An Inverse-Problem Approach to the Estimation of Despeckled and Deconvolved Images From Radio-Frequency Signals in Pulse-Echo Ultrasound

Samuel Beuret, *Member, IEEE*, Adrien Besson, *Member, IEEE*, Akihiro Sugimoto, *Member, IEEE*, and Jean-Philippe Thiran, *Senior Member, IEEE*

***Abstract*—In recent years, there has been notable progress in the development of inverse problems for image reconstruction in pulse-echo ultrasound. Inverse problems are designed to circumvent the restrictions of delay-and-sum, such as limited image resolution and diffraction artifacts, especially when low amount of data are considered. However, the radio-frequency image or tissue reflectivity function that current inverse problems seek to estimate do not possess a structure that can be easily leveraged by a regularizer, in part due to their high dynamic range. The performance of inverse-problem image reconstruction is thus impeded. In contrast, despeckled images exhibit a more exploitable structure. Therefore, we first propose an inverse problem to recover a despeckled image from single-plane-wave radio-frequency echo signals, employing total-variation norm regularization. Then, we introduce an inverse problem to estimate the tissue reflectivity function from radio-frequency echo signals, factoring in the despeckled image obtained by the first problem into a spatially-varying reflectivity prior. We show with simulated, in-vitro, and in-vivo data that the proposed despeckled image estimation technique recovers images almost free of diffraction artifacts and improves contrast with respect to delay-and-sum and non-local means despeckling. Moreover, we show with in-vitro and in-vivo data that the proposed reflectivity estimation method reduces artifacts and improves contrast with respect to a state-of-the-art inverse problem positing a uniform prior. In particular, the proposed techniques could prove beneficial for imaging with ultra-portable transducers, since these devices are likely to be limited in the amount of data they can acquire and transmit.**

Manuscript received 22 December 2023; revised 10 May 2024 and 26 July 2024; accepted 31 July 2024. Date of publication 12 August 2024; date of current version 23 August 2024. This work was supported by the Ecole polytechnique fédérale de Lausanne. The associate editor coordinating the review of this article and approving it for publication was Prof. Andreas Hauptmann. *(Corresponding author: Samuel Beuret.)*

This work involved human subjects or animals in its research. Approval of all ethical and experimental procedures and protocols was granted by the Cantonal Commission on Ethics in Human Research of the Canton of Vaud, Switzerland (Application No. 2022-0169) and performed in line with the Declaration of Helsinki.

Samuel Beuret is with the Signal Processing Laboratory (LTS5), École polytechnique fédérale de Lausanne (EPFL), 1015 Lausanne, Switzerland (e-mail: samuel.beuret@epfl.ch).

Adrien Besson is with E-Scopics, 13290 Aix-en-Provence, France (e-mail: adrien.besson@e-scopics.com).

Akihiro Sugimoto is with the National Institute of Informatics (NII), Tokyo 101-0003, Japan (e-mail: sugimoto@nii.ac.jp).

Jean-Philippe Thiran is with the Signal Processing Laboratory (LTS5), École polytechnique fédérale de Lausanne (EPFL), 1015 Lausanne, Switzerland, and also with the Department of Radiology, University Hospital Center and University of Lausanne (CHUV), 1015 Lausanne, Switzerland (e-mail: jean-philippe.thiran@epfl.ch).

Digital Object Identifier 10.1109/TCI.2024.3441234

***Index Terms*—Despeckling, inverse problems, optimization, probabilistic modeling, pulse-echo ultrasound.**

I. Introduction

ULTRAFAST ultrasound has emerged as a new paradigm in echography competing with traditional focused ultrasound. This approach relies on the emission of unfocused waves by an ultrasound transducer. Human soft tissues can therefore be imaged with a limited number of insonifications. On the one hand, this technique enables imaging at potentially very high frame-rates. These frame-rates act as the backbone behind methods such as shear-wave elastography [1] or ultrasound neurofunctional imaging [2]. On the other hand, ultrafast ultrasound can greatly reduce the amount of data that need to be acquired with respect to focused ultrasound. Thus, point-of-care imaging can be eased since portable ultrasound transducers possess limitations in terms of data storage and transmission bandwidth.

The reconstruction of radio-frequency (RF) images from the acquired echo signals is usually performed by delay-and-sum (DAS) and coherent compounding [3]. DAS possesses numerous advantages: low computational cost, conceptual simplicity and stability to noise and aberrations. However, DAS generates diffraction artifacts—side lobes, grating lobes—in the reconstructed images, and the severity of the artifacts increases when the amount of data considered is reduced. In particular, diffraction artifacts affect the diagnosis capability of images reconstructed from a low amount of data. In addition, the resolution of images reconstructed using DAS is limited by the center frequency of the ultrasound transducer, as well as its physical aperture.

To circumvent the limitations of DAS, a series of methods recast image reconstruction as an inverse problem involving a linear measurement model relating the RF measured echo signals to the unknown image of interest. In addition, a regularization term accounting for some a-priori knowledge on the unknown image is considered. The inverse-model-based image reconstruction approaches can be separated into two categories depending on if the method aims at improving resolution or not. We denote this class of techniques as *regularized beamforming*, and they are typically able to improve contrast—and, for some methods, resolution—with respect to standard beamformers such as DAS.

We denote the latter category as *regularized beamforming* methods. Two main difficulties arise in this class of problems.

Two main difficulties arise in this set of problems. First, the dimensionality of the measurement model—and thus the memory footprint of the matrix discretizing the model—is in practice very large. Second, the choice of regularizer is difficult since RF images do not admit a straightforward structure that can be easily leveraged by a regularizer in a general case. Second, the choice of regularizer is difficult since ultrasound images do not admit a straightforward structure that can be easily leveraged by a regularizer in a general case. Several measurement models have been proposed, including frequency-domain approaches [4], [5], and spatial-domain models [6], [7], [8]. Importantly, a matrix-free spatial-domain approach has been proposed in [9], where an ℓ_1 norm in a wavelet basis is used as a regularizer. Other regularizers comprise ℓ_1 and ℓ_2 norms [6], [10], a combination of space and frequency domain regularizers [7], and, more recently, regularization through a plug-and-play approach [8].

The regularized beamforming methods seen up to now seek to recover a high-quality image from measured echo signals. However, another set of techniques aims at estimating the tissue reflectivity function (TRF or *reflectivity*) itself. Thus, such methods can achieve significant resolution improvement with regards to standard beamforming approaches. However, the conditioning of the measurement models involved in these problems is significantly worse than the conditioning of models seen up to now. There is therefore a large dependency of the estimated TRF to the choice of regularizer. The latter includes, ℓ_1 norms [11], [12] or, more generally, ℓ_p norms [13].

In contrast, a different family of methods targets the estimation of the reflectivity from a high-quality RF ultrasound image, namely deconvolution methods. Such techniques suppose that the image is the result of the convolution of the TRF by a point spread function (PSF), assumed constant in a region of the medium. The first ultrasound deconvolution techniques were based on Wiener filtering [14]. More complex approaches proposed since include homomorphic deconvolution [15], blind deconvolution [16]—where the PSF is estimated alongside the TRF—, or deconvolution techniques using sparse regularizers [17], [18]. In particular, generalized Gaussian distributions (GGD) are now commonly used as reflectivity prior—or, equivalently, the TRF is commonly regularized by an ℓ_p norm. Unfortunately, the priors proposed in both deconvolution techniques and reflectivity estimation methods from echo signals often fail to take into account the large dynamic range of the reflectivity. To circumvent this issue, Corbineau et al. assume that the TRF is drawn from a GGD whose variance and shape parameter vary within the medium [19]. GGD variance and shape parameter are supposed constant in a fixed number of zones, whereas deconvolution and segmentation of the medium are performed jointly. Since they assume PSFs that are constant in the image and limited spatially, deconvolution methods are however not able to model diffraction artifacts, and thus to correct them. Consequently, deconvolution techniques are unsuited to imaging configurations using limited amount of data.

A last class of methods we can highlight comprises despeckling algorithms. In practice, their aim is to filter speckle from B-mode images to ease the analysis of ultrasound images by a practitioner. Several approaches have been proposed, including wavelet-based denoising [20], or non-local means (NLM) filtering [21]. Importantly, TV norm also lays at the basis of several despeckling techniques [22], [23], [24]. More generally, we must emphasize that despeckled images possess a more easily leverageable structure compared to the TRF, a RF image, or a B-mode image prior to despeckling. However, despeckling methods are in general not able to correct for diffraction artifacts. Therefore, the latter also affect despeckled images. The diagnosis capability of ultrasound imaging with a reduced amount of data is ultimately hindered. Furthermore, despeckled images are also used as inputs for segmentation algorithms [25]. Thus, the accuracy of image segmentation is also expected to suffer when a low number of insonifications is used.

Taking into consideration the methods discussed thus far and their limitations, we can summarize the contributions of the article as follows.

- The first contribution is a novel statistical model of data acquisition in pulse-echo ultrasound. In particular, we express measurements as the realization of a high-dimensional multivariate Gaussian random variable.
- Based on the proposed statistical model, our second contribution is a regularized inverse problem aimed at the direct reconstruction of a despeckled image from a limited number of measured RF echo signals. Specifically, we propose to identify a despeckled image of the medium with the log-scale variance of the TRF, whereas we regularize the latter with a TV norm.
- Our third contribution concerns reflectivity estimation from RF echo signals. We propose a novel non-uniform TRF prior factoring in the despeckled image estimated by the first inverse problem. Thus, our goal is to account for the high dynamic range of the reflectivity more effectively than current state-of-the-art approaches.

The article is organized as follows. Section II presents the measurement model we use to relate the TRF to the measured RF echo signals, the definition of the inverse problem for the despeckled image and the optimization of the associated loss function. Finally, the reflectivity inverse problem is introduced. We detail in Section III the experiments performed to assess the proposed methods. Section IV presents the results of the despeckled image and reflectivity estimation techniques. Ultimately, we discuss the results in Section V, whereas Section VI concludes the article.

II. Theory

A. Notation

We begin this section by briefly detailing the notation we use throughout this article. We denote vectors by bold lowercase letters, whereas matrices are represented by bold capital letters, with $\boldsymbol{I}$ the identity matrix. $\mathrm{Diag}(\boldsymbol{x})$ denotes the linear operator constructing a diagonal matrix from a vector $\boldsymbol{x}$, whereas $\mathrm{diag}(\boldsymbol{A})$ is its adjoint constructing a vector from the diagonal elements of $\boldsymbol{A}$. In the following, $\odot$ represents the element-wise product between two vectors, $|\boldsymbol{x}|^2$ is the element-wise squared modulus of $\boldsymbol{x}$, and $e^{\boldsymbol{x}}$ is the element-wise exponential of $\boldsymbol{x}$. In addition, $x_{(i,j)}$ is an alias for the $(i(N-1)+j)$-th element of $\boldsymbol{x} \in \mathbb{R}^{MN}$ or $\boldsymbol{x} \in \mathbb{C}^{MN}$. Finally, $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes a complex multivariate

centrally-symmetric Gaussian distribution with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$.

B. Measurement Model

We present herein the fundamental building block of the proposed method, namely the measurement model. We detail the imaging configuration and the fundamental assumptions behind our model of the received echo signals.

We suppose that a single, unsteered, plane wave (PW) is emitted by a linear ultrasound transducer. Furthermore, we suppose that the N^{el} piezoelectric elements of the transducer record back the echo signals generated by local impedance variations occurring within the insonified medium. Throughout this article, we assume that the transducer is aligned with the x-axis and emits waves propagating towards the positive z direction. We represent the acquired signals as a vector of complex radio-frequency (CRF) signals $\boldsymbol{m} \in \mathbb{C}^{N^{\mathrm{el}}N^{\mathrm{t}}}$, with N^{t} the number of time samples. In practice, Hilbert transform is used to obtain analytic CRF signals from the real RF signals. We justify in Section II-E our choice of working with CRF signals instead of their RF counterparts.

We consider a linear measurement model linking the measurement vector $\boldsymbol{m}$ to the TRF, a complex function factoring in local impedance variations. We rely on the model detailed in [9], [13] defined as

$$m_{(k,j)} = \int_{\boldsymbol{r}} o^{\mathrm{Tx}}(\boldsymbol{r})\, o_k^{\mathrm{Rx}}(\boldsymbol{r}) \, v_{\mathrm{pe}}\left(t_j^{\mathrm{s}} - \tau^{\mathrm{Tx}}(\boldsymbol{r}) - \tau_k^{\mathrm{Rx}}(\boldsymbol{r})\right) \gamma(\boldsymbol{r})\, \mathrm{d}\boldsymbol{r}, \quad (1)$$

where t_j^{s} denotes the j-th sampling time instant, $k = 1, \ldots, N^{\mathrm{el}}$ is the index of the piezoelectric transducer in reception and $\boldsymbol{r} = [x, z]^T$ represents an arbitrary points in the medium. In (1), τ^{Tx} and τ^{Rx} are the transmit (Tx) and receive (Rx) propagation times, defined as

$$\tau^{\mathrm{Tx}}(\boldsymbol{r}) = \frac{1}{c_0} z, \quad (2)$$

and

$$\tau_k^{\mathrm{Rx}}(\boldsymbol{r}) = \frac{1}{c_0}\sqrt{\left(x - x_k^{\mathrm{el}}\right)^2 + z^2}, \quad (3)$$

where c_0 is the speed of sound (SoS) of the medium. Functions o^{Tx} and o^{Rx} are real positive factors, whereas γ denotes the continuous TRF. Finally, v_{pe} is the pulse-echo wavelet defined in [26]. Factors o^{Tx} and o^{Rx} in (1) account for magnitude effects affecting signals scattered at point $\boldsymbol{r}$, in transmission and reception respectively. Such effects mostly include the decay of waves with distance, the focusing of the acoustic lens or the directivity of the sensors.

Several hypotheses are required to define (1). Most importantly, we suppose a first-order Born approximation. Namely, multiple scattering is neglected. In addition, we assume linear and isotropic wave propagation in the medium, large-scale variations of SoS and attenuation are neglected, and we posit that the latter does not depend on frequency. Finally, we postulate that the medium is entirely located in the far-field of the transducer's elements.

In the following, we consider that the reflectivity map is discretized according to a grid of points $\boldsymbol{r}_{(m,n)}$ with $m = 1, \ldots, N^{\mathrm{x}}$, $n = 1, \ldots, N^{\mathrm{z}}$ such that it is represented by a vector $\boldsymbol{\gamma} \in \mathbb{C}^{N^{\mathrm{x}}N^{\mathrm{z}}}$. Thus, we represent the measurement model as a matrix $\boldsymbol{A} \in \mathbb{C}^{N^{\mathrm{el}}N^{\mathrm{t}} \times N^{\mathrm{x}}N^{\mathrm{z}}}$. Importantly, we never construct matrix $\boldsymbol{A}$ explicitly and matrix-vector multiplication is implement according to (1) using linear interpolation. A similar strategy is used to implement $\boldsymbol{A}^H$. We refer the reader to [9], [13] for a detailed description of the measurement model and its implementation.

Ultimately, we model the signals as

$$\boldsymbol{m} = \boldsymbol{A}\boldsymbol{\gamma} + \boldsymbol{\nu}, \; \boldsymbol{\nu} \sim \mathcal{N}\left(\boldsymbol{0}, \sigma_\nu^2 \boldsymbol{I}\right). \quad (4)$$

In (4), vector $\boldsymbol{\nu}$ represents thermal noise, but more importantly the inaccuracies of $\boldsymbol{A}$. The neglect of multiple scattering, spatial variations of speed of sound and attenuation, as well as the estimation of v_{pe} by a Gaussian pulse introduce errors in the estimated signals. We ultimately model them by white Gaussian noise added to the signals predicted by the proposed measurement model.

We must highlight that additive Gaussian noise applied to $\boldsymbol{A}\boldsymbol{\gamma}$ is a simplistic hypothesis. The quality of the method proposed herein is therefore expected to suffer compared to a technique that would consider a more complex—most likely non-linear—measurement model factoring in multiple scattering, SoS variations, frequency-dependent attenuation or space-dependent attenuation. However, this assumption is enforced to keep the model computationally tractable. Furthermore, the Gaussian noise hypothesis—or equivalently, the use of a least-squares data term—is supposed in the vast majority of regularized beamforming and reflectivity estimation methods proposed previously [6], [7], [9], [13].

C. Inverse Problem for the Echogenicity Map

In this subsection, we introduce an inverse problem to estimate the local variance of the reflectivity in a logarithmic scale—that we interpret as despeckled image—from the measured echo signal. We also detail the statistical model underpinning the inverse problem and the assumptions it is based on, as well as the regularization strategy.

We first establish several assumptions with regard to the reflectivity. We suppose that each component γ_k, $k = 1, \ldots, N^{\mathrm{x}}N^{\mathrm{z}}$ of the discrete reflectivity vector $\boldsymbol{\gamma}$ are uncorrelated. Besides, we posit that the probability distributions from which each γ_k is drawn have zero means and variances $\mathbb{E}\{|\gamma_k|^2\} = \sigma_k^2 < \infty$. Importantly, we assume that the variance of the reflectivity σ_k^2 can be different for each k. In other terms, we suppose that the reflectivity variance varies within the medium.

The absence of correlation between the discrete reflectivity at different points of the medium is a standard hypothesis in ultrasound image deconvolution [18], [19] and inverse-problem-based pulse-echo ultrasound imaging [13]. Even if the TRF is ultimately continuous, the scale of the variations is typically inferior to grid size—especially in zones appearing as speckle in B-mode image. Thus, this fact justifies a lack of correlation

between the reflectivity at adjacent pixels. However, this argument is mitigated in the case of specular reflections. In the future, we can thus imagine a more complex model better suited for the imaging of specular reflectors.

Based on the hypotheses detailed above, the covariance matrix $\boldsymbol{\Sigma}$ of the measured signals $\boldsymbol{m}$ is then given by

$$\boldsymbol{\Sigma}(\boldsymbol{\kappa}) = \boldsymbol{A}\text{Diag}\left(e^{\boldsymbol{\kappa}}\right)\boldsymbol{A}^H + \sigma_\nu^2\boldsymbol{I}, \tag{5}$$

where $\boldsymbol{\kappa}$ denotes the element-wise logarithm of the TRF variance $\sigma_k^2, \forall k = 1, \ldots, N^{\text{x}}N^{\text{z}}$.

Importantly, we interpret $\boldsymbol{\kappa}$ as a despeckled image expressed in a logarithmic scale. We refer it as the *echogenicity* in the following. The logarithmic scale is chosen to better account for the high dynamic range of the echogenicity inherent to pulse-echo ultrasound imaging. To the best of our knowledge, this article represents the first attempt to establish an explicit link between reflectivity variance and despeckled images in ultrasound imaging. However, this claim follows from the theoretical work of Michailovich and Tannenbaum [22]. There, the reflectivity is first estimated through image deconvolution. Its squared norm—expressed in a logarithmic scale—is then filtered to estimate a despeckled image. We argue for the similarity between the approach of [22] and an estimation of the log-scale reflectivity variance.

We now present the fundamental hypothesis underpinning the proposed method. We assume that the measurement vector $\boldsymbol{m}$ is given by a single realization of a circularly-symmetric complex multivariate Gaussian variable with zero mean and covariance matrix $\boldsymbol{\Sigma}(\boldsymbol{\kappa})$. In particular, the probability density function (PDF) of the distribution is given by

$$\mathbb{P}(\boldsymbol{m}|\boldsymbol{\Sigma}) = \frac{1}{\pi|\boldsymbol{\Sigma}|}e^{-\boldsymbol{m}^H\boldsymbol{\Sigma}^{-1}\boldsymbol{m}}. \tag{6}$$

Our goal is, first, to estimate $\boldsymbol{\kappa}$. To do so, we propose to solve the following optimization problem:

$$\min_{\boldsymbol{\kappa}\in\mathbb{R}^{N^{\text{x}}N^{\text{z}}}} L(\boldsymbol{\kappa}) + \mu R(\boldsymbol{\kappa}), \tag{7}$$

with

$$L(\boldsymbol{\kappa}) = \log\left|\boldsymbol{\Sigma}(\boldsymbol{\kappa})\right| + \boldsymbol{m}^H\boldsymbol{\Sigma}^{-1}(\boldsymbol{\kappa})\boldsymbol{m}, \tag{8}$$

and $\boldsymbol{\Sigma}(\boldsymbol{\kappa})$ the covariance matrix as defined in (5). Problem (7) is a maximum-a-posteriori (MAP) estimator of $\boldsymbol{\kappa}$, where the regularizer μR is a positive functional acting as a Bayesian prior on the echogenicity map $\boldsymbol{\kappa}$, and where L denotes the negative log-likelihood that the measurement vector $\boldsymbol{m}$ is generated by the distribution defined in (6).

In this article, we propose to use the following regularizer

$$R(\boldsymbol{\kappa}) = \sum_{i}^{(N^{\text{x}}-1)(N^{\text{z}}-1)} \sqrt{[\boldsymbol{D}^{\text{x}}\boldsymbol{\kappa}]_i^2 + [\boldsymbol{D}^{\text{z}}\boldsymbol{\kappa}]_i^2} + \chi_{]-\infty,\kappa^{\max}]}(\boldsymbol{\kappa}), \tag{9}$$

where $\boldsymbol{D}^{\text{x}}$ and $\boldsymbol{D}^{\text{z}}$ denote forward finite difference operators along the x and z axes, respectively. Furthermore, χ denote the characteristic function of the interval ranging from $-\infty$ to $\kappa^{\max}$. The first part of the regularizer corresponds to an isotropic TV-norm. The second part ensures that the echogenicity remains within an expected range upper-bounded by $\kappa^{\max}$. This last term is present to alleviate potential convergence issues of the algorithm we will present in Section II-D. The strength of the prior is encoded by the regularization parameter μ. The choice of the TV norm is motivated by its previous uses as a prior on despeckled images [22], [23], [24]. More generally, the piece-wise constancy enforced by TV norm is a hypothesis regularly assumed to model the spatial distribution of physical properties of human tissues in medical imaging [27], [28], [29], [30].

D. Natural Proximal Gradient Descent

We describe in this subsection how we propose to solve the inverse problem defined in Section II-C. In particular, we detail the optimization algorithm used.

In practice, we determine the solution of (7) with an approximate natural proximal gradient descent algorithm. The basic principle of natural gradient descent is to replace the Euclidean metric implicitly assumed by gradient descent by the Kullback-Leiber (KL) divergence between the distribution parameterized by the variable the algorithm seeks to optimize—here, $\boldsymbol{\kappa}$ [31], [32]. This metric is estimated locally by the following distance

$$d(\boldsymbol{\kappa}, \hat{\boldsymbol{\kappa}}) = \sqrt{[\boldsymbol{\kappa} - \hat{\boldsymbol{\kappa}}]^T\boldsymbol{F}(\boldsymbol{\kappa})[\boldsymbol{\kappa} - \hat{\boldsymbol{\kappa}}]}, \tag{10}$$

where $\boldsymbol{F}$ denotes the Fisher information matrix. Due to the high dimensionality of the problem in our case, we approximate the Fisher information matrix by a diagonal matrix $\boldsymbol{F} \approx \text{Diag}(\boldsymbol{f})$. To handle the proposed non-smooth TV-norm regularizer, we use a proximal step similar to the one introduced for proximal Newton-type methods [33]. The proposed optimization scheme can thus be written as

$$\begin{cases} \hat{\boldsymbol{\kappa}} = \boldsymbol{\kappa}_k - \eta\text{Diag}^{-1}\left(\boldsymbol{f}(\boldsymbol{\kappa}_k)\right)\left[\nabla_{\boldsymbol{\kappa}}L(\boldsymbol{\kappa}_k)\right], \\ \boldsymbol{\kappa}_{k+1} = \text{Prox}_{\eta\mu R}^{\text{Diag}(\boldsymbol{f}(\boldsymbol{\kappa}_k))}(\hat{\boldsymbol{\kappa}}), \end{cases} \tag{11}$$

for a specific iteration k and step size η. Furthermore, the proximal step in (11) is defined as

$$\text{Prox}_{\eta\mu R}^{\text{Diag}(\boldsymbol{f})}(\hat{\boldsymbol{\kappa}}) = \arg\min_{\boldsymbol{\kappa}} \eta\mu R(\boldsymbol{\kappa}) + \frac{1}{2}[\boldsymbol{\kappa} - \hat{\boldsymbol{\kappa}}]^T\text{Diag}(\boldsymbol{f})[\boldsymbol{\kappa} - \hat{\boldsymbol{\kappa}}]. \tag{12}$$

Algorithm (11) requires the gradient of the data term L. It can be split into two terms, namely

$$\nabla_{\boldsymbol{\kappa}}\left[\log\left|\boldsymbol{\Sigma}(\boldsymbol{\kappa})\right|\right] = \text{diag}\left(\boldsymbol{A}^H\boldsymbol{\Sigma}^{-1}(\boldsymbol{\kappa})\boldsymbol{A}\right) \odot e^{\boldsymbol{\kappa}}, \tag{13}$$

and

$$\nabla_{\boldsymbol{\kappa}}\left[\boldsymbol{m}^H\boldsymbol{\Sigma}^{-1}(\boldsymbol{\kappa})\boldsymbol{m}\right] = -\left|\boldsymbol{A}^H\boldsymbol{\Sigma}^{-1}(\boldsymbol{\kappa})\boldsymbol{m}\right|^2 \odot e^{\boldsymbol{\kappa}}. \tag{14}$$

Unfortunately, the dimensions of the covariance matrix $\boldsymbol{\Sigma}(\boldsymbol{\kappa})$ prevent the direct computation of its inverse due to the prohibitive computational and memory cost it would require. Therefore (14) and (13) must be approximated.

Indeed, the computation of the gradient of the log-det term (13) is the most critical aspect of the proposed method. For example, the loss function L is also involved in the inference

of Gaussian processes and the gradient of the logarithm of determinant represent the main hindrance to their application to large datasets [34], [35]. To estimate it, we first observe that

$$\operatorname{diag}\left(\boldsymbol{A}^H\boldsymbol{\Sigma}^{-1}(\boldsymbol{\kappa})\boldsymbol{A}\right) = \mathop{\mathbb{E}}_{\hat{\boldsymbol{m}}\sim\mathcal{N}(\boldsymbol{0},\boldsymbol{\Sigma}(\boldsymbol{\kappa}))}\left|\boldsymbol{A}^H\boldsymbol{\Sigma}^{-1}(\boldsymbol{\kappa})\hat{\boldsymbol{m}}\right|^2, \quad (15)$$

following from operator $\operatorname{diag}(\cdot)$ applied to the identity

$$\begin{aligned}\mathop{\mathbb{E}}_{\hat{\boldsymbol{m}}\sim\mathcal{N}(\boldsymbol{0},\boldsymbol{\Sigma}(\boldsymbol{\kappa}))}&\left\{\left[\boldsymbol{A}^H\boldsymbol{\Sigma}^{-1}(\boldsymbol{\kappa})\hat{\boldsymbol{m}}\right]\left[\boldsymbol{A}^H\boldsymbol{\Sigma}^{-1}(\boldsymbol{\kappa})\hat{\boldsymbol{m}}\right]^H\right\}\\ &= \left[\boldsymbol{A}^H\boldsymbol{\Sigma}^{-1}(\boldsymbol{\kappa})\right]\boldsymbol{\Sigma}(\boldsymbol{\kappa})\left[\boldsymbol{A}^H\boldsymbol{\Sigma}^{-1}(\boldsymbol{\kappa})\right]^H\\ &= \boldsymbol{A}^H\boldsymbol{\Sigma}^{-1}(\boldsymbol{\kappa})\boldsymbol{A}. \end{aligned} \quad (16)$$

Therefore, we propose the subsequent Monte-Carlo estimation

$$\operatorname{diag}\left(\boldsymbol{A}^H\boldsymbol{\Sigma}^{-1}(\boldsymbol{\kappa})\boldsymbol{A}\right) \approx \frac{1}{N^{\text{est}}}\sum_{i=1}^{N^{\text{est}}}\left|\boldsymbol{A}^H\boldsymbol{\Sigma}^{-1}(\boldsymbol{\kappa})\hat{\boldsymbol{m}}_i\right|^2, \quad (17)$$

with

$$\hat{\boldsymbol{m}}_i = \boldsymbol{A}\boldsymbol{\gamma}_i + \boldsymbol{\nu}_i, \boldsymbol{\gamma}_i \sim \mathcal{N}\left(\boldsymbol{0}, \operatorname{Diag}\left(e^{\boldsymbol{\kappa}}\right)\right), \quad \boldsymbol{\nu}_i \sim \mathcal{N}\left(\boldsymbol{0}, \sigma_\nu^2\boldsymbol{I}\right), \ \forall i = 1,\dots,N^{\text{est}}, \quad (18)$$

and where N^{est} is the number of sample vectors. In particular, (17) shares strong similarities with (14). Thus, the computation of $\nabla_{\boldsymbol{\kappa}}L$ requires the solutions to $N^{\text{est}}+1$ systems of equations, one to compute (14) and N^{est} to estimate (13) according to (17).

However, instead of solving directly the systems in (17) and (14), we first rewrite them as

$$\boldsymbol{A}^H\boldsymbol{\Sigma}^{-1}(\boldsymbol{\kappa})\boldsymbol{m} = \operatorname{Diag}\left(e^{-\boldsymbol{\kappa}}\right)\left[\tfrac{1}{\sigma_\nu^2}\boldsymbol{A}^H\boldsymbol{A} + \operatorname{Diag}\left(e^{-\boldsymbol{\kappa}}\right)\right]^{-1}\tfrac{1}{\sigma_\nu^2}\boldsymbol{A}^H\boldsymbol{m}. \quad (19)$$

Then, we propose to solve the series of $N^{\text{est}}+1$ system with a block conjugate gradient (BCG) algorithm [36]. The BCG allows us to benefit from a faster convergence time than, for example, applying conjugate gradient independently for each of the $N^{\text{est}}+1$ systems. The BCG is ultimately stopped when the relative norm of the residuals becomes inferior to a stopping criterion ϵ^{BCG}. We choose to compute the solutions to the systems in the spatial domain (19)—in opposition to the signal domain in (17) and (14)—since we can find an effective diagonal preconditioner. Finally, we present our estimation of the diagonal of the Fisher information matrix in Appendix A, whereas Appendix B describes the algorithm used to compute $\operatorname{Prox}_{\eta\mu R}^{\operatorname{Diag}(\boldsymbol{f})}(\hat{\boldsymbol{\kappa}})$ (12).

The complete proximal natural gradient descent algorithm is detailed in Algorithm 1. We must emphasize that the objective function (7) is in general not convex. Thus, we cannot guarantee that the proposed scheme will converge to a global minimum. However, we noticed empirically that the algorithm converges satisfactorily if the initial estimation of $\boldsymbol{\kappa}$ is superior to the solution of (7). Therefore, we propose to set the initial estimate of $\boldsymbol{\kappa}$ as a the upper bound $\kappa^{\max}$ on $\boldsymbol{\kappa}$ imposed by the regularizer R (9). Our choice of a natural gradient approach—compared to an accelerated proximal gradient descent such as the fast iterative shrinkage-thresholding algorithm (FISTA) [37]—is motivated by its faster convergence time and the iteration stability it allows. In contrast, our experiments showed a lack of stability of FISTA in that case, which lead the nearly systematic divergence of the algorithm.

Algorithm 1: Proximal Natural Gradient Descent.

input: $\boldsymbol{m} \in \mathbb{C}^{N^{\text{el}}N^{\text{t}}}$ the CRF measurement vector, a maximum number of iterations N^{iter}
$\boldsymbol{\kappa}_0 \leftarrow \kappa^{\max}$
$k \leftarrow 0$
do
  Generate $\hat{\boldsymbol{m}}_i \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{\Sigma}(\boldsymbol{\kappa}_k)),\ \forall i = 1,\dots,N^{\text{est}}$
  $\boldsymbol{M} \leftarrow [\hat{\boldsymbol{m}}_1,\dots,\hat{\boldsymbol{m}}_{N^{\text{est}}}, \boldsymbol{m}]$
  $\boldsymbol{P} \leftarrow \frac{1}{\sigma_\nu^2}\boldsymbol{A}^H\boldsymbol{M}$
  Solve $\left[\frac{1}{\sigma_\nu^2}\boldsymbol{A}^H\boldsymbol{A} + \operatorname{Diag}(e^{-\boldsymbol{\kappa}})\right]\boldsymbol{Q} = \boldsymbol{P}$ using BCG
  $\boldsymbol{R} \leftarrow \operatorname{Diag}(e^{-\boldsymbol{\kappa}})\boldsymbol{Q}$
  $\boldsymbol{g} \leftarrow \left[\frac{1}{N^{\text{est}}}\sum_{i=1}^{N^{\text{est}}}\left|\boldsymbol{r}_i\right|^2 - \left|\boldsymbol{r}_{N^{\text{est}}+1}\right|^2\right] \odot e^{\boldsymbol{\kappa}}$
  $\boldsymbol{f} \leftarrow \left[\frac{1}{N^{\text{est}}}\sum_{i=1}^{N^{\text{est}}}\left|\boldsymbol{r}_i\right|^4 - \left[\frac{1}{N^{\text{est}}}\sum_{i=1}^{N^{\text{est}}}\left|\boldsymbol{r}_i\right|^2\right]^2\right] \odot e^{2\boldsymbol{\kappa}} + f^{\text{reg}}$
  $\hat{\boldsymbol{\kappa}} \leftarrow \boldsymbol{\kappa}_k - \eta\operatorname{Diag}^{-1}(\boldsymbol{f})\boldsymbol{g}$
  $\boldsymbol{\kappa}_{k+1} \leftarrow \operatorname{Prox}_{\eta\mu R}^{\operatorname{Diag}(\boldsymbol{f})}(\hat{\boldsymbol{\kappa}})$
  $k \leftarrow k+1$
while $k < N^{\text{iter}}$
return $\boldsymbol{\kappa}_k$

E. Reflectivity Estimation

To conclude this section, we describe herein how to estimate the reflectivity map itself—interpreted as a deconvolved image—from the measured signals, factoring in the echogenicity map obtained by solving (7). We detail the definition of the optimization problem used to do so, the statistical model it is based on, as well as the optimization algorithm proposed to solve it.

Once the log-scale local reflectivity variance $\boldsymbol{\kappa}$ is estimated, we can use it to define a prior on the discrete reflectivity map $\boldsymbol{\gamma}$. Following prior work on image deconvolution and inverse-model-based reflectivity estimation, we propose to model an element of $\boldsymbol{\gamma}$ as a single realization of a centered GGD [13], [19]. Since we consider a complex reflectivity map, we additionally impose the GGD to be centrally-symmetric. Therefore, the PDF of an element k of $\boldsymbol{\gamma}$ is given by

$$\mathbb{P}(\gamma_k) = \frac{p}{2\pi c_k^2\Gamma\left(\frac{2}{p}\right)}e^{-\left|\frac{\gamma_k}{c_k}\right|^p}. \quad (20)$$

Distribution (20) is parameterized by a scale parameter $c_k > 0$—supposed different for each pixel—and a shape parameter p—supposed constant within the medium and known a priori.

TABLE I
TRANSDUCER SPECIFICATIONS

Transducer	GE9L-D
Number of elements N^{el}	192
Pitch	230 µm
Aperture	43.93 mm
Element width w^{el}	0.9·230 µm[a]
Element height	6 mm
Elevation focus	28 mm
Center frequency	5.3 MHz
Fractional bandwidth	75%
Excitation cycles	1
Excitation frequency f_0	5.2 MHz
Sampling frequency f^{s}	20.8 MHz

[a]No official data available, guessed value according to [39]

Moreover, the variance of (20) is given by

$$\mathbb{E}\left[|\gamma_k|^2\right] = e^{\kappa_k} = c_k^2 \left[\frac{\Gamma\left(\frac{4}{p}\right)}{\Gamma\left(\frac{2}{p}\right)}\right], \tag{21}$$

with $\boldsymbol{\kappa}$ the echogenicity map solution of (7), and Γ the gamma function [38].

We assumed in Section II-B that the measurement vector $\boldsymbol{m}$ is the result of operator $\boldsymbol{A}$ acting on the reflectivity map $\boldsymbol{\gamma}$, plus a realization of white Gaussian noise with σ_ν^2 variance. The PDF of $\boldsymbol{m}$ given $\boldsymbol{\gamma}$ is thus expressed as

$$\mathbb{P}(\boldsymbol{m}|\boldsymbol{\gamma}) = \frac{1}{\pi\left[\sigma_\nu^2\right]^{N^{\text{x}}N^{\text{z}}}} e^{-\frac{1}{\sigma_\nu^2}\|\boldsymbol{m}-\boldsymbol{A}\boldsymbol{\gamma}\|^2}, \tag{22}$$

where $N^{\text{x}}N^{\text{z}}$ is the dimension of $\boldsymbol{\gamma}$. Taking into account (20), (21), and (22), we define a MAP reflectivity estimate as the solution to the following minimization problem

$$\min_{\boldsymbol{\gamma}} \frac{1}{\sigma_\nu^2}\left\|\boldsymbol{A}\boldsymbol{\gamma}-\boldsymbol{m}\right\|^2 + \left[\frac{\Gamma\left(\frac{4}{p}\right)}{\Gamma\left(\frac{2}{p}\right)}\right]^{\frac{p}{2}} \left\|\text{Diag}\left(e^{-\frac{1}{2}\boldsymbol{\kappa}}\right)\boldsymbol{\gamma}\right\|_p^p. \tag{23}$$

In this article, we investigate 4 different values of the scale parameters; $p = 1$, $p = \frac{4}{3}$, $p = \frac{3}{2}$, and $p = 2$. Values $p < 2$ are chosen to select distributions (20) presenting sparse behaviours. The Gaussian hypothesis $p = 2$ is tested for comparison purpose. In addition, we ensure that (23) remains convex by selecting $p \geq 1$, and select values of p such that the proximal operator of the regularizer in (23) admits a closed-form expression.

In particular, the proposed inverse problem automatically increases the weight of the prior in medium areas posessing a low echogenicity—with respect to a uniform prior—and reduces its importance in areas with high echogenicity. Since echo signals originating from hypoechoic zones are weak, an accurate prior is of upmost importance therein. Thus, problem (23) is especially designed to improve reflectivity reconstruction in low-echogenicity areas.

To solve problem (23), we implement FISTA according to Algorithm 2. In practice, the Lipschitz constant L is set using power iterations, and the computation of $\boldsymbol{\gamma}_{k+1}$ is detailed in Appendix C.

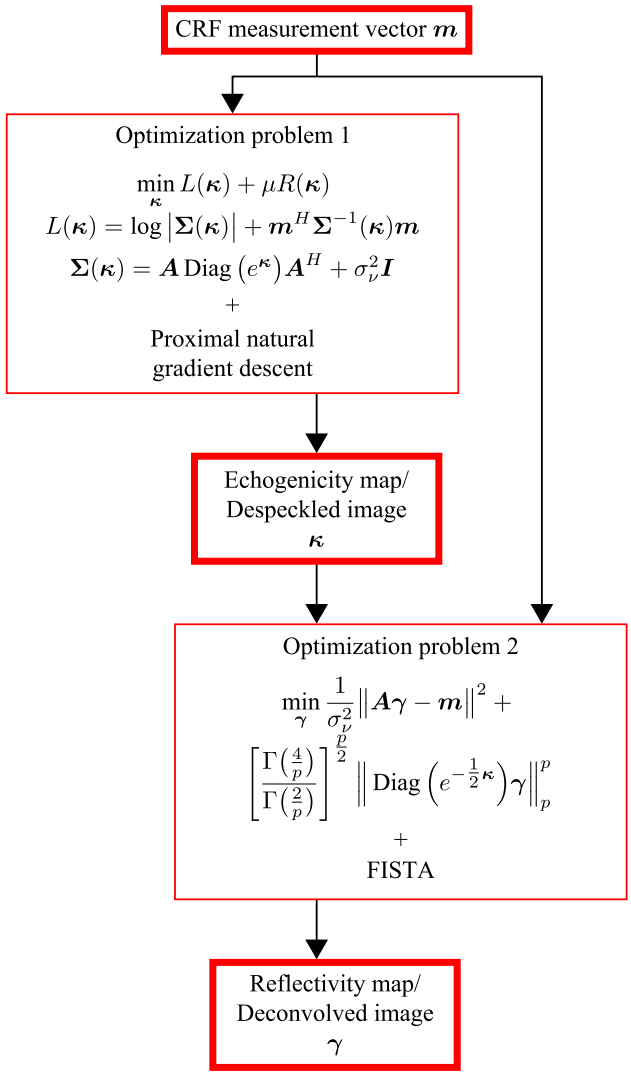

Fig. 1. Flow chart summarizing the proposed approach, where $\boldsymbol{A}$ is the discrete measurement model, σ_ν^2 is the signal noise variance, R is a bounded isotropic TV norm regularizer, μ is the regularization parameter, and p is the GGD shape parameter.

Ultimately, we display the absolute value of the reflectivity $\boldsymbol{\gamma}$ expressed in decibels, analogous to the depiction of B-mode images reconstructed with DAS. In particular, this fact justifies the use of a complex-valued vector $\boldsymbol{\gamma}$ instead of a real vector. Indeed, depicting a real reflectivity map in this manner requires applying envelope detection. However, doing so leads to artifacts since the bandwidth of the axial and lateral frequencies of the estimated reflectivity map $\boldsymbol{\gamma}$ is larger than the ones of images reconstructed using DAS.

A visual summary of the proposed approach for both echogenicity and reflectivity estimation is provided in Fig. 1.

III. EXPERIMENTS

A. Parameters

For all the experiments performed in this article, we used a GE9L-D ultrasound transducer (GE Healthcare, Chicago, Illinois, USA) linked to Verasonics Vantage 256 Ultrasound system (Verasonics, Kirkland, WA, USA). The specifications of

Algorithm 2: FISTA to Evalutate (23).

input: $\boldsymbol{m} \in \mathbb{C}^{N^{\text{el}}N^{\text{t}}}$ the CRF measurement vector,
$\boldsymbol{\kappa} \in \mathbb{R}^{N^{\text{x}}N^{\text{z}}}$ the solution of (7), a stopping criterion
$\epsilon^{\text{FISTA}} > 0$
$\boldsymbol{\gamma}_0 = \mathbf{0}$
$\boldsymbol{z}_0 = \mathbf{0}$
$t_0 = 1$
$k \leftarrow 0$
$L \leftarrow \|\frac{1}{\sigma_\nu^2}\boldsymbol{A}^H\boldsymbol{A}\|$
do
$$\hat{\boldsymbol{\gamma}} \leftarrow \left[\boldsymbol{I} - \tfrac{1}{L}\tfrac{1}{\sigma_\nu^2}\boldsymbol{A}^H\boldsymbol{A}\right]\boldsymbol{z}_k$$
$$\boldsymbol{\gamma}_{k+1} \leftarrow \arg\min_{\boldsymbol{\gamma}} \left[\frac{\Gamma(\frac{4}{p})}{\Gamma(\frac{2}{p})}\right]^{\frac{p}{2}} \left\|\text{Diag}\left(e^{-\frac{1}{2}\boldsymbol{\kappa}}\right)\boldsymbol{\gamma}\right\|_p^p + \frac{L}{2}\|\boldsymbol{\gamma} - \hat{\boldsymbol{\gamma}}\|^2$$
$$t_{k+1} \leftarrow \frac{1}{2}[1 + \sqrt{1 + 4t_k^2}]$$
$$\boldsymbol{z}_{k+1} \leftarrow \boldsymbol{\gamma}_k + \tfrac{t_k - 1}{t_{k+1}}[\boldsymbol{\gamma}_{k+1} - \boldsymbol{\gamma}_k]$$
$$k \leftarrow k + 1$$
while $\frac{1}{\sqrt{N^{\text{x}}N^{\text{z}}}}\|\boldsymbol{\gamma}_k - \boldsymbol{\gamma}_{k-1}\| > \epsilon^{\text{FISTA}}$
return $\boldsymbol{\gamma}_k$

the transducer are summarized in Table I. All the experiments required either the emission of a single unsteered PW, or the emission of a series of 115 steered PWs whose steering angles are uniformly spaced between $-28.5°$ and $28.5°$. The 115 PWs data are used solely for comparison.

In particular, almost every parameter of the measurement model $\boldsymbol{A}$ are set according to the parameters provided in Table I, whereas we assume a medium SoS c_0 of $1540\,\text{m/s}$. In addition, we suppose that v_{pe} is a Gaussian pulse whose central frequency is given in Table I, and whose the standard deviation $\sigma_{\text{pe}} = 0.16\,\mu\text{s}$ is determined to best predict the spectral power density of signals acquired on the phantom figuring uniform speckle—CIRS model 054GS general-purpose in-vitro phantom (Sun Nuclear, Melbourne, FL, USA). Similary, function o^{Rx} (1) is determined to best fit the local variance of 100 signals $\boldsymbol{m}$ acquired on the same uniform phantom.

With respect to the parameter of the inverse problem (7), we set $\sigma_\nu^2 = 10^{-5}$, $\mu = 1.910^{-5}$, and $\kappa^{\text{max}} = 34$ due to their empirical success in-vivo. In Algorithm 1, we set $N^{\text{est}} = 20$, $\eta = 8 \cdot 10^{-3}$, $f^{\text{reg}} = 10^{-5}$, $\epsilon^{\text{BCG}} = 4 \cdot 10^{-7}$, $N^{\text{iter}} = 120$, and $\epsilon^{\text{Prox}} = 10^{-7}$ as a trade-off between convergence time and the error with respect to the solution of (7). Regarding the optimization of (23) and (24) with Algorithm 2, we set the stopping tolerance to $\epsilon^{\text{FISTA}} = 10^{-7}$.

B. Simulated Data

As an initial test, we apply the proposed method to simulated data. To generate data, we rely on the linear acoustics simulator Field II [40]. We suppose a medium spanning the width of the transducer—with a $7\,\text{mm}$ margin on both sides of the geometrical model of the probe—, a depth of $60\,\text{mm}$, and a height of $5.8\,\text{mm}$.

We simulate the presence of 10 scatterers per resolution cells—unless mentioned otherwise—whose positions are drawn from a uniform distribution and whose intensities are drawn from Gaussian distribution. The variance of the Gaussian varies spatially according to a reference echogenicity map.

Simulated data allows us to test the proposed method in the absence of error caused by the most significant inaccuracies of the measurement model—namely, multiple scattering, frequency-dependent attenuation or spatial variations of SoS and attenuation. Moreover, we can compare explicitly the results of the proposed method with the ground truth echogenicity. The ground truth map will be presented in Section IV-A, alongside the results of the proposed method.

C. Phantom and In-Vivo Data

We also test the proposed method with data acquired on a CIRS model 054GS phantom. The imaged medium comprises a series of three far-field inclusions; One anechoic and two hypoechoic, 6 dB and 3 dB below the background echogenicitiy, respectively.

In addition, we apply the proposed method to in-vivo data. We consider two images of the carotid artery of a healthy volunteer (a 28 years old male), one figuring a longitudinal slice of the artery and one figuring an axial slice. Data acquisition was conducted in accordance with the Declaration of Helsinki. Moreover, the ultrasound sequence has been approved by the Cantonal Commission on Ethics in Human Research of the Canton of Vaud, Switzerland.

D. Reference Image Despeckling Method

We compare the proposed echogenicity recovery method to DAS followed by a despeckling algorithm. In particular, we selected a non-local means approach due to its use in commercial ultrasound systems [8], [21]. To correct for the spatial variations of echogenicity caused by beamforming, we normalize the image according to [39]. We first define a normalization mask as the average of 100 B-mode images acquired on a zone of the CIRS model 054GS phantom presenting a uniform echogenicity. We then subtract the normalization mask to the B-mode image of interest prior to despeckling. This procedure ensures that a medium with uniform echogenicity will indeed appear uniform in the despeckled image. Finally, we set the beamforming SoS c_0 to $1540\,\text{m/s}$, namely the same medium SoS assumed in measurement model $\boldsymbol{A}$.

A ground-truth echogenicity map is available for experiments with simulated data. However, this is not the case with in-vivo or in-vitro experiments. There, we also reconstruct high-quality images using 115 PW acquisitions to act as surrogates for ground-truth echogenicity maps.

E. Reference Reflectivity Estimation Method

To reconstruct a deconvolved image from CRF measurement vector $\boldsymbol{m}$, we propose to follow the procedure introduced in [13].

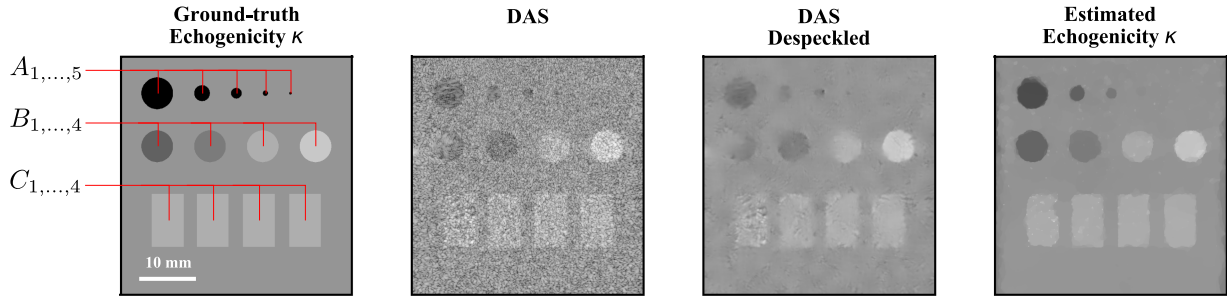

Fig. 2. From left to right: The ground-truth echogenicity of the numerical phantom used to generate simulated data, the result of delay-and-sum (DAS) applied to the data, a despeckled version of the image reconstructed with DAS, and the result of the proposed method. The zones used for the different contrast experiments are highlighted on the ground-truth echogenicity map.

It involves an inverse problem similar to (23), albeit assuming a constant echogenicity κ^{ref} in the medium. The inverse problem in this case is written as

$$\arg\min_{\boldsymbol{\gamma}} \frac{1}{\sigma_\nu^2}\left\|\boldsymbol{A}\boldsymbol{\gamma}-\boldsymbol{m}\right\|^2 + e^{-\frac{p}{2}\kappa^{\text{ref}}}\left[\frac{\Gamma\left(\frac{4}{p}\right)}{\Gamma\left(\frac{2}{p}\right)}\right]^{\frac{p}{2}}\left\|\boldsymbol{\gamma}\right\|_p^p. \quad (24)$$

As demonstrated in [13], the use of a physical model of data acquisition $\boldsymbol{A}$ alleviates the drawbacks associated with deconvolution methods assuming a constant PSF. These drawbacks include the lack of accuracy resulting from the variation of the effective PSF with position, the necessity to estimate the PSF itself, and the presence of non-local diffraction artifacts in the image. In fact, the latter represent the main hindrance to the use of deconvolution techniques in the present case of single PW imaging. Thus, the reference method based on a uniform prior (24) represents the current state-of-the-art in terms of the estimation of a deconvolved image from a low quantity of PW data.

Regarding the parameters of (24), we propose to use the same noise variance than the proposed method ($\sigma_\nu^2 = 10^{-5}$). In the phantom case, we set κ^{ref} to the median value of $\boldsymbol{\kappa}$ obtained by solving (7) in the background. In the in-vivo case, we set κ^{ref} according to the median value of $\boldsymbol{\kappa}$ in the whole medium.

IV. Results

A. Echogenicity Results With Simulated Data

We present herein the results of the proposed despeckled image recovery method. Regarding first the experiments with simulated data, we depict in Fig. 2 the results using DAS and the proposed method, as well as the numerical phantom used for the simulation. Three types of experiments are performed in the phantom. First, the top of the phantom comprises a series of five anechoic inclusions—designated as A_1 to A_5—with decreasing radii—3 mm, 1.5 mm, 1 mm, 0.5 mm, and 0.25 mm, respectively. Two hypoechoic inclusions—12 and 6 dB below the background echogenicity—are present below the anechoic inclusions. They have a radius of 3 mm and are denoted as B_1 and B_2. In addition, there are two hyperechoic inclusions—B_3 and B_4—with the same radius and whose echogenicities are 6 and 12 dB above the background one. Finally, four rectangular areas—indexed from C_1 to C_4—are present in the bottom of the phantom, with an echogenicity 6 dB above the background one. Until now, we assumed 10 scatterers per resolution cell during simulation. Experiments C_1 to C_4 are designed to test the behaviour of the proposed method when the number of scatterers is reduced, namely when an image of the medium does not present a fully resolved speckle pattern. In particular, we reduced the number of scatterers per cell by factors of 27, 9, and 3 with respect to the standard value of 10 in experiments C_1 to C_3. Experiment C_4 assumes 10 number per resolution cells for comparison purpose.

Table II presents the contrast ratios (CRs) in every cases with respect to the echogenicity of the background, for both the proposed method and DAS followed by despeckling. The CRs for each inclusion are given by

$$CR[\text{dB}] = \frac{1}{|\mathcal{K}_{\text{in}}|}\sum_{k\in\mathcal{K}_{\text{in}}}\boldsymbol{\kappa}_k^{\text{dB}} - \frac{1}{|\mathcal{K}_{\text{out}}|}\sum_{k\in\mathcal{K}_{\text{out}}}\boldsymbol{\kappa}_k^{\text{dB}}, \quad (25)$$

where $\boldsymbol{\kappa}^{\text{dB}}$ denotes either a desepckled DAS B-mode image or

$$\boldsymbol{\kappa}^{\text{dB}} = 20\log_{10}(e)\,\boldsymbol{\kappa}, \quad (26)$$

when the proposed method is tested. In (25), $\mathcal{K}_{\text{in}}$ is the set of pixels inside the inclusion and $\mathcal{K}_{\text{out}}$ is the set of pixels located at at most 3 mm from the outer boundary of the inclusion.

Regarding first the anechoic inclusions, we can see that inclusions A_1 to A_3 are visible with DAS—before and after despeckling—and the proposed approach. Inclusion A_4 is scarcely discernible with DAS, but cannot be distinguished with the proposed method, whereas inclusion A_5 is not visible for both methods. Quantitatively, the proposed approach allows for a significant contrast improvement with respect to DAS for cases A_1 to A_3. Moreover, the contrast increases with the radius of the inclusion.

As it can be observed in Fig. 2, inclusion B_1 is hidden behind a grating lobe artifact when DAS is used. This results in a poor contrast, both visually and quantitatively. The proposed method, however, successfully removes the grating lobe and contrast is improved considerably. Regarding inclusions B_2 to B_4, the proposed and reference methods achieve similar results both quantitatively and quantitatively, even if the boundaries of the inclusions are better defined with the proposed approach.

TABLE II
CONTRAST RATIOS WITH SIMULATED DATA

	A					B				C			
	1	2	3	4	5	1	2	3	4	1	2	3	4
Ground truth	-∞	-∞	-∞	-∞	-∞	-12	-6	6	12	6	6	6	6
DAS - despeckled	-6.15	-2.99	-3.96	-2.71	0.2	-1.89	-4.76	4.94	11.06	3.88	4.6	5.25	5.33
Proposed	-15.17	-10.04	-5.45	-1.02	-0.18	-10.45	-4.87	4.86	10.7	4.98	5.02	5.3	5.49

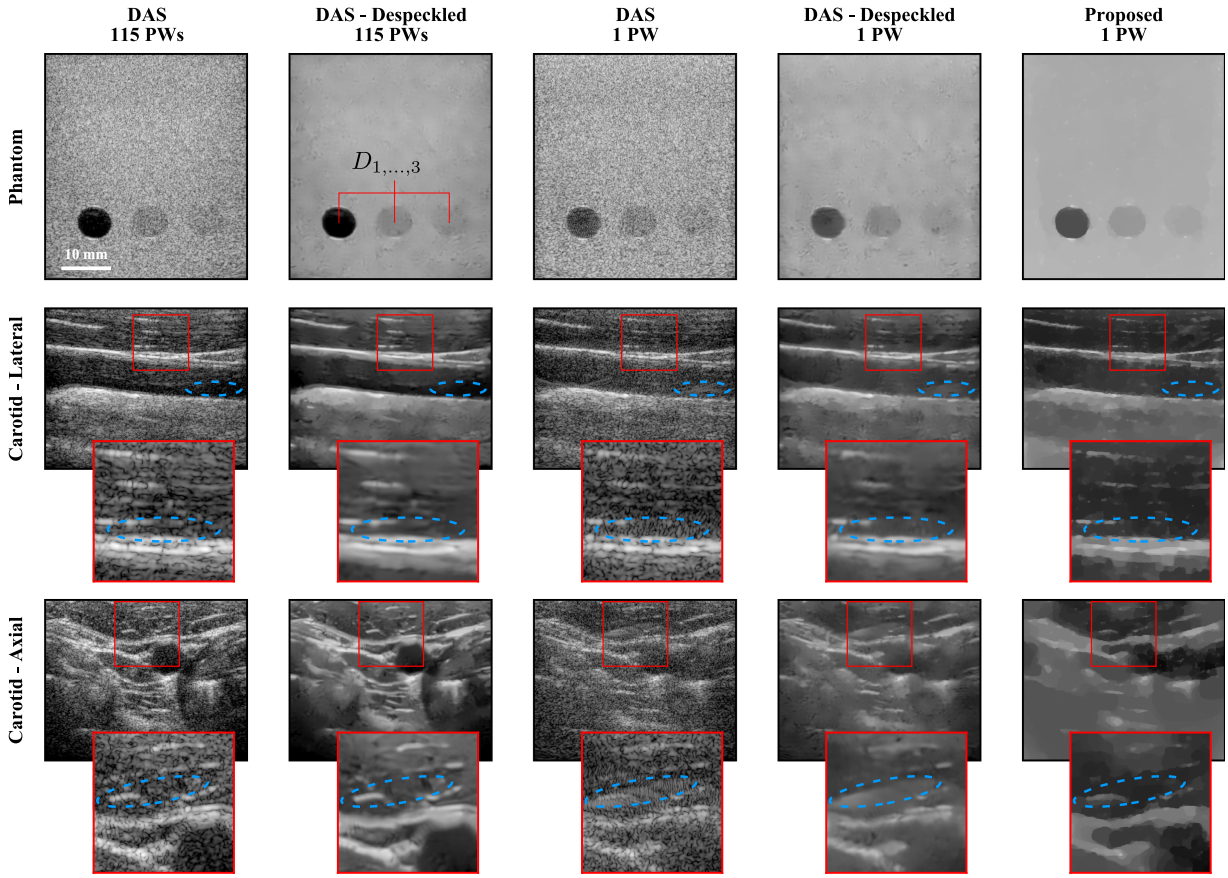

Fig. 3. Despeckled images of the phantom (top row) and in-vivo carotid artery (bottom rows). In each case, we depict the result of delay-and-sum (DAS) before and after despeckling, using both 115 plane waves (PWs) and single PW data, along with the result of the proposed approach. For the in-vivo results, we highlight two areas of interest.

Finally, we can see the contrast ratios of the four areas C_1 to C_4 diminish with both techniques when the number of scatterers per resolution cell is reduced. However, this trend is significantly alleviated by the proposed method with respect to DAS. Visually, we notice that C_1 presents echogenicity variations with DAS. Uniformity is better enforced by the proposed approach, even though some punctual structure can be observed within the area.

TABLE III
CONTRAST RATIOS WITH PHANTOM DATA

	D		
	1	2	3
Ground truth	-∞	-6	-3
DAS 115 PWs - Despeckled	-36.25	-6.53	-3.52
DAS 1 PW - Despeckled	-19.48	-6.5	-3.62
Proposed	-22.32	-5.36	-2.62

B. Echogenicity Results With In-Vitro Phantom Data

The top row of Fig. 3 depicts the results of DAS applied to in-vitro phantom data—with and without despeckling and using both 115 PWs and 1 PW—alongside the results of the proposed method. In addition, Table III provides the CRs achieved by the different method for the three inclusion present in the phantom—denoted as D_1, D_2, and D_3. The latter are computed in accordance with (25).

Regarding first the anechoic inclusion D_1, we can observe in Fig. 3 a qualitative contrast improvement with the proposed method compared to DAS with single PW data. The improvement is confirmed quantitatively since we report from Table III a contrast gain of almost 3 dB. Turning now to the hypoechoic inclusions D_2 and D_3, the CRs achieved by the proposed method are lower than both the target value and the values obtained by

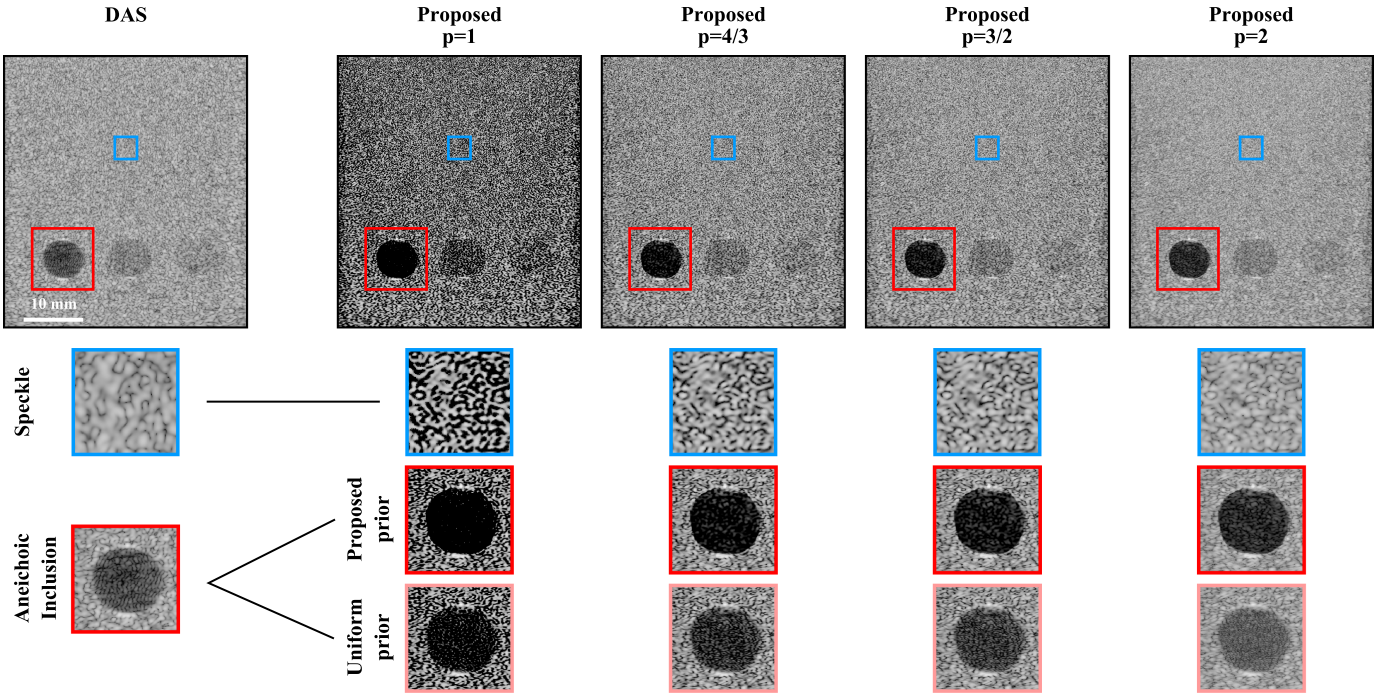

Fig. 4. Estimation of the reflectivity in the phantom. On the top, the results of delay-and-sum and the proposed method with the four ℓ_p norms are depicted. We present in the bottom of the figure areas of interest figuring a uniform speckle and the anechoic inclusion. In the latter case, the result of the reference method with a uniform prior is also displayed.

DAS, irrespectively of the number of PWs. However, we can notice that the proposed method manage to consistently recover a sharper inclusion boundary, especially in cases D_2 and D_3.

C. Echogenicity Results With In-Vivo Data

The second and third rows of Fig. 3 depict lateral and axial images of the carotid artery. In addition, we highlight for each image an area of interest.

Focusing first on the lateral image, we observe that the proposed method successfully removes diffraction artifacts. The improvement is especially noticeable inside the carotid artery, for example within the blue ellipse on the right half of the image. Significant artifacts are present in the single PW images—before and after despeckling—and they are absent in the result of the proposed method. The image enhancement is also perceptible just above the artery wall in the area of interest. Side lobe artifacts occurs in the single PW image. As it can be observe within the blue ellipse, artifacts are removed by the proposed method, leading to a result closer to the 115 PWs despeckled image. Furthermore, we can see on the area of interest that the proposed method maintains fine structures within the superficial layer of the imaged medium.

Regarding the axial image, we see in the area of interest a large artifact highlighted in blue appearing in the single PW images above the outer artery wall. As we can see on the left side of the blue ellipse, the artifact conceals in the single PW image an horizontal structure that is visible in the 115 PWs image. Once again, the proposed method removes the artifact to recover an anatomically correct result on the upper half of the image. Furthermore, the contrast inside the artery is improved compared to DAS. Unfortunately, no significant difference between the proposed method and DAS can be observed in the lower half of the image. We can attribute this lack of improvement to the physical phenomenons that are disregarded by the model $\boldsymbol{A}$—foremost, multiple scattering and SoS variations—since we expect their effect to be more severe in the far-field.

D. Reflectivity Results With In-Vitro Phantom Data

Fig. 4 presents the result of the proposed reflectivity estimation method applied to phantom data for $p = 1$, $p = \frac{4}{3}$, $p = \frac{3}{2}$, and $p = 2$, after log compression. On the left side of the figure, we depict for comparison a B-mode image obtained with DAS.

Two regions of interest are highlighted in the images. The top one figures uniform speckle, whereas the bottom one is a close-up of the anechoic inclusion. In addition to the result of the proposed method, we also display the anechoic inclusion as reconstructed by the reference reflectivity estimation method. Since the echogenicity κ^{ref} assumed by the reference method is identical to the background echogenicity recovered by the proposed echogenicity estimation technique, the reflectivities estimated by the two methods are otherwise very similar outside of the inclusions.

Focusing first on speckle, we can clearly observe the resolution improvement allowed by the inverse-problem approach. To quantify this phenomenon, the upper half of Table IV presents the axial and lateral full width at half maximum (FWHM) of the reflectivity autocorrelation within the top region of interest. We can notice that the axial FWHM is significantly improved

TABLE IV
SPECKLE RESOLUTION AND ANECHOIC CONTRAST IN THE PHANTOM

	DAS	$p=1$	$p=\frac{4}{3}$	$p=\frac{3}{2}$	$p=2$
Axial FWHM [µm]	429	178	206	211	214
Lateral FWHM [µm]	333	233	266	273	279
TCR - Reference [dB]	-18.67	-25.1	-21.05	-19.34	-14.01
TCR - Proposed [dB]		-40.83	-30.3	-28.09	-24.82

TABLE V
CONTRAST AND RESOLUTION IN VIVO

	DAS		$p=1$	$p=\frac{4}{3}$	$p=\frac{3}{2}$	$p=2$
Axial FWHM [µm]	381	Uniform	257	268	272	278
		Proposed	253	267	274	281
TCR [dB]	-11.18	Uniform	-11.26	-10.28	-9.68	-8.94
		Proposed	-12.97	-12.73	-12.3	-10.66

compared to DAS, whereas the improvement in FWHM along the lateral dimension is more marginal. The influence of p on speckle resolution is otherwise limited. We can observe a small FWHM improvement with $p=1$ with respect to the three other cases. Yet, the speckled pattern does not seem to present in Fig. 4 a finer grain size than with $p>1$. Moreover, the speckle distribution is significantly affected. In contrast, the speckle pattern recovered with $p=2$ appears like a less coarse version of the one reconstructed with DAS.

Turning now to the anechoic inclusion, we can appreciate in Fig. 4 the large contrast improvement permitted by the proposed approach compared to a uniform prior. To provide a quantitative estimation of contrast, we compute the tissue-to-clutter ratio (TCR) between the interior of the inclusion and the adjacent background [41]. Formally, the TCR is defined as

$$TCR[\text{dB}] = 20\log_{10}\left[\frac{1}{|\mathcal{K}_{\text{in}}|}\sum_{k\in\mathcal{K}_{\text{in}}}|\gamma_k|\right] - 20\log_{10}\left[\frac{1}{|\mathcal{K}_{\text{out}}|}\sum_{k\in\mathcal{K}_{\text{out}}}|\gamma_k|\right], \quad (27)$$

where $\mathcal{K}_{\text{in}}$ denotes the set of pixels inside the inclusion, and $\mathcal{K}_{\text{out}}$ is the set of pixels located at at most $3\,\text{mm}$ from the inclusion boundary. The choice of TCR rather than CR—more regularly used in the literature and employed in Section IV to compare despeckled images—stems from the fact that some pixels of $\boldsymbol{\gamma}_k$ are null with $p<2$, which would lead to a CR of $-\infty$. We present TCR values in the lower half of Table IV, for both the proposed and uniform priors. Again, the large contrast improvement observed visually is confirmed quantitatively since the improvement in TCR allowed by the proposed prior varies from 8.75 dB ($p=\frac{3}{2}$) to 15.73 dB ($p=1$). As a final remark, we can emphasize that the reflectivity estimated by the proposed method using $p=1$ is null for 99.2% of the pixels within the inclusion, whereas this number falls to 79.8% with a uniform prior. The proposed approach is therefore able to recover an almost truly anechoic inclusion.

E. Reflectivity Results With In-Vivo Data

We depict in Fig. 5 the reflectivity maps achieved by the proposed method with the two in-vivo datasets. The DAS B-mode image and the reflectivity maps obtained when a uniform prior is assumed are depicted alongside the results of the proposed method. Moreover, we highlight in Fig. 5 two zones of interest in the reflectivity maps using red ellipses.

To quantify imaging quality, we estimate the resolution of the images by measuring the axial FWHM of one the muscular fibers present in the near field of the lateral carotid image—in a zone depicted in blue in Fig. 5. To estimate contrast, we compute the TCR (27) between the interior and the exterior of the carotid artery in the lateral carotid image. The sets $\mathcal{K}_{\text{in}}$ and $\mathcal{K}_{\text{out}}$ are depicted in yellow in Fig. 5 using plain and dashed lines respectively. We present in Table V the quantitative FWHM and contrast value using both uniform and proposed priors, alongside the values obtained with DAS.

First of all, we can notice qualitatively that the model inversion allows for a significant improvement of the resolution with regard to DAS. Interestingly, the improvement is more important in areas with low echogenicity compared to areas with high echogenicity. The resolution improvement is confirmed quantitatively, since the relative improvement in axial FWHM varies from 26% to 34% compared to DAS. Similarly to the results obtained in Section IV, there is a small FWHM improvement as p decreases, for both priors. However, we can notice that the proposed prior do not provide improvement in terms of axial FWHM with respect to the reference one.

We notice an improvement of the contrast of the carotid artery between the proposed and reference priors. Quantitatively, the TCR improvement varies from 1.71 dB ($p=1$) to 2.62 dB ($p=\frac{3}{2}$). Furthermore, we observe that the reflectivity estimates suffer from substantial side-lobe-like artifacts. Their severity generally increase with p, and they are especially visible in the superficial layers of both examples and within the carotid artery in the lateral image. However, the proposed regularizer is able to alleviate artifacts significantly with respect to a uniform prior. In Fig. 5, we highlight in red two zones where the reduction in artifact level is especially visible.

F. Dependency on Noise Variance

The proposed echogenicity and reflectivity estimation techniques both rely on a set of parameters. The most important of them is arguably the noise variance σ_ν^2. To test its influence, we propose to repeat the in-vitro phantom experiments with values of σ_ν^2 ranging from 10^{-6} to 10^{-3}. We depict in Fig. 6 four examples of despeckled images, whereas Fig. 7 presents the evolution of the metrics of Table IV computed on estimated reflectivity maps.

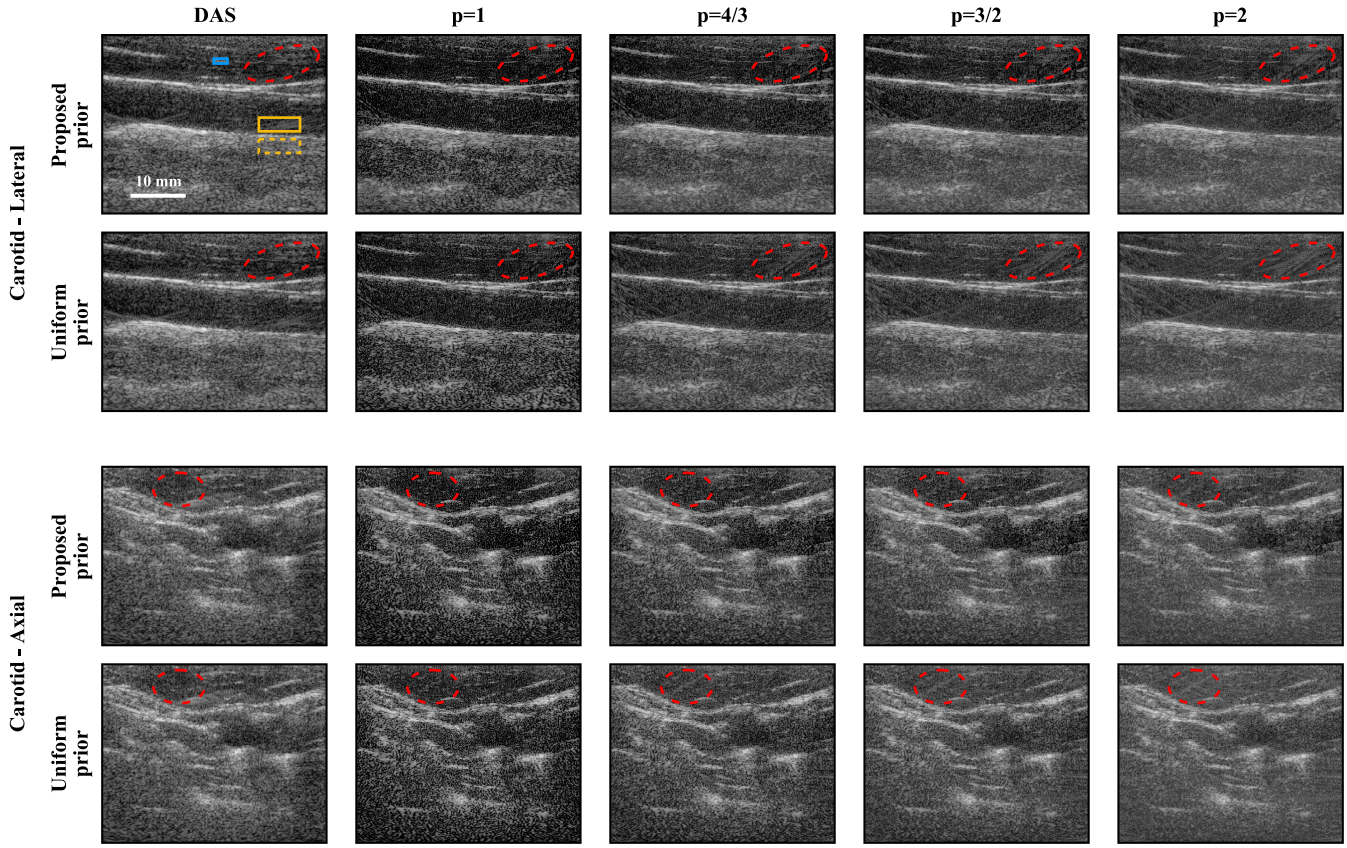

Fig. 5. In-vivo reflectivity maps estimated with single plane-wave data. We present the reflectivity map obtained by the proposed method for the two carotid images using the four ℓ_p norms. For comparison, we present the results of delay-and-sum (DAS) and the results of the reference method using a uniform prior.

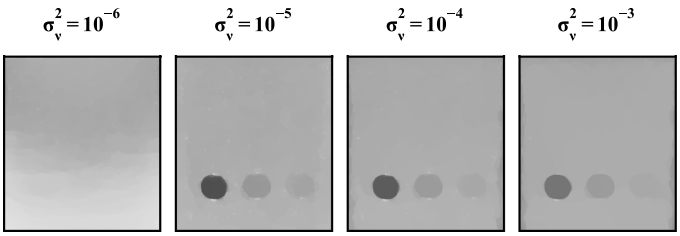

Fig. 6. Examples of phantom echogenicity maps estimated by the proposed method, for four values of the noise variance σ_ν^2.

We observe in Fig. 7 that contrast in the despeckled image generally improves as σ_ν^2 decreases. Anechoic contrast ratios of -13.03 dB, -19.34 dB, and -22.32 dB are obtained with $\sigma_\nu^2 = 10^{-3}$, $\sigma_\nu^2 = 10^{-4}$, and $\sigma_\nu^2 = 10^{-5}$, respectively. However, when the noise variance becomes too low—$\sigma_\nu^2 = 10^{-6}$—detrimental artifacts appear in the far field. We strongly suspect that these artifacts results from the inability of the method to discard multiple scattering and frequency-dependent attenuation modeling errors when σ_ν^2 is too low.

The axial speckle FWHM generally worsens as σ_ν^2 increases for $p > 1$, whereas the lateral FWHM remains approximately constant. In contrast, a reduction of speckle FWHM is observed with $p = 1$. Importantly, we do not construe this reduction as an improvement in resolution. We rather interpret this change as a decrease of the quality of autocorrelation FWHM as a resolution estimator, as reflectivity maps becomes increasingly sparse.

In opposition, the contrast of the deconvolved images generally decreases with increasing values of σ_ν^2. Also, the proposed prior constantly achieves a large contrast improvement with respect to the uniform one as long as $\sigma_\nu^2 \geq 10^{-5}$. Specifically, this threshold corresponds to the point where significant artifacts appear in the echogenicity map $\boldsymbol{\kappa}$.

V. Discussion

A. Echogenicity Estimation

As we already highlighted in Section IV-C, diffraction artifacts present in despeckled DAS images are almost entirely removed by the proposed echogenicity estimation method. We also saw that the proposed method suffers from some limitations, including a lack of contrast improvement with respect to DAS in the far field, or the oversight of some fine details. We use a TV norm to regularize the echogenicity map $\boldsymbol{\kappa}$. Therefore, the morphology of the estimated echogenicity maps differs from the structure of both the image despeckled with NLM and the results of despeckling algorithms implemented in commercial scanners. To solve this issue and alleviate the drawbacks highlighted above, we can imagine more complex regularization strategies for $\boldsymbol{\kappa}$. Instead of an explicit regularizer, plug-and-play regularization could be used. For example, a NLM algorithm

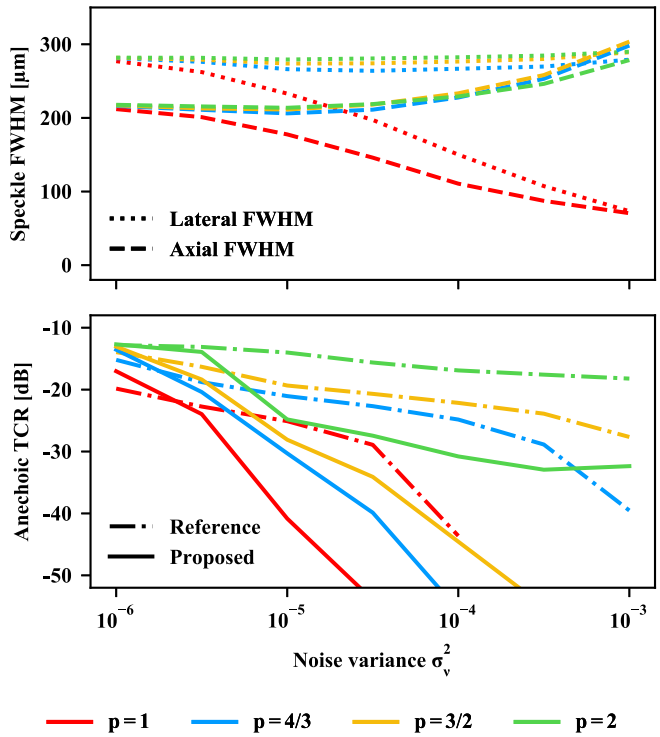

Fig. 7. Evolution of the axial and lateral speckle full width at half maximum (FWHM) in the phantom with respect to the noise variance σ_ν^2, along with the evolution of the tissue-to-clutter ratio (TCR) of the anechoic inclusion, for both the proposed and reference priors.

is used in [8] to enforce the desired image structure. A learned prior on high-quality despeckled image could also be imagined.

We focus in this article on single PW imaging. This choice is motivated by the simple implementation of operator $\boldsymbol{A}$ relatively to other imaging configurations. In addition, the low data size limits the computational burden of the method. However, the proposed model is valid for other configurations as well and considering, for example, several PWs could alleviate the negative effects of multiple scattering.

Finally, modeling errors caused by local variations of SoS and attenuation probably have a significant negative impact on the echogenicity map estimated by the proposed method. A radical approach to solve this issue would be to take into account SoS and attenuation maps in the model and estimate them with problem (7) alongside echogenicity.

B. Reflectivity Estimation

We saw in Section IV-E that the in-vivo reflectivity estimates are affected by side-lobe-like structures. We suppose that the latter originates, on the one hand, from aberration phase shifts affecting the signals—typically caused by SoS variations in the medium. On the other hand, frequency-dependent attenuation causes the pulse-echo wavelet v_{pe} to distort as waves propagate through the medium, a phenomenon disregarded by the proposed model. We believe it is another important source of artifacts in the reflectivity map. Further tests are however necessary to properly quantify the sensitivity of the proposed approach to SoS and attenuation effects—including their effect on the estimation of despeckled images. However, as detailed in Section V-A, a solution would be to directly take into account variations of SoS and attenuation in the model.

We observe that the side-lobe-like artifacts are greatly alleviated by changing the reflectivity prior—such as replacing the uniform regularizer with the proposed one or reducing the value of p. More accurate reflectivity priors could therefore further improve image quality, including, for example, taking into account a non-uniform value of p [19]. More generally, our work emphasizes the crucial role played by the prior in reflectivity estimation methods.

Finally, our proposed reflectivity estimation method is doubly affected by multiple scattering. First, the oversight of multiple scattering reduces the accuracy of the measurement model. Second, the reflectivity prior defined by the echogenicity map loses in quality. Again, our preferred approach to alleviate this issue is to consider several insonifications in the measurement model.

C. Parameter Selection

The proposed approach comprises a series of parameters that must be determined. Most parameters—N^{est}, η, f^{reg}, ϵ^{BCG}, N^{iter}, ϵ^{Prox}, and ϵ^{FISTA}—do not affect the theoretical solutions of (7) and (23). They solely enforce a trade-off between the convergence times of the optimization algorithms and the accuracy of the solutions.

However, three parameters besides p—$\kappa^{\max}$, μ and σ_ν^2—affect the solutions of (7) and (23). The influence of $\kappa^{\max}$ is extremely limited in practice—when set sufficiently large—beyond its positive effect on convergence.

In contrast, the strength μ of the TV norm regularizer must achieve a critical trade-off. If μ is too low, the despeckling capability of the proposed method is affected. In contrast, an overly large value of μ prevents the method to recover the fine details of the image—such as muscular fibers.

As detailed in Section II-B, the additive white noise $\boldsymbol{\nu}$ of variance σ_ν^2 in (4) is modelling mostly the inaccuracies of the measurement model $\boldsymbol{A}$. Consequently, we cannot determine σ_ν^2 by estimating the amount of noise only from $\boldsymbol{m}$, and we rather set σ_ν^2 from the empirical performances of the method. As seen in Section IV-F, a high value of σ_ν^2 reduces contrast in despeckled images. Fine details of the image are also increasingly overlooked, as it can be observed directly above and below the inclusions of Fig. 6. However, if σ_ν^2 is set too low, the proposed approach is not able to discard physical phenomenons neglected by the measurement model. Thus, significant artifacts are bound to appear in the despeckled images. Even though a high value of σ_ν^2 increases the contrast in deconvolved images, the improvement occurs at the cost of a reduction of resolution. Overall, we therefore recommend to set σ_ν^2 to the lowest value possible before significant artifacts affect the echogenicity map $\boldsymbol{\kappa}$.

An interrogation regarding the proposed parameters concerns their generalizability to different transducers, patients and organs. We can see that the same parameters lead to meaningful result with simulated, in-vitro and in-vivo data. This fact provides a

first positive indication for their general effectiveness for a given transducer and system. Regarding the implementation of the method with other configurations, we propose to set parameter σ_ν^2 as a fraction of the median squared modulus of the input CRF echo signals $\boldsymbol{m}$, for a series of acquisitions. We recommend a ratio of 500, even though more thorough investigations are necessary to determine this value precisely. We theoretically expect the regularization parameter μ to be independent of the imaging configuration. Indeed, the strength of the regularizer does not vary with grid size nor the amplitude of the input signals $\boldsymbol{m}$. Again, additional experiments are necessary to confidently establish this independence.

D. Applicability

We must highlight that the most important drawback of the proposed echogenicity estimation method in its current form is its computational complexity. If Algorithm 1 converges quickly in term of iterations, its computational cost per iteration is extremely high. Our current prototype implementation can therefore take up to one full day to converge with a computer comprising a Intel Core i9 CPU (Intel, Santa Clara, CA, USA), 64 GB of RAM and a Nvidia RTX 2080 Ti GPU (Nvidia, Santa Clara, CA, USA). However, we believe that our code can be significantly sped-up. Other—potentially quicker—inversion methods can also be investigated, such as approaches based on Langevin diffusion [42], [43]. We can also imagine using such a method to perform a joint estimation of reflectivity and echogenicity.

To reduce the computational burden of the proposed technique, another possibility is to replace the covariance matrix $\boldsymbol{\Sigma}$ (5) by an approximation admitting a computationally efficient inversion—for instance, a block-diagonal matrix. Such a simplification could, for example, be performed with data expressed in the frequency domain [4]. A trade-off would however be necessary between accuracy and computational cost. We also believe that self-supervised learning based on deep neural networks could offer a powerful way to harness the proposed model while achieving real-time imaging. Overall, our next line of research is to reduce the computational burden of the proposed method.

VI. Conclusion

We devise in this article a statistical model for ultrasound imaging. We show that the proposed model can improve the estimation of despeckled and deconvolved images from a limited amount of data, with respect to DAS followed by a despeckling algorithm and a state-of-the-art inverse problem for reflectivity estimation, respectively.

We introduce a novel physical interpretation of a despeckled image as the reflectivity variance. We believe this interpretation can provide meaningful insights for the design of despeckled image reconstruction methods, especially when a low number of data is considered. More widely, we consider that the proposed approach relying on the covariance matrix of measured signal can act as a powerful basis for future pulse-echo ultrasound imaging techniques, including quantitative ones.

Appendix A
Estimation of the Diagonal of the Fisher Information Matrix

We need to estimate the diagonal $\boldsymbol{f}$ of the Fisher information matrix. Let us first recall the definition of the latter. For an arbitrary probability distribution $\mathcal{G}$ associated with a PDF $\mathbb{P}(\cdot|\boldsymbol{\theta})$ and parameterized by a vector $\boldsymbol{\theta} \in \mathbb{R}^N$, the Fisher information matrix $\boldsymbol{F}(\boldsymbol{\theta}) \in \mathbb{C}^{N\times N}$ is defined as

$$\boldsymbol{F}(\boldsymbol{\theta}) = \mathop{\mathbb{E}}_{\boldsymbol{x}\sim\mathcal{G}(\boldsymbol{\theta})} \Big\{ [\nabla_{\boldsymbol{\theta}} \log\left(\mathbb{P}(\boldsymbol{x}|\boldsymbol{\theta})\right)] \; [\nabla_{\boldsymbol{\theta}} \log\left(\mathbb{P}(\boldsymbol{x}|\boldsymbol{\theta})\right)]^H \Big\}. \tag{28}$$

In our case, the probability distribution $\mathcal{G}$ is the multivariate zero-mean circularly-symmetric Gaussian distribution whose PDF is given in (6). Moreover, the parameter vector $\boldsymbol{\theta}$ is given by $\boldsymbol{\kappa} \in \mathbb{R}^{N^x N^z}$. The diagonal of the Fisher information matrix can thus be written as

$$\boldsymbol{f}(\boldsymbol{\kappa}) = \mathop{\mathbb{E}}_{\hat{\boldsymbol{m}}\sim\mathcal{N}(\boldsymbol{0},\boldsymbol{\Sigma}(\boldsymbol{\kappa}))} \Big\{ \Big| \nabla_{\boldsymbol{\kappa}} \Big[\log\left|\boldsymbol{\Sigma}(\boldsymbol{\kappa})\right| + \hat{\boldsymbol{m}}^H \boldsymbol{\Sigma}^{-1}(\boldsymbol{\kappa}) \hat{\boldsymbol{m}} \Big] \Big|^2 \Big\}. \tag{29}$$

Introducing (13) and (14) in (29) leads to

$$\begin{aligned} \boldsymbol{f}(\boldsymbol{\kappa}) = \Bigg[& \left| \operatorname{diag}\left(\boldsymbol{A}^H \boldsymbol{\Sigma}^{-1}(\boldsymbol{\kappa}) \boldsymbol{A}\right) \right|^2 \\ & - 2 \left[\mathop{\mathbb{E}}_{\hat{\boldsymbol{m}}\sim\mathcal{N}(\boldsymbol{0},\boldsymbol{\Sigma}(\boldsymbol{\kappa}))} \left\{ \left| \boldsymbol{A}^H \boldsymbol{\Sigma}^{-1}(\boldsymbol{\kappa}) \hat{\boldsymbol{m}} \right|^2 \right\} \right] \odot \\ & \operatorname{diag}\left(\boldsymbol{A}^H \boldsymbol{\Sigma}^{-1}(\boldsymbol{\kappa}) \boldsymbol{A}\right) \\ & + \mathop{\mathbb{E}}_{\hat{\boldsymbol{m}}\sim\mathcal{N}(\boldsymbol{0},\boldsymbol{\Sigma}(\boldsymbol{\kappa}))} \left\{ \left| \boldsymbol{A}^H \boldsymbol{\Sigma}^{-1}(\boldsymbol{\kappa}) \hat{\boldsymbol{m}} \right|^4 \right\} \Bigg] \odot e^{2\boldsymbol{\kappa}}, \end{aligned} \tag{30}$$

and

$$\begin{aligned} \boldsymbol{f}(\boldsymbol{\kappa}) = \Bigg[& \mathop{\mathbb{E}}_{\hat{\boldsymbol{m}}\sim\mathcal{N}(\boldsymbol{0},\boldsymbol{\Sigma}(\boldsymbol{\kappa}))} \left\{ \left| \boldsymbol{A}^H \boldsymbol{\Sigma}^{-1}(\boldsymbol{\kappa}) \hat{\boldsymbol{m}} \right|^4 \right\} \\ & - \left[\mathop{\mathbb{E}}_{\hat{\boldsymbol{m}}\sim\mathcal{N}(\boldsymbol{0},\boldsymbol{\Sigma}(\boldsymbol{\kappa}))} \left\{ \left| \boldsymbol{A}^H \boldsymbol{\Sigma}^{-1}(\boldsymbol{\kappa}) \hat{\boldsymbol{m}} \right|^2 \right\} \right]^2 \Bigg] \odot e^{2\boldsymbol{\kappa}}, \end{aligned} \tag{31}$$

after taking into account (15). Therefore, we propose to introduce a second Monte-Carlo approximation, similar in nature to the one used in (17), leading to

$$\begin{aligned} \boldsymbol{f}(\boldsymbol{\kappa}) \approx \Bigg[& \frac{1}{N^{\text{est}}} \sum_{i=1}^{N^{\text{est}}} \left| \boldsymbol{A}^H \boldsymbol{\Sigma}^{-1}(\boldsymbol{\kappa}) \hat{\boldsymbol{m}}_i \right|^4 \\ & - \left[\frac{1}{N^{\text{est}}} \sum_{i=1}^{N^{\text{est}}} \left| \boldsymbol{A}^H \boldsymbol{\Sigma}^{-1}(\boldsymbol{\kappa}) \hat{\boldsymbol{m}}_i \right|^2 \right]^2 \Bigg] \odot e^{2\boldsymbol{\kappa}}, \end{aligned} \tag{32}$$

where vectors $\boldsymbol{m}_i,\ i = 1, \ldots, N^{\text{est}}$ are generated according to (18). In particular, we can reuse in (32) the solution of the linear system of equation computed to estimate (17). In practice,

Algorithm 3: Computation of $\mathrm{Prox}_{\eta\mu R}^{\mathrm{Diag}(\boldsymbol{f})}(\hat{\boldsymbol{\kappa}})$.

input: $\hat{\boldsymbol{\kappa}} \in \mathbb{R}^{N^{\mathrm{x}}N^{\mathrm{z}}}$, $\boldsymbol{f} \in \mathbb{R}_+^{N^{\mathrm{x}}N^{\mathrm{z}}}$, $\eta\mu > 0$, an initial guess $\boldsymbol{x}_0 \in \mathbb{R}^{2(N^{\mathrm{x}}-1)(N^{\mathrm{z}}-1)}$, a stopping criterion $\epsilon^{\mathrm{Prox}} > 0$
$\boldsymbol{y} = \boldsymbol{x}_0$
$L \leftarrow [\frac{4}{\Delta x^2} + \frac{4}{\Delta z^2}]/f^{\mathrm{reg}}$
$k \leftarrow 0$
do
  $\hat{\boldsymbol{x}} \leftarrow \hat{\boldsymbol{\kappa}} + \frac{1}{L}\boldsymbol{D}\,\mathrm{Proj}_{]-\infty,\kappa^{\max}]}(\hat{\boldsymbol{\kappa}} - \mathrm{Diag}^{-1}(\boldsymbol{f})\boldsymbol{D}^T\boldsymbol{y})$
  $\boldsymbol{x}_{k+1} = \mathrm{Prox}_{\eta\mu h}(\hat{\boldsymbol{x}})$
  $t_{k+1} = \frac{1}{2}[1 + \sqrt{1 + 4t_k^2}]$
  $\boldsymbol{y} \leftarrow \boldsymbol{x}_{k+1} + \frac{t_k - 1}{t_{k+1}}[\boldsymbol{x}_{k+1} - \boldsymbol{x}_k]$
  $k \leftarrow k + 1$
while $\frac{1}{\sqrt{2(N^{\mathrm{x}}-1)(N^{\mathrm{z}}-1)}}\|\boldsymbol{x}_{k+1} - \boldsymbol{x}_k\| > \epsilon^{\mathrm{Prox}}$
return $\mathrm{Proj}_{[\kappa^{\min},\kappa^{\max}]}(\hat{\boldsymbol{\kappa}} - \mathrm{Diag}^{-1}(\boldsymbol{f})\boldsymbol{D}^T\boldsymbol{x}_k)$

however, the value of $\boldsymbol{f}$ we use in the optimization algorithm is given by (32), plus a regularization factor $f^{\mathrm{reg}} > 0$. The latter alleviates the statistical variations of estimator (32). In addition, an iterative optimization algorithm—detailed in Appendix B—is required to compute $\mathrm{Prox}_{\mu R}^{\mathrm{Diag}(\boldsymbol{f})}(\hat{\boldsymbol{\kappa}})$ (12). Its convergence speed depends on the minimum of $\boldsymbol{f}$ and a non-zero value of f^{reg} allows us to limit its running time.

Appendix B
Proximal Operator of the Total-Variation Norm

To compute the proximal operator of the regularizer μR in (7), we rely on the dual approach introduced in [37]. However, we modified the algorithm to include the diagonal matrix weighting $\mathrm{Diag}(\boldsymbol{f})$ present in the data fidelity term. The scheme in question is detailed in Algorithm 3.

Function h in Algorithm 3 is defined as

$$h(\boldsymbol{x}) = \sum_{k=1}^{(N^{\mathrm{x}}-1)(N^{\mathrm{z}}-1)} \sqrt{[x_k^{\mathrm{x}}]^2 + [x_k^{\mathrm{z}}]^2},$$

$$\text{with } \boldsymbol{x} = \begin{bmatrix}\boldsymbol{x}^{\mathrm{x}} \\ \boldsymbol{x}^{\mathrm{z}}\end{bmatrix},\ \boldsymbol{x}^{\mathrm{x,z}} \in \mathbb{R}^{(N^{\mathrm{x}}-1)(N^{\mathrm{z}}-1)}, \tag{33}$$

whereas its proximal operator is defined as

$$\mathrm{Prox}_{\mu h}(\hat{\boldsymbol{x}}) = \begin{bmatrix}\boldsymbol{x}^{\mathrm{x}} \\ \boldsymbol{x}^{\mathrm{z}}\end{bmatrix},$$

$$\text{with } \begin{bmatrix}x_k^{\mathrm{x}} \\ x_k^{\mathrm{z}}\end{bmatrix} = \frac{\max\left\{\sqrt{[\hat{x}_k^{\mathrm{x}}]^2 + [\hat{x}_k^{\mathrm{z}}]^2} - \mu, 0\right\}}{\sqrt{[\hat{x}_k^{\mathrm{x}}]^2 + [\hat{x}_k^{\mathrm{z}}]^2}} \begin{bmatrix}\hat{x}_k^{\mathrm{x}} \\ \hat{x}_k^{\mathrm{z}}\end{bmatrix}. \tag{34}$$

In addition,

$$\boldsymbol{D} = \begin{bmatrix}\boldsymbol{D}^{\mathrm{x}} \\ \boldsymbol{D}^{\mathrm{z}}\end{bmatrix},\ \text{with } [\boldsymbol{D}^{\mathrm{x}}\boldsymbol{x}]_{(i,j)} = \frac{x_{(i+1,j)} - x_{(i,j)}}{\Delta x},$$

$$[\boldsymbol{D}^{\mathrm{z}}\boldsymbol{x}]_{(i,j)} = \frac{x_{(i,j+1)} - x_{(i,j)}}{\Delta z},$$

$$i = 1, \ldots, N^{\mathrm{x}} - 1, j = 1, \ldots, N^{\mathrm{z}} - 1. \tag{35}$$

In (35), Δx and Δz denote the discretization step of $\boldsymbol{\kappa}$ along the x and z axes, respectively. Finally, we have

$$\left[\mathrm{Proj}_{]-\infty,\kappa^{\max}]}(\boldsymbol{x})\right]_k = \min\{x_k, \kappa^{\max}\}. \tag{36}$$

Appendix C
Proximal Operator of ℓ_p Norms

Algorithm 2 requires the computation of a proximal operator taking the form

$$\min_{\boldsymbol{\gamma}\in\mathbb{C}^{N^{\mathrm{x}}N^{\mathrm{z}}}} \left[\frac{\Gamma\left(\frac{4}{p}\right)}{\Gamma\left(\frac{2}{p}\right)}\right]^{\frac{p}{2}} \left\|\mathrm{Diag}\left(e^{-\frac{1}{2}\boldsymbol{\kappa}}\right)\boldsymbol{\gamma}\right\|_p^p + \frac{L}{2}\left\|\boldsymbol{\gamma} - \hat{\boldsymbol{\gamma}}\right\|^2, \tag{37}$$

with $p \geq 1$, $L > 0$. Problem (37) is separable and the solution can be computed for each element of $\boldsymbol{\gamma}$ independently through the following set of problems

$$\min_{\gamma_i} \beta_i|\gamma_i|^p + \frac{1}{2}[\gamma_i - \hat{\gamma}_i]^2,\ \forall i = 1, \ldots, N^{\mathrm{x}}N^{\mathrm{z}}, \tag{38}$$

with

$$\beta_i = \frac{1}{L}\left[\frac{\Gamma\left(\frac{4}{p}\right)}{\Gamma\left(\frac{2}{p}\right)}\right]^{\frac{p}{2}} e^{-\frac{p}{2}\kappa_i}. \tag{39}$$

According to [44], the solutions of (38) are given by

$$\gamma_i = \rho_i e^{j\angle\hat{\gamma}_i} \tag{40}$$

with $\angle\hat{\gamma}_i$ the argument of $\hat{\gamma}_i$ and ρ_i the solution of

$$\rho_i + p\beta_i\rho_i^{p-1} = |\hat{\gamma}_i|,\ \rho_i \geq 0. \tag{41}$$

If $p = 1$, the solution in the limit case $p \to 1$ writes

$$\rho_i = \max\{|\hat{\gamma}_i| - \beta, 0\}, \tag{42}$$

whereas we have

$$\rho_i = \frac{|\hat{\gamma}_i|}{1 + 2\beta_i}, \tag{43}$$

if $p = 2$. For the cases $p = \frac{3}{2}$ and $p = \frac{4}{3}$, (41) can be expressed as quadratic and cubic equations, respectively. Their solutions are thus obtained with the quadratic formula and Cardano's formula, such that

$$\rho_i = |\hat{\gamma}_i| + \frac{3\beta_i}{4}\left[\frac{3\beta_i}{2} - \sqrt{4|\hat{\gamma}_i| + \frac{9\beta_i^2}{4}}\right], \tag{44}$$

for $p = \frac{3}{2}$, and

$$\rho_i = |\hat{\gamma}_i| + C - \frac{64\beta_i^3}{81C}, \tag{45}$$

for $p = \frac{4}{3}$, with

$$C = \frac{4\beta_i}{3}\sqrt[3]{\sqrt{\frac{|\hat{\gamma}_i|^2}{4} + \frac{64\beta_i^3}{729}} - \frac{|\hat{\gamma}_i|}{2}}. \tag{46}$$

REFERENCES

[1] J.-L. Gennisson, T. Deffieux, M. Fink, and M. Tanter, "Ultrasound elastography: Principles and techniques," *Diagn. Interventional Imag.*, vol. 94, no. 5, pp. 487–495, May 2013.

[2] E. Mace, G. Montaldo, B.-F. Osmanski, I. Cohen, M. Fink, and M. Tanter, "Functional ultrasound imaging of the brain: Theory and basic principles," *IEEE Trans. Ultrasonics, Ferroelect., Freq. Control*, vol. 60, no. 3, pp. 492–506, Mar. 2013.

[3] G. Montaldo, M. Tanter, J. Bercoff, N. Benech, and M. Fink, "Coherent plane-wave compounding for very high frame rate ultrasonography and transient elastography," *IEEE Trans. Ultrasonics, Ferroelect., Freq. Control*, vol. 56, no. 3, pp. 489–506, Mar. 2009.

[4] A. Besson et al., "A sparse reconstruction framework for Fourier-based plane-wave imaging," *IEEE Trans. Ultrasonics, Ferroelect., Freq. Control*, vol. 63, no. 12, pp. 2092–2106, Dec. 2016.

[5] T. Chernyakova et al., "Fourier-domain beamforming and structure-based reconstruction for plane-wave imaging," *IEEE Trans. Ultrasonics, Ferroelect., Freq. Control*, vol. 65, no. 10, pp. 1810–1821, Oct. 2018.

[6] T. Szasz, A. Basarab, and D. Kouamé, "Beamforming through regularized inverse problems in ultrasound medical imaging," *IEEE Trans. Ultrasonics, Ferroelect., Freq. Control*, vol. 63, no. 12, pp. 2031–2044, Dec. 2016.

[7] E. Ozkan, V. Vishnevsky, and O. Goksel, "Inverse problem of ultrasound beamforming with sparsity constraints and regularization," *IEEE Trans. Ultrasonics, Ferroelect., Freq. Control*, vol. 65, no. 3, pp. 356–365, Mar. 2018.

[8] S. Goudarzi, A. Basarab, and H. Rivaz, "Inverse problem of ultrasound beamforming with denoising-based regularized solutions," *IEEE Trans. Ultrasonics, Ferroelectr., Freq. Control*, vol. 69, no. 10, pp. 2906–2916, Oct. 2022.

[9] A. Besson et al., "Ultrafast ultrasound imaging as an inverse problem: Matrix-free sparse image reconstruction," *IEEE Trans. Ultrasonics, Ferroelectr., Freq. Control*, vol. 65, no. 3, pp. 339–355, Mar. 2018.

[10] T. Szasz, A. Basarab, M.-F. Vaida, and D. Kouamé, "Elastic-net based beamforming in medical ultrasound imaging," in *Proc. IEEE 2016 13th Int. Symp. Biomed. Imag.*, 2016, pp. 477–480.

[11] G. David, J.-l. Robert, B. Zhang, and A. F. Laine, "Time domain compressive beam forming of ultrasound signals," *J. Acoustical Soc. Amer.*, vol. 137, no. 5, pp. 2773–2784, 2015.

[12] M. F. Schiffner and G. Schmitz, "Fast pulse-echo ultrasound imaging employing compressive sensing," in *Proc. IEEE 2011 Int. Ultrasonics Symp.*, 2011, pp. 688–691.

[13] A. Besson et al., "A physical model of nonstationary blur in ultrasound imaging," *IEEE Trans. Comput. Imag.*, vol. 5, no. 3, pp. 381–394, Sep. 2019.

[14] J. A. Jensen, J. Mathorne, T. Gravesen, and B. Stage, "Deconvolution of in-vivo ultrasound B-mode images," *Ultrason. Imag.*, vol. 15, no. 2, pp. 122–133, 1993.

[15] T. Taxt, "Restoration of medical ultrasound images using two-dimensional homomorphic deconvolution," *IEEE Trans. Ultrasonics, Ferroelect., Freq. Control*, vol. 42, no. 4, pp. 543–554, Jul. 1995.

[16] O. Michailovich and A. Tannenbaum, "Blind deconvolution of medical ultrasound images: A parametric inverse filtering approach," *IEEE Trans. Image Process.*, vol. 16, no. 12, pp. 3005–3019, Dec. 2007.

[17] M. Alessandrini et al., "A restoration framework for ultrasonic tissue characterization," *IEEE Trans. Ultrasonics, Ferroelect., Freq. Control*, vol. 58, no. 11, pp. 2344–2360, Nov. 2011.

[18] Z. Chen, A. Basarab, and D. Kouamé, "Compressive deconvolution in medical ultrasound imaging," *IEEE Trans. Med. Imag.*, vol. 35, no. 3, pp. 728–737, Mar. 2016.

[19] M.-C. Corbineau, D. Kouamé, E. Chouzenoux, J.-Y. Tourneret, and J.-C. Pesquet, "Preconditioned P-ULA for joint deconvolution-segmentation of ultrasound images," *IEEE Signal Process. Lett.*, vol. 26, no. 10, pp. 1456–1460, Oct. 2019.

[20] A. Khare, M. Khare, Y. Jeong, H. Kim, and M. Jeon, "Despeckling of medical ultrasound images using Daubechies complex wavelet transform," *Signal Process.*, vol. 90, no. 2, pp. 428–439, 2010.

[21] P. Sudeep et al., "Speckle reduction in medical ultrasound images using an unbiased non-local means method," *Biomed. Signal Process. Control*, vol. 28, pp. 1–8, 2016.

[22] O. V. Michailovich and A. Tannenbaum, "Despeckling of medical ultrasound images," *IEEE Trans. Ultrasonics, Ferroelectr., Freq. Control*, vol. 53, no. 1, pp. 64–78, Jan. 2006.

[23] S. Liang, F. Yang, T. Wen, Z. Yao, Q. Huang, and C. Ye, "Nonlocal total variation based on symmetric Kullback-Leibler divergence for the ultrasound image despeckling," *BMC Med. Imag.*, vol. 17, no. 1, pp. 1–12, 2017.

[24] K. Mei, B. Hu, B. Fei, and B. Qin, "Phase asymmetry ultrasound despeckling with fractional anisotropic diffusion and total variation," *IEEE Trans. Image Process.*, vol. 29, pp. 2845–2859, 2020.

[25] J. A. Noble and D. Boukerroui, "Ultrasound image segmentation: A survey," *IEEE Trans. Med. Imag.*, vol. 25, no. 8, pp. 987–1010, Aug. 2006.

[26] J. A. Jensen, "A model for the propagation and scattering of ultrasound in tissue," *J. Acoustical Soc. Amer.*, vol. 89, no. 1, pp. 182–190, 1991.

[27] M. Persson, D. Bone, and H. Elmqvist, "Total variation norm for three-dimensional iterative reconstruction in limited view angle tomography," *Phys. Med. Biol.*, vol. 46, no. 3, 2001, Art. no. 853.

[28] Z. Tian, X. Jia, K. Yuan, T. Pan, and S. B. Jiang, "Low-dose CT reconstruction via edge-preserving total variation regularization," *Phys. Med. Biol.*, vol. 56, no. 18, 2011, Art. no. 5949.

[29] J. Huang, S. Zhang, and D. Metaxas, "Efficient MR image reconstruction for compressed MR imaging," *Med. Image Anal.*, vol. 15, no. 5, pp. 670–679, 2011.

[30] U. S. Kamilov et al., "Optical tomographic image reconstruction based on beam propagation and sparse regularization," *IEEE Trans. Comput. Imag.*, vol. 2, no. 1, pp. 59–70, Mar. 2016.

[31] S. Amari and S. C. Douglas, "Why natural gradient?," in *Proc. IEEE 1998 Int. Conf. Acoust., Speech Signal Process.*, Seattle, WA, USA, 1998, vol. 2, pp. 1213–1216.

[32] J. Martens, "New insights and perspectives on the natural gradient method," *J. Mach. Learn. Res.*, vol. 21, no. 1, pp. 5776–5851, 2020.

[33] J. D. Lee, Y. Sun, and M. Saunders, "Proximal Newton-type methods for convex optimization," in *Proc. Annu. Conf. Neural Inf. Process. Syst.*. 2012, pp. 827–835.

[34] H. Liu, Y.-S. Ong, X. Shen, and J. Cai, "When Gaussian process meets big data: A review of scalable GPs," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 31, no. 11, pp. 4405–4423, Nov. 2022.

[35] K. Wang, G. Pleiss, J. Gardner, S. Tyree, K. Q. Weinberger, and A. G. Wilson, "Exact Gaussian processes on a million data points," in *Proc. Annu. Conf. Neural Inf. Process. Syst.*, 2019, vol. 32, pp. 1–12.

[36] D. P. O'Leary, "The block conjugate gradient algorithm and related methods," *Linear Algebra Appl.*, vol. 29, pp. 293–322, Feb. 1980.

[37] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM J. Imag. Sci.*, vol. 2, no. 1, pp. 183–202, Jan. 2009.

[38] M. Novey, T. Adali, and A. Roy, "A complex generalized Gaussian distribution—characterization, generation, and estimation," *IEEE Trans. Signal Process.*, vol. 58, no. 3, pp. 1427–1433, Mar. 2010.

[39] D. Perdios, M. Vonlanthen, F. Martinez, M. Arditi, and J.-P. Thiran, "CNN-based image reconstruction method for ultrafast ultrasound imaging," *IEEE Trans. Ultrasonics, Ferroelectr., Freq. Control*, vol. 69, no. 4, pp. 1154–1168, Apr. 2022.

[40] J. A. Jensen, "Field: A program for simulating ultrasound systems," in *Proc. 10th Nordic-Baltic Conf. Biomed. Imag.*, Tempere, Finland, 1997, pp. 351–353.

[41] M. V. Wijk and J. Thijssen, "Performance testing of medical ultrasound equipment: Fundamental vs. harmonic mode," *Ultrasonics*, vol. 40, no. 1–8, pp. 585–591, 2002.

[42] G. O. Roberts and O. Stramer, "Langevin diffusions and Metropolis-Hastings algorithms," *Methodol. Comput. Appl. Probability*, vol. 4, pp. 337–357, 2002.

[43] A. Durmus, E. Moulines, and M. Pereyra, "Efficient Bayesian computation by proximal Markov chain Monte Carlo: When Langevin meets Moreau," *SIAM J. Imag. Sci.*, vol. 11, no. 1, pp. 473–506, 2018.

[44] P. L. Combettes and J.-C. Pesquet, "Proximal splitting methods in signal processing," in *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*. Berlin, Germany: Springer, 2011, pp. 185–212.