will search in the space of possible solutions and look for reconstructions that resemble natural images, which also makes potential reconstruction artifacts harder to identify. Especially unexpected localized features could correspond to real pathological abnormalities or arise in the interplay of noise and recovery methods. For this reason, quantifying the uncertainty via certified confidence intervals is of increasing importance not only for MRI but for a variety of problems in imaging and computer vision [59].

At the same time, classical approaches for assessing the confidence of the solution that could allow to distinguish these two causes, abnormalities and artifacts, are not applicable. The reason is that they typically require some knowledge of the distribution of the random distortion and, consequently, of the generated image. While the measurement error is well modeled by Gaussian noise, the distribution of the resulting distortion in the reconstruction cannot be modeled as easily due to the nonlinearity of the procedure, see, *e.g.*, [22, Theorem 2] and the discussion therein. This poses a significant issue when the reconstructed image is directly employed for downstream tasks, such as automating the extraction of quantitative parameters for clinical analysis.

Imaging with confidence, a novel rigorous method tailored to UQ:

It is only recently that the first attempts to derive non-asymptotic confidence intervals for the pixel values of the reconstruction, without imposing prior knowledge have been successful (see Section 2 for a discussion). So far, most works have focused on sparsity models with respect to some known basis, as they are natural, for example, in angiography and reconstruction via variants of the LASSO [7]. In this work, we leverage these approaches to address the uncertainty quantification problem for recovery via total variation minimization, a standard recovery scheme for MR images in current industrial implementations [4] that has also conceptually inspired various machine learning-based approaches reaching state-of-the-art results such as variational networks [25]. To the best of our knowledge, for the first time, we provide rigorous confidence intervals for un-

undersampled acquired MR images and test the results with synthetic and in vivo data.

2 Related Work

There were other attempts available in the literature to tackle the aleatoric uncertainty quantification of high-dimensional parsimonious problems, such as in undersampled imaging generation. For example, a series of works developed the so-called generalized fiducial inference [26, 27, 41]. Such a method creates a relationship between the data and the parameters for which the UQ can be rigorously performed and establishes confidence intervals from such a relationship. However, high-dimensional Monte Carlo makes this method intractable in the high-dimensional imaging regime. Bayesian sparse learning techniques were also developed for undersampled imaging problems [43]. An inherent challenge with the Bayesian approach arises when inappropriate prior distributions, which fail to accurately reflect the prior physical knowledge of the problem, are employed. In such cases, the performance and reliability of the approach may deteriorate. Another line of research on UQ is conformal prediction, which generates distribution-free prediction sets for regression and classification problems [9]. However, choosing the score function in imaging problems can be challenging. Such a technique is more directly applied to epistemic uncertainty rather than to the aleatoric one [31]. Our work builds on the urgency of providing rigorous UQ to imaging problems in computer vision and image processing. Such necessity can be seen in the several workshops for UQ in computer vision and imaging that took place over the last few years dedicated to this theme [1, 5, 6].

Our uncertainty quantification extends a line of research about non-asymptotic confidence intervals for high-dimensional problems known as *debiased estimators*, which was initiated in a series of papers [11, 35, 60, 65]. The key idea is that given an estimator to retrieve noisy image retrieval from a few measurements; one can characterize the distribution of a modified estimator based on the KKT conditions of the underlying inverse problem; see Section 3.2. Such works established sharp confidence intervals for i.i.d. (sub-)Gaussian variables with Gaussian noise by assuming that the object to be retrieved is s -sparse in the canonical basis. However, such variables do not represent typical constraints imposed by the physics behind the acquisition procedure. Later, this was extended to imaging applications, where UQ was rigorously obtained for ℓ_1 model-based MRI [28, 29]. In this paper, we harness results for UQ obtained in previous works to more realistic scenarios where the images are reconstructed via the minimization of the total variation semi-norm, an idea introduced in the seminal paper [53]. This imaging retrieval technique was extensively explored in tasks such as inpainting, denoising, and deblurring; see [14, 32] and references therein. This method has been rigorously proven to yield state-of-the-art results in MRI reconstruction – see [40] and [7, Chapter 17] – and was consistently tested in several applications such as parallel imaging [12], free-breathing three-dimensional cardiac MRI [16], non-contrast-enhanced angiography [17], spectroscopy [54] and metabolite maps

of the human brain [47], to name a few. Recently, in the context of learning algorithms, the work [57] proposed a latent Bayesian model with a variational autoencoder to obtain posterior distributions for the generated images and the work [21] proposed hierarchical conditional variational autoencoders for UQ, but no theory is available for such methods. Finally, [30] developed a non-asymptotic theory for constructing confidence intervals for high-dimensional learning problems.

3 Background: Sparse Imaging Retrieval

In this section, we summarize the tools we need to construct confidence intervals for images. Throughout the paper, we consider the following high-dimensional linear model that is a discretized version of the Bloch equations, which models the MR acquisition process [42]

$$b = \mathcal{P}\mathcal{F}x^0 + \varepsilon \quad \text{with } N \gg m. \quad (1)$$

Here, $b \in \mathbb{C}^m$ is the k-space data, $\mathcal{P}\mathcal{F} \in \mathbb{C}^{m \times N}$ is a subsampled Fourier matrix, $x^0 \in \mathbb{C}^N$ is the underlying ground truth image, and $\varepsilon \sim \mathcal{CN}(0, \sigma^2 I_{N \times N})$ is complex Gaussian noise. The goal in fast undersampled MRI is to retrieve x^0 having $m < N$ k-space measurements b when $\mathcal{P}\mathcal{F}$ and σ are known.

3.1 Total Variation

As described in Sec. 2, several image modalities can be modeled by gradient-sparse images. Also, such guiding principle is behind modern deep learning architectures that learn higher-order statistics of images, *e.g.*, those based on the fields of experts framework [52]. Therefore, we extend the results to the total-variation setting to create rigorous UQ schemes for MRI. More precisely, for a given image $x^0 \in \mathbb{C}^N$, we consider the finite gradient operator D that gives $Dx^0 = (x_2^* - x_1^*, \dots, x_N^* - x_{N-1}^*)^T \in \mathbb{C}^{N-1}$ and we solve the *total variation* problem, *i.e.*,

$$\operatorname{argmin}_{x \in \mathbb{C}^N} \frac{\mu}{2} \|b - \mathcal{P}\mathcal{F}x\|_2^2 + \|x\|_{TV} \quad (2)$$

with a regularization parameter $\mu > 0$ for balancing the data fidelity term $\|b - \mathcal{P}\mathcal{F}x\|_2^2$ and the regularization term $\|x\|_{TV} := \|Dx\|_1$, which is referred to as the anisotropic formulation [7]. We use such a formulation due to the resemblance with the LASSO estimator. Still, the same techniques described here hold for other discretizations of the total variation, such as the isotropic one, *e.g.*, see [15].

3.2 Debiased Estimators for Image Retrieval

If the underlying image to be retrieved from a few samples is sparse in the canonical basis, several methods in the sparse learning literature can be used, with the LASSO being arguably the most famous one [58]. However, given a certain

noise level σ , it is provably impossible to have a practical characterization of the LASSO distribution since it depends on the support of the ground truth, which is not known in advance [22]. Moreover, a bias is introduced by the shrinkage procedure of the ℓ_1 -norm [23, Chapter 4]. Several works in the statistical literature developed strategies to debias such estimators, *e.g.*, [35, 36, 60, 65]. Roughly speaking, such works introduce a correction term to the LASSO estimator \hat{x} that aims to retrieve x^0 from $b = Ax^0 + \varepsilon$. It is given by $\hat{x}^u = \hat{x} + \frac{1}{m}MA^*(b - A\hat{x})$, where, for a given sample covariance matrix $\hat{\Sigma} = \frac{A^*A}{m}$, M should be chosen (as discussed in Sec. 4), such that $M\hat{\Sigma} \approx I_{N \times N}$. From that, the decomposition

$$\sqrt{m}(\hat{x}^u - x^0) = \frac{1}{\sqrt{m}}MA^*\varepsilon - \sqrt{m}(M\hat{\Sigma} - I_{N \times N})(\hat{x} - x^0), \quad (3)$$

where $\frac{1}{\sqrt{m}}MA^*\varepsilon \sim \mathcal{N}(0, \sigma^2 M\hat{\Sigma}M^*)$, follows. The work [36] showed that in the not realistic case of a measurement process A being described by a real Gaussian matrix with known covariance matrix Σ , the remainder term $\sqrt{m}(M\hat{\Sigma} - I_{N \times N})(\hat{x} - x^0)$ vanishes with probability converging to 1 as $m = m(N) \rightarrow \infty$ provided $m \geq Cs\sqrt{\log N}$ for a constant $C > 0$.⁶ This means that, given a certain noise level, the difference between this new estimator and the underlying image, $\hat{x}^u - x^0$, is approximately a Gaussian variable. Such distance can be quantified in a non-asymptotic way, which means that *non-asymptotic sharp confidence intervals can be derived for such an estimator*. In the case of realistic measurement operators given by a subsampled Fourier matrix, the works [28, 29] proved that the remainder term in ℓ_1 -regularization also vanishes under slightly stronger assumptions in the number of acquired k-space points as compared to the case of Gaussian measurements. From that, confidence intervals for sparse MR images can be constructed but only in the context of ℓ_1 -regularized recovery. In particular, these results do not cover variational approaches such as the total variation minimization problem Eq. (2), which continues to be the method of choice for many applications due to its simplicity and good numerical performance and which also plays a foundational role for many machine learning approaches to MRI recovery in realistic settings, as discussed in Section 2.

In this paper, we step towards closing this gap by leveraging previous debiasing approaches to quantify the uncertainty for MR images recovered from undersampled data via the variational problem Eq. (2). We introduce the estimator in Sec. 4 and establish non-asymptotic confidence intervals for every pixel of the image to be retrieved in Sec. 5. In Sec. 6, we discuss the role of noise estimation in MRI, and, finally, in Sec. 7, we numerically illustrate our findings on in vivo brain data.

4 Debaised Total Variation Estimator

Since most MR images are not sparse in the pixel domain, the standard debaised LASSO is limited in practice. Thus, inspired by [28, 29], we propose to use a

⁶ The notation $m(N)$ considers a sequence of regression problems, where both dimensions m and N are growing with the rate $\frac{s\sqrt{\log N}}{m} \rightarrow 0$.

Algorithm 1 Construction of correction matrix M

-
- 1: **Input:** Sampling mask P , regularization parameters $\lambda_1 \dots \lambda_N$
 - 2: For $i \in [N]$ solve $\hat{x}^i \leftarrow \min_{x \in \mathbb{C}^{N-1}} \frac{1}{2m} \|(\mathcal{PF})_i - (\mathcal{PF})_{-i}x\|_2^2 + \lambda_i \|x\|_1$
 - 3: Set $\hat{T}^2 = \text{diag}(\hat{\tau}_1^2, \dots, \hat{\tau}_N^2)$ with $\hat{\tau}_i^2 = \frac{1}{m} ((\mathcal{PF})_i - (\mathcal{PF})_{-i}\hat{x}^i)^* (\mathcal{PF})_i$
 - 4: For $i \in [N]$ let $\hat{c}^i \in \mathbb{C}^N$, set $\hat{c}_j^i = -\hat{x}_j^i$ for $j < i$, $\hat{c}_j^i = -\hat{x}_{j-1}^i$ for $j > i$ and $\hat{c}_i^i = 1$.
 - 5: Set $\hat{C} = (\hat{c}^1, \dots, \hat{c}^N)^T \in \mathbb{C}^{N \times N}$
 - 6: Compute $M = \hat{T}^{-2} \hat{C}$
 - 7: **Output:** Correction matrix M
-

debiased TV estimator of the form $\hat{x}^u = \hat{x} + \frac{1}{m} M(\mathcal{PF})^*(b - \mathcal{PF}\hat{x})$, where \hat{x} denotes the solution of Eq. (2). The decomposition

$$\sqrt{m}(\hat{x}^u - x^0) = \frac{1}{\sqrt{m}} M(\mathcal{PF})^* \varepsilon - \sqrt{m}(M\hat{\Sigma} - I_{N \times N})(\hat{x} - x^0). \quad (4)$$

consists of a remainder term $R := \sqrt{m}(M\hat{\Sigma} - I_{N \times N})(x^0 - \hat{x})$ and of $W := \frac{1}{\sqrt{m}} M(\mathcal{PF})^* \varepsilon$ which is complex Gaussian distributed with zero mean and covariance matrix $\frac{\sigma^2}{m} M(\mathcal{PF})^*(M(\mathcal{PF})^*)^* = \sigma^2 M\hat{\Sigma}M^*$ due to the Gaussian distribution of the noise. The remainder term vanishes provided that, for a given correction matrix M , the term $\max_{ij \in [N]} |(M\hat{\Sigma} - I_{N \times N})_{ij}|$ is small and that the solution to the variational problem Eq. (2) accurately describes the underlying image, *i.e.*, $x^0 - \hat{x}$ is small.

The so-called *nodewise LASSO*, here adapted to *complex* subsampled Fourier matrices, and described in Algorithm 1, guarantees that $\max_{ij \in [N]} |(M\hat{\Sigma} - I_{N \times N})_{ij}|$ is small [36, 60]. There, $(\mathcal{PF})_i$ denotes the i -th column of the matrix (\mathcal{PF}) and $(\mathcal{PF})_{-i}$ denotes this matrix without the i -th column. For the latter, *i.e.*, to understand when $x^0 - \hat{x}$ is small, a series of works, *e.g.*, [39, 48, 49] and [7, Theorem 17.5] show that TV minimization provably retrieves a gradient-sparse image. The above results combined lead to the construction of pixel-wise sharp confidence intervals as described in Sec. 5. A remarkable advantage of such estimators is that once the image is retrieved as a solution to a convex optimization problem, one only needs to add the term $\frac{1}{m} M(\mathcal{PF})^*(b - \mathcal{PF}\hat{x})$ to the estimator \hat{x} in order to obtain confidence intervals. The inexpensive cost of such a debiased strategy comes from the fact that, for a given subsampling pattern \mathcal{P} , the matrix M needs to be computed only once, and it is independent of the image and the noise level.

5 Confidence Intervals Implementation

The construction of confidence intervals relies on the asymptotic normality of the debiased TV estimator. If the remainder term is small, then

$$\sqrt{m}(\hat{x}^u - x^0) \sim \mathcal{CN}(0, \sigma^2 M\hat{\Sigma}M^*). \quad (5)$$

The goal is to derive a confidence interval for every pixel. We start with a *confidence region* $C_i(\alpha)$ for the complex pixel value x_i^0 that depends on the significance level $\alpha \in (0, 1)$, such that

$$\mathbb{P}(x_i^0 \in C_i(\alpha)) \geq 1 - \alpha \quad \text{for every image } x^0, \quad (6)$$

where \mathbb{P} denotes the probability due to the random noise. According to Eq. (5), the new estimator \hat{x}^u induces a complex Gaussian distribution on the difference $\hat{x}^u - x^0$. This means that it is possible to derive confidence intervals for the magnitude and the phase of the MR images. Since both the real part and the imaginary part are Gaussian distributed, the absolute value $|\hat{x}_i^u - x_i^0| = \sqrt{\Re(\hat{x}_i^u - x_i^0)^2 + \Im(\hat{x}_i^u - x_i^0)^2}$ follows a Rician distribution with zero-mean. This special case simplifies to a Rayleigh distribution with variance $\sigma \cdot \sqrt{(M\hat{\Sigma}M^*)_{ii}}$ [8]. The corresponding radius δ_i of the confidence region $C_i(\alpha)$ is given by

$$\delta_i(\alpha) = \frac{\sigma \sqrt{\log(1/\alpha)}}{\sqrt{m} \sqrt{(M\hat{\Sigma}M^*)_{ii}}}. \quad (7)$$

Therefore, $C_i(\alpha)$ is a circle centered around the debiased TV reconstruction \hat{x}_i^u with radius $\delta_i(\alpha)$, *i.e.* $C_i(\alpha) = \{z \in \mathbb{C} \mid |z - \hat{x}_i^u| \leq \delta\}$. The confidence intervals for the magnitude, which is what is typically visualized by radiologists, can be constructed due to the following inequality:

$$\mathbb{P}(|x_i^0| \in [|\hat{x}_i^u| - \delta_i(\alpha), |\hat{x}_i^u| + \delta_i(\alpha)]) \geq \mathbb{P}(x_i^0 \in C_i(\alpha)) \geq 1 - \alpha. \quad (8)$$

Hence, the $1 - \alpha$ confidence interval for the magnitude is given by $J_i(\alpha) = [|\hat{x}_i^u| - \delta_i(\alpha), |\hat{x}_i^u| + \delta_i(\alpha)]$. Please note, however, that the confidence regions are given in the complex plane. This means that we could also have CIs for the phase φ^0 of the image, which contains information, *e.g.*, on motion or electromagnetic properties of biological tissues [55]. With the same argument as in Eq. (8), and by denoting $\hat{\varphi}_i := \arcsin\left(\frac{\Re(\hat{x}_i^u)}{|\hat{x}_i^u|}\right)$ and $\gamma_i = \arccos\left(1 - \frac{\delta_i^2}{2|\hat{x}_i^u|^2}\right)$, we obtain

$$\mathbb{P}(\varphi_i^0 \in [\hat{\varphi}_i - \gamma_i, \hat{\varphi}_i + \gamma_i]) \geq 1 - \alpha. \quad (9)$$

For a detailed derivation of the confidence intervals, refer to Section 2 of the supplementary material. This paper focuses on visualizing magnitudes, so we adhere to the confidence intervals presented in Equation 8.

Rigorously, Eq. (6) should be interpreted in the following way: The random quantity in Eq. (6) is the confidence region $C_i(\alpha)$ since it is based on the debiased TV estimator, which, in turn, is reconstructed by data that is affected by random noise. In contrast, the underlying image x_i^0 is fixed. Therefore, one should consider an experiment that is conducted l times, and every time the noise is different, it leads to a (slightly) different reconstruction. Hence, also to a (slightly) different confidence interval. In the end, the confidence region will contain the underlying image in at least $1 - \alpha$ of the l cases.

Algorithm 2 Confidence intervals for MR images

- 1: **Input:** Sampling mask \mathcal{P} , k-space data b , regularization parameter μ , (estimated) noise level $\hat{\sigma}$, significance level α , correction matrix M
 - 2: solve $\hat{x} \leftarrow \min_{x \in \mathbb{C}^N} \frac{\mu}{2} \|b - \mathcal{P}\mathcal{F}x\|_2^2 + \|x\|_{TV}$
 - 3: debias reconstruction $\hat{x}^u \leftarrow \hat{x} + \frac{1}{m} M(\mathcal{P}\mathcal{F})^*(b - \mathcal{P}\mathcal{F}\hat{x})$
 - 4: For each pixel i compute radius $\delta_i \leftarrow \frac{\hat{\sigma}}{\sqrt{m} \sqrt{(M\hat{\Sigma}M^*)_{ii}}} \sqrt{\log(1/\alpha)}$ with covariance matrix $\hat{\Sigma} = (\mathcal{P}\mathcal{F})^*(\mathcal{P}\mathcal{F})/m$
 - 5: For each pixel i obtain confidence circle $C_i(\alpha) \leftarrow \{z \in \mathbb{C} \mid |z - \hat{x}_i^u| \leq \delta\}$
 - 6: For each pixel i compute its confidence interval for the magnitude $J_i(\alpha) \leftarrow [|\hat{x}_i^u| - \delta_i(\alpha), |\hat{x}_i^u| + \delta_i(\alpha)]$
 - 7: **Output:** Reconstructions \hat{x} , \hat{x}^u and confidence circles $C_i(\alpha)$ for x_i^0 and confidence intervals $J_i(\alpha)$ for the magnitude $|x_i^0|$ for every pixel i
-

Imaging with confidence, summarized in Algorithm 2, computes confidence intervals for the given pixels of an image retrieved via total variation minimization. More precisely, we propose an estimator that has an asymptotic Gaussian distribution, and from that, we derive pixel-wise uncertainty quantification for the magnitude and the phase of MR images. To visualize such a UQ procedure, we choose lines in the image, draw the confidence intervals, and measure the hit rate of how many pixels lie inside the given interval.

6 Estimating the Noise Level

Despite advances in hardware and image acquisition techniques for MRI, the problem of noise in MR images persists, making noise estimation and analysis crucial tasks in MR image reconstruction [8]. Notably, MR noise estimation increases in complexity if advanced MR reconstruction approaches (e.g. parallel imaging) are employed and, especially, if the underlying individual coil raw data is not available. This is the default case for the clinical use of MRI since usually, only the magnitude data of the image after reconstruction is stored in clinical databases for diagnostic purposes. For our conceptual work on UQ, we ensured the consistent availability of the MR raw data by obtaining it directly from the MR scanner. Therefore, we were able to utilize image regions without imaging objects to retrieve the Gaussian distributed noise characteristics directly, obtaining noise levels of 5% – 15% for the available data. For future implementation on any MR scanners, our proposed method for UQ can be directly amended to the regular image reconstruction when raw data is still available.

7 Numerical Experiments

We numerically illustrate our *imaging with confidence* method by applying it to obtain pixel-wise confidence intervals for clinically feasible MR images reconstructed from noisy measurements. In line with our discussion above, we work

with noise levels between 6% and 12.5%. We retrieve the image by minimizing the TV-regularized objective, Eq. (2), and we consider two different sampling schemes, *radial* and *spiral* sampling, which are both among the most popular for MRI applications. The code for reproducing all the experiments is available on GitHub ⁷. More information about the implementation can be found in the supplementary material in Section 3.

The radial k-space masks were generated according to the golden angle radial scheme proposed by [63]. Each mask is created by adding lines going through the center of the mask by placing a new line in an angle of 111.246° (the golden angle) to the previously added one. As for the latter, k-space masks for spiral sampling were generated by first simulating readout gradients in 2D and computing the resulting k-space trajectories. The trajectories were regridded to the Cartesian grid and transformed to boolean masks. Undersampling factors were adjusted by changing the number of spokes and fine-tuned by adjusting the maximum radius $\sqrt{(k_x^2 + k_y^2)}$. Again, individual spokes were rotated by the golden angle relative to the previous one. Each spiral spoke forms a $3/4$ rotation. The masks were fit to the size of our images, *i.e.*, 156×156 .

As we will observe below, the confidence intervals appropriately capture the difference between ground truth and reconstruction, as predicted. Note that in the plots displaying this finding, the ground truth is only added as a comparison to assess UQ but never used for designing our method.

7.1 UQ for in vivo Complex MRI Data

To assess UQ in realistic scenarios, we simulate the MRI data acquisition and reconstruction process using in vivo data as the underlying ground truth image, see Fig. 1a. We use existing in vivo data acquired with an inversion recovery spin echo (IR-SE) sequence with single-shot spiral readout as described by the works [19, 46]. Brain data of a healthy 32-year-old female were acquired on a Siemens Magnetom Skyra Connectome 3 T system with a 32-channel head coil⁸ with a matrix of 156×156 . Images with complex (magnitude and phase) data were reconstructed using an oversampled ridge-regression estimator that uses a higher order signal model that accounts for spatio-temporal field dynamics [50, 61, 62]. This high-quality complex-valued reconstruction serves as ground truth data for validating our uncertainty quantification method. Moreover, all in vivo scans were performed in accordance with local ethics guidelines and the Institutional Review Board (IRB) after obtaining the volunteer’s written informed consent.

The data acquisition process in the MRI scanner is simulated by taking the 2D Fourier transform of the image, adding Gaussian noise with noise level $\sigma = 0.1$ (see Fig. 1b) and subsampling this k-space data using the spiral and radial sampling scheme, respectively. The sampling masks are visualized in Fig. 1c for 43% spiral subsampling and in Fig. 1d for 50% radial subsampling. These subsampling rates are inspired by studies about the limitations of compressive sensing

⁷ https://github.com/HannahLaus/Project_UQ_TV.git

⁸ Siemens Healthineers, Erlangen, Germany.

approaches for MRI based on expert Radiologists’ scoring [34]. The resulting relative noise levels are $\frac{\|\varepsilon\|_2}{\|\mathcal{P}\mathcal{F}x^0\|_2} \approx 6.5\%$ for spiral and 7% for radial sampling.

Since the sampling patterns are fixed, we calculate, for both masks, the matrix M according to Algorithm 1 with $\lambda_1 = \dots = \lambda_N = 0.0035 \cdot \frac{\sqrt{m}}{\sqrt{12 \log(N)}}$ using FISTA [10] with 1000 iterations. The dimensional scaling $\frac{\sqrt{m}}{\sqrt{12 \log(N)}}$ comes from theoretical works on debiased estimators, *e.g.*, [28]. The constant 0.0035 was found via cross-validation. To reconstruct the image by the subsampled spiral and radial data, respectively, we solve Eq. (2) with the Split-Bregman algorithm [24] since it allows for complex-valued solutions. We use the parameters $\hat{\mu} = \frac{\mu}{20}$, $\lambda_1 = \lambda_2 = 0.05$. For further information on the parameters and the optimization problem, see Section 1 of the supplementary material. The parameter $\mu = \frac{\sqrt{m}}{\sigma \sqrt{12 \log N}}$ was theoretically calculated for a subsampled Fourier matrix in [28]. After this reconstruction step, we compute the debiased TV estimator according to Algorithm 2.

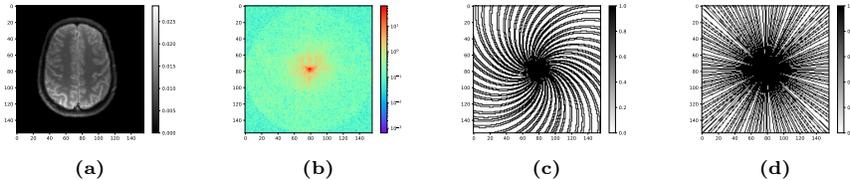


Fig. 1: Model of experimental setting 1a: Magnitude of complex ground truth data 1b: Full k-space affected by Gaussian noise with noise level $\sigma = 0.1$ 1c: Spiral subsampling mask which samples 43% of k-space 1d: Radial subsampling mask which samples 50% of k-space.

To examine the distribution of the debiased TV estimator, a Q-Q plot (see Fig. 2d) compares the quantiles of the input sample, *i.e.* every component $\frac{\sqrt{2m}\Re(\hat{x}^u - x^0)_i}{\sigma \cdot \sqrt{(M\hat{\Sigma}M^*)_{ii}}}$ with the quantiles of a standard Gaussian distribution. It indeed confirms that the input is Gaussian distributed. Since the same is true for the imaginary part, it holds that $\sqrt{m}(\hat{x}^u - x^0) \sim \mathcal{CN}(0, \sigma^2 M\hat{\Sigma}M^*)$, which means that the error of the debiased TV follows a Gaussian distribution with mean zero and covariance structure $\sigma^2 M\hat{\Sigma}M^*$. The same behavior can be observed for radial sampling. The Q-Q plot is shown in Fig. 4d.

Algorithm 2, then, constructs confidence regions for the complex pixel values from which confidence intervals for the magnitudes and phases can be derived by using the fact that the new debiased estimator follows a Gaussian distribution given by $\sqrt{m}\hat{x}^u \sim \mathcal{CN}(\sqrt{m}x^0, \sigma^2 M\hat{\Sigma}M^*)$. Here, we choose the confidence level $\alpha = 0.05$. For clarity, we show confidence intervals for all the pixels, corresponding to magnitudes, lying in a small rectangular region in the image. See Fig. 2a together with the corresponding confidence intervals in Figs. 2b and 2c in the

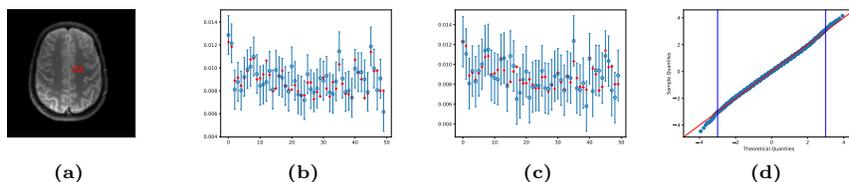


Fig. 2: Confidence intervals for 43% spiral sampling 2a: Ground truth image with rectangle for showing confidence intervals for pixels lying in this region 2b & 2c: 95% confidence intervals for the magnitude of pixels in the rectangle. The red crosses are the ground truth values, the blue circles are the corresponding estimates of the debiased TV estimator. 2b is for a reconstruction with 6.5% noise and 2c is for a reconstruction with 10% noise. 2d: Q-Q plot for input samples $\frac{\sqrt{2m}(\hat{x}^u - x^0)_i}{\sigma \cdot \sqrt{(M \hat{\Sigma} M^*)_{ii}}}$.

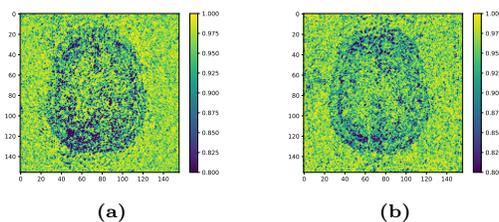


Fig. 3: Empirical probability $\hat{\mathbb{P}}_i(\alpha) = \frac{1}{100} \sum_{j=1}^{100} \mathbb{1}_{\{x_i^0 \in C_i(\alpha, \varepsilon^{(j)})\}}$ for every image pixel each time with a different realization 3a for 43% spiral undersampling and for 3b for 50% radial undersampling.

spiral undersampling setting as well as Figs. 4a, 4b and 4c for the radial case, respectively.

Interpretation of the UQ Procedure. The quality of the confidence regions is quantified by counting how many confidence regions contain the true ground truth pixel value x_i^0 on and off the support. These *hit rates* are calculated for one realization of the noise by $h = \frac{1}{N} \sum_{i=1}^N \mathbb{1}_{\{x_i^0 \in C_i\}}$ and $h_S = \frac{1}{s} \sum_{i \in \text{supp}(x^0)} \mathbb{1}_{\{x_i^0 \in C_i\}}$, where s denotes the amount of pixels that are different from zero. But in view of the interpretation of a confidence interval as described in Sec. 5, we also estimate the probability $\mathbb{P}(x_i^0 \in C_i(\alpha))$ by the empirical probability $\hat{\mathbb{P}}_i(\alpha) = \frac{1}{100} \sum_{j=1}^{100} \mathbb{1}_{\{x_i^0 \in C_i(\alpha, \varepsilon^{(j)})\}}$ for every pixel $i \in [0, 156]^2$ with $\varepsilon^{(j)}$ denoting the dependency on the j -th realization of the noise. In contrast to the hit rates, which are computed for one realization of the noise over all pixels, here we fix a pixel and observe how often the *random* confidence regions contain this pixel averaged over 100 different realizations of the noise. The image containing the empirical probability values for all pixels is shown in Fig. 3a for the spiral and in Fig. 3b for the radial case, respectively. In both cases, the maps seem to contain residual image information. In the spiral setting, the empirical probabil-

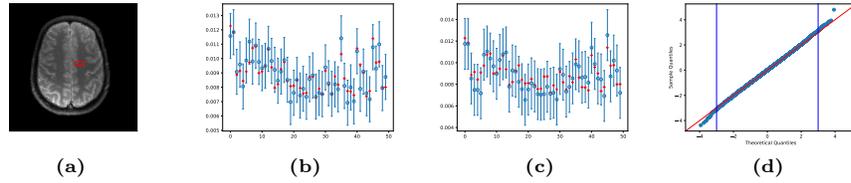


Fig. 4: Confidence intervals for 50% radial sampling 4a: Ground truth image with rectangle for showing confidence intervals for pixels lying in this region 4b & 4c: 95% confidence intervals for the magnitude of pixels in the rectangle. The red crosses are the ground truth values and the blue circles are the corresponding estimates of the debiased TV estimator. 4b is for a reconstruction with 7% noise and 4c is for a reconstruction with 10.5% noise. 4d: Q-Q plot for input samples $\frac{\sqrt{2m}\Re(\hat{x}^u - x^0)_i}{\sigma \cdot \sqrt{(M\hat{\Sigma}M^*)_{ii}}}$.

ity averaged over all support pixels is 90.64%, the one averaged over all pixels is 93.82%. For the radial case, the values are 91.95% and 94.20%, respectively.

7.2 Qualitative Evaluation

The general idea of debiasing, is that after the transformation, the reconstruction error can be decomposed as $\hat{x}^u - x^0 = R + W$, where R is a small residual term and the dominating term W is Gaussian distributed. In fact, by inspecting Fig. 5, it is possible to visualize this debiasing step. On the one hand, the difference between the biased solution of Eq. (2) (see its magnitude in Fig. 5a for spiral and in Fig. 5e for radial sampling) and the ground truth has a very structured error, *i.e.* $|\hat{x} - x^0|$ still displays much of the brain structure, which can be seen in Figs. 5b and 5f. On the other hand, the difference between the unbiased estimator (see its magnitude in Fig. 5c for spiral and Fig. 5g for radial sampling) and the ground truth in Figs. 5d and 5h is much more unstructured, and one can see that it is dominated by random noise. Moreover, as we increase the number of k-space points or we increase the noise, this difference becomes more unstructured. The distribution can be measured in the Q-Q plots in Figs. 2d and 4d, which, indeed, shows the Gaussianity of $\hat{x}^u - x^0$.

To test our findings, the same experiment as in Sec. 7.1 is conducted 100 times, each time with a different realization of the noise vector ε . In the spiral case, we do so for the subsampling rates 43% and 62%, and in the radial setting, we choose the masks according to 50%, 70%, and 90% k-space undersampling. For every mask, we select $\sigma \in \{0.1, 0.15\}$ resulting in relative noise levels $\frac{\|\varepsilon\|_2}{\|P\mathcal{F}x^0\|_2}$ between 0.06 and 0.125. In these settings the ℓ_2 and ℓ_∞ -norm of the errors $\hat{x} - x^0$ and $\hat{x}^u - x^0$, of the bias term $R := (M\hat{\Sigma} - I_{N \times N})(\hat{x} - x^0)$ and of the Gaussian term $W := \frac{1}{m}M(P\mathcal{F})^*\varepsilon$ are computed and averaged. The resulting quantities, including a more detailed analysis, can be found in Table 1 (spiral) and in Table 2 (radial) in Section 3 in the supplementary material.

With these results we numerically verify that $\|R\|_\infty$ is small as compared to $\|W\|_\infty$. Since the goal is to have *pixel-wise* confidence intervals, obtaining

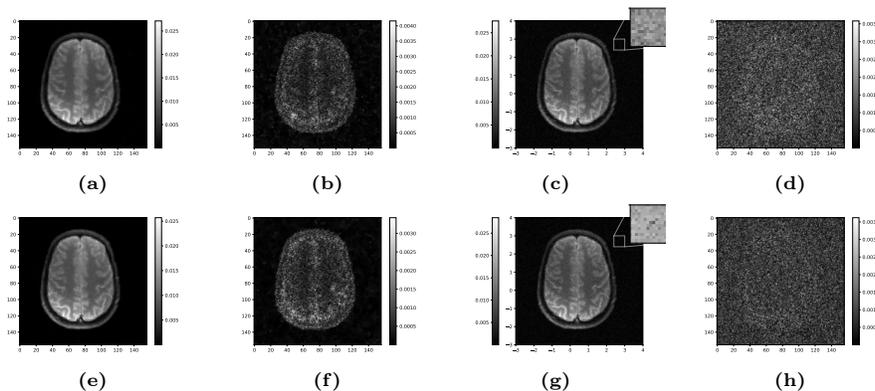


Fig. 5: Reconstruction and debiasing visualization for spiral mask with 43% subsampling and 6.5% noise in the first row and for radial mask with 50% subsampling and 7% noise in the second row 5a & 5e: TV Reconstruction 5b & 5f: Magnitude of difference between ground truth and TV estimator 5c & 5g: Debiased TV Reconstruction, the small enlargement at the top of the image visualizes the noise in the background 5d & 5h: Magnitude of difference between ground truth and debiased TV estimator.

estimators for which the residual is small in the ℓ_∞ -norm is crucial. In the conducted experiments, both the biased and the unbiased estimator have ℓ_∞ -errors compared to the ground truth that are of similar order. Despite this, and the fact that the ℓ_2 -error is even larger in the debiased case, the debiasing significantly reduces the *non-quantifiable* ℓ_∞ -error $\|\hat{x} - x^0\|_\infty$. This is the error for which no statistical information is available in closed form. Indeed, after the reconstruction given by solving Eq. (2) and applying the debiased procedure, the error is given by $\hat{x}^u - x^0 = R + W$. Since, due to our method, R is much smaller than W , which is Gaussian distributed, we can use the latter for the UQ.

The error ratio between the norm of R and of W depends on the quality of the solution of Eq. (2). The more data is available, *i.e.*, the larger m , the better the reconstruction and the better the error ratio, reflected in the confidence interval quality. A similar behavior is observed for the noise level added to the data. The norm of the Gaussian vector W depends linearly on the noise level σ . However, the reconstruction error does not seem to decrease at the same rate. This is related to the matrix difference $M\hat{\Sigma} - I_{N \times N}$, which is independent of the noise level and hence, does not decrease when the noise decreases. We also test our method on real valued brain, prostate and knee images for multiple patients, the results of those experiments are available in Section 3 of the supplementary material.

7.3 Discussion

As shown in the experiments above, our method for *imaging with confidence* is able to provide 95% confidence intervals with hit rates above 90% for each pixel.

We see such results for variational problems like Eq. (2) as particularly promising for the start of a UQ methodology for learned-based methods in MRI and in other applications where deep learning is employed. In particular, the debiasing procedure generates an estimator \hat{x}^u that follows approximately a Gaussian distribution with mean x^0 . This estimator enables us to quantify uncertainty for the MR magnitude and phase of the signal. The latter constitutes the base input data for subsequent processing steps, such as the estimation of tissue susceptibilities.

The size of the given confidence intervals is $\frac{\hat{\sigma}}{\sqrt{m}}$ up to a normalization constant. Intuitively, the more measurements we have, the smaller the range of uncertainty. Also, the dependency on $\hat{\sigma}$ comes from the fact that noisy images are more affected by uncertainty. Since, to the best of the authors' knowledge, we are the first to construct honest confidence intervals for gradient-sparse MR images, there is a lack of comparing results. Still, confidence intervals of the same size are known to be sharp, in the worst-case sense, for similar problems [13].

Limitations. One of the caveats of the *imaging with confidence* method is the estimation of the matrix M for complicated covariance structures given by other sampling schemes such as Lissajous curves [20, 37]. A more involved analysis is necessary in order to apply schemes that are more computationally efficient than the nodewise LASSO. Another crucial point is that since the size of the confidence interval depends on $\hat{\sigma}$, the method provides meaningless UQ for very noisy regimes, *i.e.* data with more than 20% noise. Finally, we acknowledge that the method should be tested in other imaging modalities, such as cardiac MRI, chemical shift imaging, or MR fingerprinting.

8 Conclusion

In this paper, we propose a theory-inspired method for the construction of confidence intervals for fast undersampled MRI. We verify experimentally with in vivo MR data that the resulting confidence intervals are valid, *i.e.* that they contain the true pixel parameter with high probability. A significant advantage of the proposed method is that its application has very modest computational requirements since it is performed without acquiring further data, no Monte Carlo integration is necessary, and it does not require a prior distribution for the underlying image class. Therefore, our method contributes to a higher level of trustworthiness in medical imaging without increasing the scan time. We see a number of promising directions for follow-up research. Firstly, we intend to explore generalizations of our approach to quantitative MRI [56].

Secondly, we intend to extend the method to neural network approaches to MRI, in particular, to variational networks [25] and to vision transformers [38, 44]. We feel that powerful methods for UQ are of particular importance here due to the black-box nature of these approaches; they are a key missing building block to make these often much superior methods [51, 64] widely available in clinical applications.

Acknowledgments

We gratefully acknowledge financial support with funds provided by the German Federal Ministry of Education and Research in the grant “*SparseMRI3D+: Compressive Sensing und Quantifizierung von Unsicherheiten für die beschleunigte multiparametrische quantitative Magnetresonanztomografie (FZK 05M20WOA)*” and funding by the European Union’s Horizon 2020 Research and Innovation Programme under Grant Agreement No. 952172.

References

1. ECCV 2022 Workshop: Workshop on Uncertainty Quantification for Computer Vision. <https://uncv2022.github.io/papers/>, accessed: 2024-02-26
2. FDA 510(k) Premarket Notification for HyperSense. <https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfpmn/pmn.cfm?ID=K162722>, accessed: 2024-02-26
3. FDA 510(k) Premarket Notification for MAGNETOM Aera and MAGNETOM Skyra. <https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfpmn/pmn.cfm?ID=K163312>, accessed: 2024-02-26
4. HyperSense: Compressed sensing and other advanced acceleration techniques. <https://www.gehealthcare.com/-/jssmedia/files/us/non-gated/mri/hypersense-booklet.pdf?rev=-1&hash=FAE260A3F4CA0A82E236DEDA5CAD39A6>, accessed: 2024-02-26
5. ICCV 2023 Workshop: Workshop on Uncertainty Quantification for Computer Vision. <https://uncv2023.github.io/papers/>, accessed: 2024-02-26
6. ICML 2021 Workshop: Workshop on Distribution-Free Uncertainty Quantification. <https://icml.cc/virtual/2021/workshop/8373>, accessed: 2024-02-26
7. Adcock, B., Hansen, A.C.: Compressive imaging: structure, sampling, learning. Cambridge University Press (2021). <https://doi.org/10.1017/9781108377447>
8. Aja-Fernández, S., Vegas-Sánchez-Ferrero, G.: Statistical Analysis of Noise in MRI: Modeling, Filtering and Estimation. Springer, Cham (2016). <https://doi.org/10.1007/978-3-319-39934-8>
9. Angelopoulos, A.N., Bates, S.: A gentle introduction to conformal prediction and distribution-free uncertainty quantification. arXiv preprint arXiv:2107.07511 (2021). <https://doi.org/10.48550/arXiv.2107.07511>
10. Beck, A., Teboulle, M.: A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM journal on imaging sciences* **2**(1), 183–202 (2009). <https://doi.org/10.1137/080716542>
11. Bellec, P.C., Zhang, C.H.: Debiasing convex regularized estimators and interval estimation in linear models. *The Annals of Statistics* **51**(2), 391 – 436 (2023). <https://doi.org/10.1214/22-AOS2243>, [10.1214/22-AOS2243](https://doi.org/10.1214/22-AOS2243)
12. Block, K.T., Uecker, M., Frahm, J.: Undersampled radial MRI with multiple coils. Iterative image reconstruction using a total variation constraint. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine* **57**(6), 1086–1098 (2007). <https://doi.org/10.1002/mrm.21236>
13. Cai, T.T., Guo, Z.: Confidence intervals for high-dimensional linear regression: Minimax rates and adaptivity. *The Annals of Statistics* **45**(2), 615 – 646 (2017). <https://doi.org/10.1214/16-AOS1461>

14. Chambolle, A., Caselles, V., Cremers, D., Novaga, M., Pock, T., et al.: An introduction to total variation for image analysis. Theoretical foundations and numerical methods for sparse recovery **9**(263-340), 227 (2010). <https://doi.org/10.1515/9783110226157.263>
15. Condat, L.: Discrete total variation: New definition and minimization. SIAM Journal on Imaging Sciences **10**(3), 1258–1290 (2017). <https://doi.org/10.1137/16M1075247>
16. Cruz, G., Atkinson, D., Buerger, C., Schaeffter, T., Prieto, C.: Accelerated motion corrected three-dimensional abdominal MRI using total variation regularized SENSE reconstruction. Magnetic resonance in medicine **75**(4), 1484–1498 (2016). <https://doi.org/10.1002/mrm.25708>
17. Cukur, T., Lustig, M., Nishimura, D.G.: Improving non-contrast-enhanced steady-state free precession angiography with compressed sensing. Magnetic Resonance in Medicine **61**(5), 1122–1131 (2009). <https://doi.org/10.1002/mrm.21907>
18. Donoho, D.: Blackboard to bedside: How high-dimensional geometry is transforming the MRI industry. Notices of the American Mathematical Society **65**(1) (2018). <https://doi.org/10.1090/noti1612>
19. Endt, S., Engel, M., Naldi, E., Assereto, R., Molendowska, M., Mueller, L., Mayrink Verdun, C., Pirkel, C.M., Palombo, M., Jones, D.K., Menzel, M.I.: In Vivo Myelin Water Quantification Using Diffusion-Relaxation Correlation MRI: A Comparison of 1D and 2D Methods. Applied Magnetic Resonance **54**(11), 1571–1588 (2023). <https://doi.org/10.1007/s00723-023-01584-1>
20. Feng, H., Gu, H., Silbersweig, D., Stern, E., Yang, Y.: Single-shot MR imaging using trapezoidal-gradient-based Lissajous trajectories. IEEE transactions on medical imaging **22**(8), 925–932 (2003). <https://doi.org/10.1109/TMI.2003.815902>
21. Fischer, P., Thomas, K., Baumgartner, C.F.: Uncertainty Estimation and Propagation in Accelerated MRI Reconstruction. In: International Workshop on Uncertainty for Safe Utilization of Machine Learning in Medical Imaging. pp. 84–94. Springer (2023). https://doi.org/10.1007/978-3-031-44336-7_9
22. Fu, W., Knight, K.: Asymptotics for lasso-type estimators. The Annals of statistics **28**(5), 1356–1378 (2000). <https://doi.org/10.1214/aos/1015957397>
23. Giraud, C.: Introduction to high-dimensional statistics. Chapman and Hall/CRC (2021)
24. Goldstein, T., Osher, S.: The split Bregman method for L1-regularized problems. SIAM journal on imaging sciences **2**(2), 323–343 (2009). <https://doi.org/10.1137/080725891>
25. Hammernik, K., Klatzer, T., Kobler, E., Recht, M.P., Sodickson, D.K., Pock, T., Knoll, F.: Learning a variational network for reconstruction of accelerated MRI data. Magnetic resonance in medicine **79**(6), 3055–3071 (2018). <https://doi.org/10.1002/mrm.26977>
26. Han, Y., Lee, T.C.: Uncertainty Quantification for Sparse Estimation of Spectral Lines. IEEE Transactions on Signal Processing **70**, 6243–6256 (2022). <https://doi.org/10.1109/TSP.2023.3235662>
27. Hannig, J., Iyer, H., Lai, R.C., Lee, T.C.: Generalized fiducial inference: A review and new results. Journal of the American Statistical Association **111**(515), 1346–1361 (2016). <https://doi.org/10.1080/01621459.2016.1165102>
28. Hoppe, F., Kraemer, F., Mayrink Verdun, C., Menzel, M., Rauhut, H.: Uncertainty quantification for sparse Fourier recovery. arXiv:2212.14864 (2022). <https://doi.org/10.48550/arXiv.2212.14864>

29. Hoppe, F., Kraemer, F., Mayrink Verdun, C., Menzel, M., Rauhut, H.: High-Dimensional Confidence Regions in Sparse MRI. In: ICASSP 2023. pp. 1–5 (2023). <https://doi.org/10.1109/ICASSP49357.2023.10096320>
30. Hoppe, F., Mayrink Verdun, C., Laus, H., Kraemer, F., Rauhut, H.: Non-Asymptotic Uncertainty Quantification in High-Dimensional Learning. arXiv preprint (2024)
31. Hüllermeier, E., Waegeman, W.: Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. *Machine Learning* **110**, 457–506 (2021). <https://doi.org/10.1007/s10994-021-05946-3>
32. Hütter, J.C., Rigollet, P.: Optimal rates for total variation denoising. In: Conference on Learning Theory. pp. 1115–1146. PMLR (2016)
33. Hyun, C.M., Kim, H.P., Lee, S.M., Lee, S., Seo, J.K.: Deep learning for undersampled MRI reconstruction. *Physics in Medicine & Biology* **63**(13), 135007 (2018). <https://doi.org/10.1088/1361-6560/aac71a9>
34. Jaspan, O., Fleysher, R., Lipton, M.: Compressed Sensing MRI: A review of the clinical literature. *The British journal of radiology* **88**, 20150487 (09 2015). <https://doi.org/10.1259/bjr.20150487>
35. Javanmard, A., Montanari, A.: Confidence intervals and hypothesis testing for high-dimensional regression. *Journal of Machine Learning Research* **15**, 2869–2909 (2014)
36. Javanmard, A., Montanari, A.: Debiasing the lasso: Optimal sample size for Gaussian designs. *The Annals of Statistics* **46**(6A) (2018). <https://doi.org/10.1214/17-AOS1630>
37. Kaethner, C., Erb, W., Ahlborg, M., Szwargulski, P., Knopp, T., Buzug, T.M.: Non-equispaced system matrix acquisition for magnetic particle imaging based on Lissajous node points. *IEEE transactions on medical imaging* **35**(11), 2476–2485 (2016). <https://doi.org/10.1109/TMI.2016.2580458>
38. Korkmaz, Y., Yurt, M., Dar, S.U.H., Özbey, M., Cukur, T.: Deep MRI reconstruction with generative vision transformers. In: Machine Learning for Medical Image Reconstruction: 4th International Workshop, MLMIR 2021, Held in Conjunction with MICCAI 2021, Strasbourg, France, October 1, 2021, Proceedings 4. pp. 54–64. Springer (2021). https://doi.org/10.1007/978-3-030-88552-6_6
39. Kraemer, F., Ward, R.: Stable and robust sampling strategies for compressive imaging. *IEEE transactions on image processing* **23**(2), 612–622 (2013). <https://doi.org/10.1109/TIP.2013.2288004>
40. Kraemer, F., Kruschel, C., Sandbichler, M.: Total variation minimization in compressed sensing. In: Compressed Sensing and its Applications: Second International MATHEON Conference 2015. pp. 333–358. Springer (2017). <https://doi.org/10.1109/CSE-EUC-DCABES.2016.233>
41. Lai, R.C., Hannig, J., Lee, T.C.: Generalized fiducial inference for ultrahigh-dimensional regression. *Journal of the American Statistical Association* **110**(510), 760–772 (2015). <https://doi.org/10.1080/01621459.2014.931237>
42. Liang, Z.P., Lauterbur, P.C.: Principles of magnetic resonance imaging. SPIE Optical Engineering Press Bellingham, WA (2000)
43. Lin, A., Song, A.H., Bilgic, B., Ba, D.: Covariance-free sparse Bayesian learning. *IEEE Transactions on Signal Processing* **70**, 3818–3831 (2022). <https://doi.org/10.1109/TSP.2022.3186185>
44. Lin, K., Heckel, R.: Vision transformers enable fast and robust accelerated MRI. In: International Conference on Medical Imaging with Deep Learning. pp. 774–795. PMLR (2022)

45. Lustig, M., Donoho, D., Pauly, J.M.: Sparse MRI: The application of compressed sensing for rapid MR imaging. *Magnetic Resonance in Medicine* **58**(6), 1182–1195 (2007). <https://doi.org/10.1002/mrm.21391>
46. Mueller, L., Rudrapatna, S.U., Tax, C.M., Wise, R., Jones, D.K.: Diffusion MRI with $b=1000\text{s/mm}^2$ at $TE<22\text{ms}$ using single-shot spiral readout and ultra-strong gradients: Implications for microstructure imaging. *Proceedings of the International Society for Magnetic Resonance in Medicine*, 2019. Presented at the ISMRM (2019)
47. Nassirpour, S., Chang, P., Avdievitch, N., Henning, A.: Compressed sensing for high-resolution nonlipid suppressed 1H FID MRSI of the human brain at 9.4 T. *Magnetic resonance in medicine* **80**(6), 2311–2325 (2018). <https://doi.org/10.1002/mrm.27225>
48. Needell, D., Ward, R.: Stable image reconstruction using total variation minimization. *SIAM Journal on Imaging Sciences* **6**(2), 1035–1058 (2013). <https://doi.org/10.1137/120868281>
49. Poon, C.: On the role of total variation in compressed sensing. *SIAM Journal on Imaging Sciences* **8**(1), 682–720 (2015). <https://doi.org/10.1137/140978569>, <https://doi.org/10.1137/140978569>
50. Pruessmann, K.P., Weiger, M., Börnert, P., Boesiger, P.: Advances in sensitivity encoding with arbitrary k-space trajectories. *Magnetic Resonance in Medicine* **46**(4), 638–651 (2001). <https://doi.org/10.1002/mrm.1241>
51. Rastogi, A., Brugnara, G., Foltyn-Dumitru, M., Mahmutoglu, M.A., Preetha, C.J., Kobler, E., Pflüger, I., Schell, M., Deike-Hofmann, K., Kessler, T., et al.: Deep-learning-based reconstruction of undersampled MRI to reduce scan times: a multicentre, retrospective, cohort study. *The Lancet Oncology* **25**(3), 400–410 (2024). [https://doi.org/10.1016/S1470-2045\(23\)00641-1](https://doi.org/10.1016/S1470-2045(23)00641-1)
52. Roth, S., Black, M.: Fields of Experts: a framework for learning image priors. In: 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05). vol. 2, pp. 860–867 vol. 2 (2005). <https://doi.org/10.1109/CVPR.2005.160>
53. Rudin, L.I., Osher, S., Fatemi, E.: Nonlinear total variation based noise removal algorithms. *Physica D: nonlinear phenomena* **60**(1-4), 259–268 (1992). [https://doi.org/10.1016/0167-2789\(92\)90242-F](https://doi.org/10.1016/0167-2789(92)90242-F)
54. Saucedo, A., Macey, P.M., Thomas, M.A.: Accelerated radial echo-planar spectroscopic imaging using golden angle view-ordering and compressed-sensing reconstruction with total variation regularization. *Magnetic Resonance in Medicine* **86**(1), 46–61 (2021). <https://doi.org/10.1002/mrm.28728>
55. Schweser, F., Deistung, A., Reichenbach, J.R.: Foundations of MRI phase imaging and processing for Quantitative Susceptibility Mapping (QSM). *Zeitschrift für medizinische Physik* **26**(1), 6–34 (2016). <https://doi.org/10.1016/j.zemedi.2015.10.002>
56. Shafieizargar, B., Byanju, R., Sijbers, J., Klein, S., den Dekker, A.J., Poot, D.H.: Systematic review of reconstruction techniques for accelerated quantitative MRI. *Magnetic Resonance in Medicine* (2023). <https://doi.org/10.1002/mrm.29721>
57. Tezcan, K.C., Karani, N., Baumgartner, C.F., Konukoglu, E.: Sampling possible reconstructions of undersampled acquisitions in MR imaging with a deep learned prior. *IEEE Transactions on Medical Imaging* **41**(7), 1885–1896 (2022). <https://doi.org/10.1109/TMI.2022.3150853>
58. Tibshirani, R.: Regression Shrinkage and Selection via the Lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **58**(1), 267–288 (1996). <https://doi.org/10.1111/j.2517-6161.1996.tb02080.x>

59. Valdenegro-Toro, M.: I find your lack of uncertainty in computer vision disturbing. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1263–1272 (2021). <https://doi.org/10.1109/CVPRW53098.2021.00139>
60. van de Geer, S., Bühlmann, P., Ritov, Y., Dezeure, R.: On asymptotically optimal confidence regions and tests for high-dimensional models. *The Annals of Statistics* **42**(3) (2014). <https://doi.org/10.1214/14-AOS1221>
61. Wilm, B.J., Barmet, C., Gross, S., Kasper, L., Vannesjo, S.J., Haeberlin, M., Dietrich, B.E., Brunner, D.O., Schmid, T., Pruessmann, K.P.: Single-shot spiral imaging enabled by an expanded encoding model: Demonstration in diffusion MRI. *Magnetic resonance in medicine* **77**(1), 83–91 (2017). <https://doi.org/10.1002/mrm.26493>
62. Wilm, B.J., Barmet, C., Pavan, M., Pruessmann, K.P.: Higher order reconstruction for MRI in the presence of spatiotemporal field perturbations. *Magnetic resonance in medicine* **65**(6), 1690–1701 (2011). <https://doi.org/10.1002/mrm.22767>
63. Winkelmann, S., Schaeffter, T., Koehler, T., Eggers, H., Doessel, O.: An optimal radial profile order based on the Golden Ratio for time-resolved MRI. *IEEE transactions on medical imaging* **26**(1), 68–76 (2006). <https://doi.org/10.1109/TMI.2006.885337>
64. Zeng, G., Guo, Y., Zhan, J., Wang, Z., Lai, Z., Du, X., Qu, X., Guo, D.: A review on deep learning MRI reconstruction without fully sampled k-space. *BMC Medical Imaging* **21**(1), 195 (2021). <https://doi.org/10.1186/s12880-021-00727-9>
65. Zhang, C.H., Zhang, S.S.: Confidence intervals for low dimensional parameters in high dimensional linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **76**(1), 217–242 (2014). <https://doi.org/10.1111/rssb.12026>