

# Mixture Manifold Networks: A Computationally Efficient Baseline for Inverse Modeling

Gregory P. Spell<sup>1</sup>, Simiao Ren<sup>1</sup>, Leslie M. Collins<sup>1</sup>, Jordan M. Malof<sup>2</sup>

<sup>1</sup>Duke University, Department of Electrical & Computer Engineering, Durham, NC USA

<sup>2</sup>University of Montana, Department of Computer Science, Missoula, MT USA

{gregory.spell, simiao.ren, leslie.collins}@duke.edu, jordan.malof@umontana.edu

## Abstract

We propose and show the efficacy of a new method to address generic inverse problems. Inverse modeling is the task whereby one seeks to determine the control parameters of a natural system that produce a given set of observed measurements. Recent work has shown impressive results using deep learning, but we note that there is a trade-off between model performance and computational time. For some applications, the computational time at inference for the best performing inverse modeling method may be overly prohibitive to its use. We present a new method that leverages multiple manifolds as a mixture of backward (e.g., inverse) models in a forward-backward model architecture. These multiple backwards models all share a common forward model, and their training is mitigated by generating training examples from the forward model. The proposed method thus has two innovations: 1) the multiple Manifold Mixture Network (MMN) architecture, and 2) the training procedure involving augmenting backward model training data using the forward model. We demonstrate the advantages of our method by comparing to several baselines on four benchmark inverse problems, and we furthermore provide analysis to motivate its design.

## Introduction

We consider the goal of solving generic inverse problems, which arise in a multitude of scientific fields. Suppose we have an observable variable  $y \in \mathbb{R}^M$ , which manifests from a mapping of a forward process over a control variable  $x \in \mathbb{R}^D$ . We write this forward mapping as:

$$y = f(x) \quad (1)$$

For many scientific endeavors, while  $y$  may be a measurement of a system, the variable of greater interest is the control variable  $x$ . This motivates the task whereby, given  $f$ , we seek to determine the inverse model,  $g = f^{-1}(y)$ , such that:

$$x = f^{-1}(y) \quad (2)$$

which we will also refer to as the backward process. For many real physical systems, this inverse process is ill-posed, in that there can be several (or sufficiently similar) settings of  $x$  which correspond to a single observed  $y$  (e.g., non-uniqueness of solutions). This "one-to-many" aspect poses

a particular challenge for deep learning approaches to learning inverse models, as conventional deep neural networks (DNNs) only predict a single output for a given input, and furthermore, the methods used to train DNNs (e.g., gradient descent) typically do not account for non-uniqueness.

A variety of approaches have been proposed to address ill-posed problems of this kind (Aster, Borchers, and Thurber 2018), however, we focus in this work on recent DNN-based approaches, termed here as deep inverse models (DIMs) (Ardizzone et al. 2019; Ren, Padilla, and Malof 2020). Recent DIMs often fall into one of two major classes that share common advantages and disadvantages: generative and iterative. Generative models address non-uniqueness by learning a conditional probability distribution  $Pr(x|y)$ , and then sampling candidate solutions from this distribution. DIMs of this kind include the (conditional) Variational Autoencoders (cVAEs) (Kingma and Welling 2013), Generative Adversarial Networks (GANs) (Salimans et al. 2016), and conditional Invertible Neural Networks (cINNs) (Kruse et al. 2019). The other class of DIM, iterative models, involve randomly sampling a large set of solution candidates (e.g., from a uniform distribution), and then iteratively refining these solutions (e.g., using gradient descent, or evolutionary approaches). DIMs of this kind include evolutionary algorithms (Zhang et al. 2020; Johnson and Rahmat-Samii 1997; Rahmat-Samii 2003) and the Neural-Adjoint (Ren, Padilla, and Malof 2020).

Generative models are computationally fast; however, they are often less accurate than Iterative methods. As reported in (Ren, Padilla, and Malof 2020), these models often produce solutions near - but not precisely at - the true solution. Because they are probabilistic, they are unlikely to sample exact solutions, even when the conditional distribution places substantial probability mass over the exact solution. By contrast, iterative models are highly accurate, but suffer from substantially greater computational burden, due to the need for an iterative (sequential) refinement procedure at test time. Sequential methods are often several orders of magnitude slower than generative counterparts. For many practical problems, such as those in scientific computing (e.g., materials design (Khatib et al. 2021), geophysics (Snieder and Trampert 1999), astronomy (Lucy 1994), oceanography (Wunsch 1996)) this burden can be tolerated, either because real-time operation is unnecessary, or

because the computational cost of the model is insignificant compared to its accuracy advantages. For many problems though, this computational burden is unacceptable, making Generative models as the only viable solution.

In this work, we propose Mixture Manifold Networks (MMNs), which achieve accuracy close to Iterative methods such as NA, but with computational burden similar to Generative models. MMN is built upon an existing model, termed the Tandem (Liu et al. 2018; Jordan and Rumelhart 1992). By contrast to Iterative and Generative models, the Tandem addresses non-uniqueness by isolating a single inverse solution for each setting of  $y$ , and then learning a direct mapping to that single solution, analogous to standard regression. As a result, the Tandem model requires no iterative refinement of its solutions, but can also directly regress onto exact solutions, potentially yielding much higher accuracy than Generative models. Recently the Tandem was found to be computationally fast, and achieve relatively high accuracy compared to other DIMs, *if all models are limited to a single solution proposal* (Ren, Padilla, and Malof 2020). One drawback of the Tandem, however, is that it can only produce a single candidate solution, and if this solution is inaccurate, then it cannot propose additional candidates. As a result, it tends to perform poorly compared to Iterative and Generative methods when a large number of solution proposals can be considered.

To improve a Tandem-esque model, we seek to endow it with the capability of producing more than one candidate solution. Since the Tandem model learns a direct mapping from each  $y$  to a single solution, we consider these solutions as comprising a manifold in the output space. The shape of this manifold will have been influenced by stochasticity inherent to the model training process, such as the random initialization of the model parameters and random order of training examples. Thus, we believe that if several such manifolds are learned, each will describe a different set of solutions for the same inputs  $y$ . While these manifolds of solutions may be similar for some  $y$ , it is possible – and likely – that certain manifolds will be better fit for different sets of  $y$  than others. With this in mind, we posit that a Tandem-esque model can be made more flexible and effective by leveraging a mixture of manifolds and choosing candidate solutions from among these manifolds. We thus name our method the *Mixture Manifold Network* (MMN).

In addition to proposing the MMN architecture, we enhance our method with a data augmentation scheme for manifold-based inverse modeling. Our second proposal comes from our observation that in the forward-backward model paradigm, the forward model is easier to learn than the backward model: the forward model does not exhibit a one-to-many challenge and generally has a more-smooth output space. This renders the forward model more capable of interpolating than the backward model. We devise a method to leverage this characteristic of the forward model to improve MMN performance.

In summary, the primary contributions of this work are:

1. *The Mixture Manifold Network (MMN) method for inverse modeling.* MMN has performance that nears that of slower iterative models, but exhibits speed similar to

generative models, thus bringing higher accuracy to more real-time applications.

2. *We demonstrate the performance of MMN on a series of benchmark inverse modeling problems.* MMN outperforms several inverse modeling baselines, and it approaches the performance of the best-performing inverse modeling method we are aware of: the Neural Adjoint (NA). MMN is dramatically faster at inference than NA.
3. *We provide analysis that lends insight into why MMN is an effective inverse modeling method.*

## Related Work

Contemporary DIM methods often fall into either a *generative* or *iterative* paradigm, with the Tandem method falling into its own separate paradigm. We discuss the Tandem more fully in the Methods section, but mention here that in this approach, a neural forward model is first trained and then its parameters are fixed before the forward model is connected to the backward model, which is then trained in end-to-end fashion. Connecting the forward and backward models thereby regulates the noisy gradient caused by one-to-manyness.

Another class of end-to-end approaches is to model the conditional posterior distribution of the control variables given the observable variables,  $p(x|y)$ . These approaches leverage variational methods, and include models such as Variational Auto-Encoders (VAEs), Mixture Density Networks (MDNs) (Bishop 1994), Invertible Neural Networks (INNs), and conditional INNs (cINNs).

VAEs use an encoder-decoder structure, with the encoder modeling  $x$  into a Gaussian-distributed random variable  $z$  conditioned on  $y$ , and the decoder going back from the latent variable  $z$  to  $\hat{x}$ . The VAE is trained by minimizing the evidence-lower-bound (ELBO) as well as a L2 MSE reconstruction loss between original input  $x$  and  $\hat{x}$ . MDNs assume a Gaussian mixture for the posterior  $p(x|y)$ , with the parameters of Gaussians (as well as the mixing proportions) predicted by a feedforward neural network, and the number of components in the mixture being selected as a hyperparameter. At inference for a particular  $y$ ,  $\hat{x}$  is sampled from the learned mixture distribution. Like VAEs, INNs also assume latent Gaussian variables  $z$ , but now *bijective* mappings are learned between  $x$  and pairs  $(y, z)$ . INNs can be trained by minimizing a supervised reconstruction loss and Maximum Mean Discrepancy (MMD) loss, or alternatively, a maximum likelihood estimate (MLE) loss used to enforce Gaussianity for the distribution of  $z$ . For proposing inverse solutions  $y$ , values of  $z$  are sampled from a Gaussian distribution, and the pairs  $(y, z)$  are mapped by the network to an estimate  $x$ . Finally, cINNs modify the INN structure in that they learn bijective mappings between  $x$  and  $z$  conditioned on  $y$ , rather than between  $x$  and the pair  $(y, z)$ .

The last class of DIM approaches we consider here are iterative methods, such as the Genetic Algorithm (GA) and Neural-Adjoint (NA). Under this paradigm, the DIM will make an initial set of solution proposals, and then from these iteratively search for better solutions. The NA uses a neural approximation of the forward process  $\hat{f}$ , and searches for op-

timal solutions  $\hat{x}$  using the gradient  $\partial \hat{f} / \partial x$ . The GA, on the other hand, is a gradient-free method, and instead searches for solutions by selecting and mutating populations of candidate solutions using evolutionary fitness rules.

Many of the methods above have recently been benchmarked for different inverse problems. Kruse (Kruse et al. 2019) benchmarked the stochastic (e.g., posterior modeling) approaches on two simple kinematic inverse problems, and Ren (Ren, Padilla, and Malof 2020) further benchmarked those methods on two additional problems. That latter study also benchmarked the Tandem (Liu et al. 2018; Jordan and Rumelhart 1992) model and their proposed NA on the same four inverse modeling tasks.

## Methods

The MMN method bears resemblance to the Tandem, particularly in the case that a single manifold is used (effectively reduces to a Tandem network). We thus first review the Tandem’s architecture and training procedure before describing our proposed MMN, and both are illustrated in Figure 1. The methods require a dataset upon which to train, consisting of sample pairs  $\{x_n, y_n\}_{n=1}^N$  from the true process  $f$ .

### Tandem Model

The Tandem architecture consists of two neural networks. One network approximates the forward process, while the other models the backward process. Supposing that the estimate,  $\hat{f}$ , of the forward process is parameterized by  $\theta_f$ , then for each  $x_n$  in the training data, the model predicts a corresponding output  $\hat{y}_n$ :

$$\hat{y}_n = \hat{f}(x_n; \theta_f) \quad (3)$$

The parameters  $\theta_f$  of the forward model are learned by optimizing a standard regression loss; here, we use the squared error loss between the true  $y_n$  and the estimate  $\hat{y}_n$ :

$$\mathcal{L}_{fwd} = \sum_{n=1}^N \left\| y_n - \hat{y}_n \right\|_2^2 \quad (4)$$

After learning the forward model  $\hat{f}$ , its parameters are fixed (e.g., “frozen”) before estimating the inverse model  $\hat{f}^{-1}$ , the second network composing the Tandem. For notational convenience, we will denote  $\hat{g} = \hat{f}^{-1}$ :

$$\hat{x}_n = \hat{f}^{-1}(y_n) = \hat{g}(y_n; \theta_g) \quad (5)$$

where  $\theta_g$  parameterizes the network  $\hat{g}$ . These parameters are learned by minimizing the following loss:

$$\mathcal{L}_g = \sum_{n=1}^N \left\| \hat{f}(\hat{g}(y_n)) - y_n \right\|_2^2 \quad (6)$$

where we reiterate that the parameters of  $\hat{f}$  have been frozen.

In addition to the loss described in 6, we employ a boundary loss proposed by (Ren, Padilla, and Malof 2020). This loss was introduced to mitigate a challenge with their Neural Adjoint in that it would frequently converge to solutions

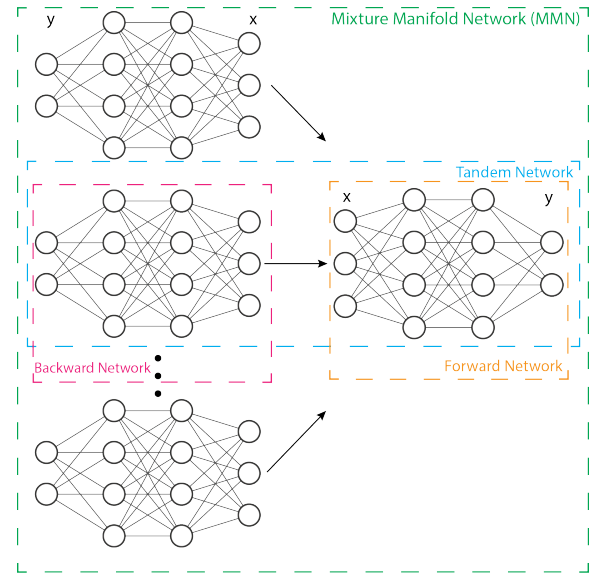


Figure 1: Model architecture illustration. The yellow box represents the forward network mapping control variable  $x$  to observation  $y$ . Pink represents the backward network mapping the inverse relationship from  $y$  to  $x$ . Connecting them forms the Tandem network, depicted in the blue box. Our proposed MMN –illustrated in the green box – can be seen as an extension of the Tandem, with multiple trained backward models corresponding to a single forward model.

outside the domain of the training data. They posit the reason for this is that outside of the training data domain (for points  $x$ ), the estimated forward model  $\hat{f}$  becomes highly inaccurate and then tends to erroneously predict that  $x$ -values outside of the space of  $\mathcal{X}$  are near the desired solutions for  $y$ . To address this undesirable behavior, an additional loss term is included to encourage the models to identify solutions within the bounds of the training data (e.g., within the space  $\mathcal{X}$ ). This loss term,  $\mathcal{L}_{bnd}$ , is formulated as:

$$\mathcal{L}_{bnd}(x) = \begin{cases} x - x_{max}, & \text{if } x \geq x_{max} \\ 0, & \text{if } x_{min} \leq x \leq x_{max} \\ x_{min} - x, & \text{if } x \leq x_{min} \end{cases} \quad (7)$$

which is equivalent to  $\mathcal{L}_{bnd} = \text{ReLU}(|x - \mu_x| - \frac{1}{2}R_x)$  for uniform distributions where  $\mu_x$  and  $R_x$  are the mean and range of the training data, respectively, and ReLU is the rectified-linear-unit commonly used as an activation function in neural networks. For data distributions without hard bounds on  $x$ , the range was defined to be the interval of 95% probability. Though the boundary loss was proposed for training the Neural Adjoint inverse models, the authors of (Ren, Padilla, and Malof 2020) demonstrate that including  $\mathcal{L}_{bnd}$  in the training of Tandem backward models also improves their performance. Thus, when training the backward component of a Tandem model, we use a combination loss from Equations 6 and 7:

$$\mathcal{L}_{bwd} = \mathcal{L}_g + \gamma \mathcal{L}_{bnd} \quad (8)$$

where the hyperparameter  $\gamma$  is used to balance the relative importance of the two losses in the combination.

### Mixture Manifold Network

We propose to extend the flexibility of a forward-backward model architecture like the Tandem by leveraging multiple backwards models associated with a single forward model. Since the outputs from each of these backward models can be considered as lying on a separate *manifold*, and since we propose to use a *mixture* of such manifolds, we have named our method the *Mixture Manifold Network*.

As in the Tandem procedure previously described, the MMN method begins by estimating a single forward model  $\hat{f}$  as an approximation of the forward process described in Equation 1. This estimate is learned by optimizing the loss in Equation 4 with predictions as in Equation 3.

Also as in the Tandem method, after learning the forward model  $\hat{f}$ , its parameters are frozen before learning subsequent pieces of the network. Whereas the Tandem model learns a single inverse model  $\hat{g} = \hat{f}^{-1}$ , our Mixture Manifold Network learns a set of  $K$  inverse models,  $\{\hat{g}^k\}_{k=1}^K$ . Crucially, each backward model  $\hat{g}^k$  is learned separately as an approximation for  $f^{-1}$  (Equation 2), using the same training data and using the combination of backward model loss and boundary loss in Equation 8. Each backward model’s predictions are can be described by:

$$\hat{x}_n^k = \hat{g}^k(y_n; \theta_g^k) \quad (9)$$

To leverage a mixture of  $K$  trained backward models, we devise a new inference scheme that is critical to the effectiveness of MMN. The scheme is a procedure whereby for each  $y_n$  MMN essentially “chooses” which of the manifolds in its mixture it will use to propose a solution  $\hat{x}_n$ ; it is illustrated in Figure 2. Consider that each backwards model in MMN has predicted a  $\hat{x}_n^k$  using Equation 9, yielding a set  $\{\hat{x}_n^k\}_{k=1}^K$ . Since the same forward model  $\hat{f}$  is tied to each backward model, the forward model can be used to assess which of the proposed solutions is the best. This can be accomplished by using the forward model to map each  $\hat{x}_n^k$  back into  $\mathcal{Y}$ -space and comparing to the original test point  $y_n$ . Whichever proposed solution yields an output closest to  $y_n$  is selected, which we formalize as:

$$\hat{x}_n = \operatorname{argmin}_{\hat{x}_n^k} \left\| \hat{f}(\hat{x}_n^k) - y_n \right\|_2^2 \quad (10)$$

### Sampling and Forward Model Generation

Whereas we have thus far described the MMN *architecture* and an inference procedure to effectively leverage that architecture, we describe here a proposed innovation to mitigate the training of efficacious backward models  $\hat{g}^k$ . Our proposal is to *augment* the data used to train the backward models comprising a MMN by leveraging the effectiveness of the forward model  $\hat{f}$ . This augmentation involves sampling points from the input space  $\mathcal{X}$  and then to use the learned forward model  $\hat{f}$  to map those points into the output space  $\mathcal{Y}$  (e.g., Equation 3). These pseudo-points which

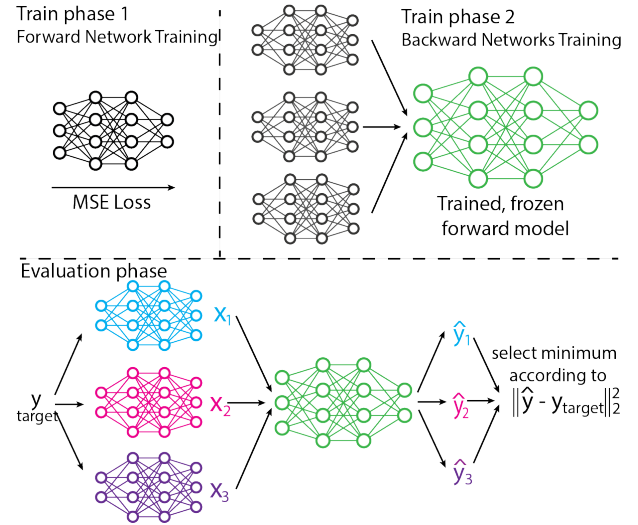


Figure 2: Illustration of the training and testing phases of MMN. The forward and backward networks are trained sequentially, with the trained forward network connected at the end of the backward network during the backward model’s training. During evaluation, the backward networks pass their outputs through the forward model and the best design is chosen according to Equation 10

have been “generated” by the forward model are then used as the training data for the backward models in the MMN. More formally, we propose to assemble a new set of training data  $\{\tilde{x}_n, \tilde{y}_n\}_{n=1}^{N'}$  with which to train the backward models  $\hat{g}^k$ . This new set of training data is generated according to:

$$\tilde{y} = \hat{f}(\tilde{x}); \tilde{x} \sim \hat{p}(x) \quad (11)$$

where  $\hat{p}(x)$  is a placeholder to indicate that the points  $\tilde{x}$  are sampled from a distribution that bears *resemblance* to the distribution describing the true data-generating process for  $x$ , but is perhaps not the same distribution. The most crucial aspect about the sampling distribution  $\hat{p}$  is that it shares the same support (e.g., bounds on  $x$ ) as the true distribution  $p$ . In some cases, we may know  $p(x)$  exactly – as in the Sine Wave and Robotic Arm benchmarks – and so we can sample directly from that distribution. In other cases, like the Meta-Material and Shell problems, generic distributions (e.g., uniform, Gaussian) with judiciously chosen bounds and means can be selected.

### Evaluation – Re-simulation Error

We measure model performance using “re-simulation” error,  $r = \mathcal{L}(f(\hat{x}), y)$ , which seeks to capture the discrepancy between the true solution  $y$  and the “re-simulated” value of  $y$  obtained by passing the estimate  $\hat{x}$  through the true forward model (also known as the simulator). We follow the convention of using mean-squared error (MSE) to measure this discrepancy. (Ren, Padilla, and Malof 2020) extended re-simulation error to account for proposal of  $T$  solutions from a model. Here,  $T$  is the number of times the simulator  $f$  is

	Sine	Arm	Meta-Mat	Shell
Dim(x)	2	4	8	8
Dim(y)	1	2	300	201
Num. Train	8000	8000	8000	40,000
Num. Val.	2000	2000	2000	10,000
Num. Test	1000	1000	1000	500

Table 1: Summary of our inverse problem datasets. Note we cover both cases that  $Dim_x > Dim_y$  and  $Dim_y > Dim_x$ .

used, which may be computationally expensive (e.g., a material design physics-based simulation).

## Benchmark Problems

We consider four benchmark tasks to demonstrate our MMN inverse modeling method; these are summarized in Table 1. We include a popular task – the Robotic Arm – which was introduced in (Kruse et al. 2019) and further benchmarked in (Ren, Padilla, and Malof 2020). That latter study also benchmarked the Sine Wave (which it introduced) and the Meta-Material (Nadell et al. 2019; Deng et al. 2021) tasks. For our fourth task, we have adopted the Shell (Peurifoy et al. 2018), which was benchmarked in (Ren et al. 2022).

We have selected these benchmarks for several reasons. First, while each demonstrates non-uniqueness of solutions, they exhibit a range of technical difficulty, with the Sine Wave being the easiest due to its low dimensional  $x$  and  $y$ , and the Meta-Material and Shell problems being harder due to their relatively higher dimensionality. Second, they encompass a breadth of application, with the Sine Wave being somewhat of a toy problem useful for close analysis, Robotic Arm being a kinematic problem with possible application for control systems, and both Meta-Material and Shell being from two different sub-fields of artificial electromagnetic material design. We believe that choosing this range of problems helps ensure that our findings for MMN generalize to other possible inverse problems. Last, each of these tasks has been benchmarked in a previous study, which helps us maintain continuity with the community.

For each inverse task, we use the same experimental designs as from previous benchmark studies, including the simulator (e.g., forward model) parameters, simulator sampling procedures, and training/testing splits. Details not included in Table 1 can be found in the studies cited above and in our supplemental material.<sup>1</sup>

## Experimental Design and Results

We show experiments and results on the above benchmark problems to demonstrate the performance of our proposed MMN method and compare it to an array of previous inverse modeling methods<sup>2</sup>. For the baselines, we find mostly consistent results with previous benchmarking studies, and we find that our MMN outperforms generative model baselines without the severe computational time of iterative models.

<sup>1</sup>See: <https://arxiv.org/abs/2211.14366>

<sup>2</sup>See: <https://github.com/gspell/MixtureManifoldNetworks>

	Sine	Arm	Meta-Mat.	Shell
MDN	3.68e-1	4.41e-3	2.65e-3	1.07e-1
INN	6.78e-1	6.26e-3	2.84e-2	1.79e-2
cINN	3.43e-1	4.17e-3	4.68e-3	7.33e-2
cVAE	8.18e-1	2.32e-2	1.20e-2	7.51e-3
Tandem	1.27e-2	2.38e-3	1.36e-3	7.41e-3
GA	8.15e-3	1.65e-4	2.66e-4	2.89e-3
NA	3.25e-3	<b>1.17e-4</b>	<b>1.97e-4</b>	<b>2.88e-3</b>
MMN	<b>1.32e-3</b>	4.33e-4	3.50e-4	3.75e-3

Table 2: Average  $T=1$  Re-Simulation Error. The best values for each benchmark problem are bolded

Thus, we argue that our MMN method provides an alternative that balances performance with inference speed.

We follow the experimental design of (Ren, Padilla, and Malof 2020) and (Kruse et al. 2019) in which the authors benchmark many of the same models, and on many of the same problems. For the Shell problem, we follow (Ren et al. 2022) where it was used as a benchmark problem for inverse models. For each benchmark problem, each model used the same training and testing data, batch size, and criteria to stop training. Model architectures and hyperparameters we selected to align with previous studies, and full implementation details can be found in the supplementary material.

For MMN experimentation, we use the the same forward and backward model architectures as used for the Tandem model for each benchmark problem. We use  $K = 6$  manifolds (backward models) for each MMN. Each of the  $K$  backward models were trained using *only* augmented data generated using the forward model sampling procedure described in the Methods section. The distributions used for this sampling differed between benchmarks, and they are outlined in the supplement. For the Sine Wave, Robotic Arm, and Meta-Material benchmarks, we sampled 40,000 points with which to train, and for the Shell benchmark, we sampled 250,000 (in the next section, we show performance as the number of generated points varies).

Our main experimental results are presented in Table 2, which shows the re-simulation error for MMN and all baseline methods for the case of  $T = 1$  solution proposals by each model. Here, we find that MMN is either the best performing method (for the Sine Wave) or is the second-best performing method to NA, which performs so well at an expense of heavy computation during evaluation. We thus find that on all four benchmark problems, MMN outperforms the Tandem model – to which it bears the most resemblance – as well as all of the probabilistic models (MDN, INN, cINN, and cVAE). Furthermore, we find that this outperformance is statistically significant, as we have conducted pairwise t-tests between the errors for MMN and each competing model. For example, the top two highest p-values across *all* benchmarks are 0.0217 and 0.00434. In our supplementary material, we present the p-values for all pairwise t-tests.

In addition to evaluation for  $T = 1$  solution proposals from each model, we can measure re-simulation error performance as a function of  $T$ , as in (Ren, Padilla, and Malof 2020). A limiting aspect of Tandem performance is its in-

## Test Resimulation Error Over Inferences

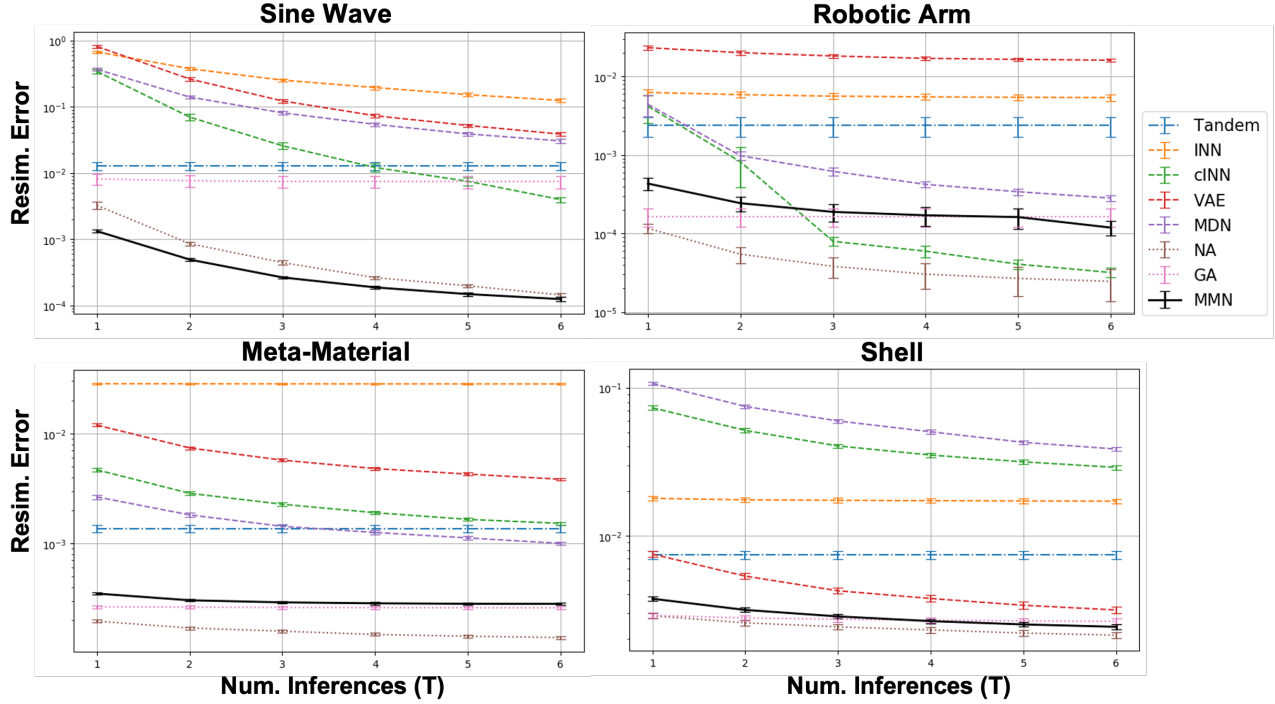


Figure 3: Re-Simulation error over Number of Inferences. We highlight our MMN method with the solid, black line, while stochastic baseline methods are shown with dashed lines and iterative baselines with dotted lines. Standard errors are shown due to their importance in statistical significance testing, following recommendations of (Bouthillier et al. 2021).

	Sine	Arm	Meta-Mat.	Shell
MDN	3.06e-2	2.83e-4	1.00e-3	3.85e-2
INN	1.25e-1	5.38e-3	2.83e-3	1.71e-2
cINN	3.97e-3	3.20e-5	1.52e-3	2.90e-2
cVAE	3.88e-2	2.61e-2	3.84e-3	3.15e-3
Tandem	1.27e-2	2.38e-3	1.36e-3	7.41e-3
GA	7.47e-3	1.64e-4	2.60e-4	2.64e-3
NA	1.44e-4	<b>2.47e-5</b>	<b>1.39e-4</b>	<b>2.13e-3</b>
MMN	<b>1.25e-4</b>	1.19e-4	2.82e-4	2.42e-3

Table 3: Average  $T=6$  Re-Simulation Error. The best values for each benchmark problem are bolded

	Sine	Arm	Meta-Mat.	Shell
MDN	0.27	0.27	0.49	0.18
INN	0.03	0.03	0.25	0.20
cINN	0.03	0.03	0.26	0.43
cVAE	0.02	0.02	0.24	0.18
Tandem	0.02	0.02	0.23	0.20
GA	28.28	28.20	97.98	71.93
NA	1.83	1.84	80.44	54.57
MMN	0.13	0.15	3.26	1.22

Table 4: Inference Time in seconds for 1,000 solutions

capability of predicting multiple solutions, but our MMN method – while a deterministic model – does admit multiple solutions. To be exact, MMN can propose up to  $T = K$  different solutions, as we take the solution from each manifold as a separate prediction. In Figure 3 we show re-simulation error as a function of  $T$  for each model on each benchmark problem. We also provide in Table 3 the re-simulation error for  $T = 6$  inferences. With multiple proposed solutions, our MMN still outperforms all but the iterative method baselines, to which MMN still remains competitive.

We have argued that a possible limitation of the iterative NA and GA methods which outperform our MMN is

that they require longer inference times, which may be prohibitive for real-time applications, or other situations when fast inference is required. The MMN has substantially faster inference times, which we show in Table 4, presenting the total inference time required in seconds for each model to propose 1000 solutions to each benchmark problem. Clearly, both the NA and GA methods require significantly more inference time than any of the other baseline methods, and this is a direct consequence of the iterative nature of their solution finding. This inference time drawback is also compounded when the dimensionality of the inverse problem is larger, as can be seen for specifically the Meta-Material and Shell benchmarks.

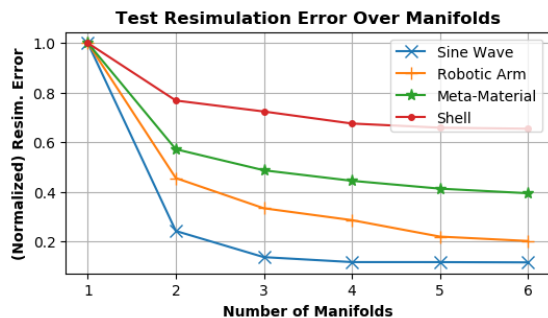


Figure 4: Resimulation error over number of manifolds for MMN on the test set for each benchmark problem. Note that resimulation error here has been normalized by the resimulation error for a single-manifold model, so as to put each benchmark’s results over number of manifolds on the same scale and facilitate viewing.

### MMN Analysis

In this section, we further characterize how our proposed MMN method performs. In particular, we examine how both of the innovations in MMN change its performance in terms of resimulation error. Our results demonstrate that each innovation individually – using multiple manifolds and training using forward-model generated points – improves performance, suggesting that both are individually beneficial.

#### On Use of Multiple Manifolds

We provide evidence that our proposed architecture for the Mixture Manifold Network – separate from using forward model point generation – is beneficial to modeling performance. To demonstrate that using multiple backwards models offers flexibility, we show the performance, in terms of  $T = 1$  resimulation error, improves as successive manifolds are included in MMN. We present this for all benchmark problems in Figure 4, where the resimulation error across number of manifolds has been normalized by single-manifold performance so that results for all benchmarks can be more-easily viewed on the same axes.

For all benchmarks, we observe a decrease in resimulation error when  $K > 1$ , and furthermore, this decrease is monotonic with increasing  $K$ . This indicates that as successive manifolds are added to MMN, the model’s overall performance increases. We believe that this experimentally demonstrates evidence for our intuition that using multiple manifolds allows MMN to “cover” different regions of  $\mathcal{X}$ -space with different manifolds, and that a well-trained forward model can discern which manifold is the most appropriate to use depending on the input point. Thus, as more manifolds are added to MMN, the forward model has more flexibility in its selection of the best manifold.

#### On Use of Sample Generation

We now demonstrate that training backward models using points that have been sampled and pushed through the forward model is also beneficial for MMN performance. Again,

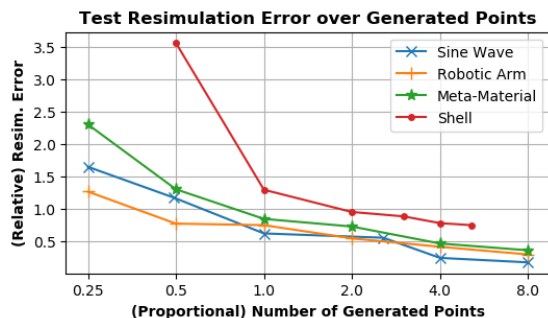


Figure 5: Resimulation error on test over number of points sampled and generated by the MMN forward model for each benchmark problem. Note that both axes have been adjusted to facilitate assessing results across all benchmarks. The x-axis, number of points sampled/generated, is in terms of proportionality to the number of real points used to train baseline models (and used to train the forward model itself). The resimulation error, shown on the y-axis, is relative to the error obtained by a (single-manifold) MMN with backward models trained on real data.

we consider performance in terms of  $T = 1$  resimulation error, and now we show that as the number of points generated by the forward model increases, so too does the performance of MMN. This result is presented, for every benchmark problem, in Figure 5. To present all benchmarks in the same space, we have again adjusted the axes of this results figure. The x-axis shows the number of generated points proportional to the number of real points used to train the baseline models, and the y-axis is relative to the resimulation error obtained by MMN with the (single) backward model trained on real data.

In addition to the resimulation error decreasing as the number of generated points increases, it is almost more important to note that for each benchmark, there is a certain number of generated points at which MMN outperforms a baseline with backward models trained on real data (e.g., the curves for all benchmarks fall below  $\bar{r} = 1$ ).

### Conclusions

In this work, we have proposed a new method for approximating solutions to inverse problems using deep neural networks, and we have assessed the performance of our method against seven modern deep inverse models on four benchmark tasks. Our method, the Mixture Manifold Network, proposes a solution from among multiple backward models that have all been trained to correspond to the same forward model, and it can furthermore at training leverage augmented data points generated from the forward model. We find that our method performs comparably to iterative methods and outperforms all other baselines. MMN, however, exhibits dramatically faster inference time than the iterative methods do. Thus, for applications where inference time may be a limiting concern (e.g., real-time control systems or rapid material design), we believe our MMN offers a promising approach to addressing such inverse problems.

## References

- Ardizzone, L.; Kruse, J.; Rother, C.; and Köthe, U. 2019. Analyzing Inverse Problems with Invertible Neural Networks. In *International Conference on Learning Representations*.
- Aster, R. C.; Borchers, B.; and Thurber, C. H. 2018. *Parameter estimation and inverse problems*. Elsevier.
- Bishop, C. M. 1994. Mixture density networks. Technical Report NCRG/94/004, Dept. of Computer Science and Applied Mathematics, Aston Univ.
- Bouthillier, X.; Delaunay, P.; Bronzi, M.; Trofimov, A.; Nichyporuk, B.; Szeto, J.; Mohammadi Sepahvand, N.; Raff, E.; Madan, K.; Voleti, V.; Ebrahimi Kahou, S.; Michalski, V.; Arbel, T.; Pal, C.; Varoquaux, G.; and Vincent, P. 2021. Accounting for Variance in Machine Learning Benchmarks. In Smola, A.; Dimakis, A.; and Stoica, I., eds., *Proceedings of Machine Learning and Systems*, volume 3, 747–769.
- Deng, Y.; Ren, S.; Fan, K.; Malof, J. M.; and Padilla, W. J. 2021. Neural-adjoint method for the inverse design of all-dielectric metasurfaces. *Optics Express*, 29(5): 7526–7534.
- Johnson, J.; and Rahmat-Samii, V. 1997. Genetic algorithms in engineering electromagnetics. *IEEE Antennas and Propagation Magazine*, 39(4): 7–21.
- Jordan, M. I.; and Rumelhart, D. E. 1992. Forward models: Supervised learning with a distal teacher. *Cognitive science*, 16(3): 307–354.
- Khatib, O.; Ren, S.; Malof, J.; and Padilla, W. J. 2021. Deep learning the electromagnetic properties of metamaterials—a comprehensive review. *Advanced Functional Materials*, 31(31): 2101748.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Kruse, J.; Ardizzone, L.; Rother, C.; and Köthe, U. 2019. Benchmarking invertible architectures on inverse problems. *Workshop on Invertible Neural Networks and Normalizing Flows, International Conference on Machine Learning*.
- Liu, D.; Tan, Y.; Khoram, E.; and Yu, Z. 2018. Training deep neural networks for the inverse design of nanophotonic structures. *Acs Photonics*, 5(4): 1365–1369.
- Lucy, L. 1994. Optimum strategies for inverse problems in statistical astronomy. *Astronomy and Astrophysics*, 289: 983–994.
- Nadell, C. C.; Huang, B.; Malof, J. M.; and Padilla, W. J. 2019. Deep learning for accelerated all-dielectric metasurface design. *Optics express*, 27(20): 27523–27535.
- Peurifoy, J.; Shen, Y.; Jing, L.; Yang, Y.; Cano-Renteria, F.; DeLacy, B. G.; Joannopoulos, J. D.; Tegmark, M.; and Soljačić, M. 2018. Nanophotonic particle simulation and inverse design using artificial neural networks. *Science advances*, 4(6): eaar4206.
- Rahmat-Samii, Y. 2003. Genetic algorithm (GA) and particle swarm optimization (PSO) in engineering electromagnetics. In *17th International Conference on Applied Electromagnetics and Communications, 2003. ICECom 2003.*, 1–5.
- Ren, S.; Mahendra, A.; Khatib, O.; Deng, Y.; Padilla, W. J.; and Malof, J. M. 2022. Inverse deep learning methods and benchmarks for artificial electromagnetic material design. *Nanoscale*, 14(10): 3958–3969.
- Ren, S.; Padilla, W.; and Malof, J. 2020. Benchmarking deep inverse models over time, and the neural-adjoint method. *Advances in Neural Information Processing Systems*, 33: 38–48.
- Salimans, T.; Goodfellow, I.; Zaremba, W.; Cheung, V.; Radford, A.; and Chen, X. 2016. Improved techniques for training gans. *Advances in neural information processing systems*, 29.
- Snieder, R.; and Trampert, J. 1999. Inverse problems in geophysics. In *Wavefield inversion*, 119–190. Springer.
- Wunsch, C. 1996. *The ocean circulation inverse problem*. Cambridge University Press.
- Zhang, T.; Liu, Q.; Dan, Y.; Yu, S.; Han, X.; Dai, J.; and Xu, K. 2020. Machine learning and evolutionary algorithm studies of graphene metamaterials for optimized plasmon-induced transparency. *Optics Express*, 28(13): 18899–18916.