
Photoacoustic Imaging with Conditional Priors from Normalizing Flows

Rafael Orozco
Georgia Institute of Technology
rorozco@gatech.edu

Ali Siahkoobi
Georgia Institute of Technology
alisk@gatech.edu

Gabrio Rizzuti
Utrecht University
g.rizzuti@uu.nl

Tristan van Leeuwen
Centrum Wiskunde & Informatica
t.van.leeuwen@cwi.nl

Felix J. Herrmann
Georgia Institute of Technology
felix.herrmann@gatech.edu

Abstract

For many ill-posed inverse problems, such as photoacoustic imaging, the uncertainty of the solution is highly affected by measurement noise and data incompleteness (due, for example, to limited aperture). For these problems, the choice of prior information is a crucial aspect of a computationally effective solution scheme. We propose a regularization scheme for photoacoustic imaging that leverages prior information learned by a generative network. We train conditional normalizing flows on pairs of photoacoustic sources (the unknowns of the problem) and the associated data in order to exploit the posterior distribution of the solution. The learned prior is combined with physics-based optimization (enforced by partial differential equations), according to the deep prior framework, in order to achieve robustness with respect to out-of-distribution data. Numerical evidence suggests the superiority of this approach with respect to non-conditional deep priors, and the ability to retrieve features of the unknowns that are typically challenging for limited-view photoacoustics.

1 Introduction

Photoacoustic imaging is a biomedical imaging technique based on the photoacoustic effect [17]. It leverages the interplay between optics and acoustics as a mean to circumvent the limitations of imaging modalities relying on single-type physics. Light beams generated by a pulsed laser can penetrate biological tissues by several centimeters, and are absorbed based on oxygen saturation or hemoglobin concentration. While optical absorption is in principle an ideal parameter for medical imaging (e.g., with respect to the detection of cancerous tissue), strong scattering imposes important limitations in its imaging resolution. Ultrasonics, on the other hand, can theoretically provide resolution of medical diagnostic value, but produce images of mechanical properties whose contrasts are not sensitive. In photoacoustics, optical and acoustic effects are combined to gain the best of both worlds. Under conditions of thermal and stress confinement, thermal energy can efficiently build up in biological tissues, which in turn undergo thermal expansion and effectively act as a spatially distributed acoustic source. In photoacoustic imaging, the actual object of interest is the induced source, as it is directly related to optical absorption and can be recovered with a relatively higher resolution than pure optical imaging, based on the acquired ultrasonic data.

Photoacoustic imaging boils down to an initial value problem. The pressure field p , function of space and time, obeys the acoustic wave equation:

$$\begin{cases} c^{-2}\partial_{tt}p - \Delta p = 0, \\ \partial_t p|_{t=0} = 0, \\ p|_{t=0} = p_0. \end{cases} \quad (1)$$

We are interested in retrieving the initial pressure distribution p_0 from acoustic data. However, due to the intrinsic limitations of optics, the viable imaging depth for human tissue ranges around 1-4 cm, so that the receiver geometry has typically a one-sided view of the object [as for planar sensors, [8]]. This restricted view poses serious challenges for photoacoustic imaging.

Due to measurement incompleteness, one has to resort to prior information. In a Bayesian setting, we posit a data likelihood $\mathbf{d} = A\mathbf{x} + \boldsymbol{\varepsilon}$, where the Gaussian white noise $\boldsymbol{\varepsilon}$ is independent from \mathbf{x} , which represents the initial distribution value p_0 appearing in equation (1). The linear operator A comprises the solution of the equation (1) followed by the restriction to the receiver position. The prior information is encapsulated by a probability density p_{prior} . The solution is estimated by solving the problem:

$$\min_{\mathbf{x}} \frac{1}{2} \|A\mathbf{x} - \mathbf{d}\|^2 - \lambda \log p_{\text{prior}}(\mathbf{x}), \quad (2)$$

for a given weighting parameter λ , to be properly tuned. For example, the compressed sensing paradigm enforces sparsity of the unknown in some transformed domain via ℓ^1 -norm regularization. This class of "hand-crafted" priors are well studied from a theoretical point of view [5], but they typically do not take advantage of domain-specific knowledge of the problem.

In recent years, many applications of deep learning to linear inverse problems have dealt with priors implicitly defined by neural networks, e.g. by pre-training a generative network on a collection of candidate solutions [4] or without training by directly exploiting the network inductive bias that favors natural-looking images [14]. Our proposal consists of a deep prior solution to photoacoustic imaging, similarly to [4] where the generative model is now trained on joint pairs of photoacoustic sources and associated measurements. The network is structured as a conditional normalizing flow [2, 15], so that it is able to generate samples from the posterior distribution of \mathbf{x} given data \mathbf{y} . For out-of-distribution data \mathbf{y} , however, the learned posterior distribution does not adequately represent the analytical posterior. We then resort to the physics-based formulation in equation (2) in combination with the learned conditional prior to compute the maximum a posteriori solution (MAP). Our approach takes advantage of the favorable trade-off between the representation error and over fit for invertible networks when compared to other generative models [3, 9].

2 Related work

The regularization scheme described in this paper follows the deep prior literature originally spun from the work of [4], which consists in reparameterizing the inverse problem unknowns by the input of a pre-trained generator. A closely related method is the deep image prior of [14]. This method must resort to early stopping strategy to avoid overfitting and only relies on the regularization effect of the network inductive bias. The complementary strengths of these two approaches can be combined in a hybrid scheme, e.g. by using image-adaptive priors [7]. Our proposal is similar to the deep prior application to magnetic resonance imaging discussed in [9], which uses invertible networks. The advantages of using normalizing flows in the deep prior context were originally advocated by [3]. We extend the normalizing flow approach by using conditional architectures, as described in [2] and [15]. Normalizing flows have been subject to a great deal of research, we refer to [10] and [11] for general overviews. The network architecture analyzed in this paper is HINT, discussed in [1] (and later generalized to the conditional setting in [2]). An application of conditional normalizing flows in variational inference was recently discussed in [12].

3 Methods

The proposed method aims at regularizing the inverse problem in equation (2) with a learned prior. The first stage of the algorithm consists of training a generator given access to a dataset of joint pairs $\mathbf{x}, \mathbf{y} \sim p_{X,Y}(\mathbf{x}, \mathbf{y})$, where $\mathbf{y} = A^T \mathbf{d}$ is time-reversed data (referred to as the adjoint solution). We




Figure 1: Samples from the posterior distribution learned in the pre-training phase. Here, we display the results for: (a) in-distribution data (e.g. test data); (b) out-of-distribution data. Due to limited aperture (the receivers are located at the position “ $z=0$ ”), the areas of highest uncertainty correspond to the vertical features of the blood vessel system.

train a normalizing flow $f_{\theta} : X \times Y \rightarrow Z_x \times Z_y$ by minimizing the Kullback-Leibler divergence between the target density $p_{X,Y}$ and the pull-back of the standard Gaussian distribution defined on the latent space $Z_x \times Z_y$ via f_{θ} [10, 11]:

$$\min_{\theta} \text{KL}(p_{X,Y} \| p_{\theta}) = \mathbb{E}_{\mathbf{x}, \mathbf{y} \sim p_{X,Y}(\mathbf{x}, \mathbf{y})} \frac{1}{2} \|f_{\theta}(\mathbf{x}, \mathbf{y})\|^2 - \log |\det J_{f_{\theta}}(\mathbf{x}, \mathbf{y})|. \quad (3)$$

Here, $J_{f_{\theta}}$ denotes the Jacobian of f_{θ} with respect to the input. Following [2], we impose a triangular structure on the normalizing flow, e.g. $f_{\theta}(\mathbf{x}, \mathbf{y}) = (f_{\theta}^{z_x}(\mathbf{x}, \mathbf{y}), f_{\theta}^{z_y}(\mathbf{y}))$. We choose the conditional architecture of the HINT invertible network, whose original (non-conditional) form was proposed in [1]. After an ideal optimization of equation (3), sampling from the posterior distribution $\mathbf{x} \sim p_{X|Y}(\mathbf{x}|\mathbf{y})$ is equivalent to computing the latent data code $\mathbf{z}_y = f_{\theta}^{z_y}(\mathbf{y})$ and evaluating $f_{\theta}^{-1}(\mathbf{z}_x, \mathbf{z}_y)$ for random \mathbf{z}_x 's. This first phase does not require repeated evaluation of the physical model A appearing in equation (2), which helps in relieving the computational burden of solving partial differential equations at scale.

When faced with out-of-distribution data, we can reasonably assume that the results previously obtained are not as reliable as for in-distribution data. The second stage of the proposed method then relies on the physics model in equation (2) and uses the pre-trained generator as an implicit deep prior. Formally, we end up solving the problem:

$$\min_{\mathbf{z}_x} \frac{1}{2} \|A f_{\theta}^{-1}(\mathbf{z}_x, \mathbf{z}_y) - \mathbf{d}\|^2 + \frac{\lambda^2}{2} \|\mathbf{z}_x\|^2, \quad \mathbf{z}_y = f_{\theta}^{z_y}(A^T \mathbf{d}). \quad (4)$$

Note that the minimization is carried over the latent variables related to the unknowns \mathbf{x} . The computational effort to evaluate A and its adjoint is manageable when single data \mathbf{d} is considered.

4 Results

The results related to the training of the generative model are collected in Figure 1. The training set consists of images of blood vessel systems \mathbf{x} and time-reversed data $\mathbf{y} = A^T \mathbf{d}$ measured at the receiver location located on top of the image[6]. The measurements are generated synthetically by solving the wave equation and adding Gaussian white noise. We qualitatively compare samples drawn from the learned posterior distribution for in- and out-of-distribution data.

In Figure 2, we compare the reconstruction results for out-of-distribution data by solving the problem in equation (2) with different choices for the prior. In particular, we consider a non-informative Gaussian prior and learned priors trained on a dataset of samples \mathbf{x} 's and joint pairs \mathbf{x}, \mathbf{y} 's (our proposal). A qualitative inspection and quality metrics indicate superior results for the conditional prior approach. Further comparison of the results regarding data fit is shown in Figure 3.

5 Discussion and conclusions

In limited-view photoacoustic imaging, the resolution of deep near-vertical structures in a blood vessel system is a fundamental challenge due to the relatively weak imprint on the measurements. Note that




Figure 2: Reconstruction results for photoacoustic imaging. We compare MAP estimates obtained with different priors: (a) ground-truth, (b) Gaussian prior, (c) learned (marginal) prior $p_X(\mathbf{x})$, (d) learned conditional prior $p_{X,Y}(\mathbf{x}, \mathbf{y})$.




Figure 3: Data fit for some of the reconstruction results shown in Figure 2: (a) Gaussian prior, (b) learned (marginal) prior $p_X(\mathbf{x})$, (c) learned conditional prior $p_{X,Y}(\mathbf{x}, \mathbf{y})$.

similar issues are common for other wave-equation-based imaging applications such as geophysical prospecting [13]. Conventional regularization methods based on sparsity in some transformed domain are typically too generic and practically unable to fill the sensitivity gap. A straightforward alternative is to use problem-specific information.

We proposed a regularization scheme for photoacoustic imaging that employs deep priors learned in an offline phase. The generative model is trained on a dataset containing pairs of candidate solutions and associated measurements, the goal is to learn the posterior distribution of the solution given some data. The primary scope of this work is to compare the regularization effect of conditional deep priors with "marginal" deep priors proposed in the recent past for many imaging applications.

The results presented in this paper suggest that the MAP estimate based on conditional deep priors is able to recover the vertical features of a blood vessel image. Our results also show that the marginal deep prior is more prone to misplace features. In our view, learned prior regularization can be beneficial for highly ill-posed problems such as limited-view photoacoustic imaging. However, we must further assess the bias imposed by this type of learned priors.

The dataset used in this work is a derivative of the ELCAP lung dataset prepared by [6]. The results presented here use a Julia implementation of invertible network architectures: <https://github.com/slimgroup/InvertibleNetworks.jl> [16]. A Julia repository to reproduce the experiments herein described can be found in <https://github.gatech.edu/orozcom3/PhotoVI.jl>.

References

- [1] Lynton Ardizzone, Jakob Kruse, Sebastian Wirkert, Daniel Rahner, Eric W. Pellegrini, Ralf S. Klessen, Lena Maier-Hein, Carsten Rother, and Ullrich Köthe. Analyzing inverse problems with invertible neural networks. *arXiv preprint arXiv:1808.04730*, 2018.
- [2] Lynton Ardizzone, Carsten Lüth, Jakob Kruse, Carsten Rother, and Ullrich Köthe. Guided image generation with conditional invertible neural networks. *arXiv preprint arXiv:1907.02392*, 2019.
- [3] Muhammad Asim, Max Daniels, Oscar Leong, Ali Ahmed, and Paul Hand. Invertible generative models for inverse problems: mitigating representation error and dataset bias. In *International Conference on Machine Learning*, pages 399–409. PMLR, 2020.
- [4] Ashish Bora, Ajil Jalal, Eric Price, and Alexandros G. Dimakis. Compressed sensing using generative models. In *International Conference on Machine Learning*, pages 537–546. PMLR, 2017.
- [5] Emmanuel J. Candes, Justin K. Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences*, 59(8):1207–1223, 2006.
- [6] Andreas Hauptmann and Ben T. Cox. Deep learning in photoacoustic tomography: current approaches and future directions. *Journal of Biomedical Optics*, 25(11):1 – 46, 2020.
- [7] Shady Abu Hussein, Tom Tirer, and Raja Giryes. Image-adaptive gan based reconstruction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 3121–3129, 2020.
- [8] Amit P. Jathoul, Jan Laufer, Olumide Ogunlade, Bradley Treeby, Ben Cox, Edward Zhang, Peter Johnson, Arnold R. Pizzey, Brian Philip, Teresa Marafioti, et al. Deep in vivo photoacoustic imaging of mammalian tissues using a tyrosinase-based genetic reporter. *Nature Photonics*, 9(4):239–246, 2015.
- [9] Varun A. Kelkar, Sayantan Bhadra, and Mark A. Anastasio. Compressible latent-space invertible networks for generative model-constrained image reconstruction. *IEEE Transactions on Computational Imaging*, 7:209–223, 2021.
- [10] Ivan Kobyzev, Simon Prince, and Marcus Brubaker. Normalizing flows: An introduction and review of current methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [11] George Papamakarios, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *arXiv preprint arXiv:1912.02762*, 2019.
- [12] Ali Siahkoohi, Gabrio Rizzuti, Mathias Louboutin, Philipp A. Witte, and Felix J. Herrmann. Preconditioned training of normalizing flows for variational inference in inverse problems. In *3rd Symposium on Advances in Approximate Bayesian Inference*, 01 2021. AABI.
- [13] Albert Tarantola. *Inverse problem theory and methods for model parameter estimation*. SIAM, 2005.
- [14] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Deep image prior. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9446–9454, 2018.
- [15] Christina Winkler, Daniel Worrall, Emiel Hoogeboom, and Max Welling. Learning likelihoods with conditional normalizing flows. *arXiv preprint arXiv:1912.00042*, 2019.
- [16] Philipp Witte, Gabrio Rizzuti, Mathias Louboutin, Ali Siahkoohi, and Felix Herrmann. InvertibleNetworks.jl: A Julia framework for invertible neural networks, November 2020.
- [17] Minghua Xu and Lihong V. Wang. Photoacoustic imaging in biomedicine. *Review of scientific instruments*, 77(4):041101, 2006.

References follow the acknowledgments. Any choice of citation style is acceptable as long as you are consistent.