

# Memory Efficient Invertible Neural Networks for 3D Photoacoustic Imaging

Rafael Orozco<sup>1</sup>, Mathias Louboutin<sup>1</sup> and Felix J. Herrmann<sup>1</sup>  
Georgia Institute of Technology

## Abstract

Photoacoustic imaging (PAI) can image high-resolution structures of clinical interest such as vascularity in cancerous tumor monitoring. When imaging human subjects, geometric restrictions force limited-view data retrieval causing imaging artifacts. Iterative physical model based approaches reduce artifacts but require prohibitively time consuming PDE solves. Machine learning (ML) has accelerated PAI by combining physical models and learned networks. However, the depth and overall power of ML methods is limited by memory intensive training. We propose using invertible neural networks (INNs) to alleviate memory pressure. We demonstrate INNs can image 3D photoacoustic volumes in the setting of limited-view, noisy, and subsampled data. The frugal constant memory usage of INNs enables us to train an arbitrary depth of learned layers on a consumer GPU with 16GB RAM.

## Problem Statement

In this work, we address the inverse problem of PAI given data  $y$ . This data is modeled with a linear forward model  $A$  which describes the propagation of a spatially varying initial acoustic wavefield  $x$ .

$$y = Ax + \varepsilon \quad \text{with} \quad \varepsilon \sim \mathcal{N}(0, \sigma^2 I) \quad (1)$$

We tackle here the realistic case of limited-view, noisy and subsampled data, which makes the inverse problem ill-posed. Multiple solutions  $x$  can explain the data  $y$ , therefore, this problem is typically solved in a variational framework where prior knowledge of the solution is incorporated by solving the minimization of the combination of an  $\ell_2$  data misfit  $L$  and a regularization term  $p_{\text{prior}}(x)$ . The success of this formulation hinges on users carefully selecting hyperparameters by hand such as the optimization step length and the prior itself, typically a multipurpose generic prior such as TV norm. On the other hand, recent machine learning work argues that we can learn the optimal step length and even a prior. We pursue this avenue but select loop-unrolled networks since this learned method also leverages the known physics of the problem encoded in the model  $A$ .

## Methods

### Loop-unrolled Networks:

Hauptmann et al. [2018] showed that known physical models can be combined with learned networks by using loop-unrolled networks. These networks emulate gradient descent where the  $i^{\text{th}}$  update is reformulated as the output of a learned network  $\Lambda_{\theta_i}$ :

$$x_{i+1}, s_{i+1} = \Lambda_{\theta_i}(x_i, s_i, \nabla_x L(A, x_i, y)) \quad (2)$$

where  $s$  is a memory variable and  $\nabla_x L$  is the gradient w.r.t the data misfit on observed data  $y$ . In our PAI case,  $A$  is a linear operator (discrete wave equation) and  $L$  is  $\ell_2$  norm data misfit so the desired gradient is  $\nabla_x L(A, x_i, y) = A^\top(Ax_i - y)$ . Thus, each gradient will require two PDE solves: one forward and one adjoint. These wave propagations embed the known physical model into the learned approach. Hauptmann et al. [2018] show that this method gives promising results on 3D PAI but note that they are limited to training shallow networks due to memory constraints. To improve on this method, we propose using an INN that is an invertible version of loop-unrolled networks.

## Invertible Neural Networks:

Our invertible neural network (INN) is based on the work of [Putzky et al., 2019], where they propose an invertible loop-unrolled method called invertible Recurrent Inference Machine (i-RIM). While typical loop-unrolled approaches have linear memory growth in depth Hauptmann et al. [2018], i-RIM has constant memory usage due to its invertible layers. A network constructed with invertible layers does not need to save intermediate activations thus trading computation cost required to recompute activations for extremely frugal memory usage. The memory gains allowed i-RIM authors to train a 480 layer model which was the state-of-the-art for the FASTMRI challenge when published Putzky et al. [2019]. For this work, we adapt i-RIM to Julia and make our code available alongside other invertible neural networks at [InvertibleNetworks.jl](#) Witte et al. [2020].

## Experiments and Results:

### Generating training dataset:

We perform supervised training on ground truth photoacoustic volumes taken from the blood vessels dataset of Bench et al. [2020]. For each 3D volume, we simulate photoacoustic data by forward propagating an initial source at the vessels and restricting the wavefield to a planar array of receivers at the top of the volume. We then subsample this data by a factor of 4 and add 10dB SNR Gaussian noise. For wave propagation, we use Devito a highly optimized DSL Louboutin et al. [2019] combined with JUDI Witte et al. [2019]. Propagating a wavefield of  $10^9$  voxels (100x300x300 volume for 1000 time steps) takes 35 seconds on a 4 core Intel Skylake CPU. We follow the prescription of Hauptmann et al. [2018] and pre-generate a set of training samples then fully train one loop-unrolled iteration  $\Lambda_{\theta_i}$  at a time. %As Hauptmann et al. [2018] explains, this training method changes our optimization from finding a end to end global minimum to instead being a greedy algorithm which finding the best possible reconstruction at each iteration.

### Results and conclusions:

After training one loop-unrolled iteration, our INN produces a better result than the baseline purely physical model based iterative method (Figure 1). For proper comparison, note that our INN result costs 2 PDE solves ( $2 \times 35$  seconds) and a forward pass on the trained INN (7 seconds), while the baseline costs 60 PDE solves ( $60 \times 35$  seconds). During INN architecture selection, we observed that the number of layers did not affect memory usage. Thus, we demonstrated that INNs provide an accurate, memory efficient method that enables fast 3D imaging for PAI.

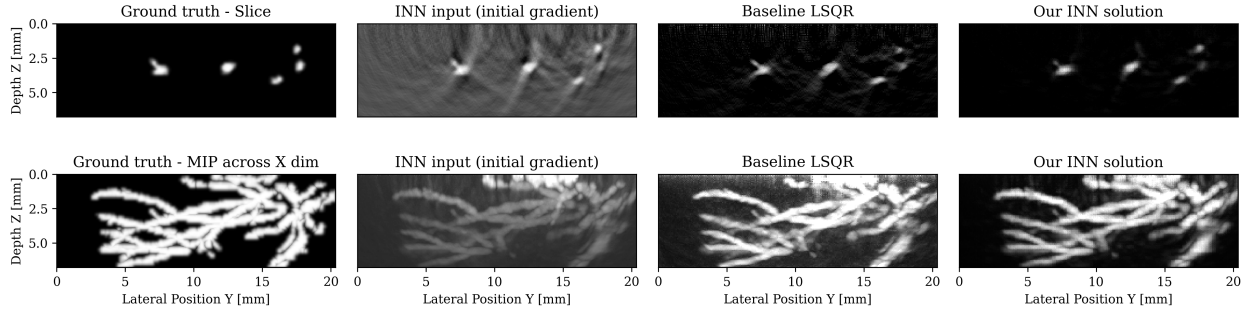


Figure 1: First row shows a slice through the 3D volume of: ground truth, the input to our INN (first gradient w.r.t data misfit on observed data  $y$ ), baseline LSQR solution after 30 iterations and our final INN reconstruction after 1 loop-unrolled iteration. Second row shows maximum intensity projection (MIP) across X dimension.

## References

- Ciaran Bench, Andreas Hauptmann, and Ben T Cox. Toward accurate quantitative photoacoustic imaging: learning vascular blood oxygen saturation in three dimensions. *Journal of Biomedical Optics*, 25(8):085003, 2020.
- Andreas Hauptmann, Felix Lucka, Marta Betcke, Nam Huynh, Jonas Adler, Ben Cox, Paul Beard, Sebastien Ourselin, and Simon Arridge. Model-based learning for accelerated, limited-view 3-d photoacoustic tomography. *IEEE transactions on medical imaging*, 37(6):1382–1393, 2018.
- Mathias Louboutin, Michael Lange, Fabio Luporini, Navjot Kukreja, Philipp A Witte, Felix J Herrmann, Paulius Velesko, and Gerard J Gorman. Devito (v3. 1.0): an embedded domain-specific language for finite differences and geophysical exploration. *Geoscientific Model Development*, 12(3):1165–1187, 2019.
- Patrick Putzky, Dimitrios Karkaloulos, Jonas Teuwen, Nikita Miriakov, Bart Bakker, Matthan Caan, and Max Welling. i-rim applied to the fastmri challenge. *arXiv preprint arXiv:1910.08952*, 2019.
- Philipp Witte, Gabrio Rizzuti, Mathias Louboutin, Ali Siahkoohi, and Felix Herrmann. Invertiblenetworks.jl: A julia framework for invertible neural networks, 2020.
- Philipp A Witte, Mathias Louboutin, Navjot Kukreja, Fabio Luporini, Michael Lange, Gerard J Gorman, and Felix J Herrmann. A large-scale framework for symbolic implementations of seismic inversion algorithms in julia. *Geophysics*, 84(3):F57–F71, 2019.