
Accurate and Robust Deep Learning Framework for Solving Wave-Based Inverse Problems in the Super-Resolution Regime

Matthew Li

Computational Science and Engineering
Massachusetts Institute of Technology
Cambridge MA 02139
mtcli@mit.edu

Laurent Demanet

Department of Mathematics
Earth Resources Lab
Massachusetts Institute of Technology
Cambridge MA 02139
laurent@math.mit.edu

Leonardo Zepeda-Núñez.

Department of Mathematics
University of Wisconsin-Madison
Madison WI 53706
zepedanunez@wisc.edu

Abstract

We propose an end-to-end deep learning framework that comprehensively solves the inverse wave scattering problem across all length scales. Our framework consists of the newly introduced wide-band butterfly network [37] coupled with a simple training procedure which dynamically injects noise during training. While our trained network provides competitive results in classical imaging regimes, most notably it also succeeds in the super-resolution regime where other comparable methods fail. This encompasses both (i) reconstruction of scatterers with sub-wavelength geometric features, and (ii) accurate imaging when two or more scatterers are separated by less than the classical diffraction limit. We demonstrate these properties are retained even in the presence of strong noise and extend to scatterers not previously seen in the training set. In addition, our network is straightforward to train requiring no restarts and has an online runtime that is an order of magnitude faster than optimization-based algorithms. We perform experiments with a variety of wave scattering mediums and we demonstrate that our proposed framework outperforms both classical inversion and competing network architectures that specialize in oscillatory wave scattering data.

1 Introduction

Wave-based inverse problems—i.e., inferring the properties of an unknown medium from indirect observations of the medium’s response to probing waves—are ubiquitous in engineering and sciences. This encounters applications in a myriad of different fields such as geophysics, astronomy, biomedical imaging, radar, spectrography, signal processing, communications, among many others [11, 54, 5, 51, 43].

All of these applications rely on the physical ability of waves to propagate information across long distances with little distortion. The amount of information that is communicated by waves is essentially proportional to its carrier frequency: higher frequencies transmit more information and thus, when used for imaging, lead to reconstructions with higher resolution. However, in order

to reduce computational cost, classical inversion methods [44, 25, 2, 52] selectively ignore later scattering events in the time series or process only the low frequency content in the data. These algorithms therefore offer fast reconstruction [13, 56, 30] of relatively low resolution images whose quality deteriorates significantly for highly heterogeneous media.

Modern techniques¹ have since been developed [3, 14, 35, 53, 47, 54] to service applications which depend on accurately imaging medium discontinuities and other fine grained material details. Unfortunately, these methods rely on high-dimensional, highly non-linear optimization programs which are susceptible to local minima. Thus, in industrial practices, domain experts are required in the loop to monitor and manually adjust the descent path. This exacerbates the already prohibitively expensive computational cost of these algorithms which require state-of-the-art computer clusters to even process the field data.

Evidently, applications requiring real time reconstructions, such as biomedical imaging, radar, continuous monitoring for CO₂ sequestration and geothermal energy, etc, remain beyond the reach of such algorithms [41].

In this context we seek to apply advances in machine learning to develop high-resolution, on-the-fly, reconstruction methods. Broadly, existing research in this area can be categorized as: (i) “physics informed” networks which minimize a PDE misfit function with a neural network ansatz [10]. While these methods offer extremely promising results they cannot yet provide real-time reconstruction. Alternatively, (ii) the mapping from the data to medium perturbations can be learned directly using a neural network [21, 20, 22, 23, 33, 38, 31].

For methods in the latter category the primary challenge lies in capturing the non-compressible long-range interactions of the inverse data-to-perturbation map efficiently. Indeed the long-range interactions often cause failures when applying typical approaches based on U-nets [49, 21, 20], as well as for generative modelling tools used for imaging such as VAEs or GANs [29, 36, 42]. Recent approaches [22, 23, 33, 38] address this issue by embedding the physics of wave propagation into the architecture of the network; specifically, taking cues from Fourier integral operators (FIOs) [27] and leveraging the butterfly algorithm [8, 7] for fast application of such operators. This builds on the use of FIOs to describe a diverse set of inverse problem modalities such as reflection seismology [46], thermoacoustic and photoacoustic tomography, radar [1, 11, 48], and single photon emission computed tomography [24]. However, these methods have several drawbacks: (i) their complexity increases super-linearly with respect to the Shannon-Nyquist scaling, (ii) they depend on stringent assumptions on symmetries and other equi- and/or in- variances, or (iii) they exhibit unstable training dynamics due to the highly-oscillatory nature of the data [58] and require intrusive initializations to obtain accurate approximations [38].

The recently introduced wide-band butterfly network (`WideBNet`) addresses these pitfalls in a unified manner [37]. Their main contribution was in extending the butterfly networks [57, 38] to assimilate wide-band data in a multi-scale fashion similar to the Cooley-Tukey algorithm [17]. This fully realizes the advantages of the butterfly algorithm in enabling a network design with training parameters and complexity which scales linearly, up-to-poly-log factors, with the dimension of the data. Additionally, the use of multi-frequency data stabilizes the inverse problem, and therefore the training dynamics of the network, and also empirically enables imaging in the super-resolution regime.

Despite the advantages of the `WideBNet` documented in [37] the stability of this network with noise remains a key gap that remains unaddressed. This is important as the mathematical super-resolution problem is known to be poorly conditioned [19]. Demonstrating robustness against noisy data is therefore key for applications, where the data acquisition is often noisy due to constraints in the exact position of the receivers or due to intrinsic statistical uncertainty.

Our contribution: We introduce the first deep learning framework capable of successfully solving the inverse scattering problem in the super-resolution regime while being robust to noise. The resulting network is able to handle multiscale features, with both sub-wavelength and wavelength level features, as well as resolve colliding scatterers separated by less than the diffraction limit. Remarkably, both of these properties hold with resilience to noise – we demonstrate robustness even with 100% noise-to-signal corruptions in the data. In addition, our network outperforms both classical and newly introduced architectures such as U-net (U-Net) [49], `SwitchNet` (`SwitchNet`) [33] and butterfly-networks (`NarrowBNet`) [57, 38] for a comprehensive range of scatterers with

¹See [16] and [55] for excellent historical reviews.

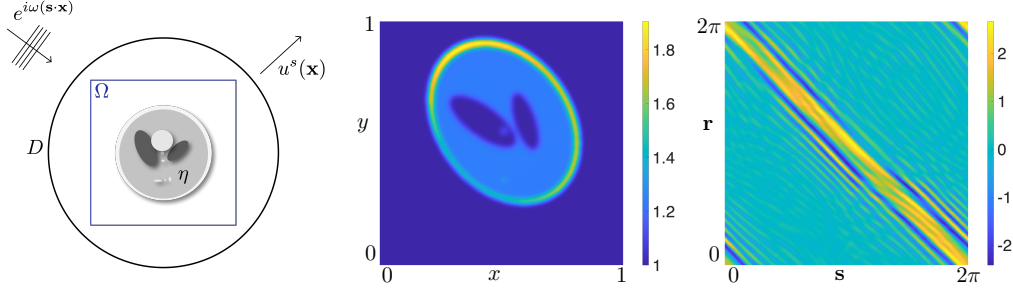


Figure 1: (left) setup for the inverse scattering problem. In the figure we probe the medium with a planewave with direction \mathbf{s} , and we sample the scattered field on the disk D , (center) slowness squared of a Shepp-logan model and (right) real part of the discrete Dirichlet-to-Neumann map.

varying wave-scattering physics. To our knowledge this is the first network to achieve these feats, thus enabling future applications in biomedical imaging, radar, and geophysical monitoring, where real-time imaging across all length scales is mission critical.

2 Physical Model

We consider the time-harmonic wave equation with constant-density acoustic physics, also called the Helmholtz equation, with frequency ω and squared slowness m , given by

$$(\Delta + \omega^2 m(\mathbf{x}))u(\mathbf{x}) = 0 \quad (1)$$

with radiating boundary conditions. We further suppose the slowness admits a scale separation into

$$m(\mathbf{x}) = m_0(\mathbf{x}) + \eta(\mathbf{x}) = 1 + \eta(\mathbf{x}), \quad (2)$$

where m_0 corresponds to the smooth background slowness, assumed known and for simplicity equal to one², and η the rough perturbation that we wish to recover. Under these assumptions, solutions to (1) can be expressed in the form

$$u(\mathbf{x}) = e^{i\omega(\mathbf{s}\cdot\mathbf{x})} + u^s(\mathbf{x}), \quad (3)$$

where $e^{i\omega(\mathbf{s}\cdot\mathbf{x})}$ is the incoming plane wave, with propagating direction \mathbf{s} , that we use to “probe” the perturbation, and $u^s(\mathbf{x})$ is the scattered field produced by the interaction of the perturbation with the impinging wave. The scattered field satisfies [16]

$$\begin{cases} (\Delta + \omega^2(1 + \eta(\mathbf{x})))u^s(\mathbf{x}) = -\omega^2\eta(\mathbf{x})e^{i\omega(\mathbf{s}\cdot\mathbf{x})} & \text{for } \mathbf{x} \in \mathbb{R}^2, \\ \lim_{|\mathbf{x}| \rightarrow \infty} |\mathbf{x}|^{1/2} \left(\frac{\partial}{\partial |\mathbf{x}|} - i\omega \right) u^s(\mathbf{x}) = 0, \end{cases} \quad (4)$$

following the setup depicted in Figure 1. For simplicity, we select the detector manifold D to be a circle of radius R that surround the domain of interest Ω as shown in Figure 1 (left). For each incoming direction $\mathbf{s} \in \mathbb{S}^1$ the data is given by sampling the scattered field with receiver elements that are located on D and indexed by $\mathbf{r} \in \mathbb{S}^1$. This yields the far-field pattern given by $\Lambda_{\mathbf{s},\mathbf{r}}^\omega = u^{sc}(R\mathbf{r})$, where \mathbf{s} denotes the incoming probing direction as defined in (4). We call $\mathcal{F}^\omega[\eta] := \Lambda_{\mathbf{s},\mathbf{r}}^\omega$ the *forward map* relating the perturbation η to its corresponding far-field pattern. In Figure 1 (right) we can observe the oscillatory behavior of a typical example of the far-field pattern for the perturbation in Figure 1 (center).

Following classical WKB asymptotic theory [45] and Fourier analysis [27] one can write

$$\Lambda_{\mathbf{s},\mathbf{r}}^\omega = (\mathcal{F}^\omega \eta)(\mathbf{x}) = \int_{\text{supp}(\eta)} a(\mathbf{x}, \mathbf{y}) e^{i\omega\phi(\mathbf{x},\mathbf{y})} \eta(\mathbf{y}) d\mathbf{y}. \quad (5)$$

where $\mathbf{x} = (\mathbf{s}, \mathbf{r})$, a and ϕ are functions. In this case a and ϕ are solutions to a transport and an eikonal equation respectively, that depend on η . Fortunately, in the Born approximation this operator can be reduced to a Fourier transform (see (13) in Appendix A).

²This assumption is only made to make the presentation more transparent.

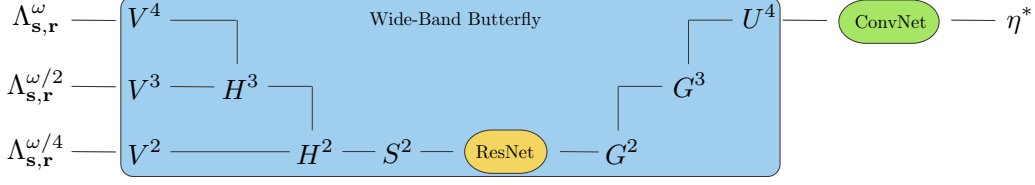


Figure 2: Sketch of the architecture of the Wide-band butterfly network for $L = 4$ used to approximate $\eta^* = (\mathcal{F})^{-1}[\Lambda_{\mathbf{s},\mathbf{r}}^\omega, \Lambda_{\mathbf{s},\mathbf{r}}^{\omega/2}, \Lambda_{\mathbf{s},\mathbf{r}}^{\omega/4}]$.

We define the *inverse map* as the map estimating η from $\Lambda_{\mathbf{s},\mathbf{r}}^\omega$, namely $(\mathcal{F}^\omega)^{-1}$. The inverse problem can be formulated as a minimization problem, i.e.,

$$(\mathcal{F}^\omega)^{-1}[\Lambda_{\mathbf{s},\mathbf{r}}^\omega] = \eta^* := \operatorname{argmin}_\mu \|\mathcal{F}^\omega[\mu] - \Lambda_{\mathbf{s},\mathbf{r}}^\omega\|^2 \quad (6)$$

for a suitable set of perturbations, and for a suitable norm. This problem is often solved using PDE-constrained optimization methods, or in the linearized case, using filtered back-projection (see Appendix A). However, this has been proven to result in unstable algorithms which can converge to non-physical local minima [28]. Thus, one key insight for designing robust algorithms, which is also incorporated in wide-band butterfly network, is to use wide-band data viz.,

$$(\mathcal{F})^{-1}[\Lambda_{\mathbf{s},\mathbf{r}}^{\omega_1}, \Lambda_{\mathbf{s},\mathbf{r}}^{\omega_2}, \dots, \Lambda_{\mathbf{s},\mathbf{r}}^{\omega_{n_\omega}}] = \eta^* := \operatorname{argmin}_\mu \sum_{i=1}^{n_\omega} w_i \|\mathcal{F}^{\omega_i}[\mu] - \Lambda_{\mathbf{s},\mathbf{r}}^{\omega_i}\|^2. \quad (7)$$

How to choose the frequencies $\{\omega_i\}_{i=1}^{n_\omega}$ and the corresponding weights $\{w_i\}_{i=1}^{n_\omega}$ remains an area of research [6]. In fact, this problem still present spurious local minima, thus regularization techniques such as recursive linearization are used [9, 47]. In a nutshell, recursive linearization (or frequency sweeps) can be understood at the algorithmic level as weights that change throughout the optimization loop: At the beginning only the low-frequency weights are not zero, and during the optimization the mass of the weights is slowly transferred to the weights involving higher frequencies. This is another key insight incorporated in to the design of the wide-band butterfly network: Instead of treating each frequency independently, processing low-frequency data provides guidance to process the higher-frequency data.

For the discrete problem we assume that Ω is discretized in $n_x \times n_y$ equispaced points in space, and that the far field data is probed following n_{sr} equispaced direction in \mathbb{S} , and that the scattered field is probed at n_{rs} equispaced points in D . For simplicity we assume that $n_x = n_y = n_{sr} = n_{rs}$.

3 Wideband Butterfly Network Architecture

The network is composed of a convolutional block and a wide-band butterfly block with a non-linear resnet in the middle following Figure 2. This construction is inspired by the filtered back-projection, which is ubiquitous in inverse problems (see Appendix A), and the butterfly factorization [40, 39]. In a nutshell, WideBNet *lifts* the linearized version of the inverse problem (see (14) in Appendix A) with the butterfly algorithm to a non-linear operator that is fully trainable. The inputs to the network are the data $\Lambda_{\mathbf{s},\mathbf{r}}^\omega$ at frequency ω , each of which is an $n_{sr} \times n_{rs}$ matrix reshaped into a tensor of dimensions $(2^L, 2^L, s^2)$. The first two dimensions of the reshaped tensor account for the geometric information in $2^L \times 2^L$ cells, and s^2 corresponds the size of the leaf nodes (see Figure 8).

Specific to the WideBNet are the local embedding blocks V^ℓ , the down-sampling layers H^ℓ , the switch layer $S^{L/2}$, the upsampling layers G^ℓ , and the interpolation layer U^L . For simplicity we describe the action of these components at the highest frequency, i.e. at the highest length scale L . For the full description we refer to the Appendix D or the original manuscript [37]. The implementation of the convolutional and resnet layers are standard and details on the dimensions are similarly left to the Appendices C and D.

- **Local embedding layer V^ℓ** takes input dimension $(2^\ell, 2^\ell, s^2)$ and processes each cell to produce a local representation as a vector of dimension r , thus outputting a tensor of dimensions $(2^\ell, 2^\ell, r)$. The first two components index the geometrical position of each cell and the last corresponds to the number of channels in the local embedding (see Figure 11 for a sketch of this operation).

- **Downsampling layer H^ℓ** merges information from contiguous cells. They downsample the number of cells by two in each axis and increase the information contained in the new aggregate cell by a factor four. If the input is $(2^\ell, 2^\ell, r)$ then the output is $(2^{\ell-1}, 2^{\ell-1}, 4r)$ (see Figure 10 for a sketch of this operation).
- **Switch layer $S^{L/2}$** was introduced in [33] to perform an all-to-all communication between the cells. The input and output dimension $(2^{L/2}, 2^{L/2}, 2^L r)$ since this layer broadcasts information in the third dimension across the first two in packages of size r^2 (see Figure 9 for a sketch of this operation).
- **Upsampling layer G^ℓ** is formally the adjoint of the downsampling layers, i.e., they split information in each cells into four contiguous cells, each with a fourth of the original aggregate cell. If the input is $(2^\ell, 2^\ell, 4r)$ then the output is $(2^{\ell+1}, 2^{\ell+1}, r)$.
- **Local sampling layer U^L** takes each cell of the input of dimension $(2^L, 2^L, r)$ which contains a local representation and they sample it in the corresponding leaf of size s^2 , thus the output is $(2^L, 2^L, s^2)$.

Observe from Figure 2 that the information is down-sampled after each H^ℓ layer to inject information at the correct length scales. This complies with Shannon-Nyquist sampling theory which describes the relation of the information content and sampling frequency.

4 Numerical Results

We benchmark the WideBNet architecture for inversion of scatterers with various geometric configurations and length scales. Our imaging domain spans $[-0.5, 0.5]^{\otimes 2}$ and is discretized over an equispaced mesh with $n = 80$ points per dimension. This corresponds to a quad-tree partitioning with $L = 4$ levels and $s = 5$ leaf nodes. We consider $n_\omega = 3$ number of frequencies at 2.5, 5.0, and 10.0 Hz in which the data was assimilated at levels $\ell = 2, 3$, and 4 respectively. We consider a homogeneous background squared slowness of $m_0 = 1$ (or equivalently a background wave-speed of $c_0 = 1$) which corresponds to 8 points per wavelength (PPW) at the highest frequency and limits the classical Rayleigh resolution to $\lambda_{\min}/2 = 4$ pixels. See Appendices C and D for details the choice of hyperparameters and implementation details, respectively.

The synthetic scattering data was generated with second-order finite differences for training and fourth-order finite differences for testing. The use of different stencils validates against overfitting of numerical dispersion artifacts. The radiating boundary conditions in (4) were implemented using perfectly matched layers (PML) spanning one wavelength with quadratic profile of intensity 80 [4]. The scattered waves were sampled at 80 equi-angular intervals on a circle of radius $r = 0.5$ where the incident waves arrived at angles aligned with the receiver geometry. The training and testing split was 21000 points and 4000 points, respectively. Note each datapoint is a tensor of dimension $n_{sr} \times n_{rs} \times n_\omega$, where $n_{sr} = n_{rs} = 80$ and $n_\omega = 3$.

The data at each frequency was normalized by subtracting and scaling by the mean and variance of the aggregate pixel intensities in the training dataset. Whereas [37] examines WideBNet with noiseless data, here we consider a multiplicative noise model similar to [33] with

$$\Lambda_{\mathbf{s}, \mathbf{r}} = \Lambda_{\mathbf{s}, \mathbf{r}}^{\text{clean}} (1 + \epsilon(\mathbf{s}, \mathbf{r})), \quad \epsilon(\mathbf{s}, \mathbf{r}) \stackrel{iid}{\sim} \mathcal{N}(0, \sigma^2). \quad (8)$$

Since each frequency is standardized independently this corresponds to multiplicative coloured noise in the time series data. We set $\sigma = 1$ to introduce 100% noise to signal ratio in our data. The noise is applied to each data pixel dynamically each epoch.

The objective function we used was the pixelwise sample loss

$$\sum_{x=\text{pixel in image}} \|(K_{\text{high}} * \eta)(x) - \text{WideBNet}[\Lambda_{s,r}^{\omega_1}, \dots, \Lambda_{s,r}^{\omega_{n_\omega}}](x)\|_2^2$$

where η is the total scattering medium and $\{\Lambda_{s,r}^{\omega_i}\}_{i=1}^{n_\omega}$ the noisy multifrequency data. The scattering medium data was smoothed using a high-pass filter K_{high} described by a Gaussian kernel with characteristic width of 0.75 grid points. This was observed to promote faster training in [37] for noiseless data. Since the support of this filter is significantly less than the Nyquist limit this processing does not pollute or eliminate sub-wavelength features in the medium.

Table 1: Mean Squared Error per pixel on testing data with $\sigma = 1$ multiplicative noise.

	WideBNet	UNet	SwitchlessBNet	NarrowBNet	SwitchNet
Blob	5.5996E-04	4.3213E-03	2.6254E-03	5.9252E-04	6.9086E-04
Square 3/5/10h	7.9643E-03	1.3671E-01	1.5380E-01	7.9946E-03	1.9018E-01
Triangle 3/5/10h	8.1558E-03	1.4777E-01	1.3805E-01	8.7101E-03	6.3816E-03
Gaussian 2h	2.5700E-03	6.2389E-02	1.1298E-01	4.7489E-03	2.4208E-03
Ten Squares 10h	3.1379E-02	1.0738E-01	4.0320E-01	9.4893E-02	8.9998E-02
Shepp Logan	7.2137E-03	3.8251E-02	3.1400E-02	1.1412E-02	8.6461E-03
GRF ($\sigma=0.04$)	3.3387E-03	1.2358E-02	2.1289E-01	9.0839E-03	9.1764E-03
GRF ($\sigma = 0.01$)	1.3585E-02	4.4441E-02	4.7032E-01	5.3533E-02	5.7688E-02

For completeness we also investigate the performance of comparable architectures such as: the popular image processing architecture U-Net where we feed in multifrequency data [49]; and SwitchNet, which similarly uses the butterfly architecture but requires superlinear $\mathcal{O}(N^{3/2})$ complexity for training stability when using single frequency data. We also consider a narrow-band implementation of our WideBNet architecture referred to as NarrowBNet which uses only 10 Hz data assimilated at level $L = 4$ and otherwise coincides with the architecture proposed in [38]. Furthermore we benchmark SwitchlessBNet, which has identical degrees to freedom as WideBNet except for the omission of the $S^{L/2}$ layer. We refer the interested reader to Appendices A and D for a comparison against more classical methods such as filtered back-projection (see Figure 7 and PDE-constrained optimization methods (see Figure 13), in which WideBNet outperforms both methods, particularly in the noisy case.

4.1 Performance across classes

Table 1 and Figure 3 depicts the performance of various architectures for different scattering media. These media comprehensively cover a wide range of wave scattering behaviour: **Randomized Shepp-Logan** is a standard biomedical imaging dataset which emulates contrasting material properties inside the human head [12]; **Gaussian Random Field** propagate waves with strong multipathing and weak coda. The parameter $\sigma = \{0.01, 0.04\}$ refers to the characteristic length scale of the fluctuations; **Blob** contains a single non-convex scatterer with wavelength level features, **Ten-Squares 10h** corresponds to ten randomly located squares each with 10 pixel sidelengths. Waves traveling in this medium exhibit strong multiple scattering; **Gaussian 2h** correspond to point scatterers at the highest resolution limits of the medium; **Squares 3/5/10h** are randomly located squares of 3, 5, and 10 pixel sidelengths. This covers multiple lengthscales ranging from subwavelength, wavelength, and superwavelength scatterers; **Triangles 3/5/10h** are similar to squares except with right triangles. For the geometric scatterers the amplitude of each scatterer is fixed at $\Delta m = 0.2$, however superposition of scatterers can lead to higher local wavespeeds. The geometric mediums contain uniformly sampled range of 2, 3, or 4 scatterers which are uniformly located in a circle of radius $r = 0.35$. All scattering mediums lie in a homogeneous background of $m_0 = 1$.

Quantitatively in Tab. 1 we observe that WideBNet generally provides the lowest average testing errors across all datasets. Notably, SwitchlessBNet, which has an identical architecture to WideBNet (as seen in Figure 2) except eliminating the switch layer, performs two orders of magnitude worse on average. This is consistent with intuition from FIO theory where the switch permutation is necessary to capture the long-range (i.e. local-to-global) interactions inherent to wave scattering data. Although SwitchNet also uses this switch permutation layer it does not outperform WideBNet despite requiring an additional million trainable weights in its architecture (see Appendix C.3 for the total number of trainable parameters of each network). The U-Net architecture performs similarly poorly indicating that the multiscale behaviour in the dataset is difficult to capture with purely convolutional filters. Using solely narrow-band (single frequency) data provides comparable behaviour on simple geometric datasets, but suffers when capturing more complex media such as the Shepp-Logan or gaussian random fields. We observe that the Blob dataset is the easiest to invert for all architectures.

Qualitatively in Figure 3 we observe that SwitchlessBNet is unable to fully resolve all of the gaussian scatterers. Moreover, although it is able to resolve the large length scale scatterers in both the square and triangle datasets, the subwavelength scatterers are notably omitted. These

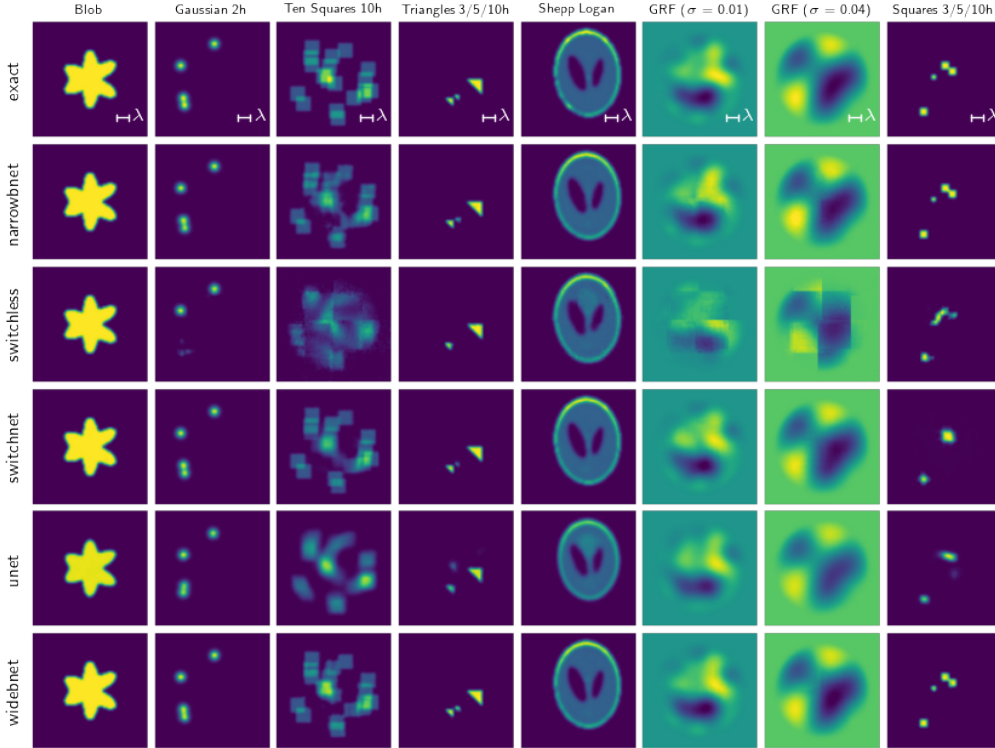


Figure 3: Comparison of performance between different architectures (rows) on various datasets (columns). The scattering data was contaminated with $\sigma = 1$ multiplicative noise following (8).

observations indicate that SwitchlessBNet struggles in the super-resolution regime. Note that the U-Net incorrectly orients the colliding point scatterers for the gaussian dataset. A similar phenomenon for the U-Net is observed for the Ten-Squares dataset which features multiple scatterers with subwavelength separation. Both NarrowBNet and SwitchNet also struggle with densely overlapping scatterers, whereas WideBNet does not. This reinforces the claim that wideband data stabilizes inversion. Also observe that NarrowBNet introduces mild artifacts in the gaussian random field with $\sigma = 0.04$, while SwitchNet images the wrong square. In all instances the Shepp-Logan phantom is well resolved with the exception of U-Net which incorrectly scales the phantom. Lastly we note that both the NarrowBNet and SwitchNet architecture are surprisingly able to resolve certain subwavelength features despite only provided single frequency data. We speculate that they utilize geometric triangulation to achieve these images.

4.1.1 Out of Distribution Generalization

We report on the unexpected ability of our trained networks to generalize **across** datasets. However we caution that this behaviour appears to be strongly dependent on the dataset – in particular, we only observe successful instances of generalization when training with the Square 3/5/10h and Triangle 3/5/10h datasets. Figure 4 documents this for the narrow- and wide-band architectures where the shape in parentheses in the column titles indicates the training set. A rough numerical comparison of extrapolative properties between datasets can be found in the heatmap shown in Figure 14 in the appendix.

Training WideBNet with the square 3/5/10 dataset appears to provide the best out of distribution performance, both qualitatively and quantitatively. For instance, observe that the wide-band network successfully captures the boundaries of the out-of-distribution blob in Figure 4. We emphasize that training points in the square dataset contain at most four scatterers – therefore this ability the infill the blob with multiple, i.e. greater than four, scatterers suggests the network has learned some generalized

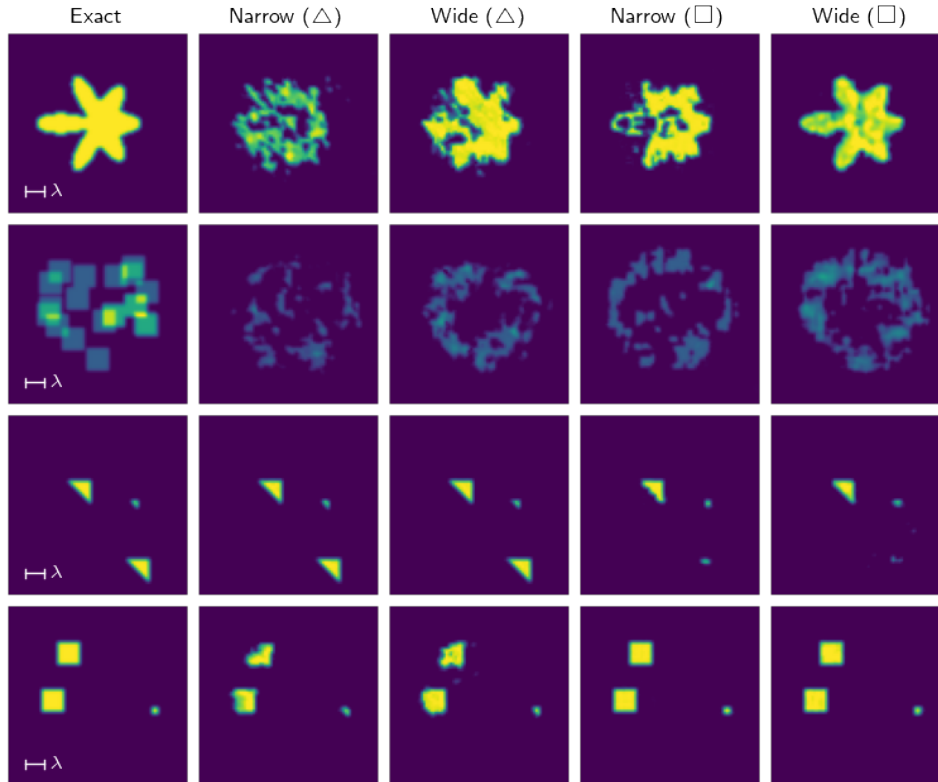


Figure 4: Example of out-of-distribution generalization properties of `NarrowBNet` and `WideBNet`. The shape in parentheses indicate the training set Δ (resp. \square) corresponds to Triangle 3/5/10 (resp. Square 3/5/10). Scattering data was polluted by $\sigma = 1$ even for the out-of-distribution test cases.

physics of wave scattering. Unexpectedly, however, this behaviour is not repeated for the Ten Squares dataset despite the fact that squares with 10 pixel sidelengths are already present in the training set.

Furthermore, note that both the squared trained `NarrowBNet` and `WideBNet` are able to localize and resolve triangular scatterers. In both instances however the large triangular scatterer on the bottom is mistakenly characterized as a small triangular scatterer; it is unclear why both networks share this deficiency. Surprisingly, triangular trained networks do not generalize as well to the square dataset even though each square can be expressed as a union of two (appropriately rotated) triangles. We leave to future work a systematic investigation of how training datasets can be designed to promote better out of distribution behaviour on trained `WideBNet` networks. However we note that these results strongly indicate that the networks learn a far more complicated signal processing algorithm than simply joint blind deconvolution with super resolution.

4.2 Super Resolution beneath the Diffraction Limit

We report on the claimed super-resolution capabilities of our trained networks: for a comparison with optimization-based methods please see Figure 13. Figure 5 depicts the testing performance on a scattering medium with three squares of 3, 5, and 10 pixel side length that are separated by a distance of Δ^3 . All datasets are polluted with $\sigma = 1$ multiplicative noise. When $\Delta \gg \lambda$ we observe that the both `NarrowBNet` and `WideBNet` successfully image all scatterers, although `NarrowBNet` struggles with the boundaries of the largest scatterer. However, when $\Delta \leq \lambda$ we note that only `WideBNet` succeeds in imaging, as claimed. `U-Net` performs poorly over the entire range of

³Due to the way the scatterers were defined in code the Δ is not exactly the separation between the *boundaries* of the scatterers.

Δ separation as it resolves the largest scatterer but ignores either one or both of the remaining small scatterers.

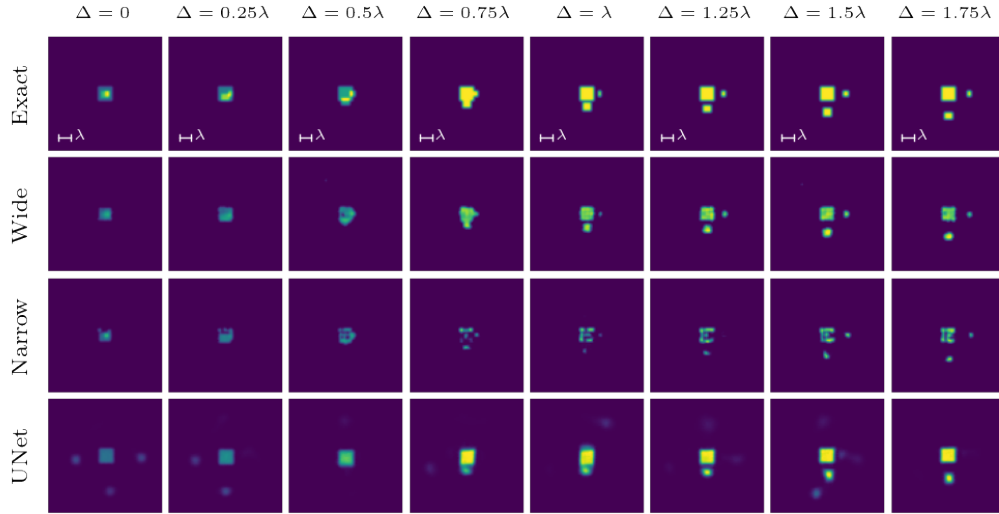


Figure 5: Comparison of imaging capabilities of different networks for square scatterers with side lengths 3, 5, and 10 pixels, which are separated by Δ . We vary Δ between sub-wavelength and wavelength separations (one wavelength $\lambda = 8$ pixels).

4.3 Stability with Noise

Lastly we consider the robustness of the learned networks to noise. Figure 6 depicts the performance of a network trained with $\sigma = 1$ noise with testing data that is polluted by different noise levels. The network demonstrates strong denoising abilities as it produces a similar images even for $\sigma = 2$, i.e. a noise level two times higher than what it saw in training. Conversely, the imaging performance also does not suffer when the input data is noiseless with $\sigma = 0$. We note that the image quality degrades gracefully with increasing σ ; while remnants of the large scatterer can still be identified, the small scatterer is not longer reliably imaged at higher noise levels.

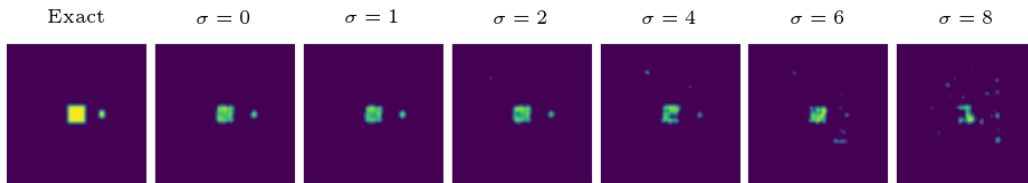


Figure 6: Degradation of image reconstruction with increasing noise-to-signal ratio σ . The network was trained with data contaminated with $\sigma = 1$ noise.

5 Conclusions

We compared WideBNet against a range of competing network architectures on a suite of benchmark scattering configurations. Both quantitative, and qualitatively, WideBNet provides better performance. In particular, WideBNet excels in the super-resolution regime whereas other trained networks fail to consistently produce acceptable images. This super-resolution behaviour was found to extend to higher levels of noise than in the training set, though this breaks down for sufficiently high noise levels. Additional mechanisms for robustness to noise remains an open problem. The WideBNet networks were also found to have surprising generalizability across datasets. However, this behaviour was

limited only to certain training datasets. The characterization of “good” training datasets, and/or strategies to promote transfer learning across datasets, is left to future work.

Acknowledgements

The authors thank Total S.A. for support. L.D. is also supported by AFOSR grant FA9550-17-1-0316. L.Z.-N. is also supported in part by the Wisconsin Alumni Research Foundation, the National Science Foundation under the grant DMS-2012292, and NSF TRIPODS award 1740707.

References

- [1] G. Ambartsoumian, R. Felea, V. P. Krishnan, C. Nolan, and E. T. Quinto. A class of singular fourier integral operators in synthetic aperture radar imaging. *Journal of Functional Analysis*, 264(1):246–269, 2013.
- [2] G. Backus and F. Gilbert. The Resolving Power of Gross Earth Data. *Geophys. J. Int.*, 16(2):169–205, 10 1968.
- [3] E. Baysal, D. D. Kosloff, and J. W. C. Sherwood. Reverse time migration. *GEOPHYSICS*, 48(11):1514–1524, 1983.
- [4] J.-P. Berenger. A perfectly matched layer for the absorption of electromagnetic waves. *J. Comput. Phys.*, 114:185–200, 1994.
- [5] B. Biondi. *3D Seismic Imaging*. Society of Exploration Geophysicists, 2006.
- [6] C. Borges, A. Gillman, and L. Greengard. High resolution inverse scattering in two dimensions using recursive linearization. *SIAM J. Imaging Sci.*, 10(2):641–664, 2017.
- [7] S. Börm, C. Börst, and J. M. Melenk. An analysis of a butterfly algorithm. *Comput. Math. Appl.*, 74(9):2125 – 2143, 2017. Advances in Mathematics of Finite Elements, honoring 90th birthday of Ivo Babuška.
- [8] E. Candès, L. Demanet, and L. Ying. A fast butterfly algorithm for the computation of fourier integral operators. *Multiscale Model. Sim.*, 7(4):1727–1750, 2009.
- [9] Y. Chen. Inverse scattering via Heisenberg’s uncertainty principle. *Inverse Probl.*, 13(2):253, 1997.
- [10] Y. Chen, L. Lu, G. E. Karniadakis, and L. Dal Negro. Physics-informed neural networks for inverse problems in nano-optics and metamaterials. *Opt. Express*, 28(8):11618–11633, Apr 2020.
- [11] M. Cheney. A mathematical tutorial on synthetic aperture radar. *SIAM Rev.*, 43(2):301–312, 2001.
- [12] M. Chung. Random shepp-logan phantom. <https://github.com/matthiaschung/Random-Shepp-Logan-Phantom>, 2018.
- [13] A. Colli, D. Prati, M. Fraquelli, S. Segato, P. P. Vescovi, F. Colombo, C. Balduini, S. Della Valle, and G. Casazza. The use of a pocket-sized ultrasound device improves physical examination: Results of an in- and outpatient cohort study. *PLOS ONE*, 10(3):1–10, 03 2015.
- [14] D. Colton and A. Kirsch. A simple method for solving inverse scattering problems in the resonance region. *Inverse Problems*, 12(4):383–393, aug 1996.
- [15] D. Colton and R. Kress. *Integral Equation Methods in Scattering Theory*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2013.
- [16] D. Colton and R. Kress. *Inverse Acoustic and Electromagnetic Scattering Theory*. Springer-Verlag New York, New York, PA, 3 edition, 2013.

- [17] J. W. Cooley and J. W. Tukey. An algorithm for the machine calculation of complex Fourier series. *Math. Comput.*, 19(90):297–301, 1965.
- [18] T. A. Davis. Algorithm 832: UMFPACK v4.3—an unsymmetric-pattern multifrontal method. *ACM Transactions on Mathematical Software*, 30(2):196–199, June 2004.
- [19] D. L. Donoho. Superresolution via sparsity constraints. *SIAM Journal on Mathematical Analysis*, 23(5):1309–1331, 1992.
- [20] Y. Fan, J. Feliu-Fabà, L. Lin, L. Ying, and L. Zepeda-Núñez. A multiscale neural network based on hierarchical nested bases. *Research in the Mathematical Sciences*, 6(2):21, Mar 2019.
- [21] Y. Fan, L. Lin, L. Ying, and L. Zepeda-Núñez. A multiscale neural network based on hierarchical matrices. *arXiv:1807.01883*.
- [22] Y. Fan and L. Ying. Solving inverse wave scattering with deep learning. *arXiv:1911.13202*.
- [23] Y. Fan and L. Ying. Solving traveltime tomography with deep learning. *arXiv:1911.11636*.
- [24] R. Felea and E. T. Quinto. The microlocal properties of the local 3-d spect operator. *SIAM Journal on Mathematical Analysis*, 43(3):1145–1157, 2011.
- [25] B. Gutenberg. Ueber erdbebenwellen. vii a. beobachtungen an registrierungen von fernbeben in göttingen und folgerung über die konstitution des erdkörpers (mit tafel). *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse*, 1914:125–176, 1914.
- [26] R. J. Hewett and L. Demanet. Pysit, 2013.
- [27] L. Hörmander. *The Analysis of Linear Partial Differential Operators. IV: Fourier Integral Operators*, volume 63 of *Classics in Mathematics*. Springer, Berlin, 2009.
- [28] P. Hähner and T. Hohage. New stability estimates for the inverse acoustic inhomogeneous medium problem and applications. *SIAM Journal on Mathematical Analysis*, 33(3):670–685, 2001.
- [29] P. Isola, J. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. pages 5967–5976, July 2017.
- [30] A. Javaherian, F. Lucka, and B. T. Cox. Refraction-corrected ray-based inversion for three-dimensional ultrasound tomography of the breast. *Inverse Problems*, 36(12):125010, dec 2020.
- [31] I. Kang, A. Goy, and G. Barbastathis. Dynamical machine learning volumetric reconstruction of objects’ interiors from limited angular views. *Light: Science & Applications*, 10(1):74, 2021.
- [32] Y. Khoo, J. Lu, and L. Ying. Solving parametric PDE problems with artificial neural networks. *arXiv preprint arXiv:1707.03351*, 2017.
- [33] Y. Khoo and L. Ying. SwitchNet: A neural network model for forward and inverse scattering problems. *SIAM J. Sci. Comput.*, 41(5):A3182–A3201, 2019.
- [34] D. Kingma and J. Ba. Adam: a method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, may 2015.
- [35] A. Kirsch and N. Grinberg. *The Factorization Method for Inverse Problems*. Oxford University Press, Oxford, first edition, 2008.
- [36] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi. Photo-realistic single image super-resolution using a generative adversarial network. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 105–114, July 2017.
- [37] M. T. Li, L. Demanet, and L. Zepeda-Núñez. Wide-band butterfly network: stable and efficient inversion via multi-frequency neural networks. *ArXiv e-prints*, [math.NA] 2011.12413, 2020.

- [38] Y. Li, X. Cheng, and J. Lu. Butterfly-Net: Optimal function representation based on convolutional neural networks. *arXiv preprint arXiv:1805.07451*, 2018.
- [39] Y. Li, H. Yang, E. Martin, K. Ho, and L. Ying. Butterfly factorization. *Multiscale Model. Sim.*, 13(2):714–732, 2015.
- [40] Y. Liu, X. Xing, H. Guo, E. Michielssen, and X. S. Ghysels, P. Li. Butterfly factorization via randomized matrix-vector multiplications. *arXiv:2002.03400*.
- [41] F. Lucka, M. Pérez-Liva, B. E. Tinverse problems in nano reeby, and B. T. Cox. High resolution 3d ultrasonic breast imaging by time-domain full waveform inversion. *ArXiv e-prints*, [math.NA] 2102.00755, 2021.
- [42] M. Mirza and S. Osindero. Conditional generative adversarial nets, 2014. cite arxiv:1411.1784.
- [43] F. Natterer. *The Mathematics of Computerized Tomography*. Society for Industrial and Applied Mathematics, 2001.
- [44] R. D. Oldham. The constitution of the interior of the earth, as revealed by earthquakes. *Quarterly Journal of the Geological Society*, 62(1-4):456–475, 1906.
- [45] F.W.J. Olver and W. Rheinbolt. *Asymptotics and Special Functions*. Elsevier Science, 2014.
- [46] T. J.P.M. Op ’t Root, C. C. Stolk, and M. V. de Hoop. Linearized inverse scattering based on seismic reverse time migration. *Journal de Mathématiques Pures et Appliquées*, 98(2):211–238, 2012.
- [47] R. G. Pratt. Seismic waveform inversion in the frequency domain; part 1: Theory and verification in a physical scale model. *Geophysics*, 64(3):888–901, 1999.
- [48] E. T. Quinto, A. Rieder, and T. Schuster. Local inversion of the sonar transform regularized by the approximate inverse. *Inverse Problems*, 27(3):035006, feb 2011.
- [49] O. Ronneberger, P. Fischer, and T. Brox. *U-Net: Convolutional Networks for Biomedical Image Segmentation*, pages 234–241. Springer International Publishing, Cham, 2015.
- [50] Y. Saad and M. H. Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing*, 7(3):856–869, July 1986.
- [51] J. F. Schenck. The role of magnetic susceptibility in magnetic resonance imaging: Mri magnetic compatibility of the first and second kinds. *Medical Physics*, 23(6):815–850, 1996.
- [52] P. Stefanov, G. Uhlmann, A. Vasy, and H. Zhou. Travel time tomography. *Acta Math. Sin.*, 35(6):1085–1114, Jun 2019.
- [53] A. Tarantola. Inversion of seismic reflection data in the acoustic approximation. *Geophysics*, 49(8):1259–1266, 1984.
- [54] J. Virieux, A. Asnaashari, R. Brossier, L. Métivier, A. Ribodetti, and W. Zhou. *6. An introduction to full waveform inversion*, pages R1–1–R1–40. 2017.
- [55] J. Virieux and S. Operto. An overview of full-waveform inversion in exploration geophysics. *Geophysics*, 74:WCC1–WCC26, 2009.
- [56] J. Wiskin, B. Malik, D. Borup, N. Pirshafiey, and J. Klock. Full wave 3d inverse scattering transmission ultrasound tomography in the presence of high contrast. *Scientific Reports*, 10(1):20166, 2020.
- [57] Z. Xu, Y. Li, and X. Cheng. Butterfly-Net2: Simplified Butterfly-Net and Fourier transform initialization. volume 107 of *Proceedings of Machine Learning Research*, pages 431–450, Princeton University, Princeton, NJ, USA, 20–24 Jul 2020. PMLR.
- [58] Z.-Q. Xu, Y. Zhang, T. Luo, Y. Xiao, and Z. Ma. Frequency principle: Fourier analysis sheds light on deep neural networks. *ArXiv e-prints*, 2019.

A Filtered Back-Projection

We can cast the inverse problem for recovering the perturbation as

$$\eta^* = \operatorname{argmin}_{\mu} \|\mathcal{F}^\omega[\mu] - \Lambda_{\mathbf{s},\mathbf{r}}\| \quad (9)$$

where $\Lambda_{\mathbf{s},\mathbf{r}}$ is the measured data. We linearize \mathcal{F}^ω to shed light on the essential difficulties of this problem. Using the classical Born approximation in (4) we obtain that

$$u^{sc}(\mathbf{x}) = \omega^2 \int_{\mathbb{R}^2} \Phi^\omega(\mathbf{x}, \mathbf{y}) \eta(\mathbf{y}) e^{i\omega(\mathbf{s} \cdot \mathbf{y})} d\mathbf{y}, \quad (10)$$

where Φ^ω is the Green's function of the two-dimensional Helmholtz equation in homogeneous media, i.e., Φ^ω satisfies

$$\begin{cases} (\Delta + \omega^2) \Phi^\omega(\mathbf{x}, \mathbf{y}) = -\delta(\mathbf{x}, \mathbf{y}) & \text{for } \mathbf{x} \in \mathbb{R}^2, \\ \lim_{|\mathbf{x}| \rightarrow \infty} |\mathbf{x}|^{1/2} \left(\frac{\partial}{\partial |\mathbf{x}|} - i\omega \right) \Phi^\omega(\mathbf{x}, \mathbf{y}) = 0. \end{cases} \quad (11)$$

Furthermore, we can use the classical far-field asymptotics of the Green's function to express

$$u^{sc}(R\mathbf{r}) = -\omega^2 \frac{e^{i\omega R}}{\sqrt{R}} \int_{\mathbb{R}^2} \eta(\mathbf{y}) e^{i\omega(\mathbf{s}-\mathbf{r}) \cdot \mathbf{y}} d\mathbf{y} + \mathcal{O}(R^{-3/2}). \quad (12)$$

Thus, up to a re-scaling and a phase change, the far-field pattern defined in (13) can be approximately written as a Fourier transform of the perturbation, i.e.,

$$\Lambda_{\mathbf{s},\mathbf{r}}^\omega \approx F^\omega \eta = -\omega^2 \frac{e^{i\omega R}}{\sqrt{R}} \int_{\mathbb{R}^2} e^{i\omega(\mathbf{s}-\mathbf{r}) \cdot \mathbf{y}} \eta(\mathbf{y}) d\mathbf{y} \quad (13)$$

is the linearized forward operator acting on the perturbation.

Solving the inverse problem (9) using the linearized operator in (13) results in an explicit solution given by the normal equation

$$\eta^* = ((F^\omega)^* F^\omega + \epsilon I)^{-1} (F^\omega)^* \Lambda_{\mathbf{s},\mathbf{r}}, \quad (14)$$

which is often also referred to as filtered back-projection [15]. The constant ϵ is a small regularization parameter that remedies the ill-conditioning of $(F^\omega)^* F^\omega$. In this case, $(F^\omega)^* F^\omega$ is translation invariant thus it is a convolutional operator, and $(F^\omega)^*$ is a Fourier transform thus highly oscillatory.

We point out that in the case that the background slowness squared is non-constant the operator in (13) becomes a FIO similar to (5) where a and ϕ depend on the background medium, which bring us back to the case covered in (9).

Finally, we used (14) to solve the inverse scattering problem for the noiseless and noisy case for different perturbations (see Figure 7). In this case the regularization parameter ϵ was chosen, after a labour intensive search, to be 10^{-3} and 1 in the noiseless and noisy cases, respectively. The linear system was solved iteratively using GMRES [50] with a tolerance of 10^{-4} and a restart of 10 steps. We can observe in Figure 7 that for the perturbations in which the assumptions of weak scattering hold, the filtered back-projection is able to localize the reflectors, albeit with some artifacts that become dominant as noise is incorporated. When the perturbation becomes delocalized or involving several scales the reconstruction deteriorates greatly in both noisy and noiseless cases. In contrast, WideBNet is completely agnostic to these issues and it provides an accurate reconstruction in the three different cases (which the network did not see during the training stage) regardless of the noise.

B Implementation Details

We provide a succinct and self-contained description of the architecture. For a more detailed exposition please see [37].

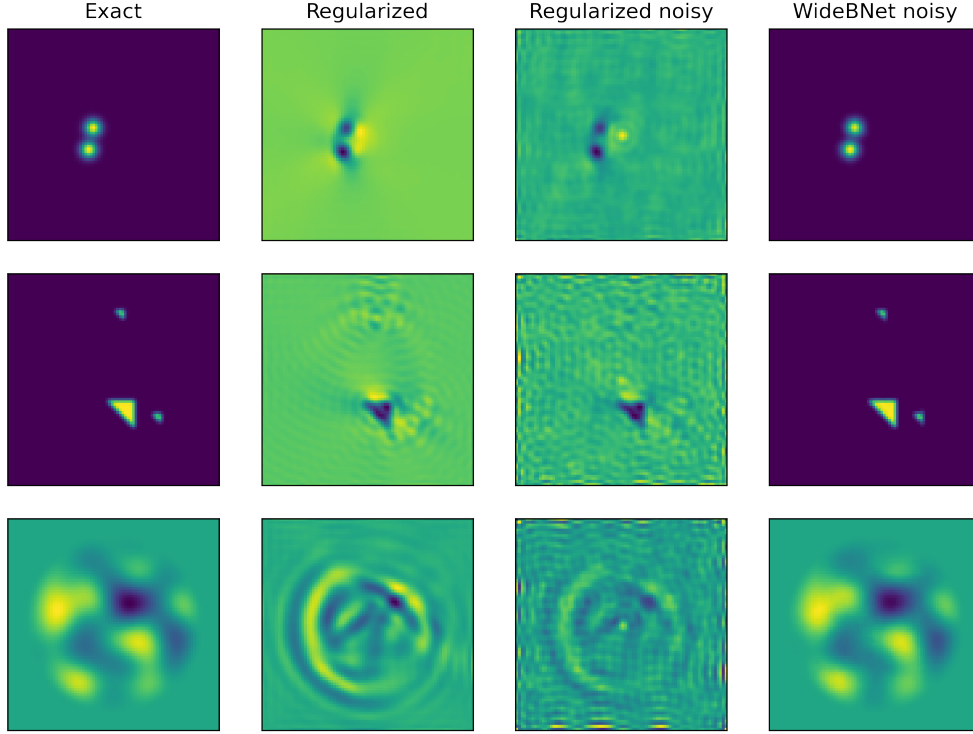


Figure 7: Comparison of reconstruction for inversion of the FIO in (14) using data at 10 Hz and WideBNet using wideband data at 2.5, 5 and 10 Hz. The two middle images have been rescaled to visually match the other two. The noise is injected to the data following (8) with $\sigma = 1$.

B.1 Input formatting

We assume the scatterers (discretized over an $n_x \times n_z$ grid) and the scattered data (an $n_{\text{src}} \times n_{\text{rev}}$ matrix for each frequency ω) are represented using complete quad-trees with L levels with leaf size s . This translates to a discretization into $n = 2^\ell s$ points for each matrix dimension, where $\ell \in [L/2, L]$, the level, indexes the size of the contiguous $2^{L-\ell} s \times 2^{L-\ell} s$ sub-matrices. The choice of L and s is informed by the wavelength scaling in each problem such that the data inside each $s \times s$ cell is non-oscillatory, i.e. contains a fixed amount of oscillations.

Following the Tensorflow convention of [height, width, channels] we reshape these quad-trees into three-tensors of size $[2^\ell, 2^\ell, s^2]$ as shown in Figure 8. The first two dimensions of the tensor contain the geometrical information, and the last contains a local representation. The data describing the local representation inside each voxel corresponds to *channels*. Following this description, we consider the slices along the height and width dimensions, i.e. the geometrical dimensions, as *cells*, e.g. a 1×1 cell of data describes slices with dimension $[1, 1, s^2]$. At the highest spatial resolution WideBNet operates on 1×1 cells, and at the lowest spatial resolution it operates on $2^{L/2} \times 2^{L/2}$ cells.

The input data $\Lambda_{s,r}^\omega \in \mathbb{C}^{n_{\text{sr}} \times n_{\text{rs}}}$ is a function of the probe frequency ω in (4). The input data to be collected from a wide-band of frequencies $\{\omega_\ell\}_{\ell=L/2}^L = \{\omega/2^{L-\ell}\}_{\ell=L/2}^L$, i.e. we use a dyadic partition containing $L/2 + 1$ intervals. With a slight abuse of notation we denote the resulting dataset as $\Lambda_{s,r}^\ell := \Lambda_{s,r}^{\omega/2^{L-\ell}} \in \mathbb{C}^{n_{\text{sr}} \times n_{\text{rev}}}$. Following the quad-tree structure we reshape each data tensor $\Lambda_{s,r}^\ell$ into a three-tensor of size $[2^\ell, 2^\ell, s^2]$. The input to WideBNet thus consists of the collection $\{\Lambda_{s,r}^\ell\}_{L/2 \leq \ell \leq L}$.

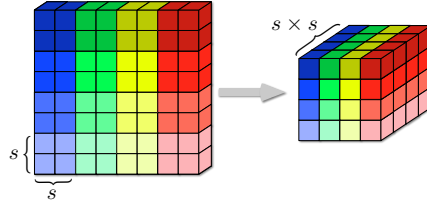


Figure 8: Sketch of the input at each level where the image of dimensions $[2^\ell s, 2^\ell s]$ is transformed to the tensorized form of size $[2^\ell, 2^\ell, s^2]$, for $\ell = 2$ and $s = 2$.

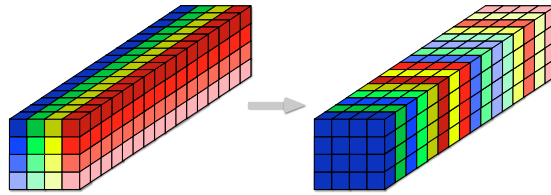


Figure 9: Sketch of the permutation of information of at the $S^{L/2}$ layer. In this case we consider $r = 1$, and $L = 4$. The tensor at the middle level is $[2^{L/2}, 2^{L/2}, r4^{L/2}] = [4, 4, 16]$. Given that the last dimension is the number of channels, we can observe that the information contained at each cells is distributed along all the other cells. This operation can be easily performed by a combination of reshape and dimension permutation.

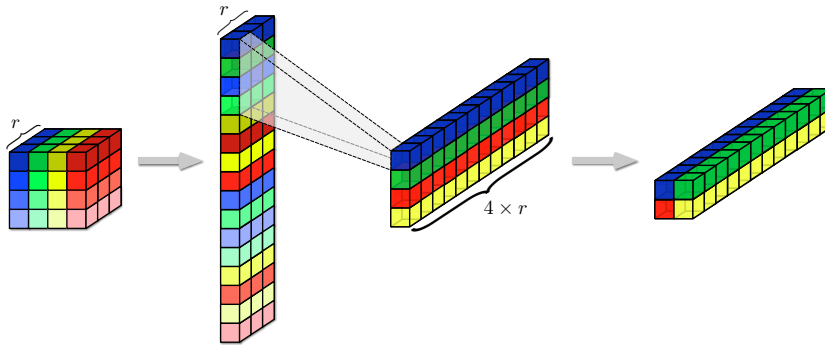


Figure 10: Sketch of the application of the H^ℓ layer, for convenience we illustrate the reshaping taking place using a Z-ordering. The layer decimates by a factor of four the number of neurons in the spatial dimensions (a factor two in each direction), while increasing four times the number of channels.

B.2 Architecture Overview

A schematic diagram of WideBNet is shown in Figure 2. At a high level, WideBNet consists of $L + 3$ specialized layers that are non-linear analogues of the butterfly factors in the butterfly factorization. In this case all the layers are linear, and we introduce non-linearities RESNET module. These layers ultimately send data into a CONVNET module, which mimic the effects of the regularized pseudo-inverse in sharpening the image/estimate as shown in (14).

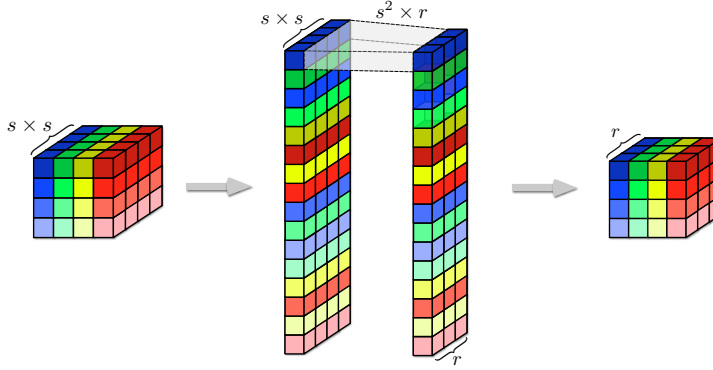


Figure 11: Sketch of the compression carried in the V^L layer, from the points contained in a leaf of size $s \times s$ to a local representation of rank r . We illustrate a reshape operation to a Z-ordering vector to illustrate local connectivity of the layer. The grey polygon represent the connections between the two layers.

The main advantage of WideBNet compared to other butterfly-based architectures, such as BNet [38] or SwitchNet [33] networks, is in how multi-frequency datasets are assimilated by exploiting the connection between spatial resolution and frequency. In particular, we stress that it is both

- (i) the connectivity/permutations inside the specialized layers $\{H^\ell\}$ and $\{G^\ell\}$ that process the wide-band data, as well as,
- (ii) the non-linearities induced by the middle $S^{L/2}$ and RESNET layer, as shown in Figure 2,

that are crucial towards achieving stable training dynamics, as well as image super-resolution, which is argued in [19].

By choosing a dyadic partition of the frequency band WideBNet exploits the inherent multiscale nature of the $\{H^\ell\}$ layers. Each $\{H^\ell\}$ layer only locally interpolates the data over voxel cells of effectively $2^{L-\ell} \times 2^{L-\ell}$, i.e. the effective length scales at this layer are of order $2^{L-\ell}$. It follows from the dispersion relation in wave-scattering that only data of bandwidth ω^ℓ are informative at this length-scale. This strategy of dyadically partitioning the bandwidth to localize spatial information is also employed by the Cooley-Tukey FFT algorithm to achieve quasi-linear time complexity [17]; in our setting this strategy affords us significant reductions in the number of trainable weights in the network, thus reducing the number of training samples and the computational complexity.

B.3 U^L and V^ℓ layers

The V^ℓ layer can be represented by a block diagonal matrix with block size $r \times s^2$. This layer takes input data (viewed as a complete quad-tree) and compresses the leaf nodes at level L , each with $s \times s$ degrees of freedom, into 1×1 cells with $\sqrt{r} \times \sqrt{r}$ degrees of freedom as depicted in Figure 11. Similarly, the U^L layer can also be represented by block diagonal matrix, albeit with block sizes of $s^2 \times r$. This layer ‘‘samples’’ the local representation back to its nominal dimensions. In both instances the compression/decompression is essentially lossless provided the number of levels L is properly adapted to the probe frequency ω . Provided that these parameters are chosen correctly, then the data is non-oscillatory (i.e. sub-wavelength) over $s \times s$ length-scales and therefore admits a low-rank representation with rank r .

We also utilize layers V^ℓ for $L/2 \leq \ell \leq L$ whose inputs are assumed to be sampled from bandwidth ω^ℓ . Each V^ℓ layer compresses the input data at level ℓ such that nodes with $2^{L-\ell}s \times 2^{L-\ell}s$ degrees of freedom are mapped into 1×1 cells with $\sqrt{r} \times \sqrt{r}$ degrees of freedom; this also has the interpretation of spatial downsampling. Note that the dyadic scaling in the definition of ω^ℓ is critical in maintaining the balance between spatial resolution and frequency.

When the input data $\Lambda_{s,r}^{\omega^\ell}$ is represented as a two-tensor of dimension $[2^\ell, 2^\ell]$, each V^ℓ layer can be implemented as a `LocallyConnected2D` layer in Tensorflow with r channels and both the kernel size and stride as $2^{L-\ell} \times 2^{L-\ell}$. The U^L layer can also be implemented as `LocallyConnected2D`

layer with rank s^2 and 1×1 kernel size and stride; the input to this layer is assumed to be of dimension $[2^\ell, 2^\ell, c]$ with c input channels.

B.4 H^ℓ and G^ℓ layers

The H^ℓ and G^ℓ layers in Figure 2 continue the theme of multiscale processing. When viewed as matrices, both H^ℓ and G^ℓ are block diagonal with block size $4^{L-\ell}r \times 4^{L-\ell}r$. Equivalently, when the input is formatted as a complete quad-tree, this implies both are *local operators* which process the nodes on the tree at length scale l to map each $2^{L-\ell}s \times 2^{L-\ell}s$ cells. Within each block there is further structure to the operators, as Figure 10 demonstrates. For each H^ℓ each sub-block has the interpretation of *aggregating* information, whereas each G^ℓ achieves the dual task of *spreading* information. We stress, however, that the action of this is entirely local within each cell. The key observation is that by permuting each node following a set pattern each operator becomes block-diagonal with block size 4^ℓ , for all $L/2 \leq \ell \leq L$.

The H^ℓ layers differ in that they process two inputs: one the output of the V^ℓ layer of dimension $[2^{L-\ell}, 2^{L-\ell}, r]$, the other the output from the previous layer of dimension $[2^\ell, 2^\ell, c]$ for some channel size c . To process the dimensions of both we first upscale each patch with redundant information to convert the data into $[2^\ell, 2^\ell, r]$. Then this is concatenated with the other dataset to form a tensor of size $[2^\ell, 2^\ell, c + r]$.

B.5 SWITCH-RESNET layer

We retain the permutation pattern of the switch layer as this is responsible for capturing the inherent non-locality of wave scattering (e.g. a point scatterer generates a diffraction pattern that is measured by all receivers in our geometry). We illustrate this pattern in Figure 9.

The input in this level serves as a condensed representation of the measured data. It is at this level that we non-linearly process the multifrequency dataset; we speculate that this also essential in facilitating the model to produce super-resolved images. We achieve this by using a residual network to refine each channel locally.

B.6 WideBNet Parameter Count

A tedious but straight-forward computation shows that the total number of parameters scales as $\mathcal{O}(N(\log N + \log^3 N))$. Note this is essentially linear in the total degrees of freedom in the data (N) up to polylogarithmic factors. Furthermore, note if naïvely L separate single channel WideBNet networks were used to compute (7) this would correspond to complexity $\mathcal{O}(N \log N^2)$; the multi-frequency assimilation only exceeds this with mild oversampling by a logarithmic factor.

Lastly, we note the effect of the partitioning of the frequencies. If all the frequencies were ingested at length scale L then the scaling becomes $\mathcal{O}(N(\log N^2 + \log N^3))$. While to leading order this presents the same asymptotic scaling, in terms of practical considerations this presents as substantial increase in the number of trainable parameters.

C Training details

C.1 Generation of Data

The scattering data was generated with 5 point-stencil second-order finite differences for training and 9 point-stencil fourth-order finite differences for testing. The radiating boundary conditions were implemented using PMLs spanning one wavelength with a quadratic profile of intensity 80 [4]. The domain of interest was discretized using a regular grid of 80×80 grid points. The scattered waves were sampled at 80 equi-angular gridpoints in D (see Figure 1) that was implemented as a circle of radius $r = 0.5$. The incident waves arrived at angles aligned with the receiver geometry. The training and testing split was 21000 points and 4000 points, respectively. For the experiments we choose three frequencies 2.5, 5 and 10Hz, and for simplicity we generated the data in the same mesh, and we use the same sampling geometry for the three frequencies. Thus each datapoint comprised by a tuple of a tensor of dimension $80 \times 80 \times 3$, and a discretization of the domain of interest Ω as a matrix of dimension 80×80 . The data was generated in Matlab using the UMFPACK [18] to

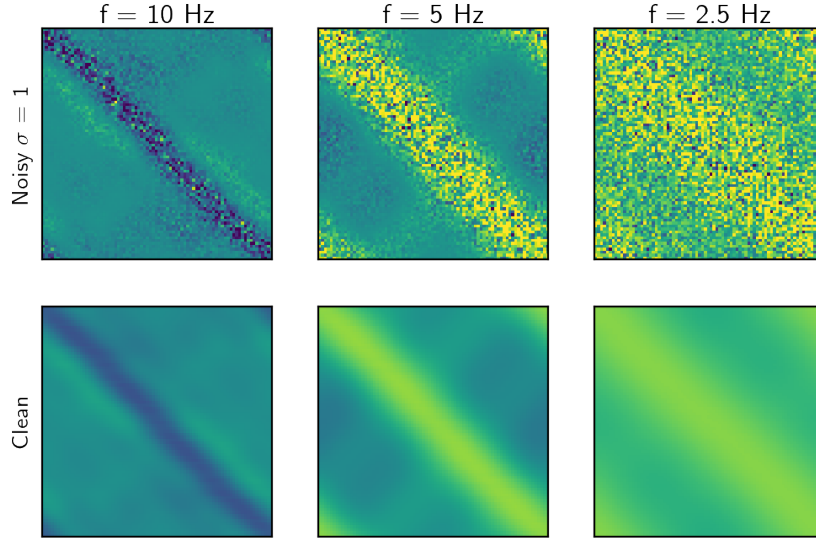


Figure 12: Depiction of effect of $\sigma = 1$ multiplicative noise (top) versus clean noiseless data (bottom).

solve the linear systems, for each dataset the generation of the data took roughly a week running on a dual-socket Intel Xeon CPU E5-2670 with 384 GB of RAM. Figure 12 (bottom) depicts one sample of the resulting far-field pattern.

C.2 Training

We implemented all the networks in Tensorflow 2.2 and we used the same optimization hyper-parameters for all the networks. The initial learning rate was set to $5e - 3$ using an exponential decay schedule with decay rate 0.95 every 2000 plateau steps with staircasing. Optimization was performed by Adam [34] using the standard parameters $\beta_1 = 0.9$ and $\beta_2 = 0.999$, which we terminate after 501 epochs. We did not observe any sensitivity to initialization reported in [57] and simply used standard glorot uniform sampling for the initial weights. No effort was taken to optimize hyper-parameters on an external validation set. Each experiment was computed with single-precision on four Nvidia GTX-1080Ti graphics cards with 128 GB of shared RAM, each training run was completed in roughly 12 hours.

The data at each frequency was normalized by subtracting and scaling by the mean and variance of the aggregate pixel intensities in the training dataset. We consider a multiplicative coloured noise model similar to [33] with

$$\Lambda_{\mathbf{s},\mathbf{r}} = \Lambda_{\mathbf{s},\mathbf{r}}^{\text{clean}}(1 + \epsilon(\mathbf{s}, \mathbf{r})), \quad \epsilon(\mathbf{s}, \mathbf{r}) \stackrel{iid}{\sim} \mathcal{N}(0, \sigma^2), \quad (15)$$

which is applied to each data pixel dynamically each epoch, where we use $\sigma = 1$ in order to introduce a 100% noise-to-signal ratio in the training data. Figure 12 depicts the result of this process.

The loss is the mean square error following

$$\sum_{x=\text{pixel in image}} \|(K_{\text{high}} * \eta)(x) - \text{WideBNet}[\Lambda_{\mathbf{s},r}^L, \dots, \Lambda_{\mathbf{s},r}^{L/2}](x)\|_2^2 \quad (16)$$

where η is the total scattering medium and $\{\Lambda_{\mathbf{s},r}^l\}$ the noisy multifrequency data. The scattering medium data was smoothed using a high-pass filter K_{high} described by a Gaussian kernel with characteristic width of 0.75 grid points. This was observed to promote faster training in [37] for noiseless data, which has the qualitative effect of smoothing the gradient. Since the support of this filter is significantly less than the Nyquist limit this processing does not pollute or eliminate sub-wavelength features in the medium.

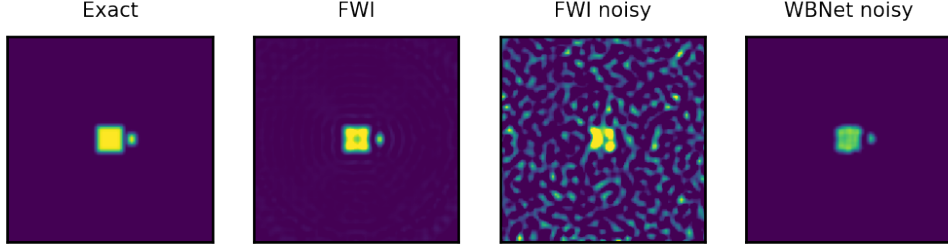


Figure 13: Comparison of reconstruction for the colliding example for $\Delta = \lambda$. From left to right, the reference perturbation to be reconstructed, the reconstruction of the perturbation using full-waveform inversion in Pysit [26] with clean data, FWI with noisy data, and WideBNet with noisy data. The computation time for the noiseless and noisy FWI is around 2 hours, and 45 minutes respectively, whereas WideBNet takes 0.1 seconds for the inference.

C.3 Hyper parameters for each architecture

- **WideBNet** : we used $N_{\text{CNN}} = 3$ convolutional layers and $N_{\text{RNN}} = 3$ residual layers and we used relu activation function throughout the network to add non-linearities. We set the number of levels $L = 4$, the leaf size $s = 5$, and the interaction rank $r = 4$. The total number of trainable parameters is 1,914,061.
- **NarrowBNet** : we used the same parameters as in WideBNet, but only processing the highest resolution. The total number of trainable parameters is 1,241,293.
- **SwitchNet** : we chose the same parameters as in [32], i.e., rank of the low-rank approximation $t = 3$, the number of partitions in x , $P_X = 8^2$, the number of partitions receiver manifold, $P_D = 4^2$, the window size of the convolution layers, $w = 10$, the number of channels in the convolution layers $\alpha = 18$, and the number of convolution layer $L = 3$. For a detailed explanation of these parameters please see [33]. The total number of trainable parameters is 3,191,416.
- **SwitchlessBNet** : we used the same parameters as in WideBNet. The only difference is the omission of the switch layer $S^{L/2}$, which does not contain any trainable parameters. The total number of trainable parameters is 1,914,061.
- **U-Net** : we considered an architecture with five levels, with a width-filter of 3, we used 4 channels at the finer level, which was doubled when going to lower levels. To make the implementation more competitive we added a periodic padding in the convolutional layers, in order to maintain the periodicity on the input and output. In addition, we used the full wide-band data, using only narrow-band data provided worse results overall. The total number of trainable parameters is 587,537.

D Comparison against other architectures

We compare the resolution with FWI implemented in Pysit [26](for a comparison against filtered back-projection please see Figure 7). We used 20 iterations to process the data at 2.5, 5 and 10 Hz, in increasing order. The optimization was performed using L-BFGS, plus line-search. We used as an example the collision data-set in which a super wavelength scatterer is roughly a wavelength away from a sub-wavelength one. In Figure 13 we can observe that if the data is clean, FWI is able to recover the perturbation, but when we add data following 6, then the reconstruction becomes much worse. Fortunately, even in this noisy setup WideBNet is able to image the perturbation.

For completeness we also provide several tables illustrating the superiority of WideBNet with respect to the other architectures considered in this manuscript. Table 2 provides the mean ℓ^2 error pixel wise during training. Tables 3 and 4 show the mean relative ℓ^2 error pixel wise during the training and testing stage for the architectures considered. Tables 3 and 4 depict the relative difficulty of learning the different dataset, where we can observe that the blob is, in general, the easiest to learn, whereas, datasets with large back-scattering, such as the Ten-squares and GRF ($\sigma = 0.01$) are the hardest one to learn.

Table 2: Comparison of mean-squared error per pixel training error between different architectures.

Dataset	WideBNet	UNet	SwitchlessBNet	NarrowBNet
Blob	4.548E-04	4.270E-03	2.364E-03	5.185E-04
Gaussian (2h)	1.641E-03	4.369E-02	4.996E-02	3.642E-03
Ten Squares (10h)	2.642E-02	1.051E-01	3.479E-01	8.022E-02
Right Triangles (3/5/10h)	3.347E-03	1.209E-01	6.175E-02	4.186E-03
Shepp Logan	6.050E-03	3.825E-02	2.545E-02	9.323E-03
GRF ($\sigma = 0.01$)	1.082E-02	3.924E-02	3.135E-01	4.579E-02
GRF ($\sigma = 0.04$)	2.978E-03	1.094E-02	1.900E-01	8.114E-03
Squares (3/5/10h)	3.412E-03	1.037E-01	6.731E-02	2.786E-03

Table 3: Comparison of relative ℓ^2 squared error per image in training between different architectures.

Dataset	WideBNet	UNet	SwitchlessBNet	NarrowBNet
Blob	4.087E-04	4.279E-03	2.273E-03	5.295E-04
Gaussian (2h)	1.573E-03	3.887E-02	4.825E-02	3.498E-03
Ten Squares (10h)	2.725E-02	1.082E-01	3.556E-01	9.290E-02
Right Triangles (3/5/10h)	3.788E-03	1.670E-01	7.279E-02	6.908E-03
Shepp Logan	6.091E-03	3.851E-02	2.565E-02	9.471E-03
GRF ($\sigma = 0.01$)	1.070E-02	3.991E-02	3.256E-01	6.104E-02
GRF ($\sigma = 0.04$)	3.259E-03	1.195E-02	2.101E-01	1.088E-02
Squares (3/5/10h)	3.441E-03	1.497E-01	8.210E-02	5.691E-03

E Generalization

In Figure 14 we provide a quantitative examination of the extrapolative properties of WideBNet models trained and tested with disparate datasets. We observe that training with the square 3/5/10h also results in low testing error on the triangle 3/5/10h and gaussian 2h.

F Super Resolution

For completeness we also provide results for super-resolution beneath the diffraction limit for two square scatterers (Figure 15), three triangles (Figure 16), and two triangles (Figure 17). Similar observations discussed in Section 4.2 hold for these examples as well – both the and U-Net architectures struggle in the subwavelength ($\Delta < \lambda$) regimes, whereas WideBNet is still able to consistent image the scatterers.

Table 4: Comparison of relative ℓ^2 squared error per image in testing between different architectures.

Dataset	WideBNet	UNet	SwitchlessBNet	NarrowBNet
Blob	5.449E-04	4.493E-03	2.598E-03	5.955E-04
Gaussian (2h)	2.299E-03	5.502E-02	1.048E-01	4.315E-03
Ten Squares (10h)	3.074E-02	1.079E-01	4.060E-01	9.098E-02
Right Triangles (3/5/10h)	7.097E-03	1.845E-01	1.197E-01	6.662E-03
Shepp Logan	7.320E-03	3.791E-02	3.209E-02	1.160E-02
GRF ($\sigma = 0.01$)	1.482E-02	4.811E-02	5.665E-01	5.964E-02
GRF ($\sigma = 0.04$)	3.800E-03	1.447E-02	2.551E-01	1.101E-02
Squares (3/5/10h)	7.085E-03	1.889E-01	1.528E-01	6.032E-03

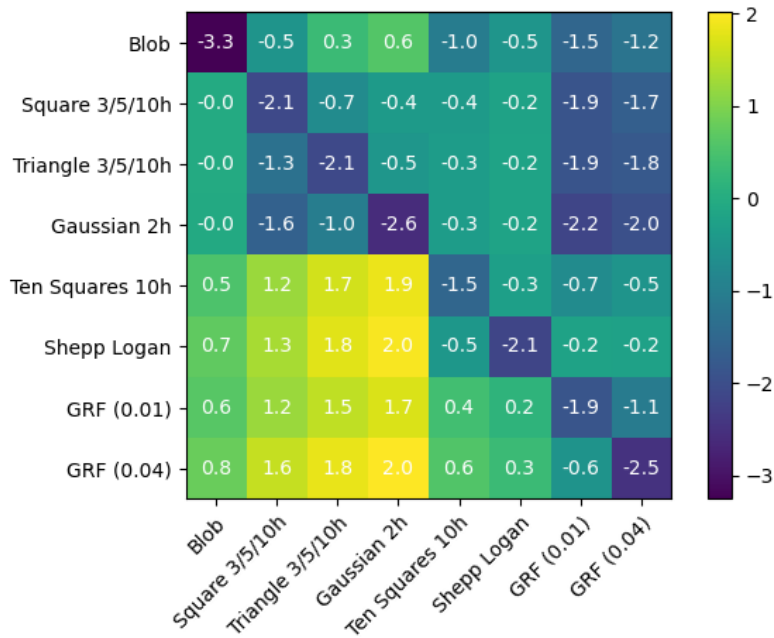


Figure 14: Generalization matrix in which we showcase the average pixelwise squared generalization error for WideBNet models trained with data sets in the ordinate axis, and tested against data sets in the coordinate axis. All numbers are in logarithmic scale (base 10).

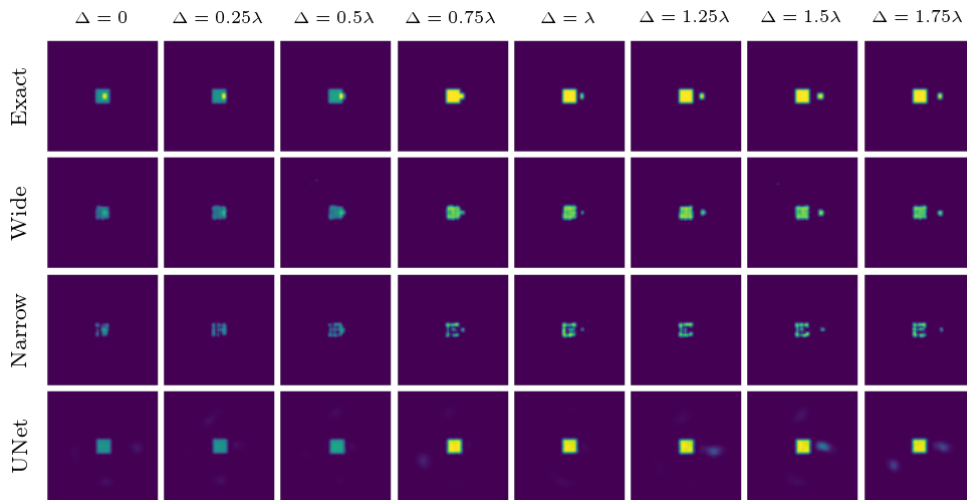


Figure 15: Comparison of imaging capabilities of different networks for square scatterers with side lengths 3 and 10 pixels, which are separated by Δ . We vary Δ between sub-wavelength and wavelength separations (one wavelength $\lambda = 8$ pixels)

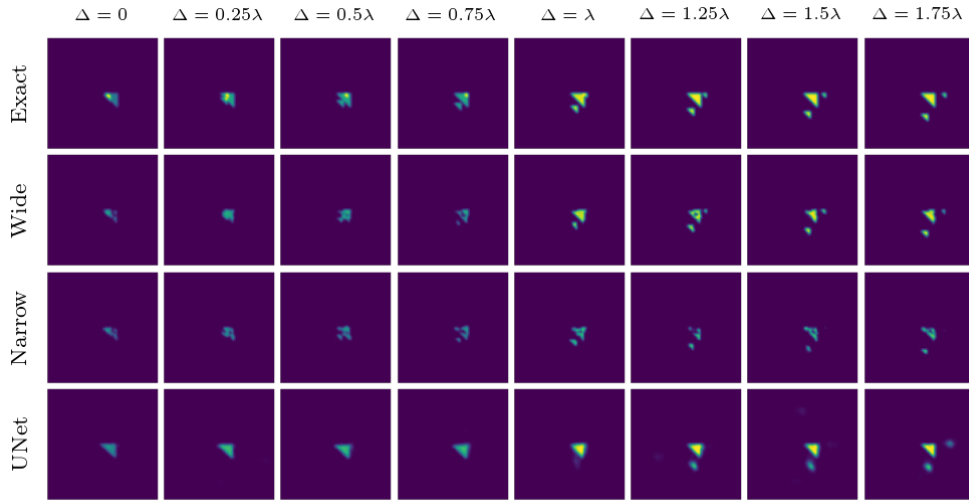


Figure 16: Comparison of imaging capabilities of different networks for triangular scatterers with side lengths 3, 5, and 10 pixels, which are separated by Δ . We vary Δ between sub-wavelength and wavelength separations (one wavelength $\lambda = 8$ pixels)

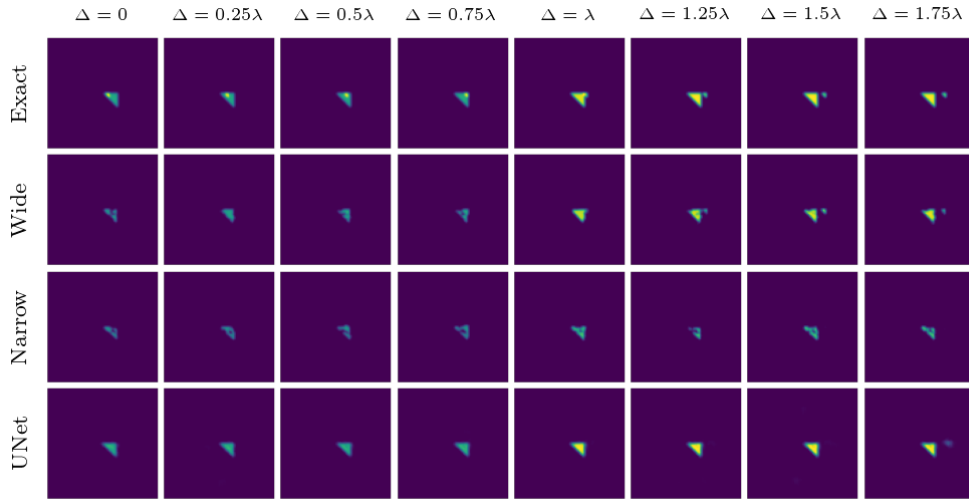


Figure 17: Comparison of imaging capabilities of different networks for triangular scatterers with side lengths 3 and 10 pixels, which are separated by Δ . We vary Δ between sub-wavelength and wavelength separations (one wavelength $\lambda = 8$ pixels)