

Specialization: Applied Mathematics

Academic year: 2025–2026

Research Internship Report

Seeing Through Sound: Assessing Physics-Informed Sound-Speed Inversion in an Acoustic Waveguide

Non-confidential report

Author: Romain Tolisano

Promotion: 2026

Academic supervisor

Laure Giovangigli
(ENSTA)

Internship supervisor

Davide Moroni
Giacomo Ignesti
Massimo Martinelli
(ISTI–CNR)

Internship completed from 18 May to 31 July 2026

ISTI–SI Laboratory

ISTI • Area della Ricerca CNR, via G. Moruzzi 1, 56124 Pisa, Italy

Acknowledgements

First and foremost, I would like to express my sincere gratitude to Davide Moroni for considering my internship application, for his responsiveness with administrative procedures, his warm welcome into the group, and his continuous guidance throughout my stay.

I also extend my thanks to Massimo Martinelli and Giacomo Ignesti for having, alongside Davide Moroni, designed a tailored research topic aligned with my interests and for agreeing to supervise me on a project outside their primary area of expertise. I am truly grateful for their support, their trust, and the substantial autonomy they granted me in conducting my work.

My thanks go as well to Said Daoudagh and the entire ISTI team for welcoming and integrating me so warmly. This internship has been a wonderful human experience that I will fondly remember.

Finally, I would like to thank Alexandre Borreda for his valuable help and attentive proofreading of this report.

Résumé

Ce stage évalue la capacité des réseaux de neurones informés par la physique (Physics-Informed Neural Networks, PINNs) à réaliser une inversion quantitative de la vitesse du son dans un guide d’ondes acoustique bidimensionnel. Nous formulons le problème de Helmholtz en champ total et en champ diffusé, puis développons intégralement en JAX une implémentation simultanée (*all-at-once*), avec un réseau à valeurs complexes représentant le champ d’onde pour chaque acquisition et un réseau de vitesse du son partagé entre les acquisitions. Une calibration systématique du problème direct montre que l’encodage par caractéristiques de Fourier et une pondération soignée des termes de perte sont déterminants pour apprendre les champs oscillants aux fréquences étudiées. Pour l’inversion, la prédiction du champ diffusé supprime la phase initiale de pré-entraînement homogène et réduit le temps d’entraînement total d’environ 10%, sans gain systématique en précision. La meilleure reconstruction combine plutôt un a priori physique sur le signe, $c \leq c_0$, avec un curriculum fréquentiel en deux étapes, et atteint une erreur normalisée par le contraste de $0,34 \pm 0,01$. Ce résultat reste dépendant de cet a priori : lorsque les bornes du matériau sont larges, la plupart des reconstructions sont moins bonnes que la carte homogène selon la même métrique, malgré une récupération locale partielle du défaut. De plus, une perte normalisée sur les traces de l’ordre de 10^{-7} peut coexister avec une erreur de pression intérieure d’environ 20%, ce qui montre qu’une faible perte ne certifie pas la reconstruction. Aucune référence classique de type inversion des formes d’onde complètes (FWI) n’est incluse. La contribution principale est donc un banc d’essai reproductible, une implémentation différentiable développée intégralement et une méthodologie d’évaluation sur plusieurs initialisations aléatoires permettant de comparer les choix de modélisation et d’optimisation, ainsi qu’un diagnostic explicite des obstacles restant à surmonter pour fiabiliser l’inversion par PINN.

Mots-clés : réseaux de neurones informés par la physique ; guides d’ondes acoustiques ; problèmes inverses ; reconstruction de la vitesse du son ; équation de Helmholtz ; inversion des formes d’onde complètes.

Abstract

This internship assesses whether Physics-Informed Neural Networks (PINNs) can perform quantitative sound-speed inversion in a two-dimensional acoustic waveguide. We formulate the Helmholtz problem in total-field and scattered-field forms and develop from scratch an all-at-once JAX implementation, with one complex wavefield network per acquisition and a sound-speed network shared across acquisitions. A systematic forward calibration shows that Fourier features and careful loss balancing are decisive for learning oscillatory wavefields at the frequencies considered. In inversion, predicting the scattered field eliminates the homogeneous warm-up and reduces end-to-end training time by about 10%, but yields no systematic accuracy gain. The best reconstruction instead combines a physics-based sign prior, $c \leq c_0$, with a two-stage frequency curriculum, reaching a contrast-normalised error of 0.34 ± 0.01 . This result remains prior-dependent: under broad material bounds, most reconstructions perform worse than the homogeneous map according to the same metric, despite partially recovering the defect locally. Moreover, a normalised trace loss of order 10^{-7} can coexist with an interior pressure error of about 20%, showing that a small loss does not certify the reconstruction. No classical full-waveform-inversion baseline is included. The main contribution is therefore a reproducible benchmark, a from-scratch differentiable implementation, and a multi-seed evaluation methodology for comparing modelling and optimisation choices, together with an explicit diagnosis of the remaining obstacles to reliable PINN-based inversion.

Keywords: physics-informed neural networks; acoustic waveguides; inverse problems; sound-speed reconstruction; Helmholtz equation; full-waveform inversion.

Contents

Acknowledgements	i
1 Introduction	1
1.1 Host Laboratory and Working Conditions	1
1.2 Context and Objectives	1
1.3 Structure of the Document	2
2 Full Waveguide Problem	3
2.1 Geometry and Notation	3
2.2 Time-Harmonic Acoustic Model	3
2.3 Transverse Modal Decomposition	4
2.4 Heterogeneous Waveguide	4
3 Scattered Field Formulation	6
3.1 Background, Total, and Scattered Fields	6
3.2 Equation for the Scattered Field	6
3.3 Equivalent Weak Formulation	6
4 Mathematical Analysis and Classical Inverse Problem	7
4.1 Well-Posedness	7
4.2 The Forward Operator	9
4.3 Inverse Problem	10
4.4 Methods and Common Challenges in the Inverse Problem	10
4.4.1 Intrinsic Limitations in Waveguide Inversion	11
5 Physics-Informed Neural Networks and Implemented Workflow	12
5.1 Concept and Positioning	12
5.2 State of the Art	12
5.3 Advantages and Limitations	13
5.4 Framework	13
5.5 PINN Formulation of the Problem	14
6 Forward PINN Study	17
6.1 Purpose and Experimental Protocol	17
6.2 Homogeneous Waveguide	17
6.2.1 Loss weighting at 600 Hz ($k_0 \approx 11.1 \text{ rad m}^{-1}$)	18
6.2.2 Frequency and adaptive weights at 1200 Hz ($k_0 \approx 22.2 \text{ rad m}^{-1}$)	19
6.3 Waveguide Containing a Sharp Defect	20
6.3.1 Total-field and scattered-field formulations	20
6.3.2 Adaptive loss weights and RAD	21
6.4 Lessons for the Inverse Problem	22
7 First Inverse Approach: Baseline PINN Formulation	23
7.1 All-at-Once Formulation and Training	23
7.2 Evaluation Protocol and Results	24
7.3 Lessons and Limitations	26
8 Second Inverse Approach: Scattered-Field Formulation	28
8.1 Motivation and Controlled Protocol	28
8.2 Results and Interpretation	29

9	Discussion	32
9.1	What the Experiments Establish	32
9.2	Methodological and Technical Contribution	33
9.3	Internship Timeline	34
9.4	What PINNs Add and Where the Field Is Moving	34
10	Conclusion and Perspectives	36
11	Technical Glossary	37
A	Proofs	38
A.1	Weak Formulation for the Discontinuous Contrast	38
A.2	Continuity of the Forward Operator	38
B	Forward-Study Reference Data	39
B.1	Homogeneous Analytic Reference	39
B.2	Heterogeneous Finite-Element Reference	39
C	Forward PINN Methods and Reproducibility	39
C.1	Field Networks and Fourier Features	40
C.2	Adaptive Loss Weights	40
C.3	Residual-Based Adaptive Distribution	41
C.4	Validation Metrics and Reproducibility	41
D	Multilayer Perceptron: Architecture and Training	42

1 Introduction

This project originates from an ENSTA Paris C++ project on ultrasound imaging, in which a Finite Element Method (FEM) was used to simulate acoustic wave propagation in a two-dimensional waveguide containing defects, and the Linear Sampling Method was applied to localise the defected regions (see [GitHub](#)). Building on that experience and a growing interest in machine-learning approaches, the natural question arose whether neural-network-based methods could be applied to the same class of problems. After discussions with the ISTI-CNR team in Pisa, the project focused on reconstructing the sound-speed field inside the waveguide using Physics-Informed Neural Networks (PINNs), a methodology that had not previously been explored by the team.

1.1 Host Laboratory and Working Conditions

The internship took place at the Institute of Information Science and Technologies (ISTI), part of the Italian National Research Council (CNR) in Pisa. ISTI organises its research activities into thematic areas and specialised laboratories. I joined the [Signals & Images Laboratory](#) (SI-Lab), which belongs to the Visual thematic area and is headed by Davide Moroni. Founded in 2002, SI-Lab is an interdisciplinary group bringing together researchers and technologists from computer science, mathematics, engineering and physics. Its activities cover signal analysis, computer vision, smart vision systems and multimedia data understanding, with applications ranging from cultural heritage and mobility to healthcare and well-being.

The day-to-day organisation of the ten-week internship was deliberately flexible. I was given substantial autonomy in setting my working hours and organising literature review, code development, numerical experiments and writing. From the sixth week onward, near-weekly meetings with my supervisors provided regular opportunities to review results, identify modelling or optimisation issues and decide on the next experiments, while preserving this freedom in the daily conduct of the work.

1.2 Context and Objectives

The initial goal was to demonstrate a successful PINN-based inversion pipeline on the waveguide problem and then extend it to medical ultrasound reconstruction. PINNs were introduced by Raissi et al. in 2017 [1] and first applied to inverse problems in 2019 [2]; the years that followed saw a wave of optimism in the research community, with numerous works reporting promising results across a variety of physics problems. This enthusiasm was naturally shared at the start of this project. The field remains highly active: new architectures, training strategies, and applications continued to appear throughout the duration of this internship, as illustrated by recent contributions in 2024–2026 [3, 4, 5, 6, 7, 8]. This ambition nonetheless proved over-optimistic: their practical behaviour on ill-posed, high-frequency problems is far from straightforward. A significant part of this work therefore consists of documenting the failure modes of PINNs on both the forward and inverse problems and incrementally designing remedies. This report does not retrace every step of the learning process but presents a matured account of the work carried out at the ISTI-CNR. The objective of this report is to identify the challenges encountered when applying PINNs to the forward and inverse waveguide problems, to evaluate the remedies investigated during the internship, and to analyse the resulting behaviour of the method within an all-at-once inverse formulation.

1.3 Structure of the Document

This document begins with the mathematical formulation of the waveguide scattering problem. The full-field model is introduced first, then rewritten as a scattered-field formulation, and finally analysed through the well-posedness of the forward problem, the definition of the forward operator, and the classical difficulties of the associated inverse problem.

The report then introduces Physics-Informed Neural Networks (PINNs), their position in the scientific machine-learning literature, and the practical advantages and limitations that motivated the numerical choices made during the internship. Before attacking the inverse problem, a forward PINN study is presented as a necessary calibration step: it identifies the architectures, constraints, and training strategies that are robust enough to be reused for inversion.

The inverse study then compares a total-field baseline with a scattered-field formulation. The discussion separates controlled comparisons from campaigns in which several ingredients changed simultaneously, and the conclusion prioritises the classical baselines and robustness tests still required.

2 Full Waveguide Problem

Before introducing the equations, it is useful to establish the physical picture. A time-harmonic acoustic wave enters the computational domain through the left artificial cross-section Γ^- and propagates along the guide. When it encounters a localised perturbation of the sound speed, part of the wave is scattered, and the resulting field is observed on the two artificial cross-sections Γ^- and Γ^+ . This section progressively formalises this setting by introducing the geometry and notation, stating the time-harmonic model with rigid-wall conditions, describing the transverse modes of the homogeneous guide, and finally formulating the heterogeneous problem with transparent Dirichlet-to-Neumann conditions on $\Gamma^\pm$.

2.1 Geometry and Notation

We consider the two-dimensional straight waveguide

$$\mathcal{G} = \mathbb{R} \times (0, H),$$

where $x \in \mathbb{R}$ denotes the longitudinal coordinate and $y \in (0, H)$ the transverse coordinate. The boundary of the guide is

$$\partial\mathcal{G} = \mathbb{R} \times \{0, H\}.$$

The outward unit normal on $\partial\mathcal{G}$ is denoted by ν . In the following, the cross-section of the guide is $\Gamma = (0, H)$.




Figure 1: Straight two-dimensional waveguide.

2.2 Time-Harmonic Acoustic Model

We work in the frequency domain with the convention $\exp(-i\omega t)$. Let

$$m(x, y) = \frac{1}{c(x, y)^2}$$

be the *squared slowness* of the medium. In the absence of volume sources, the acoustic pressure $u : \mathcal{G} \rightarrow \mathbb{C}$ satisfies the Helmholtz equation

$$\Delta u + \omega^2 m u = 0 \quad \text{in } \mathcal{G}. \quad (2.1)$$

The rigid-wall boundary condition is

$$\partial_\nu u = 0 \quad \text{on } \partial\mathcal{G}. \quad (2.2)$$

Additional conditions at infinity are required to select the physical solution. They will be stated through outgoing modal expansions. This standard acoustic model is presented in detail in Bonnet-Ben Dhia and Lunéville [9, Section 1.1.1, pp. 9–10].

2.3 Transverse Modal Decomposition

The transverse Neumann Laplacian on the cross-section $\Gamma = (0, H)$ is defined by

$$-\theta_n'' = \lambda_n \theta_n \quad \text{in } (0, H), \quad \theta_n'(0) = \theta_n'(H) = 0. \quad (2.3)$$

Its eigenpairs are explicitly given by

$$\lambda_n = \left(\frac{n\pi}{H}\right)^2, \quad \theta_0(y) = \frac{1}{\sqrt{H}}, \quad \theta_n(y) = \sqrt{\frac{2}{H}} \cos\left(\frac{n\pi y}{H}\right) \quad (n \geq 1). \quad (2.4)$$

The family $(\theta_n)_{n \geq 0}$ is an orthonormal basis of $L^2(\Gamma)$. This basis is the natural modal basis for wave propagation in a straight Neumann waveguide. The eigenpairs, separation-of-variables construction, and propagating/evanescent classification used below are recalled in [9, Section 1.1.1, pp. 10–12]. When the medium is homogeneous, $m(x, y) = m_0 > 0$, the Helmholtz equation separates in this basis. We set

$$\beta_n = \sqrt{\omega^2 m_0 - \lambda_n}, \quad \operatorname{Re} \beta_n \geq 0, \quad \operatorname{Im} \beta_n \geq 0.$$

Throughout, we exclude cut-off frequencies, namely

$$\omega^2 m_0 \notin \{\lambda_n : n \geq 0\}.$$

Consequently, every mode is either propagating or evanescent. The elementary modes are

$$u_n^\pm(x, y) = \theta_n(y) \exp(\pm i \beta_n x). \quad (2.5)$$

Modes with $\lambda_n < \omega^2 m_0$ are propagating, whereas modes with $\lambda_n > \omega^2 m_0$ are evanescent. More general solutions are expanded as superpositions of these modes; see, for example, the modal formulation of Bourgeois and Lunéville [10].

2.4 Heterogeneous Waveguide

We restrict attention to the bounded domain

$$\Omega = (-L, L) \times (0, H),$$

whose boundary decomposes as

$$\partial\Omega = \Gamma_N \cup \Gamma^+ \cup \Gamma^-,$$

where $\Gamma_N = (-L, L) \times \{0, H\}$ carries the rigid-wall Neumann condition and $\Gamma^\pm = \{\pm L\} \times (0, H)$ are the artificial vertical cross-sections on which transparent (DtN) boundary conditions will be prescribed. The medium inside Ω is heterogeneous: we introduce a compactly supported defect region $\Omega_d \subset \Omega$ and a discontinuous sound speed

$$c(x, y) = \begin{cases} c_d & (x, y) \in \Omega_d, \\ c_0 & (x, y) \in \Omega \setminus \Omega_d, \end{cases}$$

where $c_0 > 0$ is the background speed and $c_d \neq c_0$ is the local perturbation. We assume moreover that the medium is homogeneous in a neighbourhood of the artificial boundaries $\Gamma^\pm$ and in the exterior half-guides, so that $m = m_0$ there. This assumption allows the exact transparent boundary conditions to be defined through the background modal decomposition.




Figure 2: Heterogeneous waveguide $\Omega = (-L, L) \times (0, H)$. The shaded regions represent the defect subdomain Ω_d where the celerity equals c_d ; the background has celerity c_0 . The horizontal walls Γ_N carry Neumann conditions and the vertical sections $\Gamma^\pm$ carry the DtN transparent boundary conditions.

Let ν_d denote the unit normal to $\partial\Omega_d$ directed from Ω_d towards $\Omega \setminus \overline{\Omega_d}$. For a function w admitting traces from both sides of the interface, we use the conventions

$$[w]_{\partial\Omega_d} := w|_{\Omega_d} - w|_{\Omega \setminus \overline{\Omega_d}}, \quad [\partial_{\nu_d} w]_{\partial\Omega_d} := (\nabla w|_{\Omega_d} - \nabla w|_{\Omega \setminus \overline{\Omega_d}}) \cdot \nu_d.$$

The strong formulation of the heterogeneous waveguide problem reads: find $u : \Omega \rightarrow \mathbb{C}$ such that

$$\begin{cases} \Delta u + \omega^2 m u = 0 & \text{in } \Omega, \\ \partial_\nu u = 0 & \text{on } \Gamma_N, \\ \partial_\nu u = T_\omega u + g^\pm & \text{on } \Gamma^\pm, \\ [u]_{\partial\Omega_d} = 0, \quad [\partial_{\nu_d} u]_{\partial\Omega_d} = 0 & \text{on } \partial\Omega_d. \end{cases} \quad (2.6)$$

where T_ω is the Dirichlet-to-Neumann (DtN) operator encoding the outgoing wave condition at $\Gamma^\pm$ and $g^\pm$ are the incoming data (incident-field fluxes). The DtN operator is defined through the modal expansion of Section 2.3. As we take the outgoing normal derivative, either on Γ^+ or Γ^- , the DtN operator has the same definition. We will refer to $\Gamma^\pm$ as Γ . For a trace $\varphi \in H^{1/2}(\Gamma)$, writing $\varphi_n = \langle \varphi, \theta_n \rangle_{L^2(0,H)}$, the outgoing-wave radiation condition on $\Gamma^\pm$ is enforced by

$$\begin{aligned} T_\omega : H^{1/2}(\Gamma) &\longrightarrow H^{-1/2}(\Gamma) \\ \varphi &\longmapsto \sum_{n=0}^{+\infty} i\beta_n \varphi_n \theta_n \end{aligned} \quad (2.7)$$

where $\beta_n = \sqrt{\omega^2 m_0 - \lambda_n}$ with $\operatorname{Re} \beta_n \geq 0$ and $\operatorname{Im} \beta_n \geq 0$ (cf. (2.5)). Propagating modes ($\lambda_n < \omega^2 m_0$) contribute purely imaginary coefficients $i\beta_n \in i\mathbb{R}_+$, while evanescent modes ($\lambda_n > \omega^2 m_0$) contribute real negative damping coefficients. This is the exact non-dissipative DtN operator obtained from the outgoing modal expansion; see [9, Sections 3.3.3 and 3.5, pp. 47–50]. Its continuity from $H^{1/2}(\Gamma)$ to $H^{-1/2}(\Gamma)$ is established in [9, Proposition 3.6, p. 47]. In what follows, we focus on the scattered-field formulation; the analysis is analogous to the full-field problem stated above.

3 Scattered Field Formulation

3.1 Background, Total, and Scattered Fields

The background field u_0 is a solution of the Helmholtz equation in the *homogeneous* background medium of squared slowness $m_0 = 1/c_0^2$:

$$\begin{cases} \Delta u_0 + \omega^2 m_0 u_0 = 0 & \text{in } \Omega, \\ \partial_\nu u_0 = 0 & \text{on } \Gamma_N, \\ \partial_\nu u_0 = T_\omega u_0 + g^\pm & \text{on } \Gamma^\pm. \end{cases} \quad (3.1)$$

For example, consider a propagating incident mode travelling in the positive x -direction:

$$g^+ = 0 \quad g^-(y) = -2i\beta_p \theta_p(y) e^{-i\beta_p L} \quad (3.2)$$

The solution to this problem has the explicit form:

$$u_0(x, y) = \theta_p(y) e^{i\beta_p x}, \quad \beta_p = \sqrt{\omega^2 m_0 - \lambda_p} > 0, \quad (3.3)$$

for some fixed propagating-mode index $p \geq 0$. The total field $u = u_0 + u^s$ satisfies the heterogeneous problem (2.6) with the same incident data. Writing the squared slowness as

$$m = m_0 + m_s, \quad m_s(x, y) = \begin{cases} \frac{1}{c_d^2} - \frac{1}{c_0^2} & (x, y) \in \Omega_d, \\ 0 & (x, y) \in \Omega \setminus \Omega_d, \end{cases} \quad (3.4)$$

the quantity m_s is the *squared-slowness contrast*, compactly supported in the true (unknown) defect region Ω_d . More generally, we assume $m_s \in L^\infty(\Omega)$ with compact interior support.

3.2 Equation for the Scattered Field

Subtracting the equation for u_0 from that satisfied by $u = u_0 + u^s$ yields the strong formulation for the scattered field: find $u^s : \Omega \rightarrow \mathbb{C}$ such that

$$\begin{cases} \Delta u^s + \omega^2 m u^s = -\omega^2 m_s u_0 & \text{in } \Omega, \\ \partial_\nu u^s = 0 & \text{on } \Gamma_N, \\ \partial_\nu u^s = T_\omega u^s & \text{on } \Gamma^\pm, \\ [u^s]_{\partial\Omega_d} = 0, \quad [\partial_{\nu_d} u^s]_{\partial\Omega_d} = 0 & \text{on } \partial\Omega_d. \end{cases} \quad (3.5)$$

The right-hand side $f := -\omega^2 m_s u_0$ is the source term induced by the contrast m_s ; it vanishes outside Ω_d .

3.3 Equivalent Weak Formulation

The variational formulation is obtained by multiplying (3.5) by a test function $v \in H^1(\Omega)$, integrating over Ω , and integrating by parts. It is the variable-coefficient, scattered-field counterpart of the DtN formulation in [9, Section 3.5, equations (3.13)–(3.14), pp. 50–51]. The

problem-specific normal-trace argument is recorded in Appendix A.1. The Neumann condition on Γ_N is natural and the DtN conditions on $\Gamma^\pm$ contribute boundary terms through the duality bracket $\langle \cdot, \cdot \rangle_{\Gamma^\pm}$ between $H^{-1/2}(\Gamma^\pm)$ and $H^{1/2}(\Gamma^\pm)$. Throughout, this duality bracket is bilinear; the complex conjugation of the test function is therefore written explicitly.

Weak formulation. Find $u^s \in H^1(\Omega)$ such that

$$a(u^s, v) = \ell(v) \quad \forall v \in H^1(\Omega), \quad (3.6)$$

where the sesquilinear form $a : H^1(\Omega) \times H^1(\Omega) \rightarrow \mathbb{C}$ and the antilinear form $\ell : H^1(\Omega) \rightarrow \mathbb{C}$ are defined by

$$a(u, v) = \int_{\Omega} \nabla u \cdot \overline{\nabla v} \, d\Omega - \omega^2 \int_{\Omega} m u \bar{v} \, d\Omega - \langle T_{\omega} u|_{\Gamma^+}, \overline{v|_{\Gamma^+}} \rangle_{\Gamma^+} - \langle T_{\omega} u|_{\Gamma^-}, \overline{v|_{\Gamma^-}} \rangle_{\Gamma^-}, \quad (3.7)$$

$$\ell(v) = \omega^2 \int_{\Omega} m_s u_0 \bar{v} \, d\Omega. \quad (3.8)$$

For brevity, subsequent boundary expressions use $\langle T_{\omega} u, \bar{v} \rangle_{\Gamma^\pm}$ for the sum of the two duality pairings on Γ^+ and Γ^- .

4 Mathematical Analysis and Classical Inverse Problem

4.1 Well-Posedness

Throughout this section we work in the Hilbert space $V = H^1(\Omega)$ equipped with the norm

$$\|v\|_V^2 = \|\nabla v\|_{L^2(\Omega)}^2 + \|v\|_{L^2(\Omega)}^2,$$

and we denote by V' the anti-dual of V , that is, the space of continuous antilinear forms on V . We assume that $m_s \in L^\infty(\Omega; \mathbb{R})$ has compact interior support and that $m = m_0 + m_s > 0$ almost everywhere in Ω , with $m_0 > 0$ constant. The goal is to show that the weak formulation (3.6) is a Fredholm problem of index zero and, under a non-resonance condition, is uniquely solvable. The argument follows the DtN–Fredholm framework developed in [9, Sections 2.3.3–2.3.4 and 3.5, pp. 32–34 and 50–52], adapted here to a bounded variable coefficient and two artificial cross-sections.

Continuity of the sesquilinear form

Proposition 4.1 (Continuity of a). *The sesquilinear form $a : V \times V \rightarrow \mathbb{C}$ defined in (3.7) is continuous: there exists $C_a > 0$ such that*

$$|a(u, v)| \leq C_a \|u\|_V \|v\|_V \quad \forall u, v \in V.$$

Indeed, Cauchy–Schwarz estimates for the volume terms, continuity of the trace operator, and boundedness of the DtN operator [9, Proposition 3.6, p. 47] give the stated bound.

Continuity of the antilinear form

Proposition 4.2 (Continuity of ℓ). *The antilinear form $\ell : V \rightarrow \mathbb{C}$ defined in (3.8) is continuous: there exists $C_\ell > 0$ such that*

$$|\ell(v)| \leq C_\ell \|v\|_V \quad \forall v \in V.$$

Indeed, Cauchy–Schwarz gives $|\ell(v)| \leq \omega^2 \|m_s\|_{L^\infty(\Omega)} \|u_0\|_{L^2(\Omega)} \|v\|_V$.

Gårding inequality and compact perturbation

The sesquilinear form a is *not* coercive on V because of the negative-definite mass term $-\omega^2 \int_\Omega m u \bar{v} d\Omega$. However, one can recover coercivity up to a compact perturbation.

Definition 4.3 (Shifted form). For $\eta > 0$, define the shifted sesquilinear form

$$a_\eta(u, v) := a(u, v) + \eta(u, v)_{L^2(\Omega)}$$

Proposition 4.4 (Gårding inequality). *There exist constants $\alpha > 0$ and $\eta > 0$ such that*

$$\operatorname{Re} a_\eta(v, v) \geq \alpha \|v\|_V^2 \quad \forall v \in V.$$

In particular, a satisfies a Gårding inequality:

$$\operatorname{Re} a(v, v) \geq \alpha \|v\|_V^2 - \eta \|v\|_{L^2(\Omega)}^2 \quad \forall v \in V. \quad (4.1)$$

The modal sign identity for the DtN contribution is given in [9, Lemma 3.9, p. 51]. In the present variable-coefficient setting it yields

$$\operatorname{Re} a_\eta(v, v) \geq \|\nabla v\|_{L^2(\Omega)}^2 + (\eta - \omega^2 \|m\|_{L^\infty(\Omega)}) \|v\|_{L^2(\Omega)}^2,$$

so any $\eta > \omega^2 \|m\|_{L^\infty(\Omega)}$ proves the claim.

Fredholm alternative

Proposition 4.5 (Fredholm structure). *The operator $A : V \rightarrow V'$ defined by $\langle Au, v \rangle = a(u, v)$ is a Fredholm operator of index zero.*

This is the compact-perturbation argument of [9, Theorem 3.10, p. 51]: the shifted operator is an isomorphism by Lax–Milgram, while the L^2 mass term is compact by the Rellich–Kondrachov theorem.

Well-posedness result

Theorem 4.6 (Well-posedness). *Assume that $m = m_0 + m_s$ with $m_s \in L^\infty(\Omega)$, and that ω is not a resonant frequency of the waveguide, i.e., the homogeneous problem $a(u, v) = 0$ for all $v \in V$ admits only the trivial solution $u = 0$. Then the weak formulation (3.6) has a unique solution $u^s \in H^1(\Omega)$ and there exists $C > 0$ such that*

$$\|u^s\|_V \leq C \|\ell\|_{V'}.$$

The Fredholm alternative reduces well-posedness to the exclusion of resonant homogeneous solutions, as in [9, Theorems 3.10–3.11, pp. 51–52]. The exceptional frequencies are related to trapped modes and form at most a countable set [9, Corollary 4.9, p. 60].

4.2 The Forward Operator

The inverse problem consists in recovering the contrast m_s from boundary measurements of the scattered field. We introduce the forward operator that maps a contrast to the measured data.

Throughout this subsection, the background field $u_0 \in H^1(\Omega)$ (solution of (3.1)) and the frequency $\omega > 0$ are fixed. We assume that every $m_s \in \mathcal{A}$ satisfies the non-resonance condition of theorem 4.6, so that the solution operator G_{u_0} is well-defined on all of $\mathcal{A}$.

Definition 4.7 (Admissible contrasts). For fixed constants $0 < m_{\min} < m_{\max}$, the set of admissible contrasts is

$$\mathcal{A} = \{m_s \in L^\infty(\Omega; \mathbb{R}) : m_{\min} \leq m_0 + m_s \leq m_{\max} \text{ in } \Omega, \text{ supp}(m_s) \Subset \Omega\}.$$

The support condition preserves the homogeneous background medium in a neighbourhood of the artificial cross-sections, as required by the modal DtN construction.

Definition 4.8 (Forward operator). The forward operator $F_{u_0} : \mathcal{A} \rightarrow H^{1/2}(\Gamma^-) \times H^{1/2}(\Gamma^+)$ is defined by

$$F_{u_0}(m_s) = (u^s|_{\Gamma^-}, u^s|_{\Gamma^+}),$$

where $u^s \in H^1(\Omega)$ is the unique solution of (3.6) with contrast m_s .

The operator F_{u_0} decomposes as $F_{u_0} = \gamma \circ G_{u_0}$, where:

- $G_{u_0} : \mathcal{A} \rightarrow H^1(\Omega)$, $m_s \mapsto u^s$, is the *solution operator*, well-defined by theorem 4.6;
- $\gamma : H^1(\Omega) \rightarrow H^{1/2}(\Gamma)^2$, $u \mapsto (u|_{\Gamma^-}, u|_{\Gamma^+})$, denotes the pair of trace operators on $\Gamma^\pm$ (identified with Γ as in Section 2.4), which is linear and continuous.

Remark 4.9 (Nonlinearity of F_{u_0}). The forward operator F_{u_0} is *not* linear. Indeed, the sesquilinear form (3.7) depends on m_s through the mass term $-\omega^2 \int_\Omega (m_0 + m_s) u^s \bar{v} \, d\Omega$, while the right-hand side (3.8) depends on m_s through $\omega^2 \int_\Omega m_s u_0 \bar{v} \, d\Omega$. Changing m_s modifies *both* the operator and the source simultaneously. In particular, whenever m_s^1 , m_s^2 , and $m_s^1 + m_s^2$ belong to $\mathcal{A}$, the corresponding sesquilinear forms prevent additivity in general:

$$G_{u_0}(m_s^1 + m_s^2) \neq G_{u_0}(m_s^1) + G_{u_0}(m_s^2).$$

Proposition 4.10 (Continuity of F_{u_0}). When $\mathcal{A}$ is equipped with the topology induced by $L^\infty(\Omega)$, the forward operator $F_{u_0} : \mathcal{A} \rightarrow H^{1/2}(\Gamma^-) \times H^{1/2}(\Gamma^+)$ is continuous.

The proof is given in Appendix A.2. Continuity follows from a stability estimate applied to the difference of two scattered fields.

Actual sensors return finitely many numbers rather than complete boundary traces. We therefore compose F_{u_0} with a measurement operator

$$\mathcal{P} : H^{1/2}(\Gamma^-) \times H^{1/2}(\Gamma^+) \longrightarrow \mathbb{C}^{N_r},$$

where $\mathcal{P}$ samples the pair of boundary traces at the N_r receiver locations. Hence the forward map available to inversion is

$$\mathcal{P} \circ F_{u_0} : \mathcal{A} \longrightarrow \mathbb{C}^{N_r},$$

which maps each admissible contrast to the data predicted at the receivers.

4.3 Inverse Problem

The classical inverse problem is to recover the unknown contrast m_s from one or several boundary measurements of the scattered field. For acquisitions $q = 1, \dots, Q$ (different incident fields, modes or frequencies), this means solving

$$\mathcal{P}F_{u_0,q}(m_s) = d_q, \quad q = 1, \dots, Q,$$

where d_q denotes the corresponding finite-dimensional data. In practice, these data are noisy and available only for a finite acquisition set. The exact equations are therefore replaced by an optimisation problem that balances data fidelity with admissibility and regularity constraints on the contrast.

This formulation highlights the role of the forward problem in the rest of the report. Before a neural method can be trusted for inversion, it must first reproduce the scattered field accurately for a known contrast, respect the waveguide boundary conditions, and handle the oscillatory behaviour of the Helmholtz equation. The forward PINN study in Section 6 is built around this requirement.

4.4 Methods and Common Challenges in the Inverse Problem

The standard approach to this inverse problem is known as Full Waveform Inversion (FWI) [11]. FWI seeks the model parameter m that minimises a least-squares misfit between the predicted scattered field and the observed data. For a set of sources indexed by s , the objective reads

$$J_{\text{FWI}}(m) = \frac{1}{2} \sum_s \|\mathcal{P}u_s(m) - d_s\|^2, \quad (4.2)$$

where $u_s(m)$ is the wavefield obtained by solving the forward problem with model m , $\mathcal{P}$ is the measurement operator, and d_s denotes the corresponding observations. A key difficulty arising in the minimisation of (4.2) is *cycle skipping*: when the predicted and observed signals differ by more than roughly half a cycle (or their travel times by more than half a period), a pointwise least-squares misfit can associate the wrong oscillations and lead a gradient method to a local minimum [11].

An alternative formulation designed to mitigate cycle skipping is the Wavefield Reconstruction Inversion (WRI) [12]. Rather than enforcing the wave equation as a hard constraint, WRI relaxes it into a penalty term and optimises *jointly* over the wavefields $\mathbf{u}$ and the model $\mathbf{m}$. The objective introduced in [12] is

$$J_{\text{WRI}}(m, u) = \sum_{i=1}^M \left[\frac{1}{2} \|\mathcal{P}_i \mathbf{u}_i - \mathbf{d}_i\|_2^2 + \frac{\lambda^2}{2} \|A_i(\mathbf{m}) \mathbf{u}_i - \mathbf{q}_i\|_2^2 \right], \quad (4.3)$$

where M is the number of sources, $\mathcal{P}_i$ is the measurement operator for source i , $\mathbf{d}_i$ are the observed data, $A_i(\mathbf{m})$ is the discretised Helmholtz operator, $\mathbf{q}_i$ is the source term, and $\lambda > 0$ is a penalty parameter balancing data fidelity against physics compliance. This *all-at-once* formulation is structurally analogous to the PINN loss, which similarly penalises both data misfit and PDE residual simultaneously.

Both methods share two major computational challenges. First, *cycle skipping* makes optimisation sensitive to the initial model. Second, each evaluation of the forward model requires solving a large linear system defined by a finite-difference or finite-element matrix, a cost that

grows rapidly with frequency, domain size and irregularities. In FWI, an adjoint solve provides the model gradient; WRI instead enlarges the optimisation space to include the wavefield. The PINN formulation studied here asks whether neural coordinate representations and automatic-differentiation residuals provide a useful all-at-once alternative. This is a hypothesis, not an assumed cost reduction: PINN training also requires many residual and gradient evaluations, and matched forward benchmarks have found no general advantage over FEM [13].

4.4.1 Intrinsic Limitations in Waveguide Inversion

At a fixed frequency, specially tuned compact perturbations of a penetrable waveguide coefficient can have the same scattering matrix on propagating modes as the homogeneous guide [14]. The cited construction is proved completely in the monomode regime; the multimode case is supported by the continuation framework and numerical examples. This is a genuine finite-frequency non-uniqueness mechanism for distant propagating data, but it does not imply that the scattered field vanishes at finite ports, that invisibility persists across frequencies, or that a generic defect is invisible. Multiple frequencies and incident modes reduce this specific ambiguity, although they do not by themselves prove uniqueness or stability.

5 Physics-Informed Neural Networks and Implemented Workflow

5.1 Concept and Positioning

Physics-Informed Neural Networks (PINNs) were introduced by Raissi et al. [1] and use neural networks to simulate physical systems and infer their parameters. Neural networks lie at the core of deep learning: under mild conditions, they act as universal approximators and can represent any continuous function to arbitrary precision. The most widely used architecture in the PINN literature is the multilayer perceptron (MLP), a fully connected feedforward network; its formal definition and training properties are recalled in Appendix D. The central idea of PINNs is to embed the governing physics equations directly into the loss function, so that the network is trained to converge towards a continuous approximation of the true physical solution.

This paradigm differs fundamentally from classical supervised deep learning. In a standard setting, the goal is to generalise a concept from a finite, possibly sparse, dataset while avoiding overfitting. In the PINN setting, overfitting in the classical sense is not the primary concern: the network is required to satisfy the physics. As a consequence, regularisation techniques designed to prevent overfitting, such as dropout, are generally not appropriate here.

In the context of the present work, PINNs offer a natural alternative to FWI and WRI, bypassing the need for a dedicated grid-based solver by representing both the wavefield and the model parameter through a trainable network.

5.2 State of the Art

Since the introduction of PINNs, the literature has expanded rapidly, with new variants and applications appearing at a steady pace. This subsection focuses on the ideas and limitations that are directly relevant to the stated waveguide problem and to the methodological choices made in this work. Existing improvements can broadly be divided into two families: architectural modifications of the neural network and changes in the training procedure. This diversity makes PINNs a flexible framework, but also complicates the identification of robust design rules, since a method that performs well for one equation or one regime may fail for another.

The Helmholtz equation is a particularly challenging case for PINNs. Standard multilayer perceptrons are known to suffer from a spectral bias, which makes the approximation of high-frequency solutions difficult. Several approaches have been proposed to mitigate this issue, including Fourier feature embeddings and sinusoidal representations such as SIREN-type networks [15, 16]. Deep MLPs also suffer from optimisation difficulties such as vanishing gradients when the number of layers increases; residual adaptive architectures such as PirateNets have been proposed to address this type of limitation [4].

For inverse problems, the training strategy is as important as the architecture. Data-guided or data-driven pretraining can accelerate convergence by first fitting the available observations and boundary or initial conditions before introducing the full physics residual [3]. Another recurrent difficulty is the multi-objective nature of the PINN loss: the data, boundary, regularisation, and PDE terms must be weighted carefully to ensure that no component dominates the optimisation. Adaptive weighting strategies have therefore been developed to reduce gradient pathologies and improve convergence [17, 18]. However, these techniques remain problem-dependent and do not provide a universal remedy.

Most successful PINN demonstrations still concern forward problems or relatively smooth and homogeneous domains. Applications to FWI or WRI are less common, and many of them are formulated in the time domain or in seismic-imaging settings where the wavelengths are comparatively large [19, 20]. In ultrasound tomography, PINNs have shown encouraging results in some studies [21], but discontinuous sound-speed maps remain difficult: in particular, the reconstructed medium may be over-smoothed, leading to failure in strongly discontinuous cases [22]. For inverse formulations, using distinct networks for independent physical quantities, for example separate networks for the real and imaginary parts of the pressure field and another one for the wavespeed, has also been recommended [23].

Overall, the main open challenges for PINNs remain optimisation, lack of generalisation, sensitivity to architecture and hyperparameters, and the absence of complete theoretical guarantees on convergence and error control [24]. A low training loss does not necessarily imply that the learned solution is close to the true physical solution, especially when no reference solution is available. This is particularly important for inverse wave problems, where failure cases are not always explicitly reported and where the number of modelling and training choices quickly becomes large.

Although no unified theory currently describes PINN behaviour across all problems, the Neural Tangent Kernel (NTK) framework provides the main theoretical reference used in this work [25]. In the infinite-width regime, it approximates discrete parameter updates by continuous dynamics governed by a kernel, whose spectrum helps explain differences in the convergence rates of the loss components. This perspective provides theoretical motivation for the Fourier features used here [15] and has also been used to justify Self-Adaptive Weights [18], although those weights are not employed in the experiments reported in this work.

5.3 Advantages and Limitations

The main appeal of PINNs in this work is their ability to represent the solution as a continuous, differentiable function without relying on a dedicated mesh-based solver. Spatial derivatives are obtained by automatic differentiation for the classical MLP and propagated analytically through the Fourier networks used in the experiments, allowing the physical residual to be assembled directly from the governing equation. This common differentiable framework allows the wavefields and the unknown material to be inferred, in principle, within the same optimisation procedure.

These advantages come with important practical limitations. PINNs are highly sensitive to the architecture, sampling strategy, optimiser, and relative weights assigned to the different loss terms. In high-frequency Helmholtz problems, standard networks may struggle to represent oscillatory solutions, while discontinuous coefficients tend to be smoothed by the neural representation. Moreover, a small training loss does not necessarily guarantee an accurate physical solution. For this reason, the inverse problem is not addressed directly: the following section first studies the forward problem in order to identify architectures and training choices that are sufficiently reliable for the later inversion experiments.

5.4 Framework

The most widely used frameworks for implementing PINNs include DeepXDE, PyTorch, and TensorFlow. In this work, all experiments are carried out with JAX, a numerical computing

library developed by Google that combines NumPy-compatible syntax with automatic differentiation, just-in-time compilation via XLA, and transparent hardware acceleration on GPU and TPU. The choice of JAX was motivated both by prior familiarity and by the functional programming paradigm it promotes, which naturally suits the automatic-differentiation-intensive workflows required by PINNs: the physics residual, the gradient of the loss, and higher-order derivatives are all expressed as compositions of pure functions, making the code both concise and efficient. The entire implementation was built from scratch on top of raw JAX primitives, without relying on higher-level neural-network libraries such as Flax or Equinox; this deliberate choice was made primarily for pedagogical reasons, in order to gain a thorough understanding of every component of the training pipeline.

Further implementation and reproducibility details are provided in Appendix C.

5.5 PINN Formulation of the Problem

A PINN replaces a grid-based discretisation of the unknown field with a coordinate-based neural representation. For an acquisition $q = (\text{incidence}, f, p)$, where p is the incident mode, an independent network $\mathcal{N}_{\theta_q} : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ receives a spatial coordinate and returns the real and imaginary parts of the amplitude-normalised pressure:

$$(x, y) \mapsto (\hat{x}, \hat{y}) = \left(\frac{x}{L}, \frac{2y}{H} - 1 \right) \mapsto \mathcal{N}_{\theta_q}(\hat{x}, \hat{y}) = (\text{Re } \tilde{u}_q, \text{Im } \tilde{u}_q). \quad (5.1)$$

The normalised coordinates lie in $[-1, 1]^2$, which avoids presenting the first network layer with inputs of markedly different scales. Depending on the tested architecture, they are passed directly to a classical MLP or first mapped to Fourier features and then processed by a classical or modified MLP. The corresponding architectural and implementation details are given in Appendices C and D.

In the forward problem, the squared slowness $m = 1/c^2$ is prescribed. In the inverse problem, a second network $m_\phi : \Omega \rightarrow \mathbb{R}$ represents the unknown squared slowness and is shared by all acquisitions, whereas each pressure field keeps its own parameters θ_q . Its scalar output is passed through a sigmoid and an affine map so that $m_\phi(x, y)$ remains within the prescribed physical bounds. This all-at-once workflow is summarised in Figure 3.




Figure 3: Implemented PINN workflow. Each acquisition q has an independent complex-field network, while the inverse problem uses one material network shared by all acquisitions. In the forward problem, the known material replaces the lower trainable branch and the data loss is omitted; Adam minimises the applicable weighted losses.

The governing equations are not supplied as labelled training examples. Instead, automatic differentiation evaluates the spatial derivatives of the network output at collocation points, and the resulting equation and boundary residuals become loss terms. Let $m_0 = 1/c_0^2$, $k_{0,q} = \omega_q/c_0$, and let $u_{0,q}$ denote the known incident field. The dimensionless pointwise Helmholtz residual is

$$r_{\text{tot}}^q = \frac{\Delta \tilde{u}_{\text{tot},q} + \omega_q^2 m \tilde{u}_{\text{tot},q}}{k_{0,q}^2}, \quad (5.2)$$

$$r_{\text{scat}}^q = \frac{\Delta \tilde{u}_{\text{scat},q} + \omega_q^2 [m \tilde{u}_{\text{scat},q} + (m - m_0) \tilde{u}_{0,q}]}{k_{0,q}^2}. \quad (5.3)$$

The first expression is used when the network predicts the total field and the second when it predicts the scattered field.

For the inverse problem, define the non-zero transverse modal scale

$$\lambda_p^* = \begin{cases} \lambda_1, & p = 0, \\ \lambda_p, & p \geq 1. \end{cases} \quad (5.4)$$

The boundary-loss normalisation differs between the forward calibration and the inverse problem. To make this distinction explicit, let

$$a_{\text{N},q} = \begin{cases} k_{0,q}^2, & \text{forward study,} \\ \lambda_p^*, & \text{inverse problem,} \end{cases} \quad a_{\text{DtN},q} = \begin{cases} k_{0,q}^2, & \text{forward study,} \\ \beta_p^2, & \text{inverse problem.} \end{cases} \quad (5.5)$$

Thus, in the forward study, both first-order boundary residuals are divided by $k_{0,q}$ before squaring. The modal factors λ_p^* and β_p^2 are introduced only for the inverse problem, where they reflect the transverse and longitudinal scales of incident mode p .

If N_r , N_N , and N_D denote the numbers of interior, rigid-wall, and port collocation points, respectively, the implemented physics losses are

$$\mathcal{L}_{\text{PDE}}^q = \frac{1}{2N_r} \sum_{j=1}^{N_r} |r^q(\mathbf{x}_j)|^2, \quad (5.6)$$

$$\mathcal{L}_{\text{N}}^q = \frac{1}{2N_N a_{\text{N},q}} \sum_{j=1}^{N_N} (|\partial_y \tilde{u}_q(x_j, H)|^2 + |\partial_y \tilde{u}_q(x_j, 0)|^2), \quad (5.7)$$

$$\mathcal{L}_{\text{DtN}}^q = \frac{1}{N_D a_{\text{DtN},q}} \sum_{j=1}^{N_D} (|r_{\text{DtN}}^{q,-}(y_j)|^2 + |r_{\text{DtN}}^{q,+}(y_j)|^2). \quad (5.8)$$

Here $r_{\text{DtN}}^{q,\pm} = \partial_\nu \tilde{u}_q - T_{\omega_q} \tilde{u}_q - \tilde{g}_q^\pm$ for the total field, whereas $\tilde{g}_q^\pm = 0$ for the scattered field. Thus, the total-field DtN residual includes the incident contribution on the incoming boundary; the scattered-field condition is homogeneous. Modal traces are integrated by Gauss–Legendre quadrature using $3N_m$ points, with $N_m = \text{round}(2Hf/c_0) + 5$ retained modes.

In inverse experiments, an additional data term compares the predicted and observed complex pressures at the measurement points $\mathbf{x}_j^d$:

$$\mathcal{L}_{\text{data}}^q = \frac{1}{N_d} \sum_{j=1}^{N_d} |\tilde{u}_q(\mathbf{x}_j^d) - \tilde{d}_q^j|^2. \quad (5.9)$$

The three physics components are combined into the total forward objective

$$\mathcal{L}_{\text{phys}}^q = w_{\text{PDE}}\mathcal{L}_{\text{PDE}}^q + w_{\text{N}}\mathcal{L}_{\text{N}}^q + w_{\text{DtN}}\mathcal{L}_{\text{DtN}}^q. \quad (5.10)$$

For inverse training, the field objective additionally contains the data term,

$$\mathcal{L}_u^q = \mathcal{L}_{\text{phys}}^q + \alpha_t w_{\text{data}}\mathcal{L}_{\text{data}}^q, \quad (5.11)$$

with the precise ramp α_t and data-weight schedule specified in the inverse-study sections. The normalisation factors in Equation (5.5) set the scale of each residual, whereas the weights w_i in Equation (5.10) set their relative influence on the parameter update. Section 6 first compares unit weights (1, 1, 1) with the manually balanced choice $(10^3, 10^4, 10)$ in PDE/Neumann/DtN order. Adam updates the network parameters from gradients of the applicable total objective. The specific architectures, Fourier embedding, adaptive loss-weighting strategy, and residual-based adaptive sampling procedure are documented in Appendix C.

6 Forward PINN Study

6.1 Purpose and Experimental Protocol

Before addressing the inverse problem, the wavefield network must be able to solve the forward problem for a known medium. The campaign isolates the field branch of Figure 3: the material is fixed, and the data-loss and material-update branches are removed. It evaluates separately the effects of loss weighting, frequency, architecture, field formulation, adaptive loss weights, and residual-based adaptive distribution (RAD), both to select a credible inverse starting point and to quantify the forward mismatch that can remain in the coupled problem.

This is deliberately a targeted calibration study rather than an exhaustive benchmark of forward PINN solvers. The components reused in the inverse problem are compared with the same number of Adam iterations and collocation points per step; parameter counts, derivative implementations, and wall-clock costs differ across architectures. The aim is to obtain sufficiently accurate convergence at moderate cost, since otherwise classical WRI or hybrid approaches such as [26] remain preferable.

The waveguide has half-length $L = 1$ and height $H = 0.6$. A mode-0 incident wave propagates from left to right. Each network is trained for 30 000 Adam iterations with 4096 interior points, 256 abscissae evaluated on each rigid wall, and 256 ordinates evaluated on each DtN boundary. The classical, Fourier, and Fourier-modified architectures all have three hidden layers of width 128. The two Fourier variants use 64 Fourier features. The total physics objective and its normalisation are defined in Equations (5.5) and (5.10); precise architecture and adaptive-strategy details are given in Appendix C, while reference-data generation is described in Appendix B.

Unless stated otherwise, the curves show the mean and one sample standard deviation over independent random seeds, together with the individual runs. The reported errors are measured on a fixed validation mesh using the mass and stiffness matrices. The reference fields do not enter the physics loss, the optimiser, or checkpoint selection; they serve only for periodic evaluation and comparison between configurations. Consequently, the comparisons use equal iteration and collocation counts, not equal parameter counts or wall-clock time, and do not establish asymptotic convergence.

6.2 Homogeneous Waveguide

For a homogeneous medium, the analytical total-field solution is the incident mode. For the mode-0 wave travelling from left to right in this study, the reference is

$$u_{\text{ref}}(x, y) = u_0(x, y) = \theta_0(y) \exp(i\beta_0 x). \quad (6.1)$$

Every 1000 Adam steps, the current approximation $u_{\text{PINN}}^{(t)}$ is evaluated on the fixed validation mesh and compared with this analytical solution through

$$e^{(t)} = \frac{\|u_{\text{PINN}}^{(t)} - u_{\text{ref}}\|}{\|u_{\text{ref}}\|}, \quad (6.2)$$

The graphs therefore show the evolution during training of the relative L^2 and H^1 errors between the PINN approximation and the analytical field, rather than only their final values. This controlled comparison isolates optimisation and representation errors from any error caused by a material interface.

6.2.1 Loss weighting at 600 Hz ($k_0 \approx 11.1 \text{ rad m}^{-1}$)

Figure 4 shows that unit weights do not provide a fair balance between the three loss terms. The classical MLP remains at approximately 56% relative error, whereas the Fourier embeddings reduce the error below 1%. Thus, even at this frequency, equal numerical loss weights do not imply comparable effects on the parameter update.




Figure 4: Evolution during training of the relative L^2 and H^1 errors between the PINN approximation and the analytical incident-mode solution at 600 Hz, with $(w_{\text{PDE}}, w_{\text{N}}, w_{\text{DtN}}) = (1, 1, 1)$. Thin curves are individual seeds; the solid curve and shaded region are the mean and one sample standard deviation over three seeds.

A heuristic choice guided by the gradient-scale diagnostic gives the manual weights $(10^3, 10^4, 10)$, proportional to $(100, 1000, 1)$. With these weights all three architectures learn the solution, including the classical MLP, and the Fourier-modified network reaches 0.15% relative L^2 error (Figure 5). This comparison establishes that an apparent architecture failure can in fact be caused by an unbalanced multi-objective loss.




Figure 5: Evolution during training of the relative L^2 and H^1 errors between the PINN approximation and the analytical incident-mode solution at 600 Hz, with manually balanced loss weights $(w_{\text{PDE}}, w_{\text{N}}, w_{\text{DtN}}) = (10^3, 10^4, 10)$. Statistics are computed over three seeds.

6.2.2 Frequency and adaptive weights at 1200 Hz ($k_0 \approx 22.2 \text{ rad m}^{-1}$)

At twice the frequency, the same normalisation and manual weights are retained. Under this width-128, 30 000-iteration protocol, the classical MLP remains near 66% error for every tested seed, whereas both Fourier architectures remain below 1% (Figure 6). This result is consistent with the spectral bias of standard MLPs [27]: loss balancing removes one optimisation obstacle but does not make the tested classical network competitive at the larger wavenumber.

The transfer of the same weights between the two frequencies is assisted by the forward normalisation defined in Equation (5.5): the Helmholtz residual is divided by k_0^2 , while the first-order Neumann and DtN residuals are divided by k_0 . Equivalently, each boundary loss is divided by k_0^2 after squaring. The modal normalisations based on λ_p^* and β_p are introduced only for the inverse problem.




Figure 6: Evolution during training of the relative L^2 and H^1 errors between the PINN approximation and the analytical incident-mode solution at 1200 Hz. Static variants use $(w_{\text{PDE}}, w_{\text{N}}, w_{\text{DtN}}) = (10^3, 10^4, 10)$; adaptive variants update inverse-gradient-norm weights during training. Statistics are computed over five seeds.

Table 1 summarises the final errors. The Fourier-modified architecture is consistently the most accurate of the tested networks. Adaptive loss weights do not improve this homogeneous case: at 1200 Hz, they increase the mean error from 0.46% to 0.74% for the standard Fourier MLP and from 0.19% to 0.27% for the modified one. A dynamically balanced loss is therefore not automatically a better objective once reliable static weights have been found.

Configuration	Classical MLP	Fourier MLP	Modified Fourier MLP
600 Hz, unit weights	56.03 ± 5.32	0.72 ± 0.11	0.37 ± 0.09
600 Hz, manual weights	0.60 ± 0.14	0.33 ± 0.03	0.15 ± 0.02
1200 Hz, manual weights	65.65 ± 1.03	0.46 ± 0.07	0.19 ± 0.05
1200 Hz, adaptive weights	—	0.74 ± 0.20	0.27 ± 0.04

Table 1: Final relative L^2 errors in percent (mean $\pm$ sample standard deviation). Three seeds are used at 600 Hz and five at 1200 Hz.

6.3 Waveguide Containing a Sharp Defect

The heterogeneous test uses the same 1200 Hz excitation and contains a circular defect of radius 0.1, centred at (0.2, 0.2), in which the wave speed is $0.8c_0$. The reference wavefield is computed with quadratic finite elements. In contrast to the homogeneous test, the coefficient is discontinuous at the circle: the wavefield and its normal derivative remain continuous, but not its second derivatives. This is a demanding case for a globally smooth network trained with a strong-form residual.

It is also intentionally more severe than the coefficient maps that the inverse PINN will represent during optimisation. The inverse material network uses smooth activations and therefore produces continuous, smooth transitions at finite width, even when the true target has a sharp boundary. The discontinuous forward coefficient should consequently be viewed as a stress test of the wavefield solver, not as an exact reproduction of the intermediate media encountered during inversion.

6.3.1 Total-field and scattered-field formulations

The four combinations of field formulation and Fourier architecture are compared in Figure 7. Predicting the scattered field lowers the final error in all three paired runs for both architectures; with only three seeds, this is descriptive evidence rather than a general statistical conclusion. The modified MLP gives the larger mean gain. The best configuration in this comparison is the modified MLP in the scattered-field formulation, with an error of $7.98 \pm 1.29\%$.




Figure 7: Circular-defect case at 1200 Hz: total-field and scattered-field formulations with static manual weights. Statistics are computed over three seeds.

The increase from 0.19% in the homogeneous case to approximately 8% is consistent with a harder approximation problem: the defect simultaneously introduces a coefficient jump, a 25% larger local wavenumber, and scattered-wave content. Because these changes are not isolated experimentally, the increase cannot be attributed specifically to the interface regularity mismatch. It should be read as the performance of the present solver on this sharp-defect benchmark, not as a measured interface penalty or an irreducible error floor.

6.3.2 Adaptive loss weights and RAD

Starting from the modified scattered-field network, Figure 8 compares static weights, adaptive weights, RAD, and their combination. Adaptive loss weights improve every paired seed and reduce the mean error from 7.91% to 7.08%. In contrast, the tested RAD configuration degrades every paired seed, whether or not adaptive weights are enabled. Its additional sampling and refresh steps also increase the mean training time by approximately 13–17%.




Figure 8: Circular-defect case with the modified scattered-field network. Adaptive loss weighting and residual-based adaptive distribution (RAD) are compared with the static baseline over three paired seeds.

Configuration	Relative L^2 (%)	Relative H^1 (%)	Time (s)
Fourier, total field	10.66 ± 0.79	10.68 ± 0.80	88.3
Fourier, scattered field	10.18 ± 1.04	10.22 ± 1.05	86.5
Modified Fourier, total field	8.52 ± 0.93	8.53 ± 0.93	146.4
Modified Fourier, scattered field	7.98 ± 1.29	7.99 ± 1.30	146.7
Modified scattered, static weights	7.91 ± 1.25	7.92 ± 1.25	146.8
Modified scattered, adaptive weights	7.08 ± 1.37	7.09 ± 1.37	146.8
Modified scattered, RAD	9.43 ± 1.14	9.49 ± 1.14	166.2
Modified scattered, RAD and adaptive weights	11.57 ± 1.40	11.63 ± 1.40	171.1

Table 2: Final defect-case errors and measured optimiser/adaptation times over three seeds. Errors are reported as mean $\pm$ sample standard deviation; times are seed means. Times exclude JIT compilation and monitoring or reference evaluations.

This negative result is specific to the tested RAD policy, not to adaptive sampling in general. A plausible interpretation is that RAD repeatedly selects the large residuals generated near the interface, although the globally smooth network cannot efficiently remove them. More points then reinforce the same local approximation conflict. Moreover, several defect-case curves are still decreasing at iteration 30 000; the table should therefore be interpreted as a fixed-budget comparison, not as proof of an irreducible 7% error floor.

6.4 Lessons for the Inverse Problem

Two conclusions are retained as design requirements for the inverse problem. First, balancing the PDE, Neumann, and DtN loss weights is critical: unit weights can make an otherwise suitable architecture appear unable to learn the solution. Second, Fourier features are essential in the frequency range considered here. Under the tested width-128, 30 000-iteration protocol at $k_0 \simeq 22.2 \text{ rad m}^{-1}$, the classical MLP remains near 66% relative L^2 error, whereas both Fourier architectures achieve sub-percent errors. The inverse experiments therefore retain a Fourier representation and carefully balanced physics-loss weights.

Analytically propagating the Fourier networks' spatial first derivatives and Laplacians also considerably accelerates training relative to the earlier automatic-differentiation implementation. This gain makes the subsequent multi-seed studies more practical.

These choices define the retained configuration rather than universal prescriptions for PINNs. Among the tested models, the Fourier-modified network gives the lowest mean error in every forward comparison, but requires roughly two-thirds more optimiser time than the standard Fourier MLP in the defect tests and does not always have the smallest seed-to-seed variability. The scattered-field formulation gives a modest improvement in all three paired defect runs and removes the known incident field from the learned quantity. Adaptive weights improve all three defect runs but worsen the mean homogeneous results; their benefit is therefore setting-dependent. The tested RAD policy is not retained.

The approximately 7% error obtained on the sharp-defect benchmark is a limitation observed for the present global, strong-form solver and training recipe, not evidence of a universal or irreducible PINN error floor. Because the inverse material network represents smooth transitions, the same error level need not occur during inversion. Nevertheless, forward mismatch may blur the reconstructed contrast or be absorbed into the material estimate; inverse conclusions must therefore be assessed from data residuals and reconstruction error, not from the PINN loss alone. These downstream effects are evaluated in the complete inverse pipelines of Sections 7 and 8, with representative reconstructions in Figures 9 and 11.

7 First Inverse Approach: Baseline PINN Formulation

7.1 All-at-Once Formulation and Training

The inverse study activates the complete pipeline in Figure 3. It follows a WRI-like formulation: the wavefields and the medium are treated as independent unknowns, while the wave equation is imposed as a penalty rather than as a hard constraint. For each acquisition $q = (\text{incidence}, f, \text{mode})$, an independent network $u_{\theta_q} : \Omega \rightarrow \mathbb{C}$ represents the total pressure. A single network $m_\phi : \Omega \rightarrow \mathbb{R}$ represents the squared slowness shared by all acquisitions. The field and material objectives are

$$\mathcal{L}_u^q = w_{\text{PDE}} \mathcal{L}_{\text{PDE}}^q + w_{\text{N}} \mathcal{L}_{\text{N}}^q + w_{\text{DtN}} \mathcal{L}_{\text{DtN}}^q + \alpha_t w_{\text{data}} \mathcal{L}_{\text{data}}^q, \quad (7.1)$$

$$\mathcal{L}_m = \frac{1}{|\mathcal{Q}|} \sum_{q \in \mathcal{Q}} \mathcal{L}_{\text{PDE}}^q + w_{\text{TV}} \text{TV}(m_\phi), \quad (7.2)$$

where the last term is present only in the TV experiment. The data loss compares the complex pressure with 30 synthetic FEM samples on each cross-section, so the pressure networks are guided directly by the measurements.

At each Adam step, the field and material gradients are evaluated on the same newly sampled collocation cloud and at the same pre-update parameter values. The two optimiser updates are then applied. There is no differentiation through the field update: the data affect m_ϕ only indirectly, through the pressure iterates passed from one step to the next.

The baseline acquisition set contains the three propagating modes at 700 Hz, with incidence from left to right. This frequency admits modes 0–2 while avoiding the immediate vicinity of the mode-2 cut-off. Incident modes close to cut-off can produce reconstruction artefacts; avoiding them was therefore an important practical constraint during the inversion.

The simplest field baseline uses 64 Fourier features followed by three width-128 hidden layers. The modified Fourier architecture from Section 6 is tested separately; its superior forward accuracy is not assumed to transfer to inversion. The material network is a two-layer, width-64 tanh MLP whose scalar output is mapped by a sigmoid into the squared-slowness interval induced by the prescribed sound-speed bounds and initialised at the homogeneous background. Adaptive weights and TV regularisation are likewise treated as experimental variants.

Each total-field network is first trained for 5 000 Adam steps in the known homogeneous medium, without data loss. Joint inversion then uses $(w_{\text{PDE}}, w_{\text{N}}, w_{\text{DtN}}, w_{\text{data}}) = (10^3, 10^3, 10, 10^3)$. During its first 10 000 steps, α_t increases linearly from 0.1 to 1 and also scales the material update. This prevents the initially homogeneous model from changing before the pressure fields have adapted to the observations. The field and material parameters use an initial learning rate of 10^{-3} followed by a delayed cosine decay to 10^{-4} . This schedule was introduced in response to late oscillations seen in preliminary runs, but its effect is not isolated in the campaigns below.

Training is performed in single precision. Because the Helmholtz residual contains second derivatives, only a paired double-precision experiment could determine whether round-off contributes materially to the observed stagnation.

7.2 Evaluation Protocol and Results

The target is a rectangular defect $\Omega_d = [-0.2, 0.2] \times [0, 0.3]$ with $c_d = 0.8c_0$ in a background $c_0 = 340 \text{ m s}^{-1}$. Since Ω_d occupies only about 10% of Ω , a global relative L^2 error is dominated by the background and can make an almost homogeneous map appear accurate. The primary metric is therefore the contrast-normalised global error, evaluated on the same fixed uniform grid for every run,

$$E_{\Delta c} = \frac{\|c_\phi - c^*\|_{L^1(\Omega)}}{\|c^* - c_0\|_{L^1(\Omega)}}. \quad (7.3)$$

In practice, both norms are approximated by equal-weight sums on the same fixed 256×80 grid. The metric therefore has the exact reference value $E_{\Delta c} = 1$ for the trivial reconstruction $c_\phi = c_0$. Because it still mixes localisation and background artefacts, it is interpreted together with the mean absolute errors inside Ω_d and in $\Omega \setminus \Omega_d$.

These metrics use the synthetic truth only for final evaluation. Checkpoints are selected without that truth by minimising the fixed-weight pressure objective on a fixed monitoring cloud. However, the monitoring objective reuses the observed port traces employed during training and is therefore not an independent validation set. The reported mean and sample standard deviation use three optimisation seeds, varying random initialisation and stochastic collocation. They describe optimisation variability, not confidence intervals for measurement, model or reconstruction uncertainty.

Experiment	Configuration	$E_{\Delta c}$	MAE _D (ms ⁻¹)
Reference	Homogeneous reconstruction $c = c_0$	1.00	68.0
A	Fourier, fixed weights, one incidence at 700 Hz	1.80 ± 0.25	27.1 ± 8.6
A	Modified Fourier, fixed weights	2.23 ± 0.15	32.5 ± 7.2
A	Field-adaptive weights, Fourier / modified Fourier	2.85 ± 0.36 / 2.81 ± 0.12	55.0 ± 8.6 / 54.1 ± 5.1
A	Material-adaptive weights and TV, Fourier	1.89 ± 0.38	32.3 ± 11.8
B	Two incidences at 700 Hz, Fourier	2.07 ± 0.65	33.6 ± 19.3
C	One incidence at 700 Hz and 1100 Hz, Fourier	1.47 ± 0.05	16.3 ± 1.5
D	Restricted $c/c_0 \in [0.6, 1.01]$, Fourier / modified Fourier	0.98 ± 0.14 / 0.91 ± 0.22	22.7 ± 3.8 / 21.3 ± 5.5
E	Restricted range, end of the bidirectional 700 Hz stage	0.80 ± 0.09	18.2 ± 1.8
E	Final stage after parameter continuation at two frequencies	0.34 ± 0.01	6.5 ± 0.7

Table 3: Contrast-normalised and in-defect mean absolute sound-speed errors for the tested total-field variants (mean $\pm$ one sample standard deviation, three optimisation seeds). Experiments A–C allow $c/c_0 \in [0.6, 1.4]$; D–E incorporate the prior that the anomaly is not appreciably faster than the background. Lower is better for both metrics.

Experiments A–C compare architecture and weighting variants with changes to the acquisition content under the broad material bounds. The contrast-normalised metric changes the interpretation of these first campaigns. For reference, the homogeneous reconstruction has an error of 68 ms^{-1} inside Ω_d and zero error in the background. Experiment B, which adds the opposite incidence at a single frequency, yields no consistent improvement over A and increases the seed-to-seed spread. Experiment C, which adds 1100 Hz, is more useful: relative to A, the mean absolute error in the defect decreases from 27.1 ms^{-1} to 16.3 ms^{-1} and the variability of $E_{\Delta c}$ becomes small. Nevertheless, $E_{\Delta c} > 1$ because errors remain distributed throughout the background. The gradient-cosine diagnostics are consistent with an optimisation difficulty: acquisition-wise material gradients can become opposed and their alignment changes during

training. They do not, by themselves, establish that one incidence is physically redundant or quantify identifiability.

Campaign D is a controlled ablation of the admissible material range. The recorded run manifests show that the executed A and D configurations use the same acquisitions, architecture, seeds, 5 000-step warm-up, 40 000 inverse steps and learning-rate schedule; only the bounds change. For the standard Fourier model, $E_{\Delta c}$ decreases for every paired seed and its mean falls from 1.80 to 0.98. The main visible effect is to suppress background artefacts: the background mean absolute error falls from 10.8 ms^{-1} to 5.0 ms^{-1} , while the defect error decreases from 27.1 ms^{-1} to 22.7 ms^{-1} . The upper bound $1.01c_0$, rather than c_0 , is numerical slack that keeps the homogeneous initialisation strictly inside the sigmoid-parameterised interval; the physical prior is that the anomaly is not faster than the background. This prior remains strong and is justified only when faster inclusions can be excluded independently; neither the location nor the shape of the defect is supplied to the network.

Experiment E applies a curriculum. Six fields (three modes and two incidences) are first trained at 700 Hz; the inferred material and the left-incident fields are retained, the new 1100 Hz fields are pretrained on that material, and the seven left-incident fields are jointly refined. The Adam states and learning-rate schedules are restarted between stages: “continuous” therefore refers to parameter continuation, not to an uninterrupted optimiser trajectory. Within the same three runs, $E_{\Delta c}$ decreases from 0.80 ± 0.09 to 0.34 ± 0.01 ; the final mean absolute errors are 6.5 ms^{-1} in the defect and 1.84 ms^{-1} in the background. The acquisition-averaged pressure errors are also reduced to $5.81 \pm 0.28\%$ in L^2 and $6.37 \pm 0.30\%$ in H^1 . This staged decrease measures the curriculum as a whole, not the isolated effect of frequency continuation: adding 1100 Hz, retaining only left incidence, warm starting, restarting the optimiser schedule and taking additional optimisation steps remain entangled.




Figure 9: Ground truth and representative reconstructions with a common colour scale. The displayed seed is the median seed according to $E_{\Delta c}$: experiment C gives $E_{\Delta c} = 1.47$, whereas the final stage of E gives $E_{\Delta c} = 0.33$. Seed selection is used only for visualisation; all numerical comparisons use the three runs in Table 3.

7.3 Lessons and Limitations

The forward-study ranking does not transfer directly to inversion. The modified Fourier network is about 67% slower than the standard one in A and provides no consistent inverse improvement: it is worse under the broad bounds, while under the restricted bounds the difference between means is smaller than the observed seed-to-seed spread. Adaptive field weights also degrade both the fixed-weight monitoring objective and the material estimate. Because their definition changes with the weights, their moving objective is not comparable across iterations. Finally, material-adaptive weighting and TV were enabled together, so their individual effects cannot be inferred; the combined variant helps one seed but not the mean. The present evidence therefore favours the simpler Fourier network with fixed weights.

Optimisation variability is a second methodological lesson. Some configurations exhibit substantial seed-to-seed dispersion, so an isolated run may not represent the typical behaviour of the method. During the early stages of the internship, configurations were not always repeated across several random seeds. An unusually favourable or unfavourable trajectory could therefore influence the interpretation of a result and the subsequent experimental choices. The current workflow addresses this weakness through controlled multi-seed campaigns, systematic retention of run configurations and per-seed histories, and reporting of means with sample standard deviations. These changes make the comparisons more reproducible, traceable and robust to atypical optimisation trajectories. The three seeds used here still characterise optimisation variability rather than statistical generalisation.

The acquisition studies are not equal-compute comparisons. Excluding compilation and monitoring, the mean measured training times are approximately 416 seconds for A, 1 063 seconds for B, 1 271 seconds for C and 1 488 seconds cumulatively for E. The accuracy gained by C and E must therefore be weighed against costs of about three and 3.6 times the baseline, respectively. A controlled follow-up should compare configurations at both equal iteration count and equal wall-clock budget.

More fundamentally, the relaxed physics constraint allows the networks to reproduce the discrete boundary traces without recovering the correct interior field. In a representative median baseline run, the normalised trace loss reaches the order of 10^{-7} while the acquisition-averaged pressure error remains about 20%. This gap is compatible with the combination of a soft PDE constraint, sparse boundary observations and the non-convexity discussed in Section 4.4; the experiment does not identify which mechanism dominates. It nevertheless confirms the warning from Section 5: neither a small data loss nor a small PINN objective certifies the interior solution. The smooth material network also blurs the sharp interface.

The study therefore demonstrates feasibility on these development cases, rather than robust inversion or out-of-sample generalisation. The configurations were developed and evaluated on the same target geometries and contrasts; no medium was held out as a benchmark.

The evaluation also uses noiseless FEM traces generated with the same nominal waveguide model and exact geometry. All receiver samples are used for training, with no held-out boundary data. The full FEM fields are independent of the training losses, but not of study-level model selection. Moreover, no noise, geometry-error or physics-mismatch test has yet been performed. A thresholded Dice or intersection over union is deliberately not reported: for a smooth material network, it would depend strongly on an arbitrary defect threshold. Complete pressure maps, loss histories, gradient cosines and per-seed metrics are also available in the project [GitHub repository](#), under `inverse_PINN/results`, so that these questions can be investigated without overloading the present report.

These limitations define the scope of the total-field baseline. The next section uses it as the reference for comparison with the scattered-field formulation.

8 Second Inverse Approach: Scattered-Field Formulation

8.1 Motivation and Controlled Protocol

The scattered-field reformulation was motivated by preliminary total-field runs in which the incident field appeared to overwhelm the weaker defect response. In retrospect, those runs do not constitute evidence that the total-field formulation is intrinsically inferior: they used neither the present normalisation nor the same number of boundary samples, iterations and modes. In particular, some total-field runs included mode 4 at 1200 Hz, whereas the initial scattered-field tests used only modes 0–3. The mode-4 cut-off is $f_4 = 4c_0/(2H) \simeq 1133$ Hz, only about 6% below the operating frequency. Their apparently better scattered-field maps are therefore retained as motivation only and are not compared quantitatively below.

The new formulation keeps the pressure and material architectures, material parameterisation, optimiser and fixed loss weights from Section 7. Only the represented field and its governing terms change. For each acquisition the network predicts $u_s = u - u_0$ and enforces Eq. (3.5); the known incident field is subtracted from the FEM port data and the DtN condition becomes homogeneous. Each formulation uses the maximum modulus of its target over the two ports as an amplitude scale; that scale is applied consistently to its field, traces and physics residuals. This prevents the small scattered trace from disappearing merely because of its physical amplitude, but also means that raw PINN losses from the two formulations are on different scales and must not be ranked. The comparable outputs are the recovered sound speed and the reconstructed total field $u_0 + u_s$.

A near-zero initial pressure output is consistent with the homogeneous scattered-field solution. Consequently, the initial background warm-up is removed, whereas the total-field pipeline uses 5 000 warm-up steps before the same 40 000 joint inversion steps. The comparison thus evaluates the two practical pipelines, including this intended saving; it is not an equal-step ablation of the PDE substitution alone.

The controlled campaign uses a circular defect of radius 0.1 m with $c = 0.8c_0$, centred at either $x = -0.2$ (left) or $x = 0.2$ (right), and a left-incident field at 1200 Hz. Modes 0–3, 30 FEM samples per port and acquisition, and three optimisation seeds are used throughout the paired comparison; mode 4 is deliberately excluded. Results are assessed with the contrast-normalised material error $E_{\Delta c}$ from Eq. (7.3) and the sound-speed mean absolute errors inside and outside the defect. As in Section 7, the values are descriptive means and sample standard deviations over only three optimisation runs, rather than uncertainty estimates.

8.2 Results and Interpretation

Defect	Sound-speed bounds	Field	$E_{\Delta c}$	MAE_D (m s^{-1})	MAE_{bg} (m s^{-1})	Time (s)
Left	$[0.6, 1.4]c_0$	Total	4.01 ± 0.11	43.0 ± 0.5	6.00 ± 0.20	615 ± 3
Left	$[0.6, 1.4]c_0$	Scattered	3.72 ± 0.11	40.9 ± 1.6	5.54 ± 0.18	550 ± 11
Right	$[0.6, 1.4]c_0$	Total	3.14 ± 0.16	33.5 ± 1.6	4.71 ± 0.30	613 ± 3
Right	$[0.6, 1.4]c_0$	Scattered	3.20 ± 0.31	32.5 ± 2.6	4.84 ± 0.52	548 ± 7
Left	$[0.6, 1.01]c_0$	Scattered	2.61 ± 0.23	32.1 ± 3.8	3.81 ± 0.31	524 ± 3
Right	$[0.6, 1.01]c_0$	Scattered	1.97 ± 0.16	24.9 ± 1.8	2.85 ± 0.24	527 ± 5

Table 4: Circular-defect inversions at 1200 Hz (mean $\pm$ one sample standard deviation, three optimisation seeds). Only the four broad-bound rows form a direct total/scattered comparison. Time excludes compilation and monitoring and includes the total-field warm-up when present.

Under broad bounds, the scattered formulation modestly lowers all three material errors for the left defect, but the improvement is not reproduced for the right defect: its $E_{\Delta c}$ and background error are slightly higher, with differences comparable to the seed spread. There is therefore no systematic accuracy advantage. Field accuracy is only a secondary diagnostic of material reconstruction. Moreover, normalising its L^2 error by the total FEM field is insensitive to part of the defect response because the denominator contains the exactly known incident field. As a post-hoc sensitivity check, dividing the same total-field reconstruction error by the FEM scattered-field norm for both pipelines gives $42.7 \pm 2.3\%$ (total) versus $51.9 \pm 2.7\%$ (scattered) on the left, and $33.9 \pm 1.0\%$ versus $40.8 \pm 1.3\%$ on the right. This common scale relates the error to the defect response, but it is not an independent validation metric and can amplify errors when that response is weak. Its lower values for the total-field runs therefore suggest, rather than establish, better field recovery here and do not imply a better material map.

All broad-bound values satisfy $E_{\Delta c} > 1$: globally, their error exceeds that of the trivial homogeneous map under this contrast-normalised metric. This does not mean that the reconstructions contain no information. A homogeneous map has an in-defect error of 68 m s^{-1} , whereas the observed values are 32 m s^{-1} to 43 m s^{-1} ; the defect is partially recovered, but this gain is outweighed by widespread background artefacts because each circle occupies only about 2.6% of the domain. Figure 11 illustrates this distinction without selecting the best run; the corresponding ground truth is shown in Figure 10.




Figure 10: Ground-truth sound-speed map for the right-hand circular defect used in the broad-bound comparison.




Figure 11: Broad-bound reconstructions of the right-hand circle on the colour scale of Figure 10. The dashed circle is the true boundary. Median seed 0 gives $E_{\Delta c} = 3.06$ (total) and 3.18 (scattered); both underestimate the contrast and show structured background artefacts.

The left defect remains harder than its translated right counterpart for both formulations under left incidence. This reproducible asymmetry is compatible with an illumination or optimisation sensitivity, but the present experiment does not isolate a proximity effect. Restricting the scattered-field model to $c \leq 1.01c_0$ improves both positions and suppresses part of the background error; relative to the corresponding broad-bound scattered runs, this is the only executed configuration change. As in Campaign D, the 1% upper slack keeps the homogeneous initialisation inside the sigmoid range; the physical prior is $c \leq c_0$. However, $E_{\Delta c}$ remains above the homogeneous reference, and no restricted total-field campaign was run, so this comparison supports the sign prior rather than the scattered-field reformulation.

Finally, an exploratory scattered-field run used two opposite incidences and two defects of opposite contrasts. Its single seed gives $E_{\Delta c} = 2.67$, with mean absolute errors of 50.4 m s^{-1} in the defects and 7.08 m s^{-1} in the background. The inferred map recovers the signs of both contrasts but substantially underestimates their amplitudes (mean predicted celerities of approximately 318 m s^{-1} and 353 m s^{-1} , versus 272 m s^{-1} and 408 m s^{-1}). Because this run also changes the acquisitions and learning-rate schedule and has no repeated seed or total-field control, it is only a qualitative feasibility test.

The clearest repeatable benefit observed here is computational: removing the initial warm-up reduces measured training time by about 10% while leaving material accuracy of the same order. This efficiency and the explicit separation of the known background make the scattered formulation a reasonable default for subsequent exploratory work, not a claim that it solves the inverse ambiguity or is more accurate. The conclusion is limited to noiseless data generated under the same nominal waveguide model; full per-seed maps, pressure errors and loss histories are also available in the project [GitHub repository](#), under `inverse_PINN/results`.

Even under the limited campaign considered here, the partial recovery of defect locations and contrast signs is encouraging. With more time, a broader controlled study could test additional

training and acquisition configurations, including higher frequencies, more incidence directions, and combinations of both. The additional spectral content and illumination could improve spatial resolution and constrain the material more strongly, thereby producing more accurate reconstructions. This expectation still requires multi-seed validation: Section 7 showed that adding an incidence at a single frequency does not necessarily provide a systematic gain.

9 Discussion

This section separates three questions raised by the numerical study: what the experiments support, what was contributed during the internship, and what PINNs add to this class of inverse problems.

9.1 What the Experiments Establish

The experiments do not identify a universal PINN recipe. They instead establish a set of necessary choices under the tested computational budget. The forward study shows that loss balancing and Fourier features are not secondary refinements. Unit weights can make a suitable network appear unable to solve the problem, and the classical MLP fails at the larger wavenumber even after the loss components have been balanced. The modified Fourier network is the most accurate forward model tested, but its additional cost and lower inverse accuracy show that a forward ranking cannot be transferred directly to inversion. Analytic propagation of the spatial derivatives also proved important in practice because it made the repeated training campaigns substantially more affordable.

For the inverse problem, acquisition content, material priors and the training path have at least as much influence as the field architecture. Adding a second incidence at one frequency does not produce a systematic gain, whereas the additional frequency is more useful in the tested cases. A plausible physical interpretation is that the additional frequency changes both the wavenumber and the propagating modal content, thereby providing information that can break some fixed-frequency ambiguities. Reversing the incidence at the same frequency changes the illumination but not this spectral content. This interpretation is consistent with the finite-frequency non-uniqueness mechanism discussed in Section 4.4 and in [14], although the present experiments do not isolate it. Restricting the admissible sound-speed range removes many background artefacts, but this improvement relies on valid prior knowledge. The best reconstruction is obtained through a continuation strategy that combines this prior with several stages and frequencies. Its result establishes the feasibility of the complete procedure, not the isolated value of any one ingredient.

The study also exposes an important limitation of the all-at-once formulation. A network can closely fit the measured boundary traces while retaining an appreciable error in the interior wavefield and material. In a representative median baseline run, the normalised trace loss reaches the order of 10^{-7} while the acquisition-averaged interior pressure error remains about 20%. This quantitative gap provides direct evidence that neither a small data loss nor a small aggregate PINN objective certifies a correct reconstruction. The gap is particularly relevant for sparse boundary observations and a soft PDE constraint, where errors in the learned field can be absorbed into the material estimate. Independent field errors and material metrics remain necessary whenever a synthetic reference is available.

One methodological choice also remains unablated: the material gradient is computed only from the PDE residual, together with the optional TV term, and not directly from the data loss. Moreover, no differentiation is performed through the field update, so the measurements influence the material only through the pressure iterates passed from one step to the next. This block-coordinate-style design is legitimate, but it determines what drives m_ϕ and should be compared with a fully coupled joint update.

The scattered-field formulation provides a useful simplification rather than an accuracy breakthrough. Removing the known incident field eliminates the homogeneous warm-up and reduces

the measured training cost, but the paired circular-defect tests show material errors of the same order as the total-field formulation. More broadly, the observed seed-to-seed dispersion confirms that isolated runs can be misleading. Controlled multi-seed comparisons are therefore part of the method, rather than an optional statistical refinement.

Taken together, these results support the use of PINNs as a feasible experimental framework for the present waveguide problem. They do not yet demonstrate robust inversion under noise, model mismatch or unseen media. The main scientific outcome is thus a clearer hierarchy of difficulties: representation and loss scaling must be resolved first, after which acquisition design, prior information and optimisation variability become the dominant limitations. A matched FWI baseline remains necessary to assess the relative accuracy and computational cost of the PINN approach.

9.2 Methodological and Technical Contribution

The rapidly expanding PINN literature offers a very large design space of architectures, sampling rules, loss-balancing methods and optimisation schedules. An exhaustive comparison was neither possible nor useful within the internship. The work instead followed a progressive decision path: formulate the waveguide problem, generate controlled data, calibrate the forward solver, activate the inverse coupling, and then change one identifiable part of the pipeline whenever possible. The main outcome is therefore not a universally optimal configuration, but a coherent and reproducible framework in which new choices can be evaluated without restarting the study from first principles.

The first contribution is the construction of the numerical benchmark. The data pipeline generates synthetic FEM traces for prescribed frequencies, modes, incidence directions and material contrasts. Homogeneous analytical solutions and full FEM fields provide independent references for the forward tests, while the inverse study uses metrics that separate the defect error from background artefacts. This benchmark connects the mathematical waveguide formulation to repeatable learning experiments and can be extended to new acquisition configurations.

The second contribution is a from-scratch JAX implementation of the complete PINN pipeline. It supports complex frequency-domain wavefields, modal Neumann and DtN conditions, independent field networks for each acquisition, and a material network shared across all acquisitions. Total- and scattered-field formulations are treated within the same code base. The implementation also includes Fourier embeddings, the gated modified MLP, analytic spatial derivatives, adaptive strategies and consistent physical normalisation of every loss component.

The third contribution is an inverse-specific training and evaluation methodology. It combines homogeneous field pretraining, a progressive data coupling, separate field and material updates, admissible-range parameterisation, frequency continuation and fixed monitoring clouds. Run configurations, histories and per-seed results are retained systematically. This workflow emerged from the difficulties encountered during the internship and is more reliable than interpreting a single favourable or unfavourable trajectory.

Many canonical inverse-PINN examples estimate only a few global coefficients [2]. The FWI and PINN-WRI studies most relevant here already recover spatial velocity fields, but their reference implementations rely on conventional fully connected networks [19, 20]. Within the literature reviewed for this project, the present work goes further by testing Fourier embeddings and a gated architecture in a frequency-domain, WRI-like inversion, while also examining the effect of physics-loss weights. The claim is not that these are the first learned Fourier representations

used anywhere in FWI. Rather, their controlled evaluation within a pure all-at-once PINN pipeline is an identifiable extension of the classical PINN baselines.

The problem itself adds several difficulties beyond low-dimensional parameter identification: the unknown is a two-dimensional field with localised sharp contrasts, the pressure is complex and oscillatory, and several modes and acquisitions share the same material through non-local boundary conditions. Documenting which apparently promising techniques fail to improve this setting is also part of the contribution. In particular, the experiments prevent forward accuracy, adaptive weighting or a single low-loss run from being presented as evidence of reliable inversion.

9.3 Internship Timeline

The ten-week schedule is summarised in Table 5. The first half of the internship was devoted to learning, implementation and problem formulation. The second half progressively shifted towards controlled experiments, mathematical consolidation and the construction of the numerical studies reported in this document.

Period	Main activities
Week 1	Initial literature review, familiarisation with JAX, and implementation of a first PINN for the heat equation.
Week 2	Broader review of the PINN literature, with particular attention to recent architectures, training methods and known failure modes.
Week 3	First adaptation to the inverse waveguide problem, modification of the C++ data-export pipeline, and development of the initial experimental notebooks.
Week 4	Study of the mathematical setting and classical inversion methods, together with the construction of test sound-speed fields.
Week 5	Diagnosis of the first inversion failures, including inappropriate physical choices, architectural limitations and optimisation difficulties.
Weeks 6–7	First encouraging results obtained by testing alternative field formulations, architectures and training strategies.
Week 8	Consolidation of the mathematical formulation and proofs, critical review of the experimental direction, and beginning of the report.
Week 9	Construction of the controlled forward study presented in Section 6.
Week 10	Construction of the controlled inverse study presented in Sections 7 and 8.

Table 5: Internship timeline and main activities.

9.4 What PINNs Add and Where the Field Is Moving

The main appeal of PINNs in this problem is their common differentiable representation of fields, material and physics. A coordinate network provides a continuous wavefield that can be evaluated away from the training points, while automatic or analytic derivatives insert the governing equation directly into the objective. The same framework combines sparse measurements, boundary conditions, multiple acquisitions and prior constraints without deriving a separate discrete adjoint for each new network parameterisation. It also makes field-material interactions and conflicts between loss gradients directly observable.

The experiments nevertheless show that flexibility is not the same as robustness. The neural representation introduces spectral bias and smooths sharp interfaces; the soft physics constraint leaves room for inaccurate interior fields; and the multi-objective optimisation is sensitive to

weights, initialisation and training schedule. In the present two-dimensional setting, a pure PINN is therefore most valuable as a flexible research framework. The results do not establish that it is more accurate or computationally preferable to a mature numerical forward solver.

Recent FWI research increasingly reflects this distinction by combining neural and classical components. Herrmann et al. [26] retain a neural representation for the unknown material but recover the efficiency of a conventional forward solver and continuous-adjoint gradient. More recently, Chen et al. [28] introduced a physics-guided FWI method in which a learned Fourier-domain representation of the seismic data is embedded inside a physics-guided optimisation loop. These approaches use neural networks where their parameterisation or learned representation is useful, without requiring one PINN to replace every numerical component.

10 Conclusion and Perspectives

This internship investigated how physics-informed neural networks can be used to solve an inverse acoustic problem in a two-dimensional waveguide. Starting from the mathematical model, I developed a complete workflow that connects synthetic finite-element data, a differentiable PINN implementation and an evaluation protocol adapted to the reconstruction of a spatially varying sound speed. The forward experiments established the importance of loss scaling, Fourier features and efficient spatial derivatives. The inverse experiments then showed that a material field can be reconstructed within this framework, while also exposing the strong influence of the acquisition, the prior information, the training schedule and the random initialisation. The best reconstruction requires about 3.6 times the measured optimiser time of baseline campaign A, and no equal-wall-clock comparison was performed.

The main conclusion is therefore not that PINNs provide a ready-made replacement for classical inversion methods. Their value lies in the flexibility with which measurements, governing equations and prior constraints can be combined in a single differentiable model. This flexibility comes with a demanding multi-objective optimisation problem: a low training loss or a good fit to the measurements alone does not guarantee an accurate reconstruction. The most important outcome of the project is consequently the construction of a coherent and reproducible framework in which these choices can now be tested and compared reliably.

Beyond these technical results, the internship was my first sustained experience of the research process. Confronting a broad and rapidly evolving literature taught me how to delimit a question, formulate testable hypotheses and distinguish a promising result from one that is merely favourable. Iterations, unsuccessful experiments and discussions were not peripheral to the work: they progressively shaped both the methodology and my understanding of the problem. I leave the project with a more informed view of the research profession, but also with greater autonomy, critical distance and confidence when approaching an open-ended problem. The scientific exchanges and the support received throughout the internship also made it a particularly valuable human and personal experience.

The highest-priority methodological extension is to study the representation of the unknown material, which was explored less extensively than the wavefield architecture even though it represents the quantity of interest. The current two-layer, width-64 tanh MLP could be compared with convolutional or multiscale models that introduce a spatial bias better suited to localised structures and sharp interfaces. With a stronger material representation in place, a broader inverse study could then consider additional frequencies, incidence directions, defect geometries and noise levels, with systematic multi-seed comparisons. A comparison with a classical FWI baseline, using the same data, evaluation metrics and computational budget, would also be essential.

11 Technical Glossary

This glossary collects the specialised terms that are central to the formulation and interpretation of the study. Terms are ordered alphabetically.

Term	Definition
All-at-once formulation	An inverse formulation in which the wavefields and the material parameters are treated as simultaneous unknowns and linked through a penalised physical residual.
Cycle skipping	A failure mode of waveform inversion in which predicted and observed oscillations are incorrectly matched, leading the optimisation towards a local minimum.
DtN operator	Dirichlet-to-Neumann operator mapping a field trace on an artificial boundary to its outgoing normal derivative.
Fourier features	A sinusoidal coordinate embedding used to reduce the spectral bias of neural networks and improve the representation of oscillatory fields.
FWI	Full Waveform Inversion, a PDE-constrained method that reconstructs material parameters by minimising the mismatch between predicted and observed waveforms.
NTK	Neural Tangent Kernel, a theoretical framework used to analyse the training dynamics of infinitely wide neural networks.
PINN	Physics-Informed Neural Network, trained by incorporating governing equations and boundary conditions into its loss function.
RAD	Residual-based Adaptive Distribution, a sampling strategy that concentrates collocation points in regions with large physical residuals.
Scattered-field formulation	A reformulation in which the known background field is removed and the network represents only the field generated by the material contrast.
Spectral bias	The tendency of a neural network to learn low-frequency components before higher-frequency variations.
Squared slowness	The material parameter $m = 1/c^2$, equal to the inverse square of the sound speed.
TV	Total variation, a regularisation term favouring piecewise-smooth or piecewise-constant material reconstructions.
WRI	Wavefield Reconstruction Inversion, a relaxed formulation that jointly optimises the material and wavefields while penalising the wave-equation residual.

A Proofs

A.1 Weak Formulation for the Discontinuous Contrast

The corresponding total-field DtN formulation is established in [9, Section 3.5, equations (3.13)–(3.14), pp. 50–51]. We record only the normal-trace argument needed for the discontinuous contrast and the scattered-field source.

Proposition A.1 (From strong to weak formulation). *Let $u^s \in H^1(\Omega)$ satisfy the equation in (3.5) in the sense of distributions and its boundary conditions in the sense of normal traces. Then u^s satisfies the weak formulation (3.6).*

Proof. Since $m, m_s \in L^\infty(\Omega)$ and $u^s, u_0 \in L^2(\Omega)$, the equation implies

$$\operatorname{div}(\nabla u^s) = \Delta u^s = -\omega^2(mu^s + m_s u_0) \in L^2(\Omega),$$

and hence $\nabla u^s \in H(\operatorname{div}, \Omega)$. The generalised Green formula [29, Theorem 1.32] therefore applies on the whole domain. This global $H(\operatorname{div})$ regularity already encodes the absence of a normal-flux jump across the internal interface [29, Theorem 1.36].

Substituting the equation in Green's formula, using $\partial_\nu u^s = 0$ on Γ_N and $\partial_\nu u^s = T_\omega u^s$ on $\Gamma^\pm$, and rearranging gives, for every $v \in H^1(\Omega)$,

$$\int_\Omega \nabla u^s \cdot \overline{\nabla v} \, d\Omega - \omega^2 \int_\Omega m u^s \bar{v} \, d\Omega - \langle T_\omega u^s, \bar{v} \rangle_{\Gamma^\pm} = \omega^2 \int_\Omega m_s u_0 \bar{v} \, d\Omega,$$

which is (3.6). $\square$

A.2 Continuity of the Forward Operator

Proof of Proposition 4.10. Since γ is continuous, it suffices to show that G_{u_0} is continuous. Fix $m_s^2 \in \mathcal{A}$ and consider $m_s^1 \in \mathcal{A}$. Denote by $u_1^s, u_2^s \in H^1(\Omega)$ the corresponding solutions and by a_k and ℓ_k the forms (3.7)–(3.8) associated with m_s^k , for $k = 1, 2$. Set

$$\delta m = m_s^1 - m_s^2, \quad \delta u = u_1^s - u_2^s.$$

For all $u, v \in V$,

$$a_1(u, v) - a_2(u, v) = -\omega^2 \int_\Omega \delta m u \bar{v} \, d\Omega.$$

Using $a_k(u_k^s, v) = \ell_k(v)$, we obtain

$$a_2(\delta u, v) = \omega^2 \int_\Omega \delta m (u_0 + u_1^s) \bar{v} \, d\Omega \quad \forall v \in V.$$

Let $C_2 > 0$ be the stability constant of theorem 4.6 associated with the fixed contrast m_s^2 . Then

$$\|\delta u\|_V \leq q(\|u_0 + u_2^s\|_{L^2(\Omega)} + \|\delta u\|_V), \quad q = C_2 \omega^2 \|\delta m\|_{L^\infty(\Omega)}.$$

For $q < 1$, absorption gives

$$\|\delta u\|_V \leq \frac{q}{1-q} \|u_0 + u_2^s\|_{L^2(\Omega)}.$$

Thus $u_1^s \rightarrow u_2^s$ in $H^1(\Omega)$ whenever $m_s^1 \rightarrow m_s^2$ in $L^\infty(\Omega)$. Since m_s^2 was arbitrary, G_{u_0} , and therefore F_{u_0} , is continuous on $\mathcal{A}$. $\square$

B Forward-Study Reference Data

The reference fields described here are used only to evaluate the forward PINNs in Section 6. They do not enter the physics loss, the optimiser, or checkpoint selection, and no measurement noise is added in this controlled study.

B.1 Homogeneous Analytic Reference

For the homogeneous waveguide, the total field is the exact incident mode

$$u_0(x, y) = \theta_p(y) \exp(-is\beta_p x), \quad \beta_p = \sqrt{k_0^2 - (p\pi/H)^2}, \quad (\text{B.1})$$

using the temporal convention $\exp(-i\omega t)$ introduced in Section 2. Here s is the boundary-incidence flag used by the code: $s = -1$ denotes incidence from the left and gives $\exp(+i\beta_p x)$, which propagates in the positive x -direction, whereas $s = +1$ denotes incidence from the right and gives $\exp(-i\beta_p x)$. The experiments use $p = 0$ and $s = -1$. Error norms are evaluated on an independent structured triangulation with 256×80 rectangular cells, each split into two linear triangles. Its assembled P_1 mass and stiffness matrices provide the discrete L^2 and H^1 products in (C.5).

B.2 Heterogeneous Finite-Element Reference

The defect reference solves the same Helmholtz–DtN problem as the PINN on $(-1, 1) \times (0, 0.6)$ at 1200 Hz. The circular defect has centre $(0.2, 0.2)$, radius 0.1, and wave speed $c_d = 0.8c_0$, with $c_0 = 340 \text{ m s}^{-1}$. The corresponding background wavenumber is $k_0 = 22.1759 \text{ m}^{-1}$.

The reference solver uses conforming quadratic (P_2) finite elements on a mesh conforming to the circular interface and the same modal DtN boundary treatment. The export contains 89 372 complex field degrees of freedom on 44 407 quadratic triangles. The geometric P_2 mass and stiffness matrices are exported in the same degree-of-freedom ordering and are used directly for the relative errors. This makes the comparison sensitive to the whole two-dimensional field rather than only to a set of boundary sensors.

The finest FEM solution is a numerical reference rather than an exact solution. Relative to the preceding 44 209-degree-of-freedom mesh, the field changes by approximately 0.217% in the discrete L^2 norm and 2.60% in the discrete H^1 norm.

C Forward PINN Methods and Reproducibility

This appendix records the detailed implementation and reproducibility settings used in Section 6. The general PINN workflow and the loss definitions are introduced in Section 5; general MLP definitions and the modified-MLP equations are given in Appendix D. The source code, experiment configurations, and supporting notebooks are available in the project [GitHub repository](#).

C.1 Field Networks and Fourier Features

The coordinate normalisation and field-network outputs are defined in Equation (5.1). The classical baseline applies a three-layer, width-128, tanh MLP directly to $(\hat{x}, \hat{y})$.

For the Fourier variants, the coordinates are first mapped back to physical units and embedded as [27]

$$\gamma(x, y) = \left[\cos\left(B \begin{bmatrix} x \\ y \end{bmatrix}\right), \sin\left(B \begin{bmatrix} x \\ y \end{bmatrix}\right) \right], \quad B = B_0 \text{diag}(\sigma_x, \sigma_y), \quad (\text{C.1})$$

where $B_0 \in \mathbb{R}^{64 \times 2}$ has independent standard-normal entries. The resulting input dimension is 128. Only the two scales $\boldsymbol{\sigma} = (\sigma_x, \sigma_y)$ are trainable; the sampled matrix B_0 remains fixed. For incident mode p they are initialised by

$$\sigma_x^{(0)} = \sqrt{k_0^2 - (p\pi/H)^2}, \quad \sigma_y^{(0)} = (p + \mathbf{1}_{p=0})\pi/H. \quad (\text{C.2})$$

The standard Fourier MLP and the gated modified MLP [17] have 49 796 and 82 820 trainable parameters, respectively, compared with 33 666 for the classical MLP.

The field parameters use Adam with a constant learning rate of 10^{-3} . The Fourier scales use a separate learning rate, cosine-decayed from 10^{-2} to 10^{-5} over the first 15 000 iterations, after which the scales are frozen. All affine layers use Xavier-normal initialisation and zero biases; the final Fourier-network weight matrix is additionally divided by 10 at initialisation. No L-BFGS phase or gradient clipping is used. Physical first derivatives and Laplacians are propagated analytically through the Fourier networks; the classical baseline uses automatic differentiation. Consequently, timing comparisons include both architecture and derivative-implementation effects.

C.2 Adaptive Loss Weights

The three unweighted physics losses are defined in Equations (5.6)–(5.8). The unit-weight experiment uses $(1, 1, 1)$; all static-weight experiments use $(10^3, 10^4, 10)$ in PDE/Neumann/DtN order. This latter triplet was chosen heuristically, guided by an unweighted component-gradient diagnostic, to bring the three contributions to the parameter update to comparable scales.

The adaptive variant uses an inverse-gradient-norm heuristic motivated by gradient balancing methods for PINNs [17]. The weights are initialised from the values below; after the first 500 iterations, the raw gradient norm of component i is evaluated every 500 iterations on the current training batch and an exponential moving average is updated:

$$\lambda_i^{(t)} = (1 - \alpha)\lambda_i^{(t-1)} + \alpha \|\nabla_{\theta} \mathcal{L}_i\|_2, \quad w_i^{(t)} = \frac{1}{\varepsilon + \lambda_i^{(t)}}, \quad \alpha = 0.1. \quad (\text{C.3})$$

The weights are held constant between updates and are neither clipped nor renormalised. Only currently trainable parameters contribute to the norm, so the Fourier scales are excluded after they are frozen. For the homogeneous campaign, $(\lambda_{\text{PDE}}^{(0)}, \lambda_{\text{N}}^{(0)}, \lambda_{\text{DtN}}^{(0)}) = (10^{-3}, 10^{-4}, 10^{-1})$ and $\varepsilon = 10^{-5}$. For the defect campaign the corresponding values are $(10^{-2}, 10^{-3}, 1)$ and $\varepsilon = 10^{-4}$. These initialisations reproduce approximately the static weight ratios while allowing their absolute and relative values to evolve.

C.3 Residual-Based Adaptive Distribution

The tested RAD strategy follows the residual-weighted principle of Wu et al. [30]. A pool of 100 000 uniformly distributed candidates is assigned probabilities

$$p_j \propto \frac{|r(\mathbf{x}_j)|^k}{\frac{1}{N_c} \sum_{\ell=1}^{N_c} |r(\mathbf{x}_\ell)|^k} + c. \quad (\text{C.4})$$

For the defect experiment, $k = 1$ and $c = 0$. The distribution is refreshed every 500 iterations. Each subsequent PDE batch contains 3072 newly sampled uniform points and 1024 points drawn with replacement from this cached distribution, so the total interior budget remains 4096. The comparison is paired by seed with the non-RAD baseline.

C.4 Validation Metrics and Reproducibility

Training batches are resampled independently at every Adam iteration. Loss histories are evaluated on a fixed regular monitoring set; gradient diagnostics use the current training cloud, and reference errors are evaluated every 1000 iterations. If M and K denote the validation mass and stiffness matrices, the reported complex-field errors are

$$e_{L^2} = \sqrt{\frac{\mathbf{e}^* M \mathbf{e}}{\mathbf{u}^* M \mathbf{u}}}, \quad e_{H^1} = \sqrt{\frac{\mathbf{e}^* (M + K) \mathbf{e}}{\mathbf{u}^* (M + K) \mathbf{u}}}, \quad \mathbf{e} = \mathbf{u}_{\text{PINN}} - \mathbf{u}_{\text{ref}}. \quad (\text{C.5})$$

For scattered-field runs, the known incident field is added before these total-field metrics are computed. Reported training times include optimiser steps and, when enabled, RAD refreshes and adaptive-weight updates; JIT compilation and all monitoring or reference evaluations are excluded.

Quantity	Value
Implementation	Python 3.12.3; JAX 0.10.2; Optax 0.2.8; single-precision training
Accelerator	One CUDA GPU; model, driver, and CUDA version not retained
Domain	$(-1, 1) \times (0, 0.6)$
Background speed	$c_0 = 340 \text{ m s}^{-1}$
Frequencies	600 Hz and 1200 Hz
Incident field	Mode 0, left-to-right propagation
Hidden layers	3×128 , tanh, linear two-component output
Fourier features	64 frequencies (128 sine/cosine features)
Adam iterations	30 000
Collocation per step	4096 PDE; 256 coordinates on each wall/port
Static weights	$(10^3, 10^4, 10)$ in PDE/Neumann/DtN order
Random seeds	0–2; seeds 0–4 for the homogeneous 1200 Hz case

Table 6: Main forward-study implementation parameters. Complete run configurations and per-seed histories are retained with the campaign outputs.

D Multilayer Perceptron: Architecture and Training

Architecture

A feedforward neural network, or multilayer perceptron (MLP), is a parametric function $\mathcal{N}_\theta : \mathbb{R}^{n_0} \rightarrow \mathbb{R}^{n_L}$ defined by the recursive composition

$$\mathbf{z}^{(0)} = \mathbf{x}, \quad \mathbf{z}^{(l)} = \sigma(W^{(l)}\mathbf{z}^{(l-1)} + \mathbf{b}^{(l)}), \quad l = 1, \dots, L-1, \quad \mathcal{N}_\theta(\mathbf{x}) = W^{(L)}\mathbf{z}^{(L-1)} + \mathbf{b}^{(L)}, \quad (\text{D.1})$$

where $W^{(l)} \in \mathbb{R}^{n_l \times n_{l-1}}$ and $\mathbf{b}^{(l)} \in \mathbb{R}^{n_l}$ are the weight matrix and bias vector of layer l , n_l denotes the width of layer l , and $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is the activation function applied element-wise. The set of all trainable parameters is collected in $\theta = \{W^{(l)}, \mathbf{b}^{(l)}\}_{l=1}^L$. The tanh activation used here is smooth enough for the second-order strong Helmholtz residual. The field networks have linear outputs for the real and imaginary pressure components. The material network is the exception: its scalar output is passed through a sigmoid and an affine map to enforce the squared-slowness interval induced by the prescribed sound-speed bounds. The task-specific network interactions are shown in Figure 3.

Initialisation

The choice of initial parameters θ_0 has a significant effect on the optimisation trajectory. A naive initialisation can cause the tanh activations to saturate immediately, where gradients vanish. The *Xavier* initialisation [31] addresses this by drawing the entries of $W^{(l)}$ from a normal distribution scaled to preserve activation variance across layers:

$$W_{ij}^{(l)} \sim \mathcal{N}\left(0, \frac{2}{n_{l-1} + n_l}\right).$$

This scaling keeps activations in the linear regime of tanh at the start of training, where gradients are largest. The biases are initialised to zero. All networks in this work use this scheme.

Input and Output Normalisation

Stable training requires that both the inputs and the outputs be of comparable magnitude throughout the network. The input vector $\mathbf{x} \in \mathbb{R}^{n_0}$ is mapped to a normalised vector by an affine transformation calibrated from the geometry of the problem, and the outputs are similarly rescaled to a reference magnitude. Without this normalisation, components operating at very different scales cause the optimiser to make overly conservative updates in some directions while taking oversized steps in others, ultimately slowing or destabilising convergence.

Depth and the Vanishing Gradient Problem

Increasing the depth L expands the representable function class but can also make optimisation harder because products of layer Jacobians attenuate or amplify the gradient. Residual and gated architectures such as PirateNets were designed to improve this gradient flow [4]. Their benefit depends on the PDE, depth and training budget; the standard and modified architectures are therefore compared empirically rather than selected from depth alone.

Modified MLP

The modified MLP architecture, introduced by Wang et al. [17], augments the standard feedforward network with two encoder sub-networks U and V and a multiplicative gating mechanism. Given the input X , the forward pass is defined by

$$U = \sigma(XW^1 + b^1), \quad (\text{D.2})$$

$$V = \sigma(XW^2 + b^2), \quad (\text{D.3})$$

$$H^{(1)} = \sigma(XW^{z,1} + b^{z,1}), \quad (\text{D.4})$$

$$Z^{(k)} = \sigma(H^{(k)}W^{z,k} + b^{z,k}), \quad k = 1, \dots, L, \quad (\text{D.5})$$

$$H^{(k+1)} = (1 - Z^{(k)}) \odot U + Z^{(k)} \odot V, \quad k = 1, \dots, L, \quad (\text{D.6})$$

$$f_\theta(x) = H^{(L+1)}W + b, \quad (\text{D.7})$$

where $\odot$ denotes element-wise (Hadamard) multiplication. The encoder outputs U and V are injected into every hidden layer through a residual gate controlled by $Z^{(k)}$, which alleviates the vanishing-gradient problem and enables the training of deeper networks.




Figure 12: Improved fully-connected architecture (Modified MLP) for physics-informed neural networks [17]. The encoder sub-networks U and V inject residual information into every hidden layer via a multiplicative gate, which mitigates the vanishing-gradient problem and accounts for multiplicative interactions between the inputs.

References

- [1] Maziar Raissi, Paris Perdikaris, and George E. Karniadakis. Physics informed deep learning (part i): Data-driven solutions of nonlinear partial differential equations. *arXiv preprint arXiv:1711.10561*, 2017.
- [2] Maziar Raissi, Paris Perdikaris, and George E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
- [3] Wei Zhou and Y. F. Xu. Data-guided physics-informed neural networks for solving inverse problems in partial differential equations. *arXiv preprint arXiv:2407.10836*, 2024.
- [4] Sifan Wang, Bowen Li, Yuhan Chen, and Paris Perdikaris. PirateNets: Physics-informed deep learning with residual adaptive networks. *Journal of Machine Learning Research*, 25(402):1–51, 2024.
- [5] Zhenyu Li. A review of physics-informed neural networks. *Applied and Computational Engineering*, 133:164–172, 2025.
- [6] Andrea Bonfanti, Ismael Medina, Roman List, Björn Staeves, Roberto Santana, and Marco Ellero. Pinn balls: Scaling second-order methods for pinns with domain decomposition and adaptive sampling. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025.
- [7] Victorita Dolean, Daria Hrebenshchykova, Stephane Lanteri, and Victor Michel-Dansac. Neural network-driven domain decomposition for efficient solutions to the helmholtz equation. *arXiv preprint arXiv:2511.15445*, 2025.
- [8] Nickson Golooba and Woldegebriel Assefa Woldegerima. Resolving gradient pathology in physics-informed epidemiological models. *arXiv preprint arXiv:2603.23799*, 2026.
- [9] Anne-Sophie Bonnet-Ben Dhia and Éric Lunéville. Propagation dans les guides d’ondes. Course notes, ENSTA Paris, 2021. HAL Id: hal-04117183. <https://hal.science/hal-04117183v1>.
- [10] L. Bourgeois and E. Lunéville. The linear sampling method in a waveguide: a modal formulation. *Inverse Problems*, 24(1):015018, 2008.
- [11] Jean Virieux and Stéphane Operto. An overview of full-waveform inversion in exploration geophysics. *Geophysics*, 74(6):WCC1–WCC26, 2009.
- [12] Tristan van Leeuwen and Felix J. Herrmann. Mitigating local minima in full-waveform inversion by expanding the search space. *Geophysical Journal International*, 195(1):661–667, 2013.
- [13] Tamara G. Grossmann, Urszula Julia Komorowska, Jonas Latz, and Carola-Bibiane Schönlieb. Can physics-informed neural networks beat the finite element method? *IMA Journal of Applied Mathematics*, 89(1):143–174, 2024.
- [14] Antoine Bera, Anne-Sophie Bonnet-Ben Dhia, and Lucas Chesnel. A continuation method for building invisible obstacles in waveguides. *The Quarterly Journal of Mechanics and Applied Mathematics*, 74(1):83–116, 2021.

- [15] Sifan Wang, Hanwen Wang, and Paris Perdikaris. On the eigenvector bias of fourier feature networks: From regression to solving multi-scale pdes with physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 384:113938, 2021.
- [16] Vincent Sitzmann, Julien N. P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems*, 2020.
- [17] Sifan Wang, Yujun Teng, and Paris Perdikaris. Understanding and mitigating gradient flow pathologies in physics-informed neural networks. *SIAM Journal on Scientific Computing*, 43(5):A3055–A3081, 2021.
- [18] Levi D. McClenney and Ulisses M. Braga-Neto. Self-adaptive physics-informed neural networks. *Journal of Computational Physics*, 474:111722, 2023.
- [19] M. Rasht-Behesht, C. Huber, K. Shukla, and G. E. Karniadakis. Physics-informed neural networks (pinns) for wave propagation and full waveform inversions. *Journal of Geophysical Research: Solid Earth*, 127(5), 2022.
- [20] Chao Song and Tariq A. Alkhalifah. Wavefield reconstruction inversion via physics-informed neural networks. *IEEE Transactions on Geoscience and Remote Sensing*, 60:1–12, 2022.
- [21] Xilun Liu and Mohamed Almekkawy. Ultrasound computed tomography using physical-informed neural network. In *2021 IEEE International Ultrasonics Symposium (IUS)*, pages 1–4, 2021.
- [22] Cristian Assumma. Physics-informed neural networks (pinns) for sound speed estimation from multi-view ultrasound data. Master’s thesis, Politecnico di Torino, 2026.
- [23] Ehsan Haghighat, Maziar Raissi, Adrian Moure, Hector Gomez, and Ruben Juanes. A physics-informed deep learning framework for inversion and surrogate modeling in solid mechanics. *Computer Methods in Applied Mechanics and Engineering*, 2021.
- [24] Afla Ajithkumar Sophiya, Akarsh K Nair, Sepehr Maleki, and Senthil K. Krishnababu. A comprehensive analysis of pinns: Variants, applications, and challenges. *arXiv:2505.22761*, 2025.
- [25] Sifan Wang, Xinling Yu, and Paris Perdikaris. When and why PINNs fail to train: A neural tangent kernel perspective. *Journal of Computational Physics*, 449:110768, 2022.
- [26] Leon Herrmann, Tim Burchner, Felix Dietrich, and Stefan Kollmannsberger. On the use of neural networks for full waveform inversion. *Computer Methods in Applied Mechanics and Engineering*, 415:116278, 2023.
- [27] Matthew Tancik, Pratul Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems*, 33:7537–7547, 2020.
- [28] Gui Chen, Yang Liu, Haoran Zhang, and Mi Zhang. Seismic full waveform inversion via a physics-guided Fourier representation neural network. arXiv preprint arXiv:2606.30126, 2026.
- [29] Patrick Ciarlet. Ms03 – notes de cours, 2024.

- [30] Chenxi Wu, Min Zhu, Qinyang Tan, Yadhu Kartha, and Lu Lu. A comprehensive study of non-adaptive and residual-based adaptive sampling for physics-informed neural networks. *Computer Methods in Applied Mechanics and Engineering*, 403:115671, 2023.
- [31] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, 2010.