DeepGEM: Generalized Expectation-Maximization for Inverse Problems with Model Mismatch

Anonymous Author(s)
Affiliation
Address
email

Abstract

We study the problem of jointly solving an inverse problem coupled with correcting for model mismatch. Typically, inversion algorithms assume that a forward model, which relates a source to its resulting measurements, is known and fixed. Using collected indirect measurements and the forward model, the goal becomes to recover the source. When the forward model is unknown, or imperfect, artifacts due to model mismatch occur in the recovery of the source. We aim to recover the desired source with imperfect knowledge of the forward model. In this paper we propose DeepGEM, a variational Expectation-Maximization (EM) framework that can be used to solve for the unknown parameters of the forward model in an unsupervised manner. DeepGEM makes use of a normalizing flow generative network to efficiently capture complex posterior distributions, which leads to more accurate evaluation of the source's posterior distribution used in EM. We showcase the effectiveness of our DeepGEM approach by achieving strong performance on the challenging problem of blind seismic tomography, where we significantly outperform the standard method used in seismology. We also demonstrate the generality of DeepGEM by applying it to blind deconvolution.

1 Introduction

Physics-based inversion methods typically recover an unknown source from indirect measurements by assuming that the source and measurements are related via a known forward model [14, 7, 29]. For example, non-blind deconvolution algorithms often assume that a measured blurry image is related to its true sharp image via a known spatially-invariant blur kernel [9]; and traditional seismic inversion methods assume that the spatially-varying velocity of the Earth's interior is known *a priori* when solving for an earthquake's hypocenter [29]. However, these "known" forward models are generally idealized and ignore intricacies of the systems that are either hard to capture or simply unknown. Inversion algorithms with forward *model mismatch* result in biased reconstructions. For instance, bias is regularly seen in non-blind deconvolution results, where reconstruction artifacts are often present due to the use of an incorrect blur kernel [9].

Reducing model mismatch is key to reducing inversion bias and eliminating artifacts in recovered sources. When considering how learning can help, a natural first idea might be to learn a direct map from measurements to the desired source via supervised learning. However, such a "model-free" approach is generally not practical, due to the lack of available ground truth training data. For example, the blur kernel caused by handheld camera shake cannot be reproduced to get a training set of sharp-blurry pairs to train a deconvolution approach. In seismic tomography, synthetic earthquakes cannot be placed densely throughout the interior of the Earth to measure the ground response as a function of hypocenter location. An alternative approach, which we adopt, is to develop unsupervised methods that treat the true forward model as something that is unobserved and must be inferred. An additional consideration is that we must solve for the true forward model from a single dataset,

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

without knowledge of the true source that produces the measurements – this problem setting is commonly referred to as *blind inversion*.

In this paper, we propose Deep Generalized Expectation-Maximization (DeepGEM) for solving inverse problems that are plagued by model mismatch. Using the indirect measurements as input, DeepGEM jointly estimates the source and forward model that together produce the observed measurements. DeepGEM is a variational inference based framework that makes use of deep learning machinery to easily capture and optimize complex probabilistic distributions that cannot be easily integrated in analytic Expectation-Maximization (EM) solutions. Our proposed framework is generic and can be applied to blind inversion problems described by differentiable forward models. In Section 4 we showcase the effectiveness of our DeepGEM approach on the challenging problem of blind seismic tomography, where we significantly outperform methods used in seismology. We also demonstrate the generality of DeepGEM by applying it to blind deconvolution in Section 5.

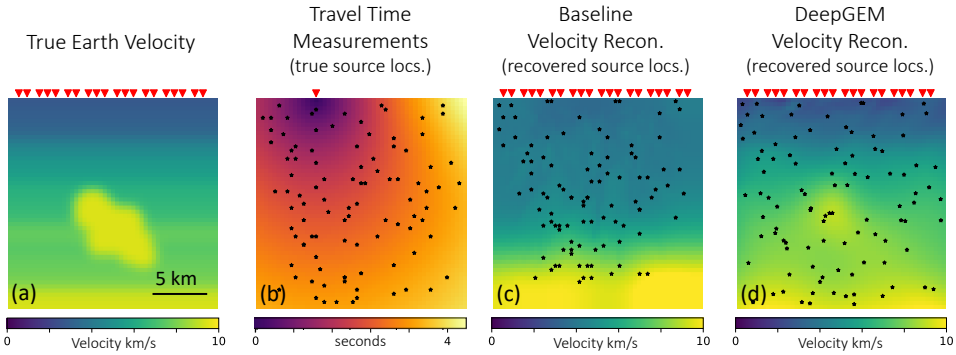

Figure 1: **DeepGEM applied to blind seismic tomography.** (a) A simulated cross section of the Earth's interior (velocity structure), along with the locations of receivers on the surface (red triangles) that collect measurements. (b) The time it takes for a wave traveling from a source below the surface to reach the specified receiver is visualized for each location in the region of interest. The overlaid dots represent the true locations of simulated earthquakes and indicate the measured travel times that constrain optimization. (c) The subsurface velocity reconstruction obtained using a baseline approach optimized with the help of a seismologist. Note that the bright anomaly is missing from this reconstruction. Overlaid dots represent the inferred earthquake locations. (d) DeepGEM reconstructed subsurface velocity and inferred earthquake locations. Note that DeepGEM is able to accurately recover the gradient of the velocity field as well as partially recover the central anomaly.

2 Background and related work

The joint optimization of a forward model with source recovery is a very challenging ill-posed problem, leading to many possible solutions that are hard to disambiguate. For example, in blind deconvolution the blurry image observed can be equivalently explained by a sharp source image convolved with an extended kernel or a blurry source image convolved with a impulse kernel; to prefer one solution over another, additional information, such as image priors, must be considered. Previous work on inversion in poorly characterized systems (e.g., model mismatch) focused on limited contexts, such as spatially-invariant blind deconvolution [20, 21, 13, 10] and CT with a simple rotational error [33]. These methods tend to be highly specialized to each application domain, and cannot be easily generalized. In contrast, we propose a flexible model-based Bayesian framework that can be applied across multiple differentiable inversion problems subject to model mismatch.

2.1 Model-based Bayesian inversion

In model-based inversion, unobserved sources x and observed measurements y are related through a forward model: $y = f(x)$. When model mismatch is suspected, one can parameterize the forward model as $f_\theta(x)$ and then solve for the true model parameters θ^*. A common approach is to solve a *maximum a posteriori* (MAP) objective: either $\text{MAP}_{\theta,x}$ or MAP_θ.

MAP$_{\theta,x}$ solves for the optimal point estimate of the pair $\{\hat{\theta}, \hat{x}\}$ that maximizes a joint objective:

$$\{\hat{\theta}, \hat{x}\} = \arg\max_{\theta,x} \left[\log p(\theta, x|y)\right] = \arg\max_{\theta,x} \left[\log p(y|\theta, x) + \log p(\theta) + \log p(x)\right]. \quad (1)$$

In practice, formal probabilistic definitions of $p(\theta)$ and $p(x)$ are often unknown and replaced with regularization terms (e.g., total variation, sparsity) [21, 20]. Although $\text{MAP}_{\theta,x}$ provides a straightforward approach to solve for θ, the energy landscape (1) is typically poorly behaved due to the ill-posed nature of the problem [21]. As a result, optimization is likely to get stuck in a (bad) local minimum.

MAP_θ attempts to smooth the energy landscape by reducing the number of parameters that must be optimized. This is done by solving for parameters of the forward model, θ, that perform best under the full volume of possible x interpretations:

$$\hat{\theta} = \arg\max_{\theta} \left[\log p(\theta|y)\right] = \arg\max_{\theta} \left[\log\left(\int_x p(y|\theta,x)p(x)dx\right) + \log p(\theta)\right]. \tag{2}$$

Since this marginalization integral is often intractable, Expectation-Maximization (EM) algorithms have long been used for solving MAP_θ efficiently [11]. EM is an iterative algorithm that alternates between: "E"-Step) calculating the posterior of x conditioned on the current estimated forward model parameters $\theta^{(t-1)}$; and M-Step) updating θ to maximize the expected value of the log likelihood:

$$\theta^{(t)} = \arg\max_{\theta} \left[\mathbb{E}_{x\sim p(x|y,\theta^{(t-1)})}\left[\log p(y|\theta,x)\right] + \log p(\theta)\right]. \tag{3}$$

The advantage of MAP_θ over $\text{MAP}_{\theta,x}$ optimization for the blind deconvolution problem was described in [21] and its success demonstrated via EM optimization in [20]. However, it is important to note that evaluating the expectation in (3) over complex distributions is often intractable. For instance, the authors of [20] were forced to restrict the posterior distribution to a Gaussian distribution. Alternatively, stochastic EM methods [5, 8] bypasses the need to evaluate the expectation directly, approximating it by sampling the distribution. In this paper, we solve MAP_θ using complex distributions parameterized by deep neural networks.

Forward model parameterization Model-based inversion requires that the parametric form of $f_\theta(x)$ is well matched with the true forward model, which is not always known. Alternatively, neural networks can be used approximate the forward model, where θ represents the network weights. This setup is flexible, in that it can be used to approximate arbitrarily complex forward models, with the downside that it is often not interpretable when the parameters have no physical meaning. In this work, we develop and make use of interpretable, physically-motivated deep neural networks to parameterize $f_\theta(\cdot)$ for the problems of blind seismic tomography and blind deconvolution.

2.2 Deep variational distributions

We are interested in solving blind inverse problems using the EM algorithm to optimize the MAP_θ objective. Optimizing in this fashion requires the use of the posterior distribution $p(x|y,\theta^{(t-1)})$ in evaluating Eq. 3. As inverse problems are often ill-posed, we expect that the posterior distribution of source x conditioned on the forward model parameters θ is likely to be complex and perhaps even multi-modal. Therefore, it is important to be able to parameterize a flexible family of distributions to best estimate this conditional posterior.

Deep Probabilistic Imaging (DPI) [31] uses a normalizing flow-based generative model, $G_\phi(\cdot)$, to solve for the uncertainty of an unknown source x given a fixed forward model $f(x)$ and measurements y. DPI solves the variational objective:

$$\hat{\phi} = \arg\min_{\phi} \text{KL}(q_\phi(x)||p(x|y)) = \arg\min_{\phi} E_{x\sim q_\phi(x)}\left[-\log p(y|x) - \log p(x) + \log q_\phi(x)\right], \tag{4}$$

where $q_\phi(x)$ is the implicit distribution defined by $G_\phi(z)$ for $z \in \mathcal{R}^{|x|} \sim \mathcal{N}(0,1)$ and $\log p(y|x) \propto ||y - f(x)||_2 + c$ when there exists *i.i.d.* Gaussian noise on the measurements, y. After inference, $q_\phi(x)$ can be efficiently sampled by evaluating $G_\phi(z)$ for a Gaussian sample z.

The DPI variational objective is equivalent to the Variational Autoencoder [16, 27] objective, except with a fixed decoder, $f(x)$. In practice, vanilla VAEs constrain the posterior to be a Gaussian distribution, relying on the reparameterization trick for tractable optimization. Alternatively, DPI uses flow-based networks to efficiently sample and directly evaluate $q_\phi(x)$ [17, 12]. DPI's use of a flow-based network allows for complicated and multi-modal posterior distributions constrained only by the space of possible invertible network architectures. Our proposed DeepGEM approach utilizes similar tools to model flexible distributions over x, while simultaneously learning the forward model parameters θ.

3 Methods

We propose a deep variational EM approach (DeepGEM) that optimizes the MAP_θ objective in Eq. 2 to recover the parameters of a forward model $f_\theta(x)$ using only measurements y. Once learned, the updated forward model can then be used to estimate the posterior distribution of the unknown source, x. DeepGEM iterates between two stages that are inspired by the standard EM algorithm: (1) an "E"-step that learns a variational distribution, $q_{\phi^{(t)}}(x)$, to approximate the posterior distribution of x given the current forward model parameters $\theta^{(t-1)}$, and (2) an M-step (refer to Eq. 3) that solves for $\theta^{(t)}$ that maximizes the expected value of the log likelihood function of θ, with respect to the posterior distribution $q_{\phi^{(t)}}(x)$ estimated in the prior "E"-step. Each step is alternated and solved to convergence.

3.1 "Expectation" step ("E"-step)

The goal of DeepGEM's "E"-step is to solve for the posterior distribution $p(x|y, \theta^{(t-1)})$ that facilitates optimizing Eq. 3. Because this posterior distribution can be very complex, and even multi-modal, we propose to use a flexible variational approach to learn the parameters ϕ of an approximate posterior distribution $q_\phi(x)$. The variational distribution $q_\phi(x)$ can then be used to evaluate Eq. 3 via efficient sampling.

Using DPI, we solve for a flexible variational distribution, $q_\phi(x)$ that well approximates the posterior distribution $p(x|y, \theta^{(t-1)})$. DPI uses a normalizing flow network, $G_\phi(z)$, with input $z \in \mathbb{R}^{|x|}$, where $x = G_\phi(z) \sim q_\phi(x)$ when $z \sim \mathcal{N}(0, \mathbb{1})$. Normalizing flow networks allow for exact computation of the log-likelihood $\log q_\phi(x)$ needed to solve

$$\begin{aligned}\phi^{(t+1)} &= \arg\min_\phi \text{KL}(q_\phi(x)||p(x|y, \theta^{(t-1)})) \\ &\approx \arg\min_\phi \frac{1}{N}\sum_{n=1}^{N}[-\log p(y|\theta^{(t-1)}, x_n) - \log p(x_n) + \log q_\phi(x_n)] \\ &\quad \text{for} \quad x_n = G_\phi(z_n), \quad z_n \sim \mathcal{N}(0, \mathbf{1}),\end{aligned} \tag{5}$$

(as derived from Eq. 4) where $\log p(x)$ is a prior on the source and $\log p(y|x, \theta^{(t)})$ is the data likelihood. When assuming the measurements y experience *i.i.d* additive Gaussian noise with standard deviation σ_y, $\log p(y|\theta^{(t)}, x_n) = \frac{1}{2\sigma_y^2}\left\|y - f_{\theta^{(t)}}(x_n))\right\|^2 + c$.

3.2 Maximization step (M-step)

The goal of DeepGEM's M-step is to use the parameterized approximate posterior distribution, $q_{\phi^{(t)}}(x)$, from the prior "E"-step to update θ, the parameters of the unknown forward model $f_\theta(\cdot)$. This is achieved by sampling from the learned normalizing flow network, $G_{\phi^{(t)}}(\cdot)$, to stochastically solve Eq. 3 :

$$\theta^{(t)} \approx \arg\max_\theta \left[\frac{1}{N}\sum_{n=1}^{N}[\log p(y|\theta, x_n)] + \log p(\theta)\right] \quad \text{for} \quad x_n = G_{\phi^{(t)}}(z_n), \;\; z_n \sim \mathcal{N}(0, \mathbb{1}), \tag{6}$$

where $p(\theta)$ is a prior on the forward model. This prior can be used to encourage the forward model parameters to remain close to an initial model $\tilde{\theta}$ by defining $\log p(\theta) \propto ||\theta - \tilde{\theta}||_2 + c$.

4 Case study: blind seismic tomography

Two fundamental seismic inverse problems are spatially localizing an earthquake's hypocenter (also referred to as source localization) and tomographic reconstruction of the Earth's subsurface [29]. These problems are interconnected: source localization relies on knowing how fast waves propagate through different regions of the Earth's interior, referred to as the Earth's *velocity* structure. However, in standard seismological practice, due to difficulties in solving these problems jointly, they are generally treated separately: source localization is performed initially using oversimplified velocity models [29], and then the tomography problem is performed by taking those best-fitting hypocenters

as ground truth [25]. This approach typically results in the need to over-regularize the inverse problem by smoothing out high-frequency information [2], and can only be improved by carefully incorporating other forms of information such as waveform-derived quantities [15, 22, 4, 23] or by performing costly experiments such as controlled explosions. In contrast, we demonstrate our DeepGEM approach on blind seismic tomography, solving for the subsurface velocity (parameterized by θ) when the source hypocenters, x, are unknown. Measurements, y, used to constrain the inverse problem are the time it takes for the first wave to propagate from its source to a receiver on the Earth's surface, referred to as a *travel time* measurement (refer to Fig. 1).

4.1 Seismic tomography background

Physics of earthquake source localization: The earthquake source location, also called the hypocenter, is the location where the earthquake nucleates [29]. The source location can be triangulated using travel times from multiple receivers. However, when there are very few receivers (< 3 in 3D, < 2 in 2D), the source localization is ill-posed and there exists multiple equally optimal solutions.

Physics of travel time tomography: Travel time tomography is a method for reconstructing the velocity structure using the absolute arrival times of earthquake waves from the earthquake to the receiver [29]. Given perfect knowledge of the earthquake locations x and receivers r at every position in the ground, the exact velocity can be computed by solving the Eikonal equation $V(r) = \|\nabla_r T(x,r)\|_2^{-1}$, where $T(x,r)$ is the travel time from an earthquake to a receiver. The solution to the Eikonal equation is physically sound; however, seismologists often use simplifications, such as straight ray tomography, for efficiency [2].

Deep Learning for travel time tomography: EikoNet [30] implicitly solves the Eikonal equation [32] by learning a mapping from a source-receiver pair (x,r) to its associated travel time, learning θ such that $f_\theta(x,r) \approx T(x,r)$ where $T(x,r)$ is the true travel time. The velocity structure can then be extracted from the learned EikoNet by solving the Eikonal equation through automatic differentiation. In particular, the computed velocity $V(s)$ at location s is $\|\nabla_s f_\theta(x,s)\|_2^{-1}$. Note that this computed velocity depends on the source location x. Ideally, the velocity should be invariant to the source location; however, in practice, this is only true when EikoNet is trained with densely-sampled (x,r) pairs. Additionally, the travel time from a source to a receiver, $T(x,r)$, should be identical to the travel time from a receiver to a source, $T(r,x)$, but remains unconstrained by EikoNet.

4.2 DeepGEM setup for blind seismic tomography

For blind seismic tomography, we parameterize the forward model using a modification of EikoNet, $f_\theta(x,r)$, with unknown source location x and known receiver location r as inputs and $y \approx T(x,r)$, the absolute travel time between x and r, as output. In order to solve Eqs. 5 and 6 for an updated forward model, we must define priors $p(x)$ and $p(\theta)$. The prior over source locations, $p(x)$, is often well defined, typically a Gaussian distribution $\mathcal{N}(\bar{x}, \sigma_x)$ with a standard deviation of $\sigma_x \sim 2$ km.

We construct a prior over the forward model, $p(\theta)$, that encourages EikoNet to learn a velocity close to $\tilde{V}(s)$. Additionally, as discussed above, there are constraints specific to seismic tomography that are not explicitly enforced through EikoNet's architecture: (1) velocity reconstruction invariance with respect to the source location, and (2) travel time symmetry between sources and receivers. We

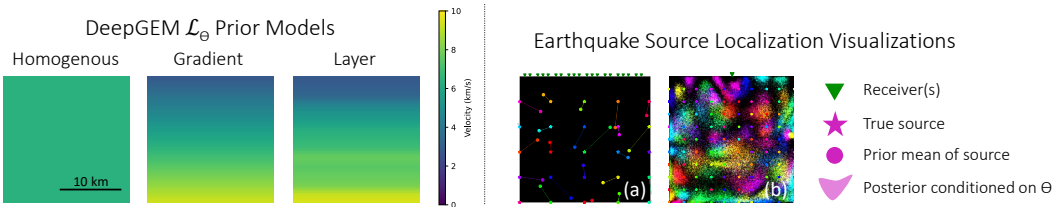

Figure 2: (Left) A visualization of the homogeneous, gradient, and layer models used in $p(\theta)$'s $\mathcal{L}_\theta$ term. (Right) Visualizations used to describe source configurations and earthquake source posteriors. (a) Visualizations of the true and initialized source locations are plotted as stars (true) connected to circles, which indicate the expected source locations according to $p(x)$. Note that the expected source locations deviate significantly from the true locations. (b) Visualizations of the learned posterior distribution, $q_\phi(x|y,\theta)$, for each source are plotted as colored histograms and overlaid with stars (of the same color) indicating the true source location.

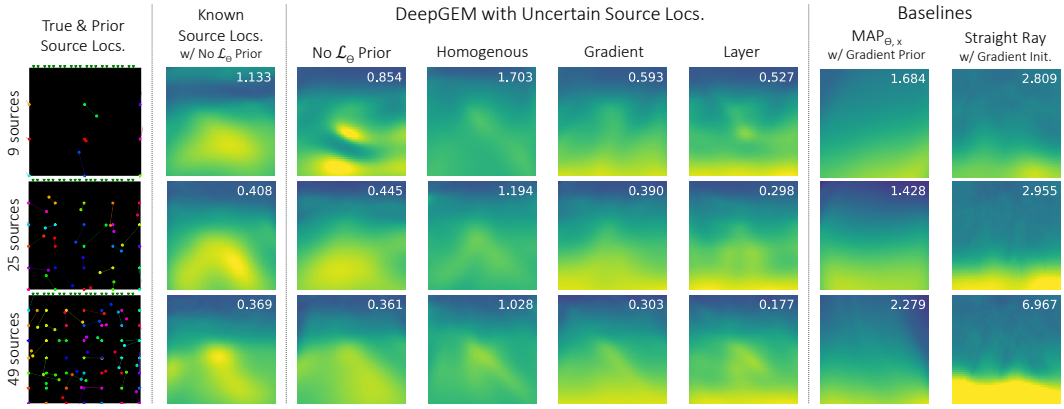

Figure 3: **DeepGEM reconstructions significantly outperform baselines, and improve with more sources and better $\mathcal{L}_\theta$ prior models.** Each row corresponds to the reconstructions obtained from a single noisy observation of travel times, where the true Earth velocity is shown in Fig. 1(a). Results shown are simulated using 20 surface receivers and a varying number of sources (9, 25, and 49) in a uniform grid. Columns 3-6 show DeepGEM results obtained using different $\mathcal{L}_\theta$ priors. Note that results improve as the $\mathcal{L}_\theta$ prior becomes closer to the true velocity structure, and as the number of sources increases. As a reference, column 2 shows the velocity reconstruction obtained using DeepGEM under fixed, perfectly known source locations. Columns 7-8 show results obtained by the baseline approaches. The velocity reconstruction MSE is included in the top right of each reconstruction. DeepGEM substantially outperforms both straight ray and $\text{MAP}_{\theta,x}$ baselines.

augment the prior $p(\theta)$ to include these constraints, implemented as soft constraints:

$$\log p(\theta) = \lambda_\theta \overbrace{\sum_{\substack{r\in\mathcal{R},\\ s\in\mathcal{S}}} ||\tilde{V}(s) - V_r(s)||_2}^{\mathcal{L}_\theta} + \lambda_V \overbrace{\sum_{\substack{r_i,r_j\in\mathcal{R},\\ s\in\mathcal{S}}} \left\|V_{r_i}(s) - V_{r_j}(s)\right\|_2}^{\mathcal{L}_V} + \lambda_T \overbrace{\sum_{\substack{r\in\mathcal{R},\\ s\in\mathcal{S}}} \|T(s,r) - T(r,s)\|_2}^{\mathcal{L}_V}. \tag{7}$$

The velocity constraint is represented through $\mathcal{L}_V$ and travel time constraint through $\mathcal{L}_T$, with corresponding hyperparameters λ_V and λ_T. $\mathcal{S}$ is a set of points sampled uniformly from the region of interest, $\mathcal{R}$ is the set of all receiver locations, and $V_r(s) = ||\nabla_s f_\theta(r,s)||_2^{-1}$.

Implementation details: In our experiments we define a realistic prior over unknown source locations as $p(x) = \mathcal{N}(\bar{x}, \sigma_x)$, where $\bar{x} \sim \mathcal{N}(x, \sigma_x)$ and $\sigma_x = 2$ km. We assume the measurements y are sampled from a Gaussian distribution with mean $\tilde{y}$ – true travel times computed using the package `eikonalfm` [28, 32] – and a realistic standard deviation of $\sigma_y = 0.2$ seconds. To simulate real world passive tomography, we assume receivers are located only along the surface (the top edge of the 2D image) of the region of interest, which is 20 km $\times$20 km, and sources can be anywhere within the region of interest.

The posterior distribution of the source locations, x, is estimated using a Real-NVP network $G_\phi(\cdot)$ with 16 affine coupling layers. An updated EikoNet (described in the supplemental material) has been modified to parameterize $f_\theta(x)$. This EikoNet is pretrained with samples from the prior $p(x)$ as input and the simulated travel time measurements as output. We use Adam as the optimizer [34] with a batch size of 32 and an E-step learning rate of 1e-3 and M-step learning rate of 5e-5. Hyperparameters $(\lambda_T, \lambda_V, \lambda_\theta)$ were empirically chosen by inspecting the loss on a grid search over hyperparameters. Results presented have been run with 10 EM iterations, each with 800 "E"-step epochs and 2000 M-step epochs. Each DeepGEM model takes $\sim$6 hours on a NVIDIA Quatro RTX 5000, for a total of $\sim$500 hours of development time.

4.3 Results

Comparison to Baseline Approaches: We compare results from DeepGEM to results obtained using a baseline run by a seismologist. This iterative baseline alternates between source inversion and straight ray tomography, and is the standard approach used for blind tomography. Further detail on this baseline is provided in the supplemental material. The gradient model shown in Fig. 2 is used to perform the initial source inversion for this baseline; nonetheless, we find that the solution quickly diverges. Therefore, we rely on the expertise of the domain expert to decide when to terminate the

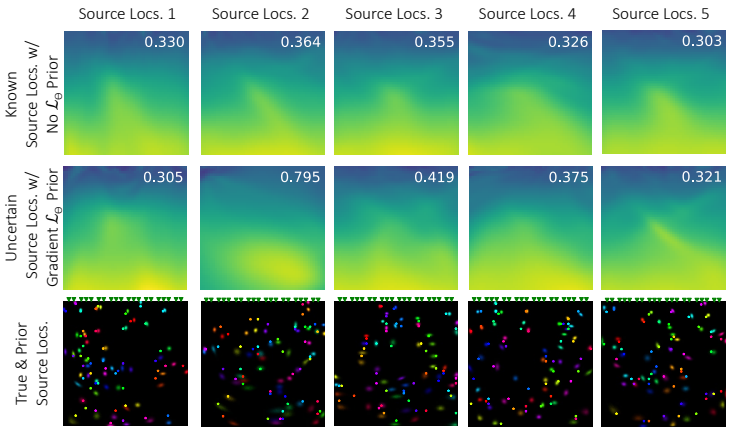

Table 1: Velocity and source localization error on random source configurations

Source Count	Velocity Error (km/s)	Source Loc. Error (m)
9	0.52 ± 0.063	62.88 ± 21.77
25	0.42 ± 0.043	36.62 ± 4.93
49	0.44 ± 0.181	33.78 ± 4.56
100	0.27 ± 0.021	30.38 ± 1.43

Figure 4: **DeepGEM consistently recovers prominent features across various source configurations.** Row 2 shows DeepGEM results obtained using different random source configurations, where the true Earth velocity is shown in Fig. 1(a). As a reference, row 1 shows the velocity reconstruction obtained using DeepGEM under fixed, perfectly known source locations. As can be seen by the table, both velocity and localization error decrease with an increasing number of sources in the region of interest. Each mean and standard deviation is computed using 5 random realizations of true source configurations.

optimization. As seen in Fig. 3, DeepGEM consistently outperforms this human-in-the-loop baseline across all source configurations. In Fig. 3 we also compare with a $\text{MAP}_{\theta,x}$ solution. $\text{MAP}_{\theta,x}$ is consistently outperformed by the DeepGEM MAP_θ approach across all source configurations.

Sensitivity to $\mathcal{L}_\theta$ Prior Choice: We evaluate DeepGEM's recovery of the true velocity structure shown in Fig. 1 using one of three different $\mathcal{L}_\theta$ priors shown in Fig. 2: homogeneous, gradient, and layer, as well as $\mathcal{L}_\theta = 0$. The homogeneous model takes on value of 6.419 km/s, the average velocity value of the true velocity structure. The gradient captures the increasing velocity as depth increases, and the layer model represents the true model without the added anomaly. As shown in Fig. 3, the mean squared error tends to decrease with the gradient and layer $\mathcal{L}_\theta$ prior models, which are closer to the true velocity structure.

Sensitivity to Source Configuration: We evaluate results with sources that are both uniformly and randomly spaced throughout the region of interest. Improved performance is expected when the number of sources is increased and/or when sources are well distributed. To better understand DeepGEM's performance we introduce an ablation, where the velocity structure is learned by training $f_\theta(\cdot)$ with access to perfect source locations x (i.e., $p(x)$ is a delta function). As is shown in Fig 3, even when training with true source locations, the anomaly is not well resolved until $\sim$49 sources.

Fig. 3 shows one realization of DeepGEM results obtained from different counts of uniformly spaced sources. As expected, the MSE between the reconstructed and true velocity structure tends to decrease as the number of sources increases. Fig. 4 shows results obtained using five randomly generated source configurations, with 49 sources each. These results demonstrate that, although the reconstructed velocity structure is somewhat sensitive to the underlying source configuration, the primary features of the true velocity can still be recovered in all cases. Velocity and source localization error obtained for different random configurations are shown in the accompanying table.

Sensitivity to Number of Receivers: In the case of a single receiver, there exists an entire ring of source locations that result in the same travel time measurement. Fig. 5 contains results from this challenging one-receiver setting using DeepGEM with/without measurement noise and with/without known source locations. Since there is only one receiver, the velocity model is able to easily overfit. However, perhaps surprisingly, artifacts are more severe when source locations are perfectly constrained. These artifacts are caused by the velocity model overfitting to noise in the travel time measurements, and are substantially reduced when noise-free travel time measurements are used. Note that the recovered source location posterior $q_\phi(x)$ obtained by the "E"-step is non-Gaussian.

Sensitivity to Velocity Structure: In Fig. 6, DeepGEM is tested on randomly generated velocity fields, each generated from a Gaussian random field (GRF) described in the supplemental material. These results show that DeepGEM works well at recovering the primary features across a variety

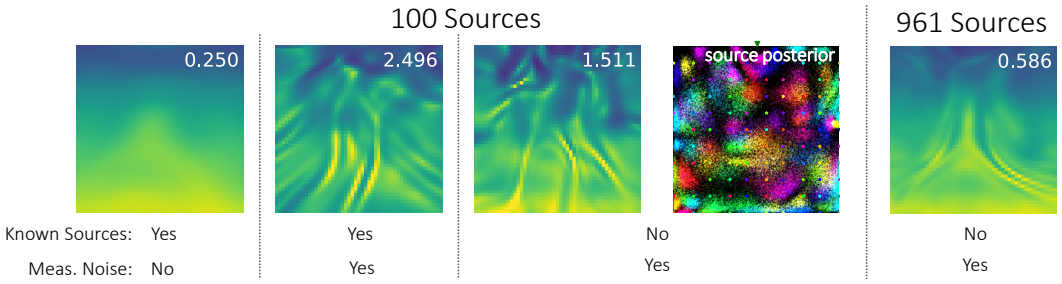

Figure 5: **DeepGEM recovery with a single receiver.** Velocity reconstructions shown in columns 3 and 5 demonstrate that DeepGEM is able to learn some of the true velocity features (see Fig. 1(a)), even when limited to measurements from a single receiver. However, these reconstructions show clear signs of overfitting to the measurement data. This is demonstrated by observing that the DeepGEM reconstruction with perfectly known source locations is significantly better with a single receiver when no measurement error is included on the measurements (comparing columns 1 and 2).

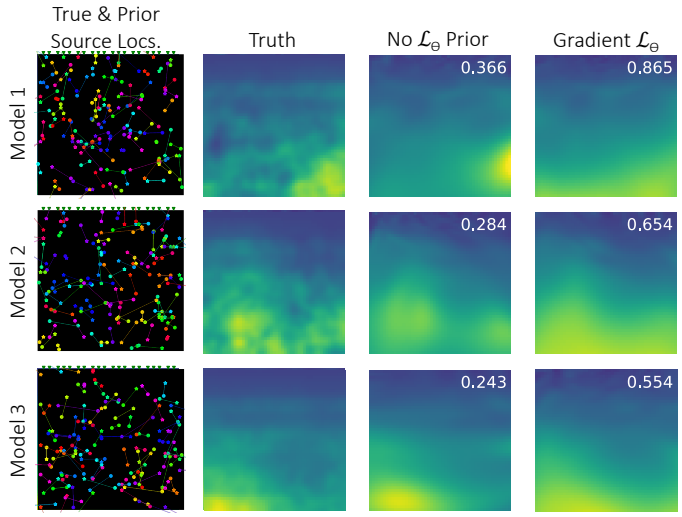

Table 2: Velocity and source localization error on random velocity fields

$\mathcal{L}_\theta$	Velocity Error (km/s)	Source Loc. Error (m)
None (& Known Source Locs.)	0.57 ± 0.25	NA
None	0.51 ± 0.30	24.40 ± 7.18
Homogenous	0.94 ± 0.45	21.37 ± 5.34
Gradient	1.07 ± 0.53	26.20 ± 6.46

Figure 6: **Performance of DeepGEM recovery on random velocity fields.** Ten random velocity fields were drawn from a GRF-based distribution and used to simulate travel time measurements with 20 receivers and 100 randomly placed sources. (left) Reconstructions obtained for 3 of these configurations are shown. (right) A table lists the mean and standard deviation of velocity and source error obtained across the ten models, each recovered using different $\mathcal{L}_\theta$ priors.

of different velocity structures. As compared to when $\mathcal{L}_\theta = 0$, the $\mathcal{L}_\theta$ gradient prior biases the reconstruction towards smoother velocity structures. The accompanying table contains error statistics for ten different randomly generated velocity fields.

5 Case study: blind deconvolution

We apply DeepGEM to the problem of blind deconvolution to further demonstrate the generality of our approach. Blind deconvolution is a classic ill-posed imaging problem that aims to reconstruct a sharp image from a blurry image with an unknown PSF [18, 13, 24, 20, 21, 19, 6, 1, 26]. Blurry images, caused by handheld camera shake, can be modeled using a single spatially-invariant blur kernel:

$$y = x * k_\theta + \varepsilon \quad \text{for} \quad \varepsilon \sim \mathcal{N}(0, \sigma), \tag{8}$$

where y is the blurry image, $*$ represents a 2D convolution, x is the true sharp image, k_θ is the spatially invariant blur kernel, and ε is additive Gaussian noise.

5.1 DeepGEM setup for blind deconvolution

For blind deconvolution, we parameterize the forward model $f_\theta(\cdot)$ using a deep network consisting of multiple convolution layers without non-linear activation, as proposed in [3]. Multiple convolutional layers without activation simply overparameterizes a linear blur kernel, which has been emperically shown to produce multiple good minima that are easier to converge to. To ensure the blur kernel is

non-negative and volume preserving, we use a Softmax layer and normalize the kernel. For an n layer network with weights θ_i for $i = 1, ..., n$, the resulting parameterized forward model is:

$$f_\theta(x) = x * \hat{k}_\theta = x * \left[\frac{\text{Softmax}(\theta_1 * \theta_2 * ...\theta_n)}{\|\text{Softmax}(\theta_1 * \theta_2 * ...\theta_n)\|_1} \right] \tag{9}$$

Implementation details: We demonstrate DeepGEM using simple Total Variation (TV) regularization in place of $\log p(x)$. We assume a Gaussian prior on the noise on the blurry measurements where $\varepsilon \sim \mathcal{N}(0, 0.01)$ as well as a sparsity prior on the reconstructed kernel through an $\ell_{0.8}$ soft constraint. The posterior distribution of the sharp image, x, is estimated using a Real-NVP network $G_\phi(\cdot)$ with 16 affine coupling layers. We use 5 convolution layers to parameterize k_θ. We use Adam as the optimizer [34] with a batch size of 64 and an E-step learning rate of 5e-4 and M-step learning rate of 1e-4. Hyperparameters, weights used for sparsity and TV priors, were empirically chosen by a grid search over hyperparameters. Results presented have been run with 10 EM iterations, each with 400 "E"-step epochs and 400 M-step epochs, which takes $\sim$1 hour on a NVIDIA Tesla V100.

5.2 Results

In Fig. 7 we show results from DeepGEM on two different blurry images [21]. The blurry image in Fig. 7(a) exhibits artifacts from the blur kernel in the form of repeated features. This is due to the blur kernel being similar to the sum of two delta point spread functions. The recovered source image is much sharper, containing clear text with minimal ringing artifacts. The reconstructed kernel roughly matches the true kernel's shape with two lobes along the same diagonal. In Fig. 7(b), ringing in the blurry image is also reduced. Although the reconstructed kernel is not located spatially at the same location as the true kernel, this does not significantly harm reconstruction; since the kernels are shift-invariant, the reconstructed image and learned kernel can both be shifted such that they reconstruct the same blurry image. Please refer to the supplemental material for more results.

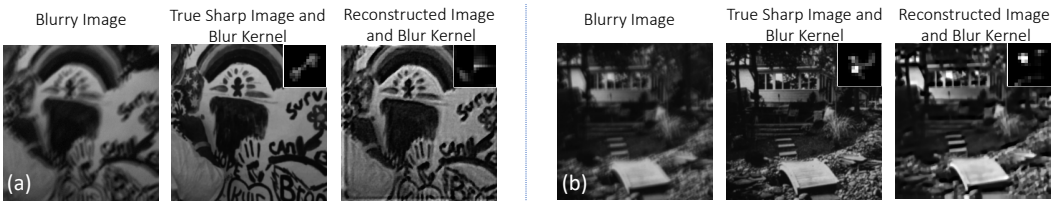

Figure 7: Blurry measured images (columns 1 and 4) generated using the true sharp image and blur kernel (shown in columns 2 and 5). Reconstructed sharp image and the corresponding inferred blur kernel from DeepGEM (columns 3 and 6).

6 Conclusions

In this paper we present DeepGEM, a deep probabilistic framework for tackling blind inverse problems through estimation of the forward model. DeepGEM achieves strong performance in the task of joint seismic tomography and earthquake source localization, substantially outperforming standard approaches currently being used in seismology on synthetic data. The proposed framework is flexible and can be applied to different applications that require estimation or fine tuning of forward model parameters. We demonstrate this flexibility by also applying the approach to the simple, but challenging, blind deconvolution problem. Future work includes applying this method to real seismic data, extending to other applications, and incorporating data-driven priors. Our results highlight the benefit of blending physically sound model-based techniques with learning machinery for inversion with model mismatch.

Broader Impacts. DeepGEM can be used to solve for a system's model mismatch, which can then help improve our understanding of complex physical systems. However, this tool is not trustworthy enough for safety-critical systems. Nonetheless, this approach can benefit society through a better understanding of fundamental science and advanced earthquake prediction models via seismic imaging.

References

[1] Raied Aljadaany, Dipan K Pal, and Marios Savvides. Douglas-rachford networks: Learning both the image prior and data fidelity terms for blind image deconvolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10235–10244, 2019.

[2] M. P. Barmin, M. H. Ritzwoller, and A. L. Levshin. *A Fast and Reliable Method for Surface Wave Tomography*, pages 1351–1375. Birkhäuser Basel, Basel, 2001.

[3] Sefi Bell-Kligler, Assaf Shocher, and Michal Irani. Blind super-resolution kernel estimation using an internal-gan. *CoRR*, abs/1909.06581, 2019.

[4] Ebru Bozdağ, Jeannot Trampert, and Jeroen Tromp. Misfit functions for full waveform inversion based on instantaneous phase and envelope measurements. *Geophysical Journal International*, 185(2):845–870, 2011.

[5] Olivier Cappé and Eric Moulines. On-line expectation-maximization algorithm for latent data models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(3):593–613, Jun 2009.

[6] Ayan Chakrabarti. A neural approach to blind motion deblurring. In *European conference on computer vision*, pages 221–235. Springer, 2016.

[7] Chengqian Che, Fujun Luan, Shuang Zhao, Kavita Bala, and Ioannis Gkioulekas. Inverse transport networks. *arXiv preprint arXiv:1809.10820*, 2018.

[8] Jianfei Chen, Jun Zhu, Yee Whye Teh, and Tong Zhang. Stochastic expectation maximization with variance reduction. In *NeurIPS*, pages 7978–7988, 2018.

[9] S. Cho, Jue Wang, and S. Lee. Handling outliers in non-blind image deconvolution. In *2011 International Conference on Computer Vision*, pages 495–502, 2011.

[10] Sunghyun Cho and Seungyong Lee. Fast motion deblurring. In *ACM SIGGRAPH Asia 2009 papers*, pages 1–8. 2009.

[11] Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):1–22, 1977.

[12] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real nvp, 2017.

[13] Rob Fergus, Barun Singh, Aaron Hertzmann, Sam T Roweis, and William T Freeman. Removing camera shake from a single photograph. In *ACM SIGGRAPH 2006 Papers*, pages 787–794. 2006.

[14] Jeffrey A Fessler. Model-based image reconstruction for mri. *IEEE signal processing magazine*, 27(4):81–89, 2010.

[15] Andreas Fichtner, Brian LN Kennett, Heiner Igel, and Hans-Peter Bunge. Theoretical background for continental-and global-scale full-waveform inversion in the time–frequency domain. *Geophysical Journal International*, 175(2):665–685, 2008.

[16] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[17] Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In *Advances in neural information processing systems*, pages 10215–10224, 2018.

[18] D. Kundur and D. Hatzinakos. Blind image deconvolution. *IEEE Signal Processing Magazine*, 13(3):43–64, 1996.

[19] Orest Kupyn, Volodymyr Budzan, Mykola Mykhailych, Dmytro Mishkin, and Jiří Matas. Deblurgan: Blind motion deblurring using conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8183–8192, 2018.

[20] A. Levin, Y. Weiss, F. Durand, and W. T. Freeman. Efficient marginal likelihood optimization in blind deconvolution. In *CVPR 2011*, pages 2657–2664, 2011.

[21] Anat Levin, Yair Weiss, Fredo Durand, and William T Freeman. Understanding and evaluating blind deconvolution algorithms. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1964–1971. IEEE, 2009.

[22] Jingrui Luo and Ru-Shan Wu. Seismic envelope inversion: reduction of local minima and noise resistance. *Geophysical Prospecting*, 63(3):597–614, 2015.

[23] Yi Luo and Gerard T Schuster. Wave-equation traveltime inversion. *Geophysics*, 56(5):645–653, 1991.

[24] D. Perrone and P. Favaro. Total variation blind deconvolution: The devil is in the details. In *2014 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2909–2916, Los Alamitos, CA, USA, jun 2014. IEEE Computer Society.

[25] Nicholas Rawlinson, S Pozgay, and S Fishwick. Seismic tomography: a window into deep earth. *Physics of the Earth and Planetary Interiors*, 178(3-4):101–135, 2010.

[26] Dongwei Ren, Kai Zhang, Qilong Wang, Qinghua Hu, and Wangmeng Zuo. Neural blind deconvolution using deep priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3341–3350, 2020.

[27] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *International conference on machine learning*, pages 1278–1286. PMLR, 2014.

[28] J. A. Sethian. Fast marching methods. *SIAM Review*, 41(2):199–235, 1999.

[29] Peter M Shearer. *Introduction to seismology*. Cambridge university press, 2019.

[30] Jonathan D. Smith, Kamyar Azizzadenesheli, and Zachary E. Ross. Eikonet: Solving the eikonal equation with deep neural networks, 2020.

[31] He Sun and Katherine L. Bouman. Deep probabilistic imaging: Uncertainty quantification and multi-modal solution characterization for computational imaging, 2020.

[32] Eran Treister and Eldad Haber. A fast marching algorithm for the factored eikonal equation. *Journal of Computational Physics*, 324:210–225, 11 2016.

[33] Mingyang Xie, Yu Sun, Jiaming Liu, Brendt Wohlberg, and Ulugbek S. Kamilov. Joint reconstruction and calibration using regularization by denoising, 2020.

[34] Aston Zhang, Zachary C. Lipton, Mu Li, and Alexander J. Smola. *Dive into Deep Learning*. 2020. `https://d2l.ai`.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [Yes]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [N/A]
   (b) Did you include complete proofs of all theoretical results? [N/A]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]