Photoacoustic Imaging with Conditional Priors from Normalizing Flows

Rafael Orozco[1], Ali Siahkoohi[1], Gabrio Rizzuti[2], Tristan van Leeuwen[2] and Felix J. Herrmann[1]
[1] Georgia Institute of Technology, [2] Utrecht University

Motivation

Photoacoustic imaging: a multi-physics medical imaging modality which takes advantage of high contrast caused by light and the high resolution of acoustic waves.

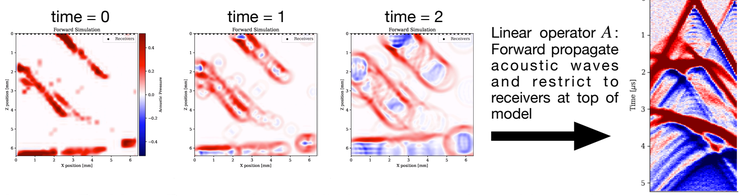

Linear operator A: Forward propagate acoustic waves and restrict to receivers at top of model

Inverse Problem: Due to noise and limited-view receivers this problem is ill-posed. The adjoint solution (time-reversal given by adjoint operator $A^\top$) contains artifacts that we will improve using a Bayesian framework.

Prior: The uncertainty of the solution is highly affected by measurement noise and data incompleteness (due 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 that leverages prior information learned by a generative network.

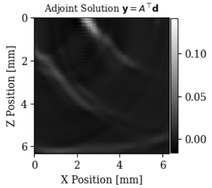

Goal: Explore the role of different priors in ill-posed photoacoustic imaging problems .

Methods

Variational Inference is a method which aims to convert posterior distribution inference $\mathbf{x} \sim p(\mathbf{x}\,|\,\mathbf{y})$ into an optimization problem. We are looking to optimize over parameters θ for some family of distributions q. We will chose this family to be a class of normalizing flows

$$q_\theta(x\,|\,y) \approx p(x\,|\,y)\,.$$

Normalizing Flows are a composition of invertible and learnable maps that transform a target density to a simpler base density (such as the standard normal). They make use of the change of variables formula for likelihood maximization training. Here we use conditional HINT [1] and train on data pairs $(\mathbf{x},\mathbf{y})$ where $\mathbf{x}$ is the ground truth image and $\mathbf{y}$ is time reversed data $\mathbf{y} = A^\top\mathbf{d}$. Our training objective is

$$\min_\theta \mathrm{KL}\left(p_{X,Y}||p_\theta\right) = \mathbb{E}_{\mathbf{x},\mathbf{y}\sim p_{X,Y}(\mathbf{x},\mathbf{y})} \frac{1}{2} \left\| f_\theta(\mathbf{x},\mathbf{y}) \right\|^2 - \log|\det \mathrm{J}_{f_\theta}(\mathbf{x},\mathbf{y})|\,.$$

Implicit deep prior: When faced with out-of-distribution data, we can assume that the previously trained flow is not as reliable as for in-distribution data. In this case, we propose a second phase in which we make use of the physics of the wave operator A and use the pre-trained generator as an implicit deep prior. This is similar to the approach in [2]. The solution to this is the maximum a posteriori (MAP). Formally, we solve

$$\min_{\mathbf{z}_x} \frac{1}{2}\|Af_\theta^{-1}(\mathbf{z}_x,\mathbf{z}_y) - \mathbf{d}\|_2^2 + \frac{\lambda^2}{2}\|\mathbf{z}_x\|_2^2\,, \quad \mathbf{z}_y = f_\theta^{\mathbf{z}_y}(A^\top\mathbf{d})\,. \qquad (1)$$

Results

Generative prior training: After training the conditional normalizing flow, we have access to samples from the posterior distribution.

Posterior Sampling from **in-distribution** data $\mathbf{d}$.

Note: high reconstruction quality in the conditional mean and the high standard deviation in deeper vessels.

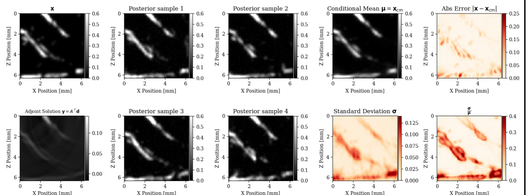

Posterior Sampling from **out-of-distribution** data $\mathbf{d}$.

Note: poor reconstruction quality in vertical vessels due to null space of forward operator.

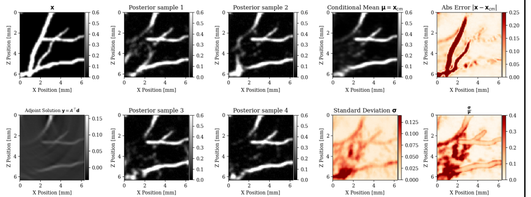

Second phase: We compare the reconstruction results for out-of-distribution data by solving the MAP problem in equation (1) with different choices for the prior. In particular, we consider a non-informative Gaussian prior, marginal learned priors and conditional priors (our proposal). A qualitative inspection and quality metrics indicate superior results for the conditional prior approach. Further comparison of the results regarding data fit is shown in right panel.

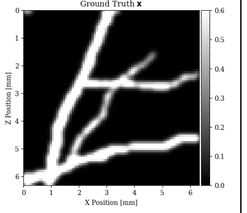

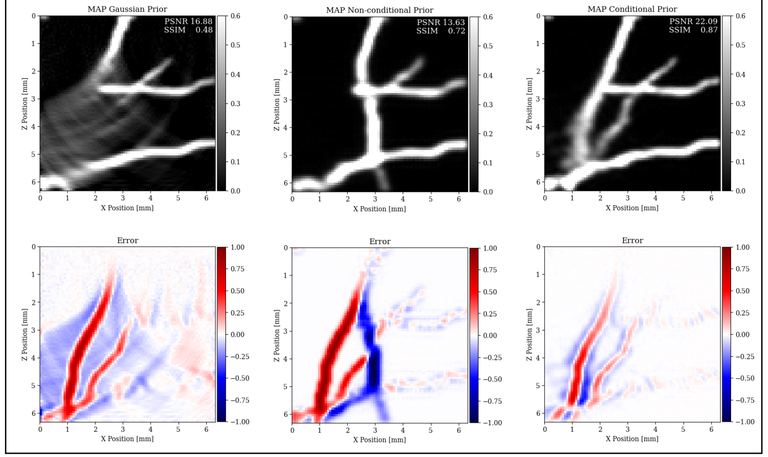

Conclusions and Future Work

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. Conventional regularization methods are typically too generic. A straightforward alternative is to use problem-specific information.

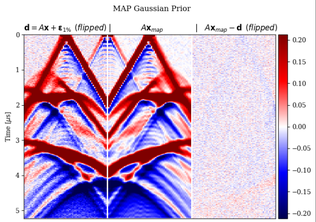

We proposed a regularization scheme that employs deep priors learned in an offline training phase. The generative model is trained on a dataset containing pairs of solutions and associated measurements, the goal is to learn the posterior distribution of the solution given some data. The primary scope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

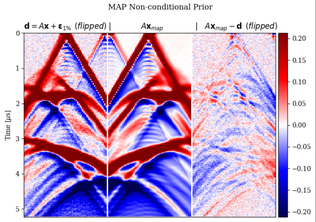

Our results suggest that the MAP estimate based on conditional deep priors is able to recover the vertical features of a blood vessel image (difficult to image since in null space of forward operator). Our results also show that the non-conditional deep prior is more prone to misplace features.

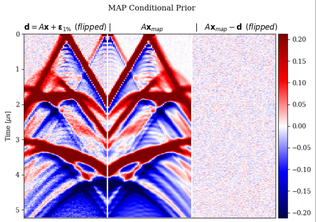

Data and Software

The dataset used in this work is a derivative of the ELCAP lung dataset prepared by [3]. The results presented here uses our Julia implementation of invertible network architectures: **https://github.com/slimgroup/InvertibleNetworks.jl** [4]

References

[1] Kruse, Jakob, et al. "HINT: Hierarchical invertible neural transport for density estimation and Bayesian inference." arXiv preprint arXiv:1905.10687 (2019).

[2] Asim, Muhammad, et al. "Invertible generative models for inverse problems: mitigating representation error and dataset bias." International Conference on Machine Learning. PMLR, 2020.

[3] Hauptmann, Andreas, and Ben T. Cox. "Deep Learning in Photoacoustic Tomography: Current approaches and future directions." *Journal of Biomedical Optics* 25.11 (2020): 112903.

[4] Philipp Witte, Gabrio Rizzuti, Mathias Louboutin, Ali Siahkoohi, and Felix Herrmann. InvertibleNetworks.jl: A Julia framework for invertible neural networks, November 2020.