

People's Democratic Republic of Algeria
Ministry of Higher Education and Scientific Research

Kasdi Merbah University – Ouargla

Faculty of Mathematics and Matter Sciences
Department of Physics



Master's Thesis

Submitted in partial fulfilment of the requirements
for the degree of **Master in Medical Physics**

**Physics-Informed Neural Networks for Light Diffusion
and Thermal Therapy in Biological Tissues**

Presented by:

BENZINE SANA and GHARBI IMANE

Defended on June 11, 2026

Examination Committee

Dr. Gadjia Mohammed Elarebi	President	Kasdi Merbah University – Ouargla
Dr. Ayat Zahia	Examiner	Kasdi Merbah University – Ouargla
Prof. Bentouila Omar	Supervisor	Kasdi Merbah University – Ouargla
Dr. Zenkhri Djamel Eddine	Co-supervisor	Kasdi Merbah University – Ouargla

Academic Year 2025–2026

Dedication

I dedicate this graduation work to the greatest source of my strength and success in life: my beloved mother and father, who have been my support in every step of my journey, never withholding their love, prayers, and encouragement.

To my dear brothers and my kind family, who have always been a source of happiness and support, and who stood by me through all circumstances.

To my respected professors and dear colleagues, who accompanied me throughout my academic journey and had a positive impact on my academic and personal development.

To everyone who contributed in one way or another to my success, I dedicate the fruit of this humble work.

Sana

Dedication

To the pure soul hovering in the heaven of my heart, to the one I deeply wished could be here to witness this milestone with me — my dear father, may Allah have mercy on him and grant him the highest dwelling in Paradise. I dedicate this success to you, the one who instilled in me the love of knowledge and perseverance.

To the fountain of giving and tenderness, my beloved mother, the pillar of my journey and my companion who supported me step by step with her prayers and endless patience.

To my brothers, sisters, and my honorable family, each by their name, who have always been my ultimate support and a source of optimism.

To my esteemed professors, who illuminated my path of knowledge throughout my academic journey and never hesitated to offer guidance and advice.

To my friends and companions, who shared with me the moments of hard work and fatigue.

To all of you, I dedicate this humble work as a fruit of effort and gratitude.

Imane

Acknowledgements

First and foremost, we thank **Almighty Allah** for His blessings, guidance, and strength that enabled us to bring this work to completion.

We would like to express our deepest gratitude to our supervisor, **Prof. Bentouila Omar**, for his invaluable guidance, continuous support, and constant encouragement throughout the course of this research.

We would also like to extend our sincere thanks to **Dr. Zenkhri Djamel Eddine** for his meaningful contribution and attentive follow-up during the preparation of this work.

Our sincere gratitude is also extended to the members of the **defense committee**:

We are particularly grateful to **Dr. Gadja Mohammed Elarebi** for his willingness to chair the defense committee. His presence and oversight lent great value to this examination, and we sincerely appreciate the time and effort he dedicated to this role.

We would also like to warmly thank **Dr. Ayat Zahia** for her kind acceptance to examine and assess this thesis. Her careful reading of this work and the constructive remarks she provided were greatly appreciated and contributed meaningfully to the quality of the final outcome.

Finally, we wish to express our heartfelt thanks to all our professors, colleagues, and everyone who supported us throughout our academic journey.

Abstract

This thesis investigates **Physics-Informed Neural Networks (PINNs)** as a unified framework for the coupled optical–thermal modelling of laser–tissue interaction: the propagation of light in biological tissue and the laser-induced temperature field underlying hyperthermic therapy. Five progressive studies share a single code base. The forward photon-diffusion problem is solved in homogeneous tissue (relative L^2 error 5.5×10^{-5}) and in a three-layer skin–fat–muscle geometry, where the analytical interface-correction term $D'(z) \partial\phi/\partial z$ proves decisive, reducing the error from $\sim 17\%$ to 1.8×10^{-4} . The inverse problem recovers the optical coefficients from noisy fluence data with sub-percent error in the single-layer case, but exposes an *identifiability gap* in the multi-layer case, where the smallest coefficient of each layer is the worst recovered (43% error for μ_a^{skin}). Finally, the diffusion solution drives a time-dependent Pennes bioheat model ($L^2 = 1.04 \times 10^{-1}$).

Beyond accuracy, the simulations expose two physical phenomena: a **Q -balance effect**, by which the absorbed power coincides in skin and tumour and the thermal selectivity stays near unity; and a **self-shielding limit** of plasmonic enhancement, where an excessive tumour absorption collapses the optical penetration depth below the tumour thickness and shifts heating toward the skin — a parametric sweep locating a modest optimum near $\mu_a^{\text{tumour}} \approx 2 \text{ cm}^{-1}$. The work shows that, even on a laptop-class budget, PINNs provide a viable framework for forward, inverse, and coupled optical–thermal modelling, and a tool to probe the physical limits of laser hyperthermia.

Keywords: Physics-Informed Neural Networks; Diffusion Approximation; Bioheat Equation; Hyperthermia; Inverse Problem; Photothermal Therapy; Self-Shielding; Biomedical Optics.

Résumé

Ce mémoire étudie les **réseaux de neurones informés par la physique (PINN)** comme cadre unifié pour la modélisation couplée optique–thermique de l’interaction laser–tissu. Cinq études progressives partagent un même code. Le problème direct de diffusion des photons est résolu en milieu homogène ($L^2 = 5,5 \times 10^{-5}$) puis dans une géométrie à trois couches peau–graisse–muscle, où le terme de correction d’interface $D'(z) \partial\phi/\partial z$ s’avère décisif (erreur réduite de $\sim 17\%$ à $1,8 \times 10^{-4}$). Le problème inverse récupère les coefficients optiques à partir de mesures bruitées avec une erreur inférieure à 1% dans le cas mono-couche, mais révèle un *défaut d’identifiabilité* en multi-couches (jusqu’à 43% pour les plus petits coefficients). Enfin, la solution optique alimente un modèle de Pennes transitoire ($L^2 = 1,04 \times 10^{-1}$).

Au-delà de la précision, deux phénomènes physiques émergent : un **effet d’équilibre de Q** (sélectivité thermique proche de l’unité malgré le contraste tissulaire) et une **limite d’auto-occultation** du renforcement plasmonique, avec un optimum modéré vers $\mu_a^{\text{tumeur}} \approx 2 \text{ cm}^{-1}$.

Mots-clés : PINN ; diffusion optique ; équation bioheat ; hyperthermie ; problème inverse ; thérapie photothermique ; auto-occultation ; optique biomédicale.

ملخص

يتناول هذا العمل الشبكات العصبية المُقيّدة بالفيزياء (PINN) كإطارٍ موحدٍ للنمذجة المقترنة بصرياً وحرارياً لتفاعل الليزر مع الأنسجة الحيوية: انتشار الضوء في الأنسجة وحقل درجة الحرارة الناتج عن الليزر في العلاج بفرط الحرارة. تتشارك خمس دراسات متدرّجة شيفرةً حاسوبيةً واحدة.

تُحلّ مسألة انتشار الفوتونات المباشرة في نسيج متجانس بدقة $L^2 = 5,5 \times 10^{-5}$ ، ثمّ في هندسة ثلاثية الطبقات (جلد--دهون--عضل) حيث يثبت حدّ التصحيح عند الواجهات أهميته الحاسمة، إذ يُخفّض الخطأ من 17% إلى $1,8 \times 10^{-4}$. وتستعيد المسألة العكسية المعاملات البصرية من قياسات مشوشة بخطأ دون 1% في الحالة أحادية الطبقة، لكنها تكشف قصور التمييز في الحالة متعددة الطبقات، حيث تُجَبّ المعاملات الصغيرة بضوضاء البيانات بأسوأ خطأ يبلغ 43% لمعامل امتصاص الجلد. وأخيراً يُغذّي الحلّ البصريّ نموذج ينس الحرارة الزمّني المعتمد على معاملات تروية متغيرة بالعمق.

كشفت المحاكاة عن ظاهرتين فيزيائيتين: تأثير توازن Q الذي يُبقي الانتقائية الحرارية قرب الواحد رغم التباين النسيجي، وقيد الحجب الذاتي في التعزيز البلازموني، مع تحديد قيمة مثلى متواضعة عند $\mu_a \approx 2 \text{ cm}^{-1}$.

الكلمات المفتاحية: الشبكات العصبية المُقيّدة بالفيزياء؛ تقريب الانتشار؛ معادلة ينس؛ فرط الحرارة؛ المسألة العكسية؛ العلاج الحراري الضوئي؛ البصريّات الحيوية.

Contents

Acknowledgements	iii
Abstract	iv
Nomenclature	xv
General Introduction	1
1 Physical Foundations of Laser–Tissue Interaction	4
Preamble	4
1.1 General Introduction	4
1.1.1 Importance of Optical–Tissue Interaction Studies	4
1.1.2 Medical Applications	5
1.2 Core Optical Parameters	6
1.2.1 Absorption Coefficient (μ_a)	6
1.2.2 Scattering Coefficient (μ_s)	7
1.2.3 Reduced Scattering Coefficient (μ'_s)	8
1.2.4 Refractive Index (n)	8
1.2.5 Anisotropy Factor (g)	9
1.3 Mathematical Models of Light Transport	10
1.3.1 The Radiative Transfer Equation	10
1.3.2 The Diffusion Approximation	11
1.3.3 Monte Carlo Simulation	12
1.4 Optical Properties of Selected Biological Tissues	12
1.4.1 The Biological (Therapeutic) Window	12
1.4.2 Dermis	12
1.4.3 Skeletal Muscle	13
1.4.4 Comprehensive Comparison at 800 nm	13
1.5 Medical Applications Based on Differential Optical Properties	14
1.6 Summary	15
2 Physics-Informed Neural Networks: Foundations, Advances, and Application to Medical Physics	16
2.1 What Is a Physics-Informed Neural Network?	16

2.1.1	Conceptual Definition	16
2.1.2	Structural Anatomy of a PINN	17
2.1.3	Key Technical Properties	18
2.2	Historical Development	18
2.3	Importance of PINNs in the Medical Physics Context	18
2.3.1	The Unique Challenges of Medical Physics Modelling	18
2.3.2	PINNs vs. Classical Methods: A Structured Comparison	19
2.3.3	Specific Medical Physics Applications	19
2.4	State of the Art	21
2.4.1	Advances in PINN Methodology	21
2.4.2	Benchmarks and Accuracy	21
2.4.3	Current Limitations and Open Problems	23
2.5	Justification for the Use of PINNs in This Thesis	23
2.6	Chapter Summary	24
3	Methodology and Resources	25
3.1	Physical Problem Formulation	25
3.1.1	The Photon Diffusion Equation	25
3.1.2	Tissue Optical Properties	26
3.2	Physics-Informed Neural Network Framework	27
3.2.1	Theoretical Foundation	27
3.2.2	Network Architecture	27
3.2.3	Training Strategy	27
3.3	Forward PINN: Single-Layer Tissue	28
3.3.1	Objective	28
3.3.2	Implementation	28
3.3.3	Evaluation Metrics	29
3.4	Forward PINN: Three-Layer Tissue	29
3.4.1	Objective	29
3.4.2	Conservative PDE and Smooth Coefficient Approximation	29
3.4.3	Interface-Aware Collocation	29
3.4.4	Reference Solution (Conservative FDM)	29
3.4.5	Error Analysis	30
3.5	Inverse PINN: Optical Property Recovery	30
3.5.1	Objective and Medical Motivation	30
3.5.2	Inverse Problem Formulation	30
3.5.3	Synthetic Measurement Generation	31
3.5.4	Trainable Variable Tracking	31
3.5.5	Recovery Accuracy	31
3.6	Computational Resources	31
3.7	Workflow Overview	31
3.8	Validation Strategy and Summary	31

4	Optical Results: Forward, Inverse, and Sensitivity Analysis	33
4.1	Forward PINN: Single-Layer Homogeneous Tissue	33
4.1.1	Problem Setup Recap	33
4.1.2	Fluence Rate Profile and Accuracy	34
4.1.3	Training Loss Convergence	35
4.1.4	Physical Validity Assessment	35
4.2	Forward PINN: Three-Layer Heterogeneous Tissue	36
4.2.1	Problem Setup Recap	36
4.2.2	Fluence Profile and Continuity	36
4.2.3	The Decisive Role of the Interface Correction Term	38
4.2.4	Quantitative Comparison: Single-Layer vs Three-Layer	38
4.3	Inverse Problem Setup	38
4.3.1	From Forward to Inverse	38
4.3.2	Forward Problem and Synthetic Measurements	39
4.3.3	PINN with Trainable Optical Parameters	39
4.4	Single-Layer Inverse Recovery	39
4.4.1	Convergence of the Trainable Variables	39
4.4.2	Recovery Accuracy under 2% Noise	39
4.5	Multi-Layer Inverse Recovery: Six Parameters at Once	41
4.5.1	Experimental Setup	41
4.5.2	Recovery Results	41
4.5.3	The Identifiability Gap — A Physical Diagnosis	43
4.6	Sensitivity Study	44
4.6.1	Experiment 1: Effect of the Number of Measurements	44
4.6.2	Experiment 2: Effect of the Noise Amplitude	45
4.6.3	Experiment 3: Effect of the Sampling Geometry	45
4.6.4	Summary Figure	46
4.7	Discussion	46
4.7.1	Accuracy Assessment in the Medical Physics Context	46
4.7.2	What the Inverse Experiment Demonstrates	47
4.7.3	Limits Exposed by the Multi-Layer Experiment	47
4.7.4	Practical Guidance from the Sensitivity Study	47
4.7.5	Comparison with Existing PINN Literature	47
4.8	Chapter Summary	48
5	Thermal Results: Bioheat PINN and Photothermal Therapy Exploration	49
5.1	Introduction and Clinical Motivation	50
5.2	Physical Model and Problem Parameters	50
5.2.1	Governing Equation: Pennes Bioheat Equation	50
5.2.2	Tissue Parameters	50
5.2.3	PINN Architecture and Training Configuration	51
5.3	Training Convergence Analysis	51

5.4	Temperature Field Analysis	52
5.4.1	Space-Time Temperature Maps	52
5.5	Therapeutic Efficacy and Selectivity Analysis	54
5.6	Error Analysis and Accuracy Assessment	55
5.6.1	Comparison Across the Studies	56
5.7	Clinical Discussion	56
5.8	Limitations	57
5.9	Strategy 1: Surface-Cooling Protocols	58
5.9.1	Clinical Motivation	58
5.9.2	Four Cooling Protocols	58
5.9.3	Numerical Results	58
5.9.4	Physical Interpretation	59
5.10	Strategy 2: Plasmonic Enhancement (PPTT)	60
5.10.1	Modelling AuNPs as an Effective Increase of μ_a	60
5.10.2	Counter-Intuitive Result	60
5.10.3	The Self-Shielding Mechanism	62
5.11	Strategy 3: Sweet-Spot Search	62
5.11.1	Parametric Sweep	62
5.11.2	Sweet-Spot Results	63
5.11.3	Reading the Sweet-Spot Curve	63
5.12	Discussion	64
5.13	Chapter Summary	65
General Conclusion		66
References		69
A Selected Code Listings		73
A.1	Forward PINN: Single-Layer Diffusion	73
A.2	Conservative Multi-Layer Formulation	74
A.3	Inverse PINN with Trainable Optical Variables	75
A.4	Bioheat PINN with Output Transform	75
A.5	Photothermal Exploration with FDM Solver	76
A.6	Reproducibility Notes	76

List of Figures

1.1	Schematic of the four fundamental phenomena occurring when a laser beam is incident on a biological tissue: specular reflection at the surface, absorption by chromophores, multiple scattering, and transmission.	5
2.1	Structural anatomy of a PINN. The network maps input coordinates $\mathbf{x} = (z, t)$ to the solution $\hat{u}$. Automatic differentiation computes exact derivatives, which enter the PDE, boundary, and data loss terms. The total weighted loss is minimised by a two-phase Adam + L-BFGS optimizer, whose gradient updates the network weights $\boldsymbol{\theta}$	17
3.1	Methodological pipeline common to all three PINN studies. The dashed data term $\mathcal{L}_{\text{data}}$ applies only to the inverse problem.	32
4.1	Single-layer forward PINN — four-panel summary. (Top-left) PINN vs. FDM fluence profiles; the two curves are visually indistinguishable, with a relative ℓ^2 error of 5.52×10^{-5} . (Top-right) Absolute error on a logarithmic scale. (Bottom-left) Gaussian source function $S(z)$ centred at $z_0 = 0.1$ cm. (Bottom-right) Convergence of the PDE residual and the boundary losses over training.	34
4.2	Multi-layer PINN solution: (top-left) fluence profile across the three layers; (top-right) $\mu_a(z)$ and $D(z)$ as smoothed step functions; (bottom-left) absolute error; (bottom-right) training loss.	37
4.3	Close-up of the two tissue interfaces. The PINN reproduces the slope change required by flux continuity.	37
4.4	Single-layer inverse recovery: recovered fluence profile against the FDM ground truth and the noisy measurements (top-left); convergence histories of the two trainable variables (top-right, bottom-left); and decomposition of the training loss into PDE, BC and data components (bottom-right).	40
4.5	Multi-layer inverse recovery: fluence profile (top) and convergence trajectories of the six trainable parameters (middle row: μ_a for skin, fat, muscle; bottom row: μ'_s for the same). Note that the absorption coefficients μ_a^{skin} (43% error) and μ_a^{fat} (28% error) are the worst-recovered, reflecting the noise-masking effect on small parameters.	42

4.6	Summary of the sensitivity analysis: effect of n (left), η (centre), and sampling distribution (right).	46
5.1	Bioheat PINN — six-panel results. Top row: $T(z, t)$ heat maps for PINN (left) and FDM reference (centre); pointwise error map (right). Middle row: spatial temperature profiles at six representative time instants (left); laser pulse shape $I(t)$ showing the rectangular envelope between $t_{\text{on}} = 5$ s and $t_{\text{off}} = 35$ s (right). Bottom row: temperature versus time at four tissue depths (skin, tumour, near-tumour muscle, deep muscle) (left); training loss convergence (right). Dashed cyan vertical lines mark the tissue interfaces at $z = 0.3$ and 0.5 cm; dashed green horizontal lines mark t_{on} and t_{off} .	53
5.2	Therapeutic analysis (bioheat PINN). Left: peak temperature profile $T^{\text{max}}(z)$ showing the broad plateau spanning skin and tumour regions at $T \approx 55^\circ\text{C}$. Centre: layer-averaged temperature evolution $\langle T \rangle(t)$ for skin, tumour, and muscle, illustrating the Q -balance effect: skin and tumour reach nearly identical peak temperatures. Right: time above the 43°C therapeutic threshold, showing ~ 25 s of effective treatment in the tumour layer.	54
5.3	Heatmaps of $T(z, t)$ under the four cooling protocols. The cool spot at the skin surface is visible in protocols B and D before the pulse and in C and D after the pulse, but its influence does not propagate differentially into the tumour region.	59
5.4	Comparison of the standard ($\mu_a^{\text{tumour}} = 0.80 \text{ cm}^{-1}$) and plasmonic-enhanced ($\mu_a^{\text{tumour}} = 80 \text{ cm}^{-1}$, $\times 100$) treatments. Top row: $T(z, t)$ heatmaps showing the migration of the hot region from the tumour centre (standard) to the skin–tumour interface (nano). Middle: peak temperature profile $T^{\text{max}}(z)$, with the maximum shifted toward the skin under plasmonic enhancement. Bottom-left: fluence profile $\phi(z)$ on a logarithmic scale — ϕ collapses by fourteen orders of magnitude across the tumour, illustrating the self-shielding mechanism. Bottom-right: bar chart of the key dosimetric metrics.	61
5.5	Sweet-spot analysis: peak temperatures (left), selectivity (centre), and optical penetration depth (right) as functions of μ_a^{tumour} . The selectivity peaks at a modest enhancement ($\mu_a^{\text{tumour}} \approx 2 \text{ cm}^{-1}$) before being eroded by self-shielding for larger values. The penetration depth δ crosses the tumour thickness (≈ 2 mm) near the same value, marking the transition between the deep-penetration regime and the self-shielding regime.	63

List of Tables

1.1	Refractive indices of selected media and biological tissues at $\lambda = 589 \text{ nm}$. . .	9
1.2	Representative anisotropy factors for selected biological tissues.	10
1.3	Optical properties of dermis at selected wavelengths [20].	13
1.4	Optical properties of skeletal muscle at selected wavelengths [20].	13
1.5	Comparative optical properties of dermis, muscle, and liver at $\lambda = 800 \text{ nm}$ (representative literature values).	14
1.6	Laser medical applications and their spectral rationale.	14
2.1	Comparison of PINNs with classical numerical methods for biomedical physics problems.	20
2.2	State-of-the-art PINN accuracy on biomedical and transport problems re- lated to the present thesis.	22
3.1	Optical properties of the three-layer tissue model (visible/NIR range).	26
3.2	Neural network architectures employed in each study.	28
4.1	Quantitative error metrics for the single-layer forward PINN.	35
4.2	Optical properties of the three tissue layers.	36
4.3	Effect of the interface correction term $D'(z)$ on multi-layer PINN accuracy. .	38
4.4	Single-layer inverse recovery, $n = 25$, $\eta = 2\%$	41
4.5	Initial guesses and true values for the multi-layer inverse problem.	41
4.6	Multi-layer inverse recovery, $n = 35$, $\eta = 2\%$	42
4.7	Empirical identifiability rule extracted from the multi-layer recovery experi- ment.	44
4.8	Effect of the number of measurements n on parameter recovery (noise: 2 %; distribution: uniform).	44
4.9	Effect of the noise level η on recovery accuracy ($n = 25$, uniform sampling). .	45
4.10	Effect of the sampling geometry on recovery accuracy ($n = 20$, $\eta = 2\%$). . .	45
5.1	Physical parameters of the three-layer tissue model used in the bioheat problem.	51
5.2	Training history summary for the bioheat PINN.	52
5.3	Peak temperatures and time above 43°C at representative tissue depths. . .	53
5.4	Accuracy metrics vs. FDM reference (bioheat PINN).	55
5.5	Accuracy comparison across all four thesis studies.	56

5.6	Effect of the four cooling protocols on peak temperatures and selectivity. Baseline optics ($\mu_a^{\text{tumour}} = 0.80 \text{ cm}^{-1}$), single pulse $t \in [5, 35] \text{ s}$, $I_0 = 12 \text{ W cm}^{-2}$.	59
5.7	Standard versus plasmonic-enhanced PPTT. Identical pulse parameters; only μ_a^{tumour} differs.	60
5.8	Sweet-spot search over μ_a^{tumour}	63

Nomenclature

Acronyms

Acronym	Definition
AD	Automatic Differentiation
AuNP	Gold Nanoparticle
BC	Boundary Condition
CW	Continuous Wave
DOT	Diffuse Optical Tomography
FDM	Finite-Difference Method
FEM	Finite-Element Method
fNIRS	Functional Near-Infrared Spectroscopy
IC	Initial Condition
L-BFGS	Limited-memory Broyden–Fletcher–Goldfarb–Shanno
MC	Monte Carlo
MRI	Magnetic Resonance Imaging
NIR	Near-Infrared
NN	Neural Network
ODE	Ordinary Differential Equation
PDE	Partial Differential Equation
PDT	Photodynamic Therapy
PINN	Physics-Informed Neural Network
PPTT	Plasmonic Photothermal Therapy
RTE	Radiative Transfer Equation

Latin Symbols

Symbol	Quantity	Unit
c	Specific heat capacity (tissue)	$\text{J kg}^{-1} \text{K}^{-1}$
c_b	Specific heat capacity (blood)	$\text{J kg}^{-1} \text{K}^{-1}$
D	Photon diffusion coefficient	cm
g	Anisotropy factor	dimensionless
I_0	Laser irradiance	W cm^{-2}
k	Thermal conductivity	$\text{W m}^{-1} \text{K}^{-1}$
$\mathcal{L}$	Total PINN loss	dimensionless
n	Refractive index	dimensionless
N	Number of grid / collocation points	dimensionless
Q_{laser}	Volumetric optical heat source	W cm^{-3}
$S(z)$	Optical source term	W cm^{-3}
t	Time	s
T	Temperature	$^{\circ}\text{C}$
T_a	Arterial blood temperature	$^{\circ}\text{C}$
T_{body}	Baseline body temperature	$^{\circ}\text{C}$
z	Depth coordinate	cm

Greek Symbols

Symbol	Quantity	Unit
α	Thermal diffusivity	$\text{m}^2 \text{s}^{-1}$
δ	Optical penetration depth	cm
θ	Neural-network parameters (weights)	—
λ	Wavelength	nm
λ_i	Loss weighting coefficient	—
μ_a	Absorption coefficient	cm^{-1}
μ_s	Scattering coefficient	cm^{-1}
μ'_s	Reduced scattering coefficient	cm^{-1}
μ_t	Total attenuation coefficient	cm^{-1}
μ'_t	Transport coefficient	cm^{-1}
ρ	Tissue density	kg m^{-3}
ρ_b	Blood density	kg m^{-3}
ρc	Volumetric heat capacity	$\text{J m}^{-3} \text{K}^{-1}$
τ	Thermal relaxation time	s
$\phi(z)$	Fluence rate	W cm^{-2}
ω_b	Blood perfusion rate	s^{-1}

Subscripts

Subscript	Region
skin	Cutaneous layer
fat	Subcutaneous adipose layer (forward optical problem only)
tumour	Tumour layer (bioheat problem)
muscle	Skeletal-muscle layer

General Introduction

Context and Motivation

The optical properties of biological tissues and the temperature distributions induced by laser radiation in them lie at the heart of many established and emerging medical technologies. From non-invasive diagnostic modalities such as diffuse optical tomography [1, 2] and functional near-infrared spectroscopy [3], to therapeutic interventions ranging from photodynamic therapy and laser dermatology [4] to plasmonic photothermal ablation of deep-seated tumours [5, 6], the same two coupled questions recur: *How does light propagate through a heterogeneous, scattering and absorbing biological medium?* and *What is the spatio-temporal temperature field that results from the absorbed optical energy?*

The conventional answer to the first question is the steady-state photon diffusion equation, a simplified form of the radiative transfer equation [7] valid in the strongly scattering regime characteristic of soft tissues in the near-infrared “therapeutic window” [8]. The conventional answer to the second question is Pennes’ bioheat equation [9], in which conductive heat transfer is augmented by a perfusion term modelling the cooling action of capillary blood flow. Both equations have been solved for decades using finite-difference (FDM), finite-element (FEM), and Monte Carlo (MC) methods, and these classical techniques remain the workhorses of clinical dosimetry codes [10].

In the last few years, however, the rapid maturation of scientific machine learning [11, 12] has introduced a fundamentally different modelling paradigm: **Physics-Informed Neural Networks (PINNs)**. Building on early work on neural-network solvers for differential equations [13] and pioneered in their modern form by Raissi, Perdikaris and Karniadakis in 2019 [14], PINNs reformulate the solution of a PDE as the training of a neural network whose loss function embeds the differential equation itself, evaluated by automatic differentiation. This single change of perspective brings three properties that are particularly attractive in medical physics: a unified treatment of forward and inverse problems, a mesh-free continuous representation of the solution, and a natural framework for assimilating sparse, noisy clinical measurements.

Problem Statement

Despite the rapidly growing PINN literature, several gaps remain when one tries to apply the methodology to laser–tissue interaction in a master-thesis-scale study. The objectives

are typically:

1. to obtain an accurate forward solution of the diffusion equation for multi-layer tissues, including the often-overlooked interface correction term,
2. to recover unknown optical parameters from sparse and noisy fluence measurements (the inverse problem),
3. to couple the optical solution to a time-dependent bioheat model for hyperthermic dosimetry,
4. and to use the resulting digital model to interrogate physical limits of laser therapy such as therapeutic selectivity and self-shielding by plasmonic enhancement.

Most published PINN works address one of these objectives in isolation, on cluster-grade hardware, and with limited medical-physics interpretation. The present thesis attempts to address **all four** within a single coherent framework, on a laptop-class computational budget, with explicit attention to the physical meaning of the results.

Objectives of the Thesis

The thesis pursues five interrelated objectives:

1. **Methodological:** implement, validate, and document a PINN framework for the photon diffusion equation in 1D, both homogeneous and three-layer, with conservative formulation and smooth coefficient transitions.
2. **Inverse:** formulate and solve the corresponding inverse problem of optical-parameter recovery, and quantify its sensitivity to measurement noise, sample size, and measurement geometry.
3. **Thermal:** extend the framework to the space-time Pennes bioheat equation with depth-varying tissue properties and blood perfusion, modelling a rectangular laser pulse delivered to a buried tumour.
4. **Physical interpretation:** use the validated PINN and a complementary FDM solver to expose two physical phenomena — the Q -balance effect and the self-shielding limit of plasmonic photothermal therapy — and quantify them via systematic parametric studies.
5. **Reproducibility:** deliver a complete, executable research package (notebooks, scripts, and figures) that other students may use as a starting point for their own work.

Identifiability gap in multi-layer inverse problems. A quantitative numerical demonstration that the smallest coefficient in each tissue layer becomes effectively unidentifiable when the noise floor of the measurements exceeds its sensitivity contribution to the fluence, leading to relative errors of 28–43% on the small absorption coefficients of skin and fat under 2% noise.

Q -balance effect in surface hyperthermia. The identification, both numerically and analytically, of a physical regime in which the volumetric optical power coincides between skin and tumour ($Q_{\text{skin}} \approx Q_{\text{tumour}}$), explaining a counter-intuitive zero thermal selectivity despite strong tissue contrast.

Self-shielding limit of plasmonic photothermal therapy. A parametric sweep over the tumour absorption coefficient showing that aggressive plasmonic enhancement ($\mu_a \times 100$) degrades rather than improves selectivity, because the optical penetration depth δ collapses below the tumour thickness; a modest sweet-spot is identified.

These contributions emerge not from algorithmic novelty in the PINN itself, but from the use of the PINN as a numerical microscope for laser–tissue physics.

Thesis Outline

The manuscript is organised in six chapters:

- **Chapter 1: Physical Foundations of Laser–Tissue Interaction.** Introduces the optical properties of tissues, the radiative transfer equation, its diffusion approximation, and the optical parameters of the biological layers used in this work.
- **Chapter 2: Physics-Informed Neural Networks — State of the Art.** Reviews the PINN methodology, its historical development, and its applications in medical physics, with explicit justification for its choice in the present thesis.
- **Chapter 3: Methodology and Resources.** Specifies the five PINN/FDM studies, the conservative formulation, the interface-aware collocation strategy, the inverse-problem formulation, the bioheat coupling, and the computational environment.
- **Chapter 4: Optical Results — Forward, Inverse, and Sensitivity Analysis (Studies 1–3).** Presents the two forward photon-diffusion solutions (single- and three-layer geometries), the inverse recovery of optical parameters from noisy measurements, and a 15-experiment sensitivity study. The chapter exposes the identifiability gap by which the small coefficients in each tissue layer are recovered with the largest relative errors.
- **Chapter 5: Thermal Results — Bioheat PINN and Photothermal Therapy Exploration (Studies 4–5).** Couples the validated optical solver to the time-dependent Pennes bioheat equation and uses the resulting coupled model to investigate surface cooling, plasmonic enhancement, and a parametric sweep over μ_a^{tumour} . The chapter exposes the Q -balance effect and the self-shielding limit of plasmonic photothermal therapy.
- **General Conclusion.** Synthesises the methodological and physical findings, and outlines avenues for future work.

Chapter 1

Physical Foundations of Laser–Tissue Interaction

Preamble

Understanding the interaction between light and biological tissues constitutes the cornerstone of modern optical medical applications [10, 15]. Since the invention of the laser in 1960, the medical field has witnessed a genuine revolution in the diagnostic and therapeutic use of light. The success of these applications depends on an accurate understanding of the optical properties of different tissues and how they interact with electromagnetic radiation [8, 16].

This chapter presents the physical and mathematical foundations governing light propagation in biological tissues. Section 1.1 introduces the fundamental phenomena of light–tissue interaction and their clinical relevance. Sections 1.2.1–1.2.5 define and analyse the five core optical parameters. Section 1.3 develops the mathematical models used to describe light transport, from the exact Radiative Transfer Equation to the Diffusion Approximation that underpins the present work. Section 1.4 provides a characterisation of biologically important tissues — dermis and skeletal muscle (two of the three modelled layers), with liver included as a strongly absorbing comparison case — together with a comparative analysis. The chapter closes with a summary in Section 1.6 that motivates the modelling choices made in subsequent chapters.

1.1 General Introduction

1.1.1 Importance of Optical–Tissue Interaction Studies

The study of light–tissue interaction forms the physical foundation of a broad range of modern medical technologies. The complexity of this field arises from the inherently heterogeneous nature of biological tissues, whose structural and chemical composition varies over length scales from nanometres (protein fibres) to millimetres (organ layers), leading to spatially varying and wavelength-dependent optical properties.

When light impinges on the surface of a biological tissue, four fundamental physical phenomena occur simultaneously:

1. **Specular surface reflection.** At the air–tissue interface, a fraction of the incident light is reflected according to the Fresnel equations. For normal incidence, this fraction is typically 4–7% and is determined by the tissue refractive index.
2. **Absorption.** Biomolecules called *chromophores* — primarily haemoglobin, melanin, water, and lipids — absorb photon energy and convert it into thermal, chemical, or re-emitted optical energy.
3. **Scattering.** Microscopic refractive-index inhomogeneities (organelles, collagen fibres, cell membranes) redirect photons without removing them, producing the characteristic diffuse glow of illuminated tissue.
4. **Transmission.** The fraction of light that traverses the tissue with neither significant absorption nor scattering constitutes the transmitted component, exploited in transillumination imaging.




Figure 1.1: Schematic of the four fundamental phenomena occurring when a laser beam is incident on a biological tissue: specular reflection at the surface, absorption by chromophores, multiple scattering, and transmission.

1.1.2 Medical Applications

Quantitative knowledge of these phenomena underpins a wide range of clinical technologies:

- **Optical diagnostics.** Optical Coherence Tomography (OCT), Photoacoustic Imaging (PAI), and near-infrared spectroscopy (NIRS) exploit tissue-specific scattering and absorption signatures to produce structural and functional images without ionising radiation.

- **Laser therapy.** Optical surgery, Photodynamic Therapy (PDT), and Laser-Induced Thermotherapy (LITT) deliver controlled doses of light energy to ablate, sensitise, or thermally destroy target tissues.
- **Physiological monitoring.** Pulse oximetry and functional near-infrared spectroscopy (fNIRS) track haemodynamic changes in real time through the non-invasive measurement of transmitted or reflected near-infrared light.

In all these applications, the accurate prediction of the *fluence rate distribution* $\phi(\mathbf{r})$ inside the tissue is the central computational challenge — and the primary objective of this thesis.

1.2 Core Optical Parameters

The optical behaviour of a biological tissue is completely characterised by five parameters: the absorption coefficient μ_a , the scattering coefficient μ_s , the reduced scattering coefficient μ'_s , the refractive index n , and the anisotropy factor g . This section defines each parameter, discusses its physical origin, and establishes its clinical significance.

1.2.1 Absorption Coefficient (μ_a)

Definition and Mathematical Description The *absorption coefficient* μ_a (cm^{-1}) quantifies the probability per unit path length that a photon is absorbed by the medium. Under the assumption of a homogeneous, non-scattering medium, the Beer–Lambert law describes the depth-dependent intensity:

$$I(z) = I_0 \exp(-\mu_a z), \quad (1.1)$$

where I_0 is the surface irradiance and z is the depth below the surface. The inverse $\ell_a = 1/\mu_a$ is the *absorption mean free path*, i.e. the average distance travelled by a photon before being absorbed.

Principal Chromophores in Biological Tissues The wavelength dependence of μ_a is governed by the spectral absorption cross-sections of the principal tissue chromophores:

1. **Haemoglobin.** The dominant absorber in the visible range (400–600 nm). It exists in two forms with distinct spectra: oxygenated (HbO_2) and deoxygenated (Hb). The Soret absorption band near 415 nm is extremely strong; secondary maxima appear at 540 nm and 576 nm for HbO_2 . At 800 nm (the *isosbestic point*), $\mu_a(\text{Hb}) = \mu_a(\text{HbO}_2)$, which is exploited by pulse oximeters.
2. **Melanin.** Concentrated in the epidermal basal layer. Its absorption decreases monotonically with wavelength, approximately as $\mu_a^{\text{mel}} \propto \lambda^{-3}$, providing strong ultraviolet shielding and moderate visible absorption. Melanin content determines skin phototype and strongly influences laser treatment parameters for dermatology.

3. **Water.** The dominant tissue constituent (60–80% by mass), water exhibits low absorption in the visible and near-infrared (NIR) below 1 400 nm, with strong absorption bands centred at 1 450 nm, 1 940 nm, and 3 000 nm. CO₂ and Er:YAG lasers operate at the 10 600 nm and 2 940 nm water absorption peaks, respectively, for precision tissue ablation.
4. **Lipids.** Present in adipose tissue and myelin sheaths, lipids exhibit characteristic absorption peaks at 1 210 nm and 1 720 nm, enabling spectroscopic quantification of body fat composition.

Clinical Significance

- **Wavelength selection.** Knowledge of the absorption spectrum guides the choice of laser wavelength to maximise therapeutic effect in the target chromophore while minimising collateral damage.
- **Selective photothermolysis.** The principle of selective photothermolysis (Anderson and Parrish, 1983) exploits the differential absorption between a target structure and surrounding tissue to achieve thermally confined destruction.
- **Penetration depth.** Together with μ'_s , the absorption coefficient determines the effective penetration depth (Eq. (1.9)), which sets the maximum treatable depth for any given laser system.

1.2.2 Scattering Coefficient (μ_s)

Definition The *scattering coefficient* μ_s (cm⁻¹) is the probability per unit path length that a photon undergoes a scattering event. The associated scattering mean free path $\ell_s = 1/\mu_s$ is the average distance between successive scattering events.

Physical Origins in Biological Tissues Two scattering regimes are relevant:

1. **Rayleigh scattering.** Dominates when the scatterer size $d \ll \lambda$. The scattering cross-section scales as λ^{-4} , making it negligible in the NIR. Sub-cellular structures such as ribosomes ($d \approx 25$ nm) contribute to Rayleigh-type scattering.
2. **Mie scattering.** Dominates when $d \approx \lambda$. It is the prevailing mechanism in biological tissues, caused by mitochondria (0.5–3 μ m), cell nuclei (5–10 μ m), and collagen fibrils (50–300 nm diameter). The Mie scattering cross-section decreases more slowly with wavelength than Rayleigh, approximately as λ^{-b} with $b \approx 0.5$ –2 depending on tissue type.

1.2.3 Reduced Scattering Coefficient (μ'_s)

Definition and Physical Meaning Because biological tissue scattering is strongly *forward-directed* (anisotropy factor $g \approx 0.9$), a photon must undergo multiple scattering events before its direction is effectively randomised. The *reduced scattering coefficient* captures this collective effect:

$$\mu'_s = \mu_s(1 - g). \quad (1.2)$$

The associated *transport mean free path* $\ell^* = 1/\mu'_s$ is the average distance after which the photon direction is fully randomised. It replaces ℓ_s as the relevant length scale in the diffusion regime.

Practical Importance

- μ'_s is the parameter that enters the Diffusion Approximation (Section 1.3.2), making it directly measurable by time-resolved or frequency-domain optical techniques.
- Typical values in soft tissues lie in the range $\mu'_s = 5\text{--}30\text{ cm}^{-1}$ in the NIR, two to three orders of magnitude larger than μ_a .
- This strong inequality $\mu'_s \gg \mu_a$ is the physical condition that justifies the Diffusion Approximation and, consequently, the photon diffusion equation solved numerically in this thesis.

1.2.4 Refractive Index (n)

Definition The *refractive index* is the ratio of the speed of light in vacuum c to its phase velocity v in the medium:

$$n = \frac{c}{v}. \quad (1.3)$$

At tissue boundaries, both surface reflection (Fresnel equations) and refraction (Snell's law, $n_1 \sin \theta_1 = n_2 \sin \theta_2$) depend on n , which directly affects the Robin-type boundary condition used in the diffusion model.

Representative Values Table 1.1 lists refractive indices for biologically relevant media at $\lambda = 589\text{ nm}$.

The mismatch between the refractive indices of air ($n = 1.00$) and tissue ($n \approx 1.4$) produces the *effective reflection parameter* A that appears in the Robin boundary condition of the photon diffusion equation (Chapter 3, Eq. (3.1)):

$$A = \frac{1 + R_{\text{eff}}}{1 - R_{\text{eff}}}, \quad (1.4)$$

where R_{eff} is the angle-averaged internal reflection coefficient computed from the Fresnel equations [17].

Table 1.1: Refractive indices of selected media and biological tissues at $\lambda = 589 \text{ nm}$.

Medium / Tissue	Refractive index n
Air	1.000
Water	1.330
Blood	1.33–1.35
Cerebrospinal fluid	1.335
Cornea	1.376
Dermis (skin)	1.40–1.42
Brain grey matter	1.36–1.37
Brain white matter	1.37–1.38
Skeletal muscle	1.37–1.39
Liver	1.37–1.38

1.2.5 Anisotropy Factor (g)

Definition The *anisotropy factor* is the mean cosine of the single-scattering deflection angle θ :

$$g = \langle \cos \theta \rangle = \int_0^\pi p(\theta) \cos \theta \sin \theta d\theta, \quad (1.5)$$

where $p(\theta)$ is the *phase function* describing the angular distribution of scattered photons. The parameter g ranges from -1 (complete backscattering) through 0 (isotropic) to $+1$ (complete forward scattering).

The Henyey–Greenstein Phase Function The Henyey–Greenstein (HG) phase function is the standard model for biological tissue scattering [18]:

$$p(\cos \theta) = \frac{1 - g^2}{4\pi(1 + g^2 - 2g \cos \theta)^{3/2}}. \quad (1.6)$$

It is analytically tractable, parameterised by a single value g , and reproduces the strongly forward-peaked scattering of tissue adequately for diffuse-regime modelling. Its limitations (it does not reproduce the sharp backward lobe seen in some tissues) have motivated extended models such as the two-term HG function, but the single-term form is sufficient for the diffusion approximation studied here.

Representative Values Table 1.2 lists typical anisotropy factors for tissues studied in this thesis.

Table 1.2: Representative anisotropy factors for selected biological tissues.

Tissue	g	Dominant scatterer
Dermis (skin)	0.80–0.90	Collagen fibrils
Brain grey matter	0.85–0.95	Cell membranes, mitochondria
Whole blood	0.990–0.995	Red blood cells
Skeletal muscle	0.90–0.95	Myofibrils
Liver	0.85–0.95	Hepatocyte organelles

1.3 Mathematical Models of Light Transport

1.3.1 The Radiative Transfer Equation

Full Form

The Radiative Transfer Equation (RTE) is the most complete and physically rigorous description of photon transport in absorbing and scattering media. In its time-dependent form it reads:

$$\frac{1}{c} \frac{\partial L}{\partial t} + \hat{s} \cdot \nabla L = -\mu_t L + \mu_s \int_{4\pi} p(\hat{s}', \hat{s}) L d\Omega' + S, \quad (1.7)$$

where $L(\mathbf{r}, \hat{s}, t)$ is the radiance [$\text{W m}^{-2} \text{sr}^{-1}$] at position $\mathbf{r}$ in direction $\hat{s}$ at time t , $\mu_t = \mu_a + \mu_s$ is the attenuation coefficient, $d\Omega'$ is the differential solid angle, and $S(\mathbf{r}, \hat{s}, t)$ is the source term [$\text{W m}^{-3} \text{sr}^{-1}$]. The five terms of Eq. (1.7) represent, from left to right: the temporal rate of change of radiance, the streaming (advection) of photons, extinction by absorption and out-scattering, gain by in-scattering from all directions, and the source.

Challenges of Direct Solution

Solving the RTE exactly is computationally prohibitive for routine medical physics calculations:

- The equation is a 7-dimensional integro-differential equation in $(x, y, z, \theta, \phi, t, \lambda)$ space.
- The integral operator couples all directions, requiring iterative or moment-expansion schemes.
- Complex tissue geometries demand unstructured meshes with $\mathcal{O}(10^6\text{--}10^9)$ degrees of freedom.

These limitations motivate the use of the Diffusion Approximation for the problems addressed in this thesis, and Monte Carlo simulation as the independent validation reference.

1.3.2 The Diffusion Approximation

Conditions of validity. The Diffusion Approximation (DA) to the RTE is valid when scattering dominates ($\mu'_s \gg \mu_a$, typically $\mu'_s/\mu_a \sim 10\text{--}100$ in soft tissues in the NIR), the observation point is several transport mean free paths from the source ($r \gg \ell^* = 1/\mu'_s$), and the angular distribution of the radiance is quasi-isotropic. These conditions break down close to the source (ballistic regime), near boundaries, and in strongly absorbing tissues, where Monte Carlo simulation provides the reference.

Diffusion equation. Expanding the radiance in spherical harmonics and retaining only the first two terms (the P_1 approximation) reduces Eq. (1.7) to the steady-state *photon diffusion equation* for the fluence rate $\phi(\mathbf{r}) = \int_{4\pi} L(\mathbf{r}, \hat{s}) d\Omega$ [W cm⁻²]:

$$-\nabla \cdot [D(\mathbf{r}) \nabla \phi(\mathbf{r})] + \mu_a(\mathbf{r}) \phi(\mathbf{r}) = S(\mathbf{r}), \quad (1.8)$$

with diffusion coefficient $D = 1/[3(\mu_a + \mu'_s)]$ — the standard steady-state form, which is the one adopted in all the simulations of this thesis. A well-known alternative, the absorption-independent Furutsu–Yamada coefficient $D_{FY} = 1/(3\mu'_s)$ [19], is preferable for time-domain transport or when μ_a is not negligible compared to μ'_s ; the two coincide in the diffusive limit $\mu_a \ll \mu'_s$, and their tissue-dependent difference is quantified in Section 1.4.4.

Boundary conditions. At the tissue surface, the Robin (Type III) condition $\phi + 2AD \partial_n \phi = 0$ [17] accounts for refractive-index mismatch with the external medium through the effective reflection parameter A , while a Dirichlet condition $\phi(z_{\max}) = 0$ is imposed at the distal boundary of a thick slab where the fluence has decayed to negligible levels.

Analytical solution and penetration depth. For a homogeneous, unbounded medium and an isotropic point source of power P [W], Eq. (1.8) yields the Green's function $\phi(r) = (P/4\pi Dr) \exp(-\mu_{\text{eff}} r)$, with

$$\mu_{\text{eff}} = \sqrt{3\mu_a(\mu_a + \mu'_s)}, \quad \delta = 1/\mu_{\text{eff}} = \sqrt{D/\mu_a}. \quad (1.9)$$

The penetration depth δ is a key clinical-design parameter: it specifies the depth at which the fluence has fallen to $1/e \approx 37\%$ of its surface value and directly controls the accessible treatment volume.

Limitations. The DA breaks down near the source (where $r < 1/\mu'_s$ and the P_1 truncation loses accuracy), in highly absorbing tissues where $\mu_a/\mu'_s \gtrsim 0.1$ produces errors up to 30% [1], and at layer interfaces, which require the conservative formulation with explicit flux continuity discussed in Chapter 3.

1.3.3 Monte Carlo Simulation

Monte Carlo (MC) simulation of photon transport stochastically samples the RTE by tracking individual photon packets through the tissue. Each packet is launched from the source with unit weight, propagated a distance $s = -\ln(\xi)/\mu_t$ drawn from an exponential distribution (where ξ is a uniform random number), scattered according to the HG phase function, and absorbed by reducing its weight by the albedo $a = \mu_s/\mu_t$ at each step. Packets that reach a weight below a threshold are terminated by Russian Roulette.

MC is exact in the limit of an infinite number of photon packets, assumes no approximations about the angular distribution of radiance, and handles arbitrary geometries and strongly absorbing media. Its principal limitation is statistical: achieving a 1% relative standard deviation in the fluence requires of order 10^6 – 10^8 packets, with run times ranging from minutes to hours on a single CPU. In this thesis, MC simulations provide the reference data against which the PINN warm-start solution $\phi_c + \phi_d$ is validated (Chapter 4).

1.4 Optical Properties of Selected Biological Tissues

1.4.1 The Biological (Therapeutic) Window

The spectral range 600–1300 nm is known as the *biological window* because it represents the interval where absorption by both haemoglobin (decreasing above 600 nm) and water (increasing above 1300 nm) is simultaneously minimised. This window provides:

- the maximum achievable penetration depth for any given tissue,
- spectral sensitivity to haemoglobin oxygenation state (exploited by pulse oximetry and fNIRS), and
- a favourable signal-to-background ratio for diffuse optical imaging.

The three laser parameters studied in this thesis ($\lambda = 800$ nm) fall near the centre of this window.

1.4.2 Dermis

Anatomical Structure The skin consists of three layers: the *epidermis* (0.05–0.1 mm, melanin-rich), the *dermis* (1–4 mm, collagen-rich, vascularised), and the *hypodermis* (fatty, variable thickness). In this thesis the dermis is modelled as a homogeneous slab of optical properties representative of the 800 nm NIR window.

Optical Properties Table 1.3 summarises the wavelength-resolved optical properties of dermis from the literature [20, 10].

Fixed parameters: $n = 1.37$, $g = 0.82$.

Table 1.3: Optical properties of dermis at selected wavelengths [20].

λ (nm)	μ_a (cm ⁻¹)	μ'_s (cm ⁻¹)	μ_{eff} (cm ⁻¹)	δ (mm)
630	0.120	27.78	3.16	3.16
700	0.080	25.35	2.48	4.03
800	0.400	36.00	6.61	1.51
900	0.060	19.44	1.95	5.13
1000	0.100	17.50	2.30	4.35

Clinical Relevance Dermis is the primary target tissue for a wide range of laser procedures including vascular lesion treatment (haemoglobin absorption at 532 nm and 585 nm), laser hair removal (melanin absorption at 694–1 064 nm), and photorejuvenation. Its high $\mu'_s = 36 \text{ cm}^{-1}$ at 800 nm makes it the most challenging tissue for PINN convergence (steepest fluence gradients), as confirmed in Chapter 4.

1.4.3 Skeletal Muscle

Optical Properties Table 1.4 provides the optical properties of skeletal muscle.

Table 1.4: Optical properties of skeletal muscle at selected wavelengths [20].

λ (nm)	μ_a (cm ⁻¹)	μ'_s (cm ⁻¹)	μ_{eff} (cm ⁻¹)	δ (mm)
630	0.070	12.98	1.72	5.82
700	0.050	11.57	1.42	7.04
800	0.100	15.00	2.13	4.70
900	0.080	13.20	1.78	5.62
1000	0.110	11.80	1.97	5.07

Fixed parameters: $n = 1.37$, $g = 0.90$.

Clinical Relevance Muscle is encountered in deep-tissue laser procedures (interstitial laser coagulation, LITT) and in the photon propagation path for transcutaneous diagnostics such as near-infrared muscle oximetry. Its moderate $\mu'_s = 15 \text{ cm}^{-1}$ and low $\mu_a = 0.1 \text{ cm}^{-1}$ at 800 nm produce the largest penetration depth among the three tissues studied here ($\delta = 4.70 \text{ mm}$), which is reflected in the widest spatial domain required for the PINN simulation ($z_{\text{max}} = 3.29 \text{ cm}$, Chapter 3).

1.4.4 Comprehensive Comparison at 800 nm

Table 1.5 compares the three tissues at the common simulation wavelength of 800 nm. These are representative literature values; the specific (idealised) coefficients adopted in the PINN simulations are listed in Chapter 3, Table 3.1.

Three key observations emerge from Table 1.5:

Table 1.5: Comparative optical properties of dermis, muscle, and liver at $\lambda = 800$ nm (representative literature values).

Tissue	μ_a (cm^{-1})	μ'_s (cm^{-1})	n	g	μ_{eff} (cm^{-1})	δ (mm)	D_{FY} (cm)
Dermis	0.40	36.00	1.37	0.82	6.609	1.51	0.00926
Muscle	0.10	15.00	1.37	0.90	2.128	4.70	0.02222
Liver	2.50	21.00	1.38	0.94	13.276	0.75	0.01587

Penetration depth spans one order of magnitude. The ratio $\delta_{\text{muscle}}/\delta_{\text{liver}} \approx 6.3$ illustrates how strongly the accessible treatment depth depends on tissue type: a low-absorption tissue such as muscle transmits light several times deeper than a strongly absorbing tissue such as liver, which directly constrains the spatial domain required to model each case.

Strong absorbers are numerically demanding. The effective attenuation $\mu_{\text{eff}}^{\text{liver}} = 13.3$ is about $6\times$ larger than that of muscle. Tissues with such large μ_{eff} produce very steep fluence gradients and a correspondingly large dynamic range in the PDE residual, which makes them markedly harder to resolve numerically and motivates the careful loss balancing and interface-aware collocation described in Chapter 3.

The Furutsu–Yamada correction is tissue-dependent. The correction $\Delta D/D$ is $+1.1\%$ for dermis (small, $\mu_a \ll \mu'_s$), $+0.7\%$ for muscle, but $+11.9\%$ for liver. This last value is non-negligible and its inclusion is mandatory for physically accurate simulation of liver LITT.

1.5 Medical Applications Based on Differential Optical Properties

Table 1.6 relates the principal laser applications to the underlying spectral rationale.

Table 1.6: Laser medical applications and their spectral rationale.

Application	Wavelength (nm)	Physical basis
fNIRS / DOT	700–900	Deep penetration; differential Hb/HbO ₂ absorption
Vascular lesions	532, 585–595, 1 064	Strong Hb absorption
Laser hair removal	694, 755, 800, 1 064	Melanin absorption; follicle sparing
Optical coherence tomography	800–1 300	High scattering, low absorption, coherence gating
Photodynamic therapy	630–690	Photosensitiser absorption + ¹ O ₂ generation
Laser tissue ablation (water)	2 940, 10 600	Er:YAG and CO ₂ ; strong water absorption
Liver LITT	800–1 064	Moderate absorption in biological window

1.6 Summary

This chapter has established the physical and mathematical foundations of laser-light propagation in biological tissues. Tissue optical behaviour is governed by five wavelength-dependent parameters — μ_a , μ_s , μ'_s , n , g — of which the reduced scattering coefficient μ'_s and the absorption coefficient μ_a enter the governing PDE directly. The Radiative Transfer Equation provides the exact description of photon transport but is computationally intractable for routine simulations; the Diffusion Approximation, valid when $\mu'_s \gg \mu_a$ and at distances beyond one transport mean free path from the source, reduces the problem to a tractable elliptic PDE. The photon diffusion equation with the standard (steady-state) diffusion coefficient and the Robin boundary condition constitutes the governing model solved numerically throughout this thesis. All simulations are performed at 800 nm, located at the centre of the biological window (600–1300 nm), where the two layered tissues studied — dermis and muscle — have well-characterised optical properties suitable for the three-layer skin–fat–muscle model used in the subsequent chapters.

Chapter 2

Physics-Informed Neural Networks: Foundations, Advances, and Application to Medical Physics

Modern biomedical research increasingly demands computational tools that are simultaneously *data-efficient*, *physically consistent*, and *capable of solving ill-posed inverse problems* under noisy, sparse measurement conditions. This chapter introduces Physics-Informed Neural Networks (PINNs) as a paradigm that naturally satisfies all three requirements. We trace the intellectual lineage of PINNs from early neural network theory through to the landmark 2019 formulation, survey the current state of the art with emphasis on medical and biological applications, and provide an explicit justification for their adoption in the present thesis.

2.1 What Is a Physics-Informed Neural Network?

2.1.1 Conceptual Definition

A Physics-Informed Neural Network is a deep neural network whose training loss function incorporates, as penalty terms, the residuals of one or more governing physical equations (PDEs or ODEs), initial conditions, and boundary conditions, so that the trained network approximates the solution of the physical problem rather than fitting data alone.

In conventional deep learning, a network $\mathcal{N}_\theta$ is trained to minimise a purely data-driven loss:

$$\mathcal{L}_{\text{data}} = \frac{1}{N_d} \sum_{i=1}^{N_d} \left\| \mathcal{N}_\theta(\mathbf{x}_i) - \mathbf{y}_i \right\|^2. \quad (2.1)$$

A PINN replaces or augments this with the composite loss:

$$\boxed{\mathcal{L}(\theta) = \underbrace{\lambda_r \mathcal{L}_{\text{PDE}}}_{\text{physics}} + \underbrace{\lambda_b \mathcal{L}_{\text{BC}}}_{\text{boundaries}} + \underbrace{\lambda_i \mathcal{L}_{\text{IC}}}_{\text{initial cond.}} + \underbrace{\lambda_d \mathcal{L}_{\text{data}}}_{\text{observations}}} \quad (2.2)$$

where each term is a mean-squared residual:

$$\mathcal{L}_{\text{PDE}} = \frac{1}{N_r} \sum_{i=1}^{N_r} \left| \mathcal{F}[\mathcal{N}_\theta](\mathbf{x}_i^r) \right|^2, \quad (2.3)$$

$$\mathcal{L}_{\text{BC}} = \frac{1}{N_b} \sum_{j=1}^{N_b} \left| \mathcal{B}[\mathcal{N}_\theta](\mathbf{x}_j^b) \right|^2, \quad (2.4)$$

$$\mathcal{L}_{\text{IC}} = \frac{1}{N_0} \sum_{k=1}^{N_0} \left| \mathcal{N}_\theta(\mathbf{x}_k^0) - u_0(\mathbf{x}_k^0) \right|^2. \quad (2.5)$$

Here $\mathcal{F}[\cdot]$ is the PDE operator acting on the network output, computed via *automatic differentiation* (AD), which is exact to machine precision. The collocation points $\{\mathbf{x}_i^r\}$, $\{\mathbf{x}_j^b\}$, $\{\mathbf{x}_k^0\}$ are sampled from the domain interior, boundaries, and initial time, respectively. No structured mesh is required.

2.1.2 Structural Anatomy of a PINN

Figure 2.1 shows the complete architecture of a PINN solving a space-time PDE.




Figure 2.1: Structural anatomy of a PINN. The network maps input coordinates $\mathbf{x} = (z, t)$ to the solution $\hat{u}$. Automatic differentiation computes exact derivatives, which enter the PDE, boundary, and data loss terms. The total weighted loss is minimised by a two-phase Adam + L-BFGS optimizer, whose gradient updates the network weights θ .

2.1.3 Key Technical Properties

1. **Mesh-free:** collocation points are scattered, not gridded; geometry change requires no re-meshing.
2. **Physics-constrained:** the trained network is guaranteed to satisfy the governing PDE at every sampled point; extrapolation remains physically consistent.
3. **Inverse-capable:** physical parameters (μ_a , k , ρc , ...) become additional trainable variables; no separate inversion algorithm is needed.
4. **Differentiable everywhere:** the output is an analytic function; derived quantities (heat flux, fluence gradient) are computed exactly without post-processing.
5. **Unified framework:** forward, inverse, and data-assimilation problems are solved within a single optimisation loop.

2.2 Historical Development

The emergence of PINNs is the culmination of three intellectual streams: the approximation theory of neural networks, the numerical solution of differential equations, and the renaissance of deep learning. The conceptual foundation was laid by [21], formalised into the universal approximation theorems of [22], and applied to PDEs by [23] and [13], who used neural networks parameterising the solution under collocation constraints on simple boundary-value problems. After two decades of relative dormancy, the deep-learning revival catalysed by AlexNet [24] and the maturation of automatic-differentiation frameworks (TensorFlow 2015, PyTorch 2017) enabled the scaling of these ideas to nonlinear, multi-dimensional problems.

The 2019 turning point. The landmark paper by [14] transformed PINNs from a niche curiosity into a mainstream computational tool. Its key contributions were the unification of forward and inverse problems within a single framework, the scalability afforded by GPU-accelerated automatic differentiation, the universality demonstrated across fluid mechanics, quantum mechanics, and solid mechanics, and the data efficiency that allowed sub-percent accuracy with very few noisy observations. Within six years, the paper accumulated over 15,000 citations and catalysed an entirely new sub-field of *scientific machine learning*, with subsequent work addressing specific failure modes of the early formulation (Section 2.4).

2.3 Importance of PINNs in the Medical Physics Context

2.3.1 The Unique Challenges of Medical Physics Modelling

Medical physics presents a particularly demanding computational environment:

- **Heterogeneity:** biological tissues are spatially inhomogeneous, with optical, thermal, and mechanical properties varying sharply across layer boundaries (skin, fat, muscle, tumour).
- **Sparse data:** *in vivo* measurements are limited; inserting dense sensor arrays is clinically impractical and ethically restricted.
- **Ill-posed inverse problems:** recovering tissue optical properties from boundary measurements (e.g. diffuse optical tomography) is severely underdetermined.
- **Coupled physics:** light transport (diffusion equation), heat transfer (bioheat equation), and biological response (cell damage) must be solved simultaneously with shared, spatially variable coefficients.
- **Real-time requirements:** treatment planning for laser surgery or thermal ablation demands fast forward-model evaluations to enable intraoperative adaptation.

Traditional numerical methods — Finite Element Methods (FEM), Finite Difference Methods (FDM), Monte Carlo — address many of these challenges but have inherent limitations: they require structured or high-quality unstructured meshes; they cannot directly solve inverse problems (a separate optimisation algorithm must be coupled); they produce no analytic solution (derivatives require numerical differentiation); and they must be re-run from scratch when a parameter changes. PINNs address all four limitations simultaneously.

2.3.2 PINNs vs. Classical Methods: A Structured Comparison

Table 2.1 provides a detailed feature matrix comparing PINNs with the three most common classical approaches in biomedical simulation.

2.3.3 Specific Medical Physics Applications

The following subsections review the four major areas where PINNs have demonstrated clinical or pre-clinical value.

Diffuse Optical Tomography (DOT)

DOT reconstructs the three-dimensional map of $\mu_a(\mathbf{r})$ and $\mu'_s(\mathbf{r})$ inside tissue from boundary measurements of near-infrared (NIR) light. This is an ill-posed inverse problem: the number of unknowns far exceeds the number of measurements, and the forward model (diffusion equation) is expensive to evaluate repeatedly. PINNs can in principle solve the forward model and the inversion simultaneously, eliminating the inner–outer loop structure of classical iterative reconstructors; a representative example is the recovery of optical properties and fields in nano-optical inverse problems with PINNs [25].

Table 2.1: Comparison of PINNs with classical numerical methods for biomedical physics problems.

Feature	<i>FDM</i>	<i>FEM</i>	<i>Monte Carlo</i>	<i>PINN</i>
Mesh required	✓	✓	✗	✗
Handles discontinuous coeff.	Partial	✓	✓	Partial [†]
Direct inverse solving	✓	✓	✓	✓
Analytic derivatives	✓	✓	✓	✓
Data assimilation	Limited	Limited	✓	✓
Parametric reuse [‡]	✓	✓	✓	✓
GPU scalability	Limited	Limited	✓	✓
Statistical noise modelling	✓	✓	✓	✓
Interpretability	High	High	Moderate	Low–Moderate
Accuracy (routine)	High	High	Very high	Moderate–High
Implementation complexity	Low	Medium	Low	Medium

[†] via smooth sigmoid coefficient approximation (used in this thesis).

[‡] without re-solving after minor parameter changes (operator learning).

Laser Therapy and Thermal Ablation Planning

Laser Interstitial Thermal Therapy (LITT) and High-Intensity Focused Ultrasound (HIFU) rely on accurate bioheat models to predict temperature distributions and avoid collateral tissue damage. The governing equation is the Pennes bioheat equation:

$$\rho c \frac{\partial T}{\partial t} = \nabla \cdot (k \nabla T) - w_b \rho_b c_b (T - T_a) + Q_{\text{laser}}(z, t), \quad (2.6)$$

where w_b is blood perfusion rate, ρ_b and c_b are blood density and specific heat, T_a is arterial blood temperature, and $Q_{\text{laser}} = \mu_a(z) \phi(z) I(t)$ is the laser-induced heat source. PINNs offer real-time forward evaluation after a single training phase, and can in principle recover thermal parameters such as k , ρc , or w_b from temperature measurements, in the spirit of PINN formulations of heat-transfer problems with unknown coefficients [26]. The present thesis implements a full PINN solution of Eq. (2.6) for a three-layer tissue with a buried tumour (the bioheat problem, Chapter 5).

Other Biomedical Applications

PINNs have also been applied to electrical impedance tomography, reconstructing tissue conductivity from surface electrodes with sub-5% error on 2D phantoms [27], and to cardiovascular flow, assimilating 4D-flow MRI data to recover pressure fields that are not directly measurable [28] — the latter being directly analogous to optical-property recovery, in which a physical field is reconstructed from indirect observations.

2.4 State of the Art

2.4.1 Advances in PINN Methodology

Since 2019 the PINN literature has expanded rapidly, with several comprehensive reviews now documenting the field [11, 12, 29]. Several methodological advances bear directly on the present work. Adaptive loss-weighting schemes, such as the Neural-Tangent-Kernel balancing of Wang et al. [30], address the gradient imbalance between PDE and boundary terms that is a primary cause of training failure. Domain-decomposition (XPINN) methods [31] and Fourier-feature embeddings [32] improve the resolution of sharp gradients near interfaces, a recurring difficulty in multi-layer tissue. Causal training [33] enforces temporal ordering in time-dependent problems and is adopted in spirit by the bioheat study of this thesis. Finally, the two-phase Adam + L-BFGS schedule [34], used throughout this work, has become the de facto standard: L-BFGS exploits curvature near the minimum to gain 2–3 orders of magnitude in final accuracy over Adam alone.

2.4.2 Benchmarks and Accuracy

Table 2.2 summarises representative PINN accuracy results from the recent literature for problems closely related to the present thesis.

Table 2.2: State-of-the-art PINN accuracy on biomedical and transport problems related to the present thesis.

Problem	Reference	Error metric	Value
1D diffusion equation (homogeneous)	Lu et al. [35]	Rel. ℓ^2	10^{-4} – 10^{-5}
2-domain diffusion (interface)	Jagtap & Karniadakis [31]	Rel. ℓ^2	$\sim 10^{-3}$
Navier–Stokes (blood flow)	Raissi et al. [28]	Rel. ℓ^2	$< 1\%$
EIT conductivity reconstruction	Khoo & Ying [27]	Mean abs. err.	$< 5\%$
Inverse optical scattering (ε recovery)	Chen et al. [25]	Field L^2 err.	2.8–5%
Heat-transfer benchmark (forward/inverse)	Cai et al. [26]	Rel. ℓ^2	$\sim 10^{-3}$
Present thesis, the single-layer forward problem	<i>This work</i>	Rel. ℓ^2	5.5×10^{-5}
Present thesis, the three-layer forward problem	<i>This work</i>	Rel. ℓ^2	1.8×10^{-4}
Present thesis, the inverse problem	<i>This work</i>	Param. error (1-layer)	$< 1.1\%$
Present thesis, the bioheat problem	<i>This work</i>	Rel. ℓ^2	1.04×10^{-1}

2.4.3 Current Limitations and Open Problems

Despite rapid progress, the PINN community acknowledges several unresolved challenges:

1. **Training instability:** PINN loss landscapes are non-convex; failure modes (spectral bias, gradient pathology) are still incompletely understood.
2. **Scalability to 3D:** the number of collocation points needed grows polynomially with dimension; 3D time-dependent problems with $\sim 10^6$ points remain computationally demanding.
3. **Sharp discontinuities:** the smooth sigmoid approximation used in this thesis (and widely in the literature) introduces a finite transition width; methods for exact interface handling (immersed boundary PINNs) are an active research topic.
4. **Theoretical guarantees:** convergence proofs and error bounds for PINNs exist only for specific problem classes; general *a priori* estimates remain an open mathematical question.
5. **Computational cost vs. FEM:** for *forward-only* problems with simple geometry, FEM is still faster than PINN training. PINNs are most advantageous for inverse, data-assimilation, or parametric multi-query scenarios.

2.5 Justification for the Use of PINNs in This Thesis

The present thesis addresses four distinct but interrelated problems in medical optics and thermal therapy: forward solution of the photon diffusion equation in homogeneous tissue (P1) and in a heterogeneous three-layer geometry (P2), inverse recovery of optical coefficients from sparse and noisy fluence measurements (P3), and the coupled solution of the Pennes bioheat equation for laser hyperthermia of a buried tumour with spatially variable perfusion (P4). These four problems collectively constitute an ideal use case for PINNs for three reasons.

Forward problems. The optical and thermal forward solutions (P1, P2, P4) benefit from the mesh-free, analytically differentiable representation that a PINN provides. The same architecture handles piecewise-constant coefficients (after sigmoid smoothing) and two-dimensional space-time domains without re-meshing, and the output transform $T = T_{\text{body}} + T_{\text{scale}} y_{NN}$ enforces the initial condition exactly. Compared with FDM, which requires careful CFL stability analysis and a separate solver for each new geometry, the PINN trades a one-time training cost (~ 5 minutes on a T4 GPU for P4) for a continuous, queryable solution.

Inverse problem. P3 is where PINNs offer the largest qualitative advantage. The unknown coefficients μ_a and μ'_s are declared as trainable variables and updated jointly with the network weights by the same Adam + L-BFGS optimiser; no separate adjoint code or iterative inversion loop is required. The PDE residual acts as an implicit Tikhonov regulariser,

attenuating the propagation of measurement noise to the recovered parameters (sub-unity propagation slopes in our sensitivity study).

Why not traditional methods. A direct FDM-only pipeline would require independent implementations for each study: a tridiagonal solver for P1, a flux-aware variant for P2, a Levenberg–Marquardt loop wrapped around FDM for P3, and an explicit conservative time-stepping scheme for P4. The PINN framework absorbs all four into a single DeepXDE codebase, which is decisive at the master-thesis scale and aligned with the trend toward reproducible scientific machine learning. Crucially, FDM remains the reference solver against which every PINN result in this thesis is validated; the proposed approach is therefore *complementary* to classical methods rather than a replacement.

2.6 Chapter Summary

PINNs embed physical laws directly into the neural-network loss function via automatic differentiation, producing solutions that are simultaneously data-informed and physics-constrained. The modern framework, formalised by Raissi, Perdikaris, and Karniadakis in 2019 [14], addresses five critical challenges of medical-physics modelling — heterogeneous tissue properties, sparse data, ill-posed inverse problems, coupled multi-physics, and real-time dosimetry requirements — simultaneously. State-of-the-art PINNs reach 10^{-5} – 10^{-3} relative errors for forward problems and below-noise parameter recovery for inverse problems, accuracy levels that the present thesis matches or improves upon (Chapter 4). Chapter 3 now presents the detailed methodology — PDE formulations, network architectures, and training strategies — that implements this vision across the five studies of this work.

Chapter 3

Methodology and Resources

This chapter presents the methodological framework adopted in this thesis to model the propagation of light in biological tissues using Physics-Informed Neural Networks (PINNs). The optical work is structured as three progressive PINN studies, each building on the previous: (i) a forward PINN for a single-layer homogeneous tissue, (ii) a forward PINN for a three-layer heterogeneous tissue, and (iii) an inverse PINN for recovering optical properties from noisy measurements. The two thermal studies — the coupled bioheat PINN and the photothermal-therapy exploration — share the same framework and are presented together with their results in Chapter 5. All implementations were carried out in Python with GPU acceleration.

3.1 Physical Problem Formulation

3.1.1 The Photon Diffusion Equation

Light transport in biological tissue is governed, under the diffusion approximation, by the steady-state photon diffusion equation (PDE) [7, 10]:

$$-\nabla \cdot [D(\mathbf{r}) \nabla \phi(\mathbf{r})] + \mu_a(\mathbf{r}) \phi(\mathbf{r}) = S(\mathbf{r}), \quad (3.1)$$

where $\phi(\mathbf{r})$ is the photon fluence rate (W cm^{-2}), μ_a is the absorption coefficient (cm^{-1}), μ'_s is the reduced scattering coefficient (cm^{-1}), $D = [3(\mu_a + \mu'_s)]^{-1}$ is the diffusion coefficient (cm), and $S(\mathbf{r})$ is the source term (W cm^{-3}).

One-Dimensional Single-Layer Case

For a slab geometry with constant optical properties, Eq. (3.1) reduces to the one-dimensional form:

$$-D \frac{d^2 \phi(z)}{dz^2} + \mu_a \phi(z) = S(z), \quad z \in [0, L], \quad (3.2)$$

with Dirichlet boundary conditions:

$$\phi(0) = \phi_0 = 1.0 \text{ (W cm}^{-2}\text{)}, \quad \phi(L) = 0. \quad (3.3)$$

The source $S(z)$ models a focused laser beam as a Gaussian distribution:

$$S(z) = S_0 \exp\left[-\left(\frac{z - z_0}{\sigma_{\text{src}}}\right)^2\right], \quad (3.4)$$

where S_0 is the source amplitude, z_0 is the peak depth, and σ_{src} is the beam width.

One-Dimensional Multi-Layer Case (Conservative Form)

When the tissue is stratified into layers with distinct optical properties, D and μ_a become piecewise functions of depth. The PDE must be written in its conservative form to ensure flux continuity at layer interfaces:

$$-\frac{d}{dz}\left[D(z) \frac{d\phi}{dz}\right] + \mu_a(z) \phi(z) = S(z). \quad (3.5)$$

Expanding the first term yields the explicit form used in the numerical implementation:

$$-D(z) \frac{d^2\phi}{dz^2} - \frac{dD}{dz} \frac{d\phi}{dz} + \mu_a(z) \phi(z) = S(z). \quad (3.6)$$

The second term, $-\frac{dD}{dz} \frac{d\phi}{dz}$, is the *interface correction*; neglecting it introduces errors of approximately 17% near tissue interfaces and must be retained for physical accuracy.

3.1.2 Tissue Optical Properties

The three-layer tissue model consists of skin, subcutaneous fat, and muscle, with optical parameters taken from the literature [10, 20]. These are summarized in Table 3.1.

Table 3.1: Optical properties of the three-layer tissue model (visible/NIR range).

Layer	Domain (cm)	μ_a (cm ⁻¹)	μ'_s (cm ⁻¹)	D (cm)
Skin	[0.0, 0.2]	0.200	20.0	0.01650
Fat	[0.2, 0.7]	0.050	12.0	0.02766
Muscle	[0.7, 1.5]	0.500	8.0	0.03922

To maintain the differentiability required by automatic differentiation inside the PINN, the sharp discontinuities in $D(z)$ and $\mu_a(z)$ are replaced by smooth sigmoid approximations:

$$D(z) = D_1 + (D_2 - D_1) \sigma[k(z - z_1)] + (D_3 - D_2) \sigma[k(z - z_2)], \quad (3.7)$$

where $\sigma(x) = (1 + e^{-x})^{-1}$ is the sigmoid function and $k = 100 \text{ cm}^{-1}$ controls the sharpness of the transition at interfaces $z_1 = 0.2 \text{ cm}$ and $z_2 = 0.7 \text{ cm}$. An analogous expression is used for $\mu_a(z)$.

3.2 Physics-Informed Neural Network Framework

3.2.1 Theoretical Foundation

A PINN approximates the solution to a PDE by embedding the differential equation directly into the loss function of a neural network [14]. Let the network $\mathcal{N}$ map a spatial coordinate $z \in \mathbb{R}$ to the predicted fluence rate $\hat{\phi}(z; \boldsymbol{\theta})$, where $\boldsymbol{\theta}$ are the trainable parameters (weights and biases). The composite loss function is:

$$\mathcal{L}(\boldsymbol{\theta}) = \lambda_{\text{PDE}} \mathcal{L}_{\text{PDE}} + \lambda_{\text{BC}} \mathcal{L}_{\text{BC}} \left[+ \lambda_{\text{data}} \mathcal{L}_{\text{data}} \right], \quad (3.8)$$

where the individual terms are:

$$\mathcal{L}_{\text{PDE}} = \frac{1}{N_r} \sum_{i=1}^{N_r} \left| \mathcal{F}[\hat{\phi}](z_i) \right|^2, \quad (3.9)$$

$$\mathcal{L}_{\text{BC}} = \left| \hat{\phi}(0) - \phi_0 \right|^2 + \left| \hat{\phi}(L) \right|^2, \quad (3.10)$$

$$\mathcal{L}_{\text{data}} = \frac{1}{N_d} \sum_{j=1}^{N_d} \left| \hat{\phi}(z_j) - \phi_j^{\text{obs}} \right|^2. \quad (3.11)$$

Here $\mathcal{F}[\hat{\phi}]$ denotes the PDE residual evaluated at collocation points $\{z_i\}_{i=1}^{N_r}$; ϕ_j^{obs} are experimental measurements used only in the inverse problem; and λ_{PDE} , λ_{BC} , λ_{data} are loss-weighting hyperparameters. Partial derivatives are computed by *automatic differentiation* (AD), which is exact to machine precision and avoids the numerical diffusion introduced by finite differences.

3.2.2 Network Architecture

A fully connected feedforward network (FNN) was used in all three studies. The architecture is described by:

$$\mathbf{a}^{(l)} = \tanh(\mathbf{W}^{(l)} \mathbf{a}^{(l-1)} + \mathbf{b}^{(l)}), \quad l = 1, \dots, n_h, \quad (3.12)$$

with a linear output layer. The hyperbolic tangent activation is chosen because it is infinitely differentiable (required by the second-order PDE) and avoids the vanishing-gradient issues of sigmoid for deeper networks. Weights are initialized with the Glorot (Xavier) uniform scheme [36]. The architectures used are listed in Table 3.2.

3.2.3 Training Strategy

A two-phase training strategy was adopted for all studies:

1. **Phase 1 — Adam optimizer:** 10,000–15,000 iterations with a learning rate of 10^{-3} . Adam performs global exploration of the loss landscape and brings the solution close

Table 3.2: Neural network architectures employed in each study.

Study	Architecture	Params
Forward, single-layer	[1, 32, 32, 32, 32, 1]	$\sim 3,553$
Forward, multi-layer	[1, 64, 64, 64, 64, 64, 1]	$\sim 20,609$
Inverse, single/multi	[1, 32, 32, 32, 32, 1]	$\sim 3,553$

to the minimum rapidly.

2. **Phase 2 — L-BFGS:** Applied after Adam convergence. L-BFGS is a quasi-Newton second-order method that exploits curvature information to achieve superlinear convergence near the minimum, typically improving accuracy by 2–3 orders of magnitude.

Loss weights $[\lambda_{\text{PDE}}, \lambda_{\text{BC}}]$ were set to $[1, 100]$ in all forward problems, and to $[1, 100, 100, 100]$ in the inverse problem (PDE, left BC, right BC, observational data), reflecting the importance of boundary and data constraints.

3.3 Forward PINN: Single-Layer Tissue

3.3.1 Objective

The first study solves Eq. (3.2) for a homogeneous tissue slab of thickness $L = 1.0$ cm with optical parameters $\mu_a = 0.1 \text{ cm}^{-1}$, $\mu'_s = 10.0 \text{ cm}^{-1}$ (hence $D = 0.0330$ cm). This constitutes the *forward* problem: all physical parameters are known, and the goal is to reconstruct the fluence rate profile $\phi(z)$.

3.3.2 Implementation

Collocation sampling. The domain $[0, 1]$ is sampled at $N_r = 200$ interior collocation points drawn via Latin Hypercube Sampling (LHS) as implemented in DeepXDE, plus $N_b = 2$ boundary points. A test set of 100 points is reserved for accuracy evaluation.

PDE residual. The residual at each collocation point z_i is:

$$\mathcal{F}[\hat{\phi}](z_i) = -D \hat{\phi}''(z_i) + \mu_a \hat{\phi}(z_i) - S(z_i), \quad (3.13)$$

where $\hat{\phi}''$ is computed by applying AD twice.

Reference solution. A Finite Difference Method (FDM) with central differences on a uniform grid of $N = 500$ nodes provides a reference solution for validation. The matrix system $\mathbf{A}\phi = \mathbf{b}$ is assembled as:

$$A_{i,i-1} = -\frac{D}{\Delta z^2}, \quad A_{i,i} = \frac{2D}{\Delta z^2} + \mu_a, \quad A_{i,i+1} = -\frac{D}{\Delta z^2}, \quad b_i = S(z_i). \quad (3.14)$$

3.3.3 Evaluation Metrics

The PINN solution is assessed using the relative ℓ^2 error:

$$\varepsilon_{\text{rel}} = \frac{\|\hat{\phi} - \phi_{\text{ref}}\|_2}{\|\phi_{\text{ref}}\|_2}, \quad (3.15)$$

and the point-wise absolute error $|\hat{\phi}(z) - \phi_{\text{ref}}(z)|$ plotted against depth.

3.4 Forward PINN: Three-Layer Tissue

3.4.1 Objective

This study extends the single-layer model to a realistic three-layer tissue structure (skin–fat–muscle) with optical properties given in Table 3.1. The domain extends to $L = 1.5$ cm with interfaces at $z_1 = 0.2$ cm and $z_2 = 0.7$ cm.

3.4.2 Conservative PDE and Smooth Coefficient Approximation

The PDE residual is computed using the expanded conservative form (Eq. (3.6)). The derivative of D is obtained analytically from Eq. (3.7):

$$\frac{dD}{dz} = (D_2 - D_1) k \sigma_1(1 - \sigma_1) + (D_3 - D_2) k \sigma_2(1 - \sigma_2), \quad (3.16)$$

where $\sigma_l = \sigma[k(z - z_l)]$. Using the analytical derivative is preferred over differentiating the sigmoid numerically because it avoids compounding AD errors when k is large.

3.4.3 Interface-Aware Collocation

Sharp gradients near interfaces require a targeted collocation strategy. In addition to $N_r = 500$ randomly sampled interior points, 80 anchor points are placed in each ± 0.05 cm neighborhood of each interface (160 anchor points in total). This ensures the PDE residual is adequately penalized in the regions of highest spatial variation.

3.4.4 Reference Solution (Conservative FDM)

The reference FDM scheme uses a cell-centered evaluation of D at the half-nodes $z_{i\pm 1/2}$ to respect the conservative form:

$$A_{i,i-1} = -\frac{D_{i-1/2}}{\Delta z^2}, \quad A_{i,i} = \frac{D_{i+1/2} + D_{i-1/2}}{\Delta z^2} + \mu_a(z_i), \quad A_{i,i+1} = -\frac{D_{i+1/2}}{\Delta z^2}. \quad (3.17)$$

This scheme is second-order accurate and automatically satisfies the flux continuity condition $[D d\phi/dz]$ across interfaces without explicit interface conditions.

3.4.5 Error Analysis

Beyond the global ℓ^2 error, the layer-by-layer maximum absolute error is reported to quantify the interface resolution quality:

$$\varepsilon_{\max}^{(\ell)} = \max_{z \in \Omega_\ell} |\hat{\phi}(z) - \phi_{\text{ref}}(z)|, \quad \ell \in \{\text{skin, fat, muscle}\}. \quad (3.18)$$

3.5 Inverse PINN: Optical Property Recovery

3.5.1 Objective and Medical Motivation

The inverse problem is the clinically most relevant component of this thesis. Rather than computing $\phi(z)$ from known μ_a and μ'_s , the inverse PINN *recovers* the optical properties from sparse, noisy measurements of the fluence rate — a capability that directly mimics optical coherence tomography (OCT) or diffuse optical tomography (DOT) [1, 2]. Tumorous tissues are distinguishable from healthy tissue by their altered optical coefficients: malignant tissue typically exhibits elevated μ_a due to increased vascularization and higher μ'_s due to structural cellular changes.

3.5.2 Inverse Problem Formulation

The optical coefficients μ_a and μ'_s are promoted to *trainable variables* alongside the network weights. This transforms the optimization into a joint problem over $\{\boldsymbol{\theta}, \mu_a, \mu'_s\}$:

$$\min_{\boldsymbol{\theta}, \mu_a, \mu'_s} \mathcal{L}(\boldsymbol{\theta}, \mu_a, \mu'_s) = \lambda_{\text{PDE}} \mathcal{L}_{\text{PDE}} + \lambda_{\text{BC}} \mathcal{L}_{\text{BC}} + \lambda_{\text{data}} \mathcal{L}_{\text{data}}. \quad (3.19)$$

The three loss terms balance three sources of information: the PDE enforces physical consistency; the BCs constrain the boundary values; and the data term forces agreement with $N_d = 25$ sensor measurements.

Single-Layer Inverse Problem

Two parameters (μ_a, μ'_s) are recovered for a single homogeneous layer. Ground-truth values are $\mu_a^* = 0.15 \text{ cm}^{-1}$ and $\mu_s^* = 12.0 \text{ cm}^{-1}$. Initial guesses are $\mu_a^{(0)} = 0.5 \text{ cm}^{-1}$ and $\mu_s^{(0)} = 5.0 \text{ cm}^{-1}$, deliberately far from the true values to test algorithmic robustness.

Three-Layer Inverse Problem

Six parameters $(\mu_a^{(l)}, \mu_s^{(l)})$ for $l = 1, 2, 3$ are recovered simultaneously. This is a significantly harder problem due to the non-convexity of the loss landscape and parameter cross-sensitivity.

3.5.3 Synthetic Measurement Generation

Experimental measurements are synthesized from the FDM reference solution with additive Gaussian noise to simulate detector noise:

$$\phi_j^{\text{obs}} = \phi_{\text{FDM}}(z_j) + \eta |\phi_{\text{FDM}}(z_j)| \xi_j, \quad \xi_j \sim \mathcal{N}(0, 1), \quad (3.20)$$

where $\eta = 0.02$ (2% noise level). This level is representative of signal-to-noise ratios achievable in clinical near-infrared spectroscopy systems [3]. Sensors are placed at $N_d = 25$ uniformly distributed depths in $[0, L]$.

3.5.4 Trainable Variable Tracking

The evolution of the estimated optical parameters during training is logged every 500 iterations using the `VariableValue` callback in DeepXDE. This allows analysis of convergence trajectories and diagnosis of training instabilities.

3.5.5 Recovery Accuracy

Parameter recovery accuracy is quantified by the relative percentage error:

$$\varepsilon_{\mu_a} = \frac{|\hat{\mu}_a - \mu_a^*|}{\mu_a^*} \times 100 \%, \quad \varepsilon_{\mu'_s} = \frac{|\hat{\mu}'_s - \mu'^*_s|}{\mu'^*_s} \times 100 \%. \quad (3.21)$$

3.6 Computational Resources

All experiments were run on a free-tier NVIDIA T4 GPU through Google Colaboratory, using PyTorch [37] as the automatic-differentiation backend, DeepXDE [35] as the PINN framework, and NumPy/SciPy [38, 39] for the FDM reference solvers. This laptop-class, cloud-based setup is deliberately chosen to keep the work reproducible without specialised hardware — a relevant consideration for research in resource-limited settings. Indicative training times range from ~ 15 s for the single-layer forward problem to ~ 5 min for the two-dimensional bioheat problem. Full reproducibility details (library versions, fixed random seeds, and run instructions) are collected in Appendix A.

3.7 Workflow Overview

Figure 3.1 illustrates the complete methodological pipeline shared by all three studies.

3.8 Validation Strategy and Summary

A consistent three-level validation strategy is applied across all studies: (i) *FDM cross-validation*, in which the PINN solution is compared against a dedicated FDM reference




Figure 3.1: Methodological pipeline common to all three PINN studies. The dashed data term $\mathcal{L}_{\text{data}}$ applies only to the inverse problem.

(standard scheme for single-layer, conservative scheme for multi-layer); (ii) *physical consistency checks*, requiring boundary conditions to be satisfied to within machine precision, positivity and monotonic decay of the fluence profile, and flux continuity at layer interfaces; and (iii) *parameter recovery* for the inverse problem, where recovered optical parameters are compared directly to the synthetic ground truth.

The methodology presented in this chapter progresses from a homogeneous forward model to a heterogeneous three-layer model and finally to a clinically motivated inverse model capable of recovering optical tissue properties from sparse, noisy data. All implementations share a common DeepXDE/PyTorch software architecture, a two-phase Adam + L-BFGS training strategy, and the validation procedure above. The complete framework is reproducible on a publicly accessible cloud GPU platform, making it transferable to other medical-physics groups without specialised hardware requirements. The next two chapters present the numerical and physical results obtained with this methodology.

Chapter 4

Optical Results: Forward, Inverse, and Sensitivity Analysis

Preamble

This chapter consolidates the numerical results obtained for the three optical Physics-Informed Neural Network (PINN) studies: the single-layer homogeneous tissue, the three-layer heterogeneous skin-fat-muscle tissue, and the inverse optical-parameter recovery problem. Each study is assessed through three complementary lenses — agreement between the PINN prediction and the finite-difference (FDM) reference, the absolute error distribution across the domain, and the convergence of the training loss — and the chapter is closed by a systematic sensitivity study that probes the robustness of the inverse PINN to measurement noise, sample size, and sampling geometry. All experiments were executed on an NVIDIA T4 GPU through Google Colaboratory, using PyTorch as the automatic-differentiation backend and DeepXDE as the PINN framework.

The chapter is organised in three parts. Part A (Sections 4.1–4.2) presents the two forward problems. Part B (Sections 4.3–4.5) addresses the inverse problem in both single- and multi-layer geometries. Part C (Section 4.6) contains the sensitivity study. The discussion (Section 4.7) and chapter summary (Section 4.8) close the chapter.

4.1 Forward PINN: Single-Layer Homogeneous Tissue

4.1.1 Problem Setup Recap

The single-layer problem solves the steady-state photon-diffusion equation on $z \in [0, 1.0]$ cm with constant optical parameters $\mu_a = 0.1 \text{ cm}^{-1}$ and $\mu'_s = 10.0 \text{ cm}^{-1}$, yielding $D = [3(\mu_a + \mu'_s)]^{-1} = 0.0330 \text{ cm}$. A Gaussian light source centred at $z_0 = 0.1 \text{ cm}$ with width $\sigma_{\text{src}} = 0.05 \text{ cm}$ drives the fluence field. The network architecture is $[1, 32, 32, 32, 32, 1]$ with tanh activation, trained first with Adam for 10,000 iterations and then refined with L-BFGS.

4.1.2 Fluence Rate Profile and Accuracy

Figure 4.1 presents a four-panel summary of the single-layer forward problem results: the PINN prediction overlaid on the FDM reference (top-left), the absolute error on a logarithmic scale (top-right), the Gaussian source $S(z)$ (bottom-left), and the training-loss convergence (bottom-right). The PINN and FDM curves are visually indistinguishable across the entire domain, and the key quantitative metrics are compiled in Table 4.1.




Figure 4.1: Single-layer forward PINN — four-panel summary. (Top-left) PINN vs. FDM fluence profiles; the two curves are visually indistinguishable, with a relative ℓ^2 error of 5.52×10^{-5} . (Top-right) Absolute error on a logarithmic scale. (Bottom-left) Gaussian source function $S(z)$ centred at $z_0 = 0.1$ cm. (Bottom-right) Convergence of the PDE residual and the boundary losses over training.

The fluence profile exhibits the expected bell shape near the surface: it rises sharply from $\phi(0) = 1.0 \text{ W cm}^{-2}$ to a maximum of approximately 2.796 W cm^{-2} at $z \approx 0.129$ cm, coinciding with the Gaussian source peak, and then decays monotonically to zero at the right boundary. The sharp rise in fluence below the surface exceeding the incident intensity is a well-known phenomenon in tissue optics called *subsurface fluence buildup*, driven by intense multiple backward scattering combined with internal reflection at the air–tissue interface, which accumulates photons near the boundary before global attenuation dominates.

Error distribution. The log-scale absolute-error plot reveals errors ranging from below 10^{-7} to a maximum of about $1.68 \times 10^{-4} \text{ W cm}^{-2}$, with the largest deviations concentrated

Table 4.1: Quantitative error metrics for the single-layer forward PINN.

Metric	Value
Relative ℓ^2 error ε_{rel}	5.52×10^{-5}
Peak absolute error ε_{max}	$1.68 \times 10^{-4} \text{ W cm}^{-2}$
Mean absolute error $\langle \varepsilon \rangle$	$7.02 \times 10^{-5} \text{ W cm}^{-2}$
Boundary residual at $z = 0$	$1.67 \times 10^{-6} \text{ W cm}^{-2}$
Boundary residual at $z = L$	$7.51 \times 10^{-6} \text{ W cm}^{-2}$
Fluence maximum ϕ_{max}	$\approx 2.796 \text{ W cm}^{-2}$ at $z \approx 0.129 \text{ cm}$
Network parameters	3,553
Adam iterations	10,000
Training time (T4 GPU)	$\approx 10 \text{ s}$

in the central region of the profile around $z \approx 0.4 \text{ cm}$ and $z \approx 0.7 \text{ cm}$ where the gradient is smallest and the collocation penalty is weakest. The error near the boundaries is several orders of magnitude lower ($\sim 10^{-6}$), confirming that the heavily weighted boundary loss ($\lambda_{\text{BC}} = 100$) enforces the Dirichlet conditions effectively. The characteristic oscillatory pattern visible in the error curve is a hallmark of PINN approximations: spectral-like oscillations between collocation points where the residual is explicitly penalised, with local minima at the training points themselves.

Boundary accuracy. The left boundary residual $|\hat{\phi}(0) - 1.0| = 1.67 \times 10^{-6} \text{ W cm}^{-2}$ and the right boundary residual $|\hat{\phi}(L)| = 7.51 \times 10^{-6} \text{ W cm}^{-2}$ are both at the sub-ppm level relative to the characteristic fluence scale, demonstrating that the loss-weight strategy ($\lambda_{\text{BC}}/\lambda_{\text{PDE}} = 100$) effectively prioritises boundary enforcement.

4.1.3 Training Loss Convergence

The convergence history (Figure 4.1, bottom-right) shows the expected two-phase behaviour: the total loss drops by about two orders of magnitude within the first 2,000 Adam iterations; the heavily weighted boundary losses converge faster than the PDE residual, plateauing near 10^{-7} – 10^{-8} while the PDE loss settles around 10^{-3} ; and the subsequent L-BFGS phase reduces the PDE loss by a further 1–2 orders of magnitude, accounting for the final accuracy reported in Table 4.1. No loss spikes are observed, confirming the stability of the two-phase strategy for this problem.

4.1.4 Physical Validity Assessment

Beyond numerical metrics, the solution is verified against the known physical constraints of photon diffusion in tissue:

- **Non-negativity:** $\hat{\phi}(z) \geq 0$ for all $z \in [0, L]$.

- **Monotonic decay beyond source:** for $z > z_0$, ϕ decreases monotonically, governed by $\mu_{\text{eff}} = \sqrt{3\mu_a(\mu_a + \mu'_s)} \approx 1.74 \text{ cm}^{-1}$.
- **Boundary conditions** satisfied to sub-ppm precision.
- **Source consistency:** the fluence peak occurs at $z \approx 0.129 \text{ cm}$, downstream of the source centroid $z_0 = 0.10 \text{ cm}$, consistent with diffusive broadening.

The single-layer forward problem thus demonstrates that a compact PINN with 3,553 parameters and two-phase training can solve the 1D photon diffusion equation in homogeneous tissue with sub-0.01% accuracy ($\varepsilon_{\text{rel}} = 5.52 \times 10^{-5}$) — validating both the architecture and the hyperparameters as a foundation for the more challenging multi-layer problem.

4.2 Forward PINN: Three-Layer Heterogeneous Tissue

4.2.1 Problem Setup Recap

The three-layer forward problem extends the domain to $L = 1.5 \text{ cm}$ with three distinct tissue layers — skin ($0 \leq z \leq 0.2 \text{ cm}$), fat ($0.2 < z \leq 0.7 \text{ cm}$), and muscle ($0.7 < z \leq 1.5 \text{ cm}$) — whose optical properties are summarised in Table 4.2.

Table 4.2: Optical properties of the three tissue layers.

Layer	$\mu_a \text{ (cm}^{-1}\text{)}$	$\mu'_s \text{ (cm}^{-1}\text{)}$	$D \text{ (cm)}$
Skin	0.20	20.0	0.0165
Fat	0.05	12.0	0.0277
Muscle	0.50	8.0	0.0392

The conservative PDE formulation, with sigmoid-smoothed coefficients ($k = 100$) and the analytic interface term $D'(z) \partial\phi/\partial z$, is used throughout (see Section 3 for the full derivation). Interface-aware collocation places 80 additional anchor points within $\pm 0.05 \text{ cm}$ of each interface.

4.2.2 Fluence Profile and Continuity

Figure 4.2 shows the four-panel summary for the three-layer forward problem. The fluence is continuous across both interfaces, and the slope of ϕ changes by exactly the factor required by the flux-continuity condition $D_1\phi'_1 = D_2\phi'_2$, confirming that the conservative formulation has been correctly enforced.

A zoomed view of the two interfaces (Figure 4.3) confirms visually the slope discontinuity expected from the heterogeneous $D(z)$.




Figure 4.2: Multi-layer PINN solution: (top-left) fluence profile across the three layers; (top-right) $\mu_a(z)$ and $D(z)$ as smoothed step functions; (bottom-left) absolute error; (bottom-right) training loss.




Figure 4.3: Close-up of the two tissue interfaces. The PINN reproduces the slope change required by flux continuity.

4.2.3 The Decisive Role of the Interface Correction Term

The most consequential result of the three-layer forward problem concerns the analytic inclusion of $D'(z)$. Table 4.3 compares the errors with and without this term.

Table 4.3: Effect of the interface correction term $D'(z)$ on multi-layer PINN accuracy.

Formulation	Relative ℓ^2 error	Status
Non-conservative ($D'(z)$ omitted)	~ 0.17 (17%)	Fails
Conservative (with $D'(z)$)	1.83×10^{-4}	Excellent

Omitting $D'(z)$ increases the error by nearly three orders of magnitude (from 1.83×10^{-4} to 0.17). This is not a numerical curiosity but a fundamental requirement of the conservative formulation: at the interface, the discontinuity in $D(z)$ generates an effective *advective* term that must be present in the residual.

4.2.4 Quantitative Comparison: Single-Layer vs Three-Layer

Moving from the single-layer to the three-layer problem raises the relative ℓ^2 error by roughly half an order of magnitude (from 5.52×10^{-5} to 1.83×10^{-4}), as expected from the three distinct coefficients, two interior interfaces, and the flux-continuity requirement. The larger network (20,609 vs 3,553 trainable parameters) and the longer Adam schedule (15,000 vs 10,000 iterations) partially offset this added difficulty, and the global relative error remains below 0.02% — a level of accuracy clinically acceptable for tissue optical characterisation. The full four-study accuracy comparison is collected later in Table 5.5.

4.3 Inverse Problem Setup

4.3.1 From Forward to Inverse

The forward problems above establish that a PINN can reproduce the steady-state fluence field with finite-difference accuracy on both homogeneous and three-layer geometries. The present section inverts this perspective: given a *sparse set of measured* fluence values along the depth axis, the trainable variables of the network are extended to include the optical coefficients themselves, and the same machinery is asked to recover them.

This question — the recovery of μ_a and μ'_s from external optical measurements — underlies a large class of biomedical imaging modalities, including diffuse optical tomography (DOT), functional near-infrared spectroscopy (fNIRS), and quantitative spectroscopy in PDT dosimetry. Classical reconstructions require iterative inversion of a forward model and careful regularisation; PINNs propose to absorb both steps into a single training loop.

4.3.2 Forward Problem and Synthetic Measurements

The inverse experiment uses the same homogeneous-tissue diffusion model as the single-layer forward problem:

$$-D \frac{d^2 \phi}{dz^2} + \mu_a \phi(z) = S(z), \quad z \in [0, L], \quad (4.1)$$

with $D = 1/[3(\mu_a + \mu'_s)]$ and a Gaussian source centred at $z_0 = 0.1$ cm. The ground-truth parameters are $\mu_a^* = 0.15 \text{ cm}^{-1}$ and $\mu'_s{}^* = 12.0 \text{ cm}^{-1}$. A high-resolution FDM solver generates the noise-free fluence $\phi^*(z)$. From this profile, n measurement points $\{(z_i, \phi^*(z_i))\}_{i=1}^n$ are selected and corrupted with multiplicative Gaussian noise:

$$\phi_i^{\text{obs}} = \phi^*(z_i) (1 + \eta \xi_i), \quad \xi_i \sim \mathcal{N}(0, 1), \quad (4.2)$$

where η is the relative noise amplitude. The reference configuration uses $n = 25$ uniformly distributed points and $\eta = 0.02$ (2 % noise).

4.3.3 PINN with Trainable Optical Parameters

The neural-network architecture is identical to that of the single-layer forward problem: $[1, 32, 32, 32, 32, 1]$ with tanh activations and Glorot initialisation. The novelty resides in the loss function, which now combines the PDE residual, the boundary conditions, and a data-fit term:

$$\mathcal{L}(\theta, \mu_a, \mu'_s) = \lambda_{\text{PDE}} \mathcal{L}_{\text{PDE}} + \lambda_{\text{BC}} \mathcal{L}_{\text{BC}} + \lambda_{\text{data}} \mathcal{L}_{\text{data}}, \quad (4.3)$$

with weights $[1, 100, 100, 100]$ for $[\text{PDE}, \text{BC}_L, \text{BC}_R, \text{data}]$ respectively. Crucially, μ_a and μ'_s are declared as `dde.Variable` objects, are initialised *far* from their true values ($\mu_a^{(0)} = 0.5$, $\mu'_s{}^{(0)} = 5.0$), and are updated jointly with the network weights θ by Adam and L-BFGS. The diffusion coefficient D is computed inside the PDE residual at every iteration from the current values of (μ_a, μ'_s) and therefore evolves with them.

The recovered values are read at the end of training and the relative errors are

$$e_p = \frac{|p_{\text{pred}} - p^*|}{p^*} \times 100\%, \quad p \in \{\mu_a, \mu'_s\}. \quad (4.4)$$

4.4 Single-Layer Inverse Recovery

4.4.1 Convergence of the Trainable Variables

Figure 4.4 shows the simultaneous convergence of the network and the two physical parameters during training. Despite the deliberately poor initialisation, the trajectories of μ_a and μ'_s are smooth and monotonic, settling on their target values well before the L-BFGS phase.

4.4.2 Recovery Accuracy under 2% Noise

The numerical recovery is summarised in Table 4.4.




Figure 4.4: Single-layer inverse recovery: recovered fluence profile against the FDM ground truth and the noisy measurements (top-left); convergence histories of the two trainable variables (top-right, bottom-left); and decomposition of the training loss into PDE, BC and data components (bottom-right).

Table 4.4: Single-layer inverse recovery, $n = 25$, $\eta = 2\%$.

Parameter	Initial	True	Predicted	Error
μ_a (cm ⁻¹)	0.50	0.15	0.1501	0.07%
μ'_s (cm ⁻¹)	5.00	12.0	11.8747	1.04%

Both relative errors lie *below* the noise amplitude itself, a behaviour characteristic of physics-constrained inverse problems: the differential operator acts as an implicit regulariser, smoothing out high-frequency components of the noise. This is the same mechanism that gives PINNs their robustness in data-scarce or noisy regimes [11].

4.5 Multi-Layer Inverse Recovery: Six Parameters at Once

4.5.1 Experimental Setup

The natural next step is to repeat the experiment in the three-layer forward geometry with six trainable parameters — (μ_a, μ'_s) for skin, fat, and muscle — initialised at the values reported in Table 4.5.

Table 4.5: Initial guesses and true values for the multi-layer inverse problem.

Layer	μ_a (cm ⁻¹)		μ'_s (cm ⁻¹)	
	Initial	True	Initial	True
Skin	0.50	0.20	10.0	20.0
Fat	0.20	0.05	20.0	12.0
Muscle	0.30	0.50	15.0	8.0

The 35 measurement points are distributed unevenly across layers, with denser sampling in the skin and at the interfaces. The network is enlarged to $[1, 64 \times 5, 1]$ to match the complexity of the three-layer forward problem, and the PDE is implemented in conservative form including the analytical $D'(z)$ term.

4.5.2 Recovery Results

The recovered parameters, after Adam (20,000 iterations) and L-BFGS, are reported in Table 4.6.

Table 4.6: Multi-layer inverse recovery, $n = 35$, $\eta = 2\%$.

Layer / parameter	True	Predicted	Error (%)	Status
Skin μ_a	0.20	0.2865	43.25	failed
Skin μ'_s	20.0	20.521	2.60	excellent
Fat μ_a	0.05	0.0359	28.28	poor
Fat μ'_s	12.0	13.556	12.97	moderate
Muscle μ_a	0.50	0.4362	12.75	moderate
Muscle μ'_s	8.0	9.296	16.20	moderate




Figure 4.5: Multi-layer inverse recovery: fluence profile (top) and convergence trajectories of the six trainable parameters (middle row: μ_a for skin, fat, muscle; bottom row: μ'_s for the same). Note that the absorption coefficients μ_a^{skin} (43% error) and μ_a^{fat} (28% error) are the worst-recovered, reflecting the noise-masking effect on small parameters.

The recovered parameters exhibit a clear pattern: scattering coefficients μ'_s are recovered with markedly better accuracy in the superficial skin layer (2.6%) than in the deeper layers (13–16%), while the absorption coefficients μ_a show the opposite trend, with the largest

error in the skin (43%) and progressively better recovery in deeper layers. The failure mode is therefore not concentrated on a single parameter, but rather on *the small parameter in each layer*: $\mu_a^{\text{skin}} = 0.20$ and $\mu_a^{\text{fat}} = 0.05 \text{ cm}^{-1}$ are the two worst-recovered quantities, while $\mu_s^{\text{skin}} = 20$ is the best.

4.5.3 The Identifiability Gap — A Physical Diagnosis

The asymmetric quality of recovery is not a numerical artefact but a manifestation of an intrinsic property of the inverse problem. The sensitivity of the fluence $\phi(z)$ to a relative perturbation of a coefficient p is bounded by the absolute magnitude of p itself: if $p \rightarrow 0$, then $\partial\phi/\partial p$ becomes very small in absolute terms, and a noise floor $\eta|\phi|$ on the measurements (here $\eta = 2\%$) masks any information that the data could otherwise carry about p .

For the parameters in our multi-layer problem, this leads to two distinct identifiability regimes:

Small absorption coefficients. The values $\mu_a^{\text{skin}} = 0.20$ and $\mu_a^{\text{fat}} = 0.05$ are small compared with the noise-induced perturbations of the local fluence: a 2% noise on ϕ translates into a fractional uncertainty of order $\eta \cdot (\mu'_s/\mu_a) \approx 0.02 \times 100 = 200\%$ on the recovered μ_a when $\mu'_s/\mu_a \gg 1$, before the PDE constraint moderates it. The skin and fat μ_a recoveries (43% and 28% error) are direct manifestations of this phenomenon.

Deep-region asymptotic insensitivity. A second mechanism affects the deep muscle layer. There the fluence approaches zero by several orders of magnitude, and the local form of the diffusion equation reduces to

$$-D \phi''(z) + \mu_a \phi(z) \approx 0, \quad \phi(z) \rightarrow 0, \quad (4.5)$$

which becomes asymptotically insensitive to D : any value of D (hence of μ'_s) satisfies the equation to numerical precision when ϕ itself is negligible. This explains the moderate-to-poor recovery of μ_s^{muscle} (16.2% error), which is worse than the corresponding skin value (2.6%) despite being a larger quantity in absolute terms.

Synthesis. Both mechanisms manifest the same general principle: the identifiability of a coefficient p is controlled by the ratio between the change in observable that a perturbation of p would induce, and the measurement noise floor. The PINN, far from concealing this, *exposes* the limit through the divergent trajectories of the trainable variables of the affected parameters.

This identifiability landscape mirrors a well-known difficulty in diffuse optical imaging [1, 40], where the recovery of optical properties requires either (i) multiple source–detector geometries that probe different combinations of regions, (ii) spectral information at several wavelengths, or (iii) time-domain measurements. In the present 1D, single-source setting none of these regularisations is available, and the partial failure modes are the expected outcome. The qualitative rule extracted from this analysis is summarised in Table 4.7.

Table 4.7: Empirical identifiability rule extracted from the multi-layer recovery experiment.

Local regime	Large coefficient	Small coefficient
Near the source (ϕ large)	excellent	poor (masked by noise)
Intermediate (ϕ moderate)	good	moderate
Deep ($\phi \rightarrow 0$)	moderate	moderate to poor

4.6 Sensitivity Study

To go beyond the single inverse experiment and probe the *robustness* of the recovery, a systematic sensitivity study was conducted on the single-layer problem. Three control variables were varied independently — the number of measurements n , the noise amplitude η , and the spatial distribution of the measurement points — and the inverse PINN was retrained from scratch for each configuration. Fifteen experiments were carried out in total.

4.6.1 Experiment 1: Effect of the Number of Measurements

The noise level was fixed at $\eta = 2\%$ and the inverse PINN was retrained for $n \in \{5, 10, 15, 20, 30, 50\}$ uniformly distributed measurement points. The resulting errors are reported in Table 4.8.

 Table 4.8: Effect of the number of measurements n on parameter recovery (noise: 2 %; distribution: uniform).

n	e_{μ_a} (%)	$e_{\mu'_s}$ (%)
5	2.45	5.44
10	7.63	4.17
15	2.10	1.85
20	1.50	0.48
30	1.95	1.75
50	3.10	1.29

The error curve is non-monotonic. The minimum is reached around $n \approx 20$, and increasing n beyond that value *degrades* the recovery. This counter-intuitive behaviour is consistent with the standard signal-processing observation that dense sampling of a noisy quantity invites over-fitting: as n grows, the data term in Eq. (4.3) dominates the loss, and the PINN follows individual noise realisations rather than the underlying smooth signal. A sweet-spot of $n \in [15, 25]$ provides the best trade-off between coverage and noise filtering. This result reinforces a design principle for DOT instrumentation: *more sensors are not always better* if their noise floor is comparable to the signal magnitude.

4.6.2 Experiment 2: Effect of the Noise Amplitude

Fixing $n = 25$, the noise level was varied across $\eta \in \{0, 1, 2, 5, 10\}\%$. The errors are summarised in Table 4.9.

Table 4.9: Effect of the noise level η on recovery accuracy ($n = 25$, uniform sampling).

η (%)	e_{μ_a} (%)	$e_{\mu'_s}$ (%)
0	0.00	0.01
1	0.45	0.78
2	0.72	1.45
5	1.97	3.66
10	3.85	7.13

A linear regression of the error against the noise level gives the slopes

$$\frac{\Delta e_{\mu_a}}{\Delta \eta} \approx 0.39, \quad \frac{\Delta e_{\mu'_s}}{\Delta \eta} \approx 0.72. \quad (4.6)$$

Both coefficients are smaller than unity, which means that the inverse PINN *attenuates* the propagation of noise from the data to the recovered parameters: a 10% noise on the measurements is reduced to roughly 4% error on μ_a and 7% on μ'_s . This natural filtering is again a consequence of the physical constraint embedded in the loss function. The qualitative interpretation is that the PDE residual acts as a Tikhonov-like prior, penalising unphysical high-frequency oscillations that the data alone would otherwise admit.

4.6.3 Experiment 3: Effect of the Sampling Geometry

With $n = 20$ and $\eta = 2\%$ held fixed, four sampling geometries were tested: *uniform* (equispaced), *near-source* (concentrated in the first half of the domain), *near-boundary* (split between the two boundaries), and *random* (uniformly drawn). The resulting errors are reported in Table 4.10.

Table 4.10: Effect of the sampling geometry on recovery accuracy ($n = 20$, $\eta = 2\%$).

Distribution	e_{μ_a} (%)	$e_{\mu'_s}$ (%)
Uniform	1.50	0.48
Near boundary	0.95	0.59
Near source	2.10	2.82
Random	0.54	0.51

The randomised distribution outperforms the uniform one by roughly a factor of three on μ_a while maintaining comparable accuracy on μ'_s . The mechanism is again related to noise filtering: a uniform sampling correlates spatially with the residual periodicity of the pseudo-random noise generator, whereas the random sampling breaks that correlation. From a clinical standpoint, this result suggests that *jittering* the position of the optical detectors — as long as it is documented — is preferable to placing them on a regular grid.

4.6.4 Summary Figure

The three sensitivity studies are summarised graphically in Figure 4.6.




Figure 4.6: Summary of the sensitivity analysis: effect of n (left), η (centre), and sampling distribution (right).

4.7 Discussion

4.7.1 Accuracy Assessment in the Medical Physics Context

The forward results demonstrate that PINNs can solve the photon-diffusion equation in biological tissue with accuracy competitive with conventional FDM solvers, while offering unique advantages:

No mesh required. The PINN operates on scattered collocation points rather than a structured grid. For the FDM reference used here, 500–1,000 uniformly spaced nodes were required; the PINN achieves equivalent accuracy with only 200–500 interior collocation points.

Differentiability. The PINN solution is an analytic function over the entire domain, enabling exact computation of derived quantities (fluence gradient, radiance, absorbed power density $\mu_a\phi$) without finite-difference post-processing.

Physics embedding. Unlike black-box neural networks trained purely on data, the PINN embeds the PDE directly into the loss function, guaranteeing that the solution obeys the

governing physics at every point in the domain — crucial for applications where conservation laws must be respected.

4.7.2 What the Inverse Experiment Demonstrates

The single-layer experiment (Section 4.4) shows that an unmodified PINN can recover two coupled optical parameters with sub-percent accuracy from noisy fluence measurements, even when initialised more than a factor of two away from the truth. The data-fidelity loss alone would admit a continuum of (μ_a, μ'_s) pairs consistent with the measurements; it is the simultaneous enforcement of the PDE residual that selects the physically meaningful pair.

4.7.3 Limits Exposed by the Multi-Layer Experiment

The multi-layer experiment (Section 4.5) nonetheless shows that the PINN is not a universal black box: when the underlying inverse problem is mathematically ill-posed, no amount of additional training compensates for the missing information. Deep-tissue μ'_s recovery from a single source–detector configuration is fundamentally limited, and the PINN faithfully reflects that limitation in the divergent trajectory of its trainable variable. Two practical mitigations follow:

1. Augmenting the dataset with measurements from additional source locations or wavelengths, so that the deep layers also experience non-negligible fluence.
2. Introducing layer-specific Tikhonov-like regularisers that anchor $\mu'_s{}^{\text{deep}}$ near a prior estimate when the data are uninformative.

4.7.4 Practical Guidance from the Sensitivity Study

The three controlled experiments converge on a coherent set of design rules for PINN-based optical recovery:

- **Number of measurements.** A modest $n \in [15, 25]$ uniform samples suffice; over-sampling a noisy quantity is counter-productive.
- **Noise tolerance.** The recovery is robust up to noise levels typical of clinical optical sensors ($\eta \leq 5\%$), with sub-noise error in the recovered parameters.
- **Sampling pattern.** Randomised sensor placements outperform regular grids of the same size, provided the positions are recorded accurately.

4.7.5 Comparison with Existing PINN Literature

The accuracy achieved here ($\varepsilon_{\text{rel}} \approx 5.5 \times 10^{-5}$ for the single-layer case) is consistent with results reported by Raissi et al. [14] for second-order linear elliptic PDEs and by Lu et al. [35] for the DeepXDE benchmark suite. For the multi-layer problem, the $\sim 1.8 \times 10^{-4}$

relative error is comparable to — and in fact better than — published PINN results for interface-coupled PDEs, which typically achieve 10^{-3} – 10^{-4} for two-domain problems with moderate contrast ratios. The present multi-layer result is notable because the optical-property contrast across interfaces (a factor of 10 in μ_a) is significantly larger than in most published benchmarks. The identifiability analysis aligns with the diffuse-optical-imaging literature [1], in which the limited identifiability of small parameters against noise is a textbook result.

4.8 Chapter Summary

This chapter has presented the optical results of the thesis across three studies:

1. The **single-layer forward problem** achieves a relative ℓ^2 error of 5.52×10^{-5} for the single-layer photon diffusion problem in ≈ 10 s on a T4 GPU, with boundary conditions satisfied to sub-ppm level.
2. The **three-layer forward problem** achieves a global relative error of 1.83×10^{-4} for the three-layer skin–fat–muscle tissue model, with accurate flux continuity at both interfaces. Omitting the $D'(z)$ interface term degrades the error by nearly three orders of magnitude (to $\sim 17\%$), demonstrating that the conservative formulation is non-negotiable.
3. The **inverse problem** recovers μ_a and μ'_s with sub-percent accuracy from 25 noisy measurements on the single-layer problem ($\varepsilon_{\mu_a} = 0.07\%$, $\varepsilon_{\mu'_s} = 1.04\%$), but exposes a clear identifiability gap on the multi-layer problem: the *smallest* coefficient in each layer is recovered with the largest relative error, with the worst case being μ_a^{skin} at 43% error. The mechanism is the masking of small-parameter sensitivity by the measurement noise floor, not the deep-region asymptotic insensitivity that one might have anticipated a priori.
4. The 15-experiment **sensitivity study** indicates that the inverse PINN is robust to noise (sub-unity propagation slopes), that an optimum number of measurements exists in the range $n \in [15, 25]$, and that random sampling is preferable to regular sampling.

These results firmly establish the PINN methodology as a viable and accurate computational tool for photon-diffusion modelling in biological tissues, and provide the validated foundation on which Chapter 5 couples the optical solution to the time-dependent bioheat equation.

Chapter 5

Thermal Results: Bioheat PINN and Photothermal Therapy Exploration

Preamble

This chapter consolidates the thermal half of the thesis. It is organised in two complementary parts. Part A (Sections 5.1–5.8) presents the bioheat problem, the time-dependent coupled optical–thermal problem in which the optical fluence field obtained in Chapter 4 drives the heat source term of the Pennes bioheat equation, and the resulting space–time temperature field $T(z, t)$ is solved by a PINN over the full domain $\Omega = [0, 1] \text{ cm} \times [0, 60] \text{ s}$. The PINN solution is validated against a conservative finite-difference reference and is shown to reproduce the clinically relevant features — peak temperature, therapeutic time window, and spatial selectivity — with engineering accuracy. The most unexpected outcome of Part A is that the thermal selectivity between skin and tumour is essentially unity, despite the four-fold contrast in absorption and the order-of-magnitude contrast in perfusion; a Q -balance argument is given to explain this counter-intuitive result.

Part B (Sections 5.9–5.13) then *uses* the validated model rather than benchmarking it. Three engineering strategies for restoring therapeutic selectivity — surface cooling in four protocols, plasmonic enhancement of the tumour absorption by $\times 100$ (modelling gold-nanoparticle photothermal therapy), and a parametric sweep over μ_a^{tumour} in search of a sweet-spot — are investigated through the same coupled solver, now run in the more economical FDM mode. Each strategy uncovers a physical limit, and the most striking of them — the *self-shielding* of the tumour by aggressive plasmonic enhancement — is identified quantitatively and discussed in the context of clinical PPTT practice.

A unified discussion (Section 5.12) and a chapter summary (Section 5.13) close the chapter and prepare the General Conclusion.

Part A — Bioheat PINN

5.1 Introduction and Clinical Motivation

Laser Interstitial Thermal Therapy (LITT) is a minimally invasive oncological procedure in which a fibre-optic probe delivers near-infrared radiation directly into a tumour, generating localised hyperthermia ($T > 43^\circ\text{C}$) that induces irreversible cell necrosis while sparing the surrounding healthy tissue [41, 42]. The therapeutic window is narrow: temperatures below 43°C are insufficient for sustained cell death, while temperatures above $\sim 55^\circ\text{C}$ in the skin produce thermal burns. Accurate forward simulation of the temperature field $T(z, t)$ is therefore a prerequisite for safe and effective treatment planning.

The bioheat problem extends the PINN framework developed for the three optical studies to the *time-dependent coupled problem*: the optical fluence field $\phi(z)$ (computed by the diffusion PINN) drives the heat source term in the Pennes bioheat equation, which the PINN then solves over the full space-time domain $\Omega = [0, 1] \text{ cm} \times [0, 60] \text{ s}$. This constitutes the first fully PINN-based optical–thermal pipeline for a three-layer tissue model in the present thesis.

5.2 Physical Model and Problem Parameters

5.2.1 Governing Equation: Pennes Bioheat Equation

The temperature field $T(z, t)$ satisfies the 1D Pennes bioheat equation [9]:

$$\rho(z) c(z) \frac{\partial T}{\partial t} = \frac{\partial}{\partial z} \left[k(z) \frac{\partial T}{\partial z} \right] - w_b(z) \rho_b c_b [T - T_a] + Q_{\text{las}}(z, t), \quad (5.1)$$

with the laser heat source:

$$Q_{\text{las}}(z, t) = \mu_a(z) \phi(z) I(t) \xi_I, \quad (5.2)$$

where $\phi(z)$ is the pre-computed photon fluence from the diffusion PINN (normalised), $\xi_I = 10^{-4}$ is a unit conversion factor, and $I(t)$ is the laser pulse function implemented as a smooth sigmoid envelope:

$$I(t) = I_0 \cdot \sigma[p_s(t - t_{\text{on}})] \cdot \sigma[-p_s(t - t_{\text{off}})], \quad \sigma(x) = \frac{1}{1 + e^{-x}}, \quad (5.3)$$

with $I_0 = 1.2 \times 10^5 \text{ W m}^{-2}$, $t_{\text{on}} = 5 \text{ s}$, $t_{\text{off}} = 35 \text{ s}$, and steepness $p_s = 20 \text{ s}^{-1}$.

5.2.2 Tissue Parameters

The three-layer geometry places the tumour below the skin and above the deeper muscle (Table 5.1).

Table 5.1: Physical parameters of the three-layer tissue model used in the bioheat problem.

Layer	Domain (cm)	μ_a (cm ⁻¹)	μ'_s (cm ⁻¹)	k (W m ⁻¹ K ⁻¹)	ρc (MJ m ⁻³ K ⁻¹)	w_b (10 ⁻³ kg m ⁻³ s ⁻¹)
Skin	[0.0, 0.3]	0.20	20.0	0.37	3.70	20.0
Tumour	(0.3, 0.5]	0.80	15.0	0.55	4.20	0.3
Muscle	(0.5, 1.0]	0.30	10.0	0.49	3.60	8.0

Blood: $\rho_b = 1060 \text{ kg m}^{-3}$, $c_b = 3770 \text{ J kg}^{-1}\text{K}^{-1}$, $T_a = 37^\circ\text{C}$.

The tumour layer is assigned a drastically reduced perfusion rate $w_b^{\text{tumour}} = 0.3 \times 10^{-3}$ vs. $w_b^{\text{skin}} = 20.0 \times 10^{-3} \text{ kg m}^{-3}\text{s}^{-1}$ — a ratio of **1:67**. This reflects the chaotic, poorly organised vasculature of malignant tissue (the Warburg effect) and is the primary physiological mechanism that enables selective tumour heating: the tumour cannot dissipate absorbed laser energy through perfusion, while the skin cools rapidly due to its abundant blood supply. This feature was deliberately encoded to test whether the PINN correctly propagates its physical consequences into the simulated temperature field.

5.2.3 PINN Architecture and Training Configuration

The space-time input (z, t) is mapped to temperature $T(z, t)$ by a six-hidden-layer fully connected network ([2, 64, 64, 64, 64, 64, 1], $\tanh$ activation, Glorot initialisation). An output transform enforces the physiological temperature scale:

$$\hat{T}(z, t; \boldsymbol{\theta}) = T_{\text{body}} + T_{\text{scale}} \cdot \mathcal{N}(z, t; \boldsymbol{\theta}), \quad T_{\text{body}} = 37^\circ\text{C}, \quad T_{\text{scale}} = 25^\circ\text{C}. \quad (5.4)$$

This transform constrains the network output to the physiologically plausible range $[37, 62]^\circ\text{C}$, which accelerates convergence by reducing the effective search space. The total loss function is:

$$\mathcal{L} = 100 \mathcal{L}_{\text{PDE}} + 10 \mathcal{L}_{\text{IC}} + 10 \mathcal{L}_{\text{BC-L}} + 10 \mathcal{L}_{\text{BC-R}}, \quad (5.5)$$

with 4,000 domain collocation points, 250 boundary points, 250 initial condition points, and 1,980 interface/temporal anchor points clustered near the tissue boundaries ($z = 0.3, 0.5 \text{ cm}$) and the pulse transitions ($t \approx t_{\text{on}}, t_{\text{off}}$).

5.3 Training Convergence Analysis

Training proceeded in two phases (Table 5.2). Adam ran for 25,000 iterations, reducing the total loss by more than three orders of magnitude to a best value of 6.94×10^{-1} ; the descent is non-monotonic, with oscillations characteristic of the time-dependent problem, in which the optimiser must balance accuracy across the full 60 s window. By component, the initial-condition loss falls fastest, the two Dirichlet boundary losses follow, and the PDE

residual remains the dominant challenge, stabilising around 10^{-1} after Adam.

Table 5.2: Training history summary for the bioheat PINN.

Phase	Iter.	Best step	Best train loss	Time (s)
Adam	25,000	17,500	6.94×10^{-1}	396
L-BFGS	7,582	32,582	2.66×10^{-2}	262
Total	32,582	—	2.66×10^{-2}	658

The subsequent L-BFGS phase improves the loss by a further factor of ~ 50 (to 2.66×10^{-2}) in 7,582 iterations, for a total training cost of ≈ 11 min on a T4 GPU — a one-time investment that then enables sub-millisecond evaluation over the whole (z, t) domain. The FDM reference is comparably fast but is non-differentiable in the optical parameters, and therefore unsuitable for the inverse and parametric studies that motivate this thesis.

5.4 Temperature Field Analysis

5.4.1 Space-Time Temperature Maps

Figure 5.1 presents the complete six-panel summary of the bioheat problem.

The PINN and FDM heat maps (top-left and top-centre panels of Figure 5.1) are visually indistinguishable in their qualitative structure: both show a hot core developing in the tumour region ($0.3 < z < 0.5$ cm) between $t = 5$ and 35 s, with the peak temperature reaching $\sim 55.2^\circ\text{C}$, and then cooling after the laser is switched off at t_{off} . The PINN correctly reproduces four physical features: the spatial localisation of the temperature maximum at $z \approx 0.35$ cm (downshifted from the optical-fluence peak by the higher $\mu_a^{\text{tumour}} = 0.80 \text{ cm}^{-1}$ that concentrates heat deposition in the tumour); the temporal dynamics, with the plateau reached around $t = 30$ s and cooling commencing at $t_{\text{off}} = 35$ s; the Dirichlet boundaries $T(0, t) = T(L, t) = 37^\circ\text{C}$ maintained throughout; and the thermal diffusion broadening of the heat plume that advances into both skin and muscle layers.

Spatial profiles. The middle-left panel shows $T(z)$ at six representative time instants ($t = 2, 5, 15, 35, 45, 55$ s), corresponding to the pre-heating baseline ($T(z, 0) = 37^\circ\text{C}$, confirming the initial condition), the laser onset, the heating plateau, and the post-irradiation cooling phase. At $t = 15$ s the tumour reaches $\approx 50^\circ\text{C}$ and the skin is already elevated ($\approx 46^\circ\text{C}$); at $t = 35$ s the tumour peaks at $\approx 55.2^\circ\text{C}$; and the tumour remains above the therapeutic threshold for a further ~ 12 s after the laser is switched off, extending the effective treatment window.

Temperature evolution at four depths. Table 5.3 reports the peak temperatures and time-above- 43°C at four key depths.




Figure 5.1: Bioheat PINN — six-panel results. **Top row:** $T(z, t)$ heat maps for PINN (left) and FDM reference (centre); pointwise error map (right). **Middle row:** spatial temperature profiles at six representative time instants (left); laser pulse shape $I(t)$ showing the rectangular envelope between $t_{\text{on}} = 5\text{ s}$ and $t_{\text{off}} = 35\text{ s}$ (right). **Bottom row:** temperature evolution versus time at four tissue depths (skin, tumour, near-tumour muscle, deep muscle) (left); training loss convergence (right). Dashed cyan vertical lines mark the tissue interfaces at $z = 0.3$ and 0.5 cm ; dashed green horizontal lines mark t_{on} and t_{off} .

Table 5.3: Peak temperatures and time above 43°C at representative tissue depths.

Location	z (cm)	T_{peak} ($^{\circ}\text{C}$)	$t_{>43}$ (s)	Clinical status
Skin	0.1	≈ 48.5	~ 25	Above burn threshold
Tumour centre	0.4	≈ 55.2	≈ 47	Effective therapy
Muscle	0.6	≈ 41.5	0	Below threshold
Deep muscle	0.85	≈ 38.5	0	Physiological

The tumour temperature rises steeply from $t_{\text{on}} = 5$ s, peaks at 55.2°C at $t \approx 35$ s, and remains above 43°C for a total of 47.2 s — greatly exceeding the 20 s minimum recommended for sustained hyperthermic cell kill [43, 44].

5.5 Therapeutic Efficacy and Selectivity Analysis

Figure 5.2 presents the therapeutic analysis. The peak-temperature profile $T_{\text{max}}(z)$ (left panel) shows that the PINN and FDM are virtually superimposed across the entire domain. The temperature peak at $z \approx 0.35$ cm ($T = 55.1^\circ\text{C}$ for both solutions) lies within the tumour layer ($0.3 < z < 0.5$ cm). However, the 43°C threshold is exceeded from $z = 0$ up to $z \approx 0.55$ cm, indicating that the thermal lesion extends slightly into the proximal skin and muscle — expected for a surface-positioned laser source with diffusive penetration of $\phi(z)$. The 45°C burn threshold is crossed from $z = 0$ up to $z \approx 0.5$ cm, encompassing the entire skin layer and the full tumour — a clinically concerning finding discussed in Section 5.7.




Figure 5.2: Therapeutic analysis (bioheat PINN). **Left:** peak temperature profile $T_{\text{max}}(z)$ showing the broad plateau spanning skin and tumour regions at $T \approx 55^\circ\text{C}$. **Centre:** layer-averaged temperature evolution $\langle T \rangle(t)$ for skin, tumour, and muscle, illustrating the Q -balance effect: skin and tumour reach nearly identical peak temperatures. **Right:** time above the 43°C therapeutic threshold, showing ~ 25 s of effective treatment in the tumour layer.

Layer-averaged dynamics. The centre panel of Figure 5.2 provides the clearest view of the therapy dynamics. The tumour-averaged temperature (red) rises faster and reaches a higher plateau than both skin and muscle, owing to its reduced perfusion ($w_b^{\text{tumour}}/w_b^{\text{skin}} = 1/67$): the tumour peaks at $\approx 54^\circ\text{C}$ and remains above the therapeutic threshold for ≈ 10 s post-off; the skin peaks at $\approx 50^\circ\text{C}$ and cools more rapidly thanks to the high perfusion and the Dirichlet boundary at $z = 0$; and the muscle reaches only $\approx 41^\circ\text{C}$ on average, below the hyperthermic threshold, protected by both distance from the source and adequate perfusion.

Selectivity assessment. The thermal selectivity ratio, defined as $\mathcal{S} = T_{\text{tumour}}^{\text{peak}}/T_{\text{skin}}^{\text{peak}}$, is

$$\mathcal{S} = \frac{55.2^\circ\text{C}}{55.1^\circ\text{C}} \approx 1.00. \quad (5.6)$$

A selectivity of $\mathcal{S} \approx 1.00$ indicates that the tumour and skin peak temperatures are essentially equal — the laser heats both tissues almost identically. This is clinically hazardous: while the tumour is therapeutically heated, the skin simultaneously reaches 55.1°C , well above the burn threshold of $\sim 45^\circ\text{C}$. The simulation correctly identifies this as a dosimetry failure mode arising from the specific geometry (a shallow tumour at $z = 0.3\text{--}0.5\text{ cm}$ combined with the non-negligible skin absorption $\mu_a^{\text{skin}} = 0.20\text{ cm}^{-1}$, which causes substantial heat deposition at the surface despite the higher tumour absorption $\mu_a^{\text{tumour}} = 0.80\text{ cm}^{-1}$). Deeper tumours or alternative irradiation geometries (interstitial fibre placement) would improve selectivity dramatically.

5.6 Error Analysis and Accuracy Assessment

The quantitative accuracy metrics are summarised in Table 5.4.

Table 5.4: Accuracy metrics vs. FDM reference (bioheat PINN).

Metric	PINN	FDM
T_{max} ($^\circ\text{C}$)	55.2	55.2
$T_{\text{max skin}}$ ($^\circ\text{C}$)	55.1	—
$T_{\text{max muscle}}$ ($^\circ\text{C}$)	47.9	—
Max pointwise error ε_{max} ($^\circ\text{C}$)	10.56	
Relative ℓ^2 error ε_{rel}	1.04×10^{-1}	
Time above 43°C (tumour)	47.2 s	
Selectivity $\mathcal{S}$	≈ 1.00	

The relative ℓ^2 error of $\approx 10.4\%$ is substantially higher than those achieved in the two forward problems (5.5×10^{-5} and 1.8×10^{-4} , respectively); the increase is expected and physically motivated by three factors. First, the bioheat problem is genuinely two-dimensional in (z, t) , with a temporal window of 60 s that is $60\times$ larger than the spatial domain, and the PINN must represent a solution that changes qualitatively from rest to active heating to cooling within a single network. Second, the laser pulse transitions from 0 to $1.2 \times 10^5\text{ W m}^{-2}$ in $\sim 0.2\text{ s}$ (sharpness $p_s = 20$), and the PINN must resolve this near-discontinuous edge across the full spatial domain. Third, the error heatmap (Figure 5.1, top-right) reveals that the 10.56°C maximum is spatially localised near the tumour core at intermediate times ($t \approx 20\text{--}40\text{ s}$) — outside this region the errors are significantly smaller, so the global ℓ^2 metric is dominated by this localised hot spot rather than being uniformly high.

Clinical relevance. Despite the $\approx 10\%$ ℓ^2 error, the PINN correctly predicts the peak temperature to within 0.1°C , the therapeutic window (47.2s above 43°C), the qualitative spatial and temporal structure of the temperature field, and the clinical safety warning (skin burn). For treatment-planning decisions — which threshold is exceeded, when, and where — the PINN provides actionable information despite the non-negligible ℓ^2 error. This distinction between global accuracy and clinical relevance is an important practical consideration in medical-physics PINN applications.

5.6.1 Comparison Across the Studies

Table 5.5 places the bioheat error in the context of all four studies.

Table 5.5: Accuracy comparison across all four thesis studies.

PDE	Type	ε_{rel}	Training time
Diffusion	Forward, 1-layer	5.52×10^{-5}	~ 15 s
Diffusion	Forward, 3-layer	1.83×10^{-4}	~ 40 s
Diffusion	Inverse (1-layer)	$\sim 1\%$ (param.)	~ 50 s
Bioheat	Forward, 3-layer	1.04×10^{-1}	~ 300 s

All times on NVIDIA T4 GPU (Google Colab).

The monotonic increase in ε_{rel} from the single-layer forward problem to the bioheat problem reflects the increasing problem complexity: from a 1D steady-state linear PDE (P1) to a 2D transient nonlinear PDE with coupled physics (P4). This scaling is quantitatively consistent with published PINN benchmarks for similar problem transitions [45, 35].

5.7 Clinical Discussion

The simulation satisfies four therapeutic-efficacy criteria for the tumour: the peak temperature (55.2°C) exceeds the 43°C hyperthermic threshold, the duration above threshold (47.2s) exceeds the recommended 20s minimum, the deep muscle is protected ($T_{\text{muscle}}^{\text{max}} = 47.9^\circ\text{C}$, average $\approx 41^\circ\text{C}$), and the tumour remains hot for ~ 10 s after the laser is switched off. However, two safety issues coexist with these therapeutic benefits: the skin peak (55.1°C) exceeds the burn threshold by $\approx 10^\circ\text{C}$, and the selectivity $\mathcal{S} \approx 1.00$ indicates no thermal contrast between the tumour and the skin — a regime that, in clinical practice, would mandate a dose reduction.

Root cause. The low selectivity has three co-occurring physical causes: (i) the tumour is relatively superficial ($z_{\text{ts}} = 0.3\text{ cm}$), so the photon fluence reaching the tumour is still $\approx 70\%$ of its surface maximum; (ii) the skin absorption ($\mu_a^{\text{skin}} = 0.20\text{ cm}^{-1}$) is not negligible, leading to substantial energy deposition before the tumour; and (iii) the 67:1 perfusion contrast between skin and tumour, while significant, is overwhelmed by the large thermal

load deposited at the surface for the chosen irradiance ($I_0 = 1.2 \times 10^5 \text{ W m}^{-2}$). This finding directly motivates the engineering-mitigation study of Part B.

Dosimetry recommendations. The validated PINN can be re-evaluated in milliseconds for new parameter combinations, enabling four immediately actionable strategies: (a) reducing I_0 to $\sim 0.8 \times 10^5 \text{ W m}^{-2}$ to bring the skin below the burn threshold while keeping the tumour above the therapeutic one; (b) targeting tumours at $z \geq 0.5 \text{ cm}$ to reduce the skin fluence to $\sim 40\%$ of its surface value; (c) applying surface cooling (modelled by changing the left Dirichlet boundary from T_{body} to a cryogen temperature); and (d) fractionated irradiation (e.g. two 15 s pulses with a 10 s interval), exploiting the CEM43 isoeffect [43, 44] to maintain thermal dose while allowing skin recovery. The cooling strategy is investigated quantitatively in Section 5.9.

5.8 Limitations

Four limitations of the current implementation are explicitly acknowledged. (i) The $\approx 10.4\%$ relative ℓ^2 error, while clinically informative, is two orders of magnitude higher than the steady-state studies; causal training [33], adaptive collocation, or wider networks would likely improve it. (ii) The 1D slab geometry cannot capture the radial dependence of a real Gaussian-beam applicator; a 2D axisymmetric extension would be needed for patient-specific planning. (iii) The Dirichlet boundary at $T_{\text{body}} = 37^\circ\text{C}$ ignores the finite thermal resistance of underlying anatomy, which a Robin (convective) condition would address. (iv) The optical fluence is pre-computed and held fixed, so temperature-induced changes in μ_a (significant above 60°C through coagulation) are not modelled — a fully coupled (ϕ, T) PINN would close this loop.

Part B — Photothermal Exploration

Section 5.1 closes with an unexpected observation: in the baseline laser-hyperthermia scenario, the peak tumour temperature and the peak skin temperature coincide to within 0.05°C , yielding a thermal selectivity indistinguishable from unity. The chapter attributes this to a Q -balance effect by which the volumetric absorbed power matches between layers. The natural next question is whether this limitation can be *overcome* by engineering interventions.

The present chapter therefore takes a different posture from the preceding ones: instead of validating the model against a reference, it *uses* the validated model to explore three candidate strategies that the medical-physics literature has proposed for boosting therapeutic selectivity:

1. **Surface cooling** of the skin during, before, or after the laser pulse (four protocols),
2. **Plasmonic enhancement** of the tumour absorption by $\times 100$ (modelling Gold Nanoparticles — AuNPs — in PPTT), and

3. A **parametric sweep** over the tumour absorption coefficient to identify the optimum enhancement factor.

For computational efficiency — each protocol or configuration would otherwise require an independent PINN training of ~ 5 minutes on a GPU — the explorations are carried out with the validated FDM solver of Section 5.1 (reference L^2 against the PINN = 1.04×10^{-1}). The *model* is the same; only the solver changes.

5.9 Strategy 1: Surface-Cooling Protocols

5.9.1 Clinical Motivation

Cryogen-spray cooling and water-bolus contact cooling are routinely used in dermatological laser therapy and in hyperthermia treatments of superficial lesions to limit epidermal injury. The intuition is straightforward: by holding the skin surface at a sub-physiological temperature, the heat generated in the dermis is partly removed before it can damage the epidermis. The question we ask is how much *tumour selectivity* this surface cooling can buy.

5.9.2 Four Cooling Protocols

Four time-dependent surface-temperature boundary conditions are compared:

- **A. None (baseline).** $T(0, t) = T_{\text{body}}$ throughout.
- **B. Pre-cooling only.** $T(0, t) = T_{\text{cool}} = 20^\circ\text{C}$ for $t < t_{\text{on}}$, then released back to T_{body} .
- **C. Post-cooling only.** Body temperature during the pulse, cooled to 20°C after t_{off} .
- **D. Pre and post cooling.** Cooled both before and after the pulse, released during it.

The transitions are implemented as smooth sigmoids ($\sigma_c = 15 \text{ s}^{-1}$) so that the boundary condition is differentiable everywhere — a requirement that would matter if the same protocol were used inside a PINN.

5.9.3 Numerical Results

The peak temperatures and the resulting selectivities are summarised in Table 5.6.

Table 5.6: Effect of the four cooling protocols on peak temperatures and selectivity. Baseline optics ($\mu_a^{\text{tumour}} = 0.80 \text{ cm}^{-1}$), single pulse $t \in [5, 35] \text{ s}$, $I_0 = 12 \text{ W cm}^{-2}$.

Protocol	$T_{\text{tumour}}^{\text{max}} (\text{°C})$	$T_{\text{skin}}^{\text{max}} (\text{°C})$	Selectivity
A. None (baseline)	55.2	55.1	1.000
B. Pre-cooling	54.7	54.7	1.000
C. Post-cooling	55.2	55.1	1.000
D. Pre + Post	54.7	54.7	1.000

The selectivity is essentially unchanged across all four protocols. The cooling does reduce the peak temperatures slightly (by less than half a degree), but it does so *symmetrically* for the skin and the tumour. The qualitative information is reinforced by Figure 5.3, which displays the $T(z, t)$ heatmaps side by side.




Figure 5.3: Heatmaps of $T(z, t)$ under the four cooling protocols. The cool spot at the skin surface is visible in protocols B and D before the pulse and in C and D after the pulse, but its influence does not propagate differentially into the tumour region.

5.9.4 Physical Interpretation

The independence of selectivity with respect to the cooling protocol is best understood through the thermal diffusion time. Defining $\alpha = k/(\rho c)$ and the characteristic diffusion time $\tau(z) = z^2/\alpha$, one finds:

$$\tau_{\text{skin}}(z = 0.15) \approx 22.5 \text{ s}, \quad \tau_{\text{tumour}}(z = 0.40) \approx 160 \text{ s}. \quad (5.7)$$

The pulse duration (30 s) is much longer than τ_{skin} : heat injected anywhere in the skin equilibrates with the cooled boundary on a sub-pulse timescale. The same pulse, however, is much shorter than τ_{tumour} : the tumour interior is essentially adiabatic with respect to the surface during the pulse. The cooling therefore acts *symmetrically* on the skin and the tumour, since the conductive path from $z = 0$ to the tumour passes through the skin, and any cooling that removes heat from the skin also removes it from the tumour via the same conduction path.

This is a non-trivial physical statement: *surface cooling cannot, on its own, achieve thermal selectivity* in the present geometry and pulse-length regime. The selectivity has to be sought in the source term, not in the boundary condition.

5.10 Strategy 2: Plasmonic Enhancement (PPTT)

5.10.1 Modelling AuNPs as an Effective Increase of μ_a

Gold nanoparticles deposited preferentially in tumour tissue through the enhanced permeability and retention (EPR) effect can raise the local absorption coefficient at the plasmon-resonance wavelength by 1–2 orders of magnitude [5, 6]. We model this by replacing μ_a^{tumour} in the optical solver with $100 \times \mu_a^{\text{tumour}} = 80 \text{ cm}^{-1}$ while leaving the skin and muscle parameters untouched. The thermal parameters are unchanged: only the optical map changes.

The diffusion solver of the three-layer forward problem is rerun with this new map (at higher spatial resolution to resolve the very thin absorption layer that ensues), and the resulting fluence is fed into the bioheat solver.

5.10.2 Counter-Intuitive Result

The outcome is summarised in Table 5.7 and visualised in Figure 5.4.

Table 5.7: Standard versus plasmonic-enhanced PPTT. Identical pulse parameters; only μ_a^{tumour} differs.

Quantity	Standard	+ AuNP $\times 100$	Change
$\mu_a^{\text{tumour}} \text{ (cm}^{-1}\text{)}$	0.80	80.0	$\times 100$
$\delta \text{ in tumour (mm)}$	1.62	0.07	$\div 23$
$T_{\text{tumour}}^{\text{max}} \text{ (}^\circ\text{C)}$	55.2	70.1	+14.9
$T_{\text{skin}}^{\text{max}} \text{ (}^\circ\text{C)}$	55.1	73.9	+18.8
Selectivity	1.000	0.948	−5.2%




Figure 5.4: Comparison of the standard ($\mu_a^{\text{tumour}} = 0.80 \text{ cm}^{-1}$) and plasmonic-enhanced ($\mu_a^{\text{tumour}} = 80 \text{ cm}^{-1}$, $\times 100$) treatments. **Top row:** $T(z, t)$ heatmaps showing the migration of the hot region from the tumour centre (standard) to the skin–tumour interface (nano). **Middle:** peak temperature profile $T^{\max}(z)$, with the maximum shifted toward the skin under plasmonic enhancement. **Bottom-left:** fluence profile $\phi(z)$ on a logarithmic scale — ϕ collapses by fourteen orders of magnitude across the tumour, illustrating the self-shielding mechanism. **Bottom-right:** bar chart of the key dosimetric metrics.

The plasmonic enhancement *worsens* the selectivity, and it does so for a non-obvious reason. The skin temperature rises by nearly 19 °C, even though the skin optical properties have not been touched.

5.10.3 The Self-Shielding Mechanism

The diagnosis is contained in the second row of Table 5.7: the optical penetration depth

$$\delta = \frac{1}{\sqrt{3\mu_a(\mu_a + \mu'_s)}} \quad (5.8)$$

collapses from 1.62 mm to 0.07 mm = 70 μm . This is *far smaller* than the 2 mm thickness of the modelled tumour. Physically:

1. The light enters the skin, is partially scattered and absorbed, and arrives at the tumour surface ($z = z_{ts} = 0.3 \text{ cm}$) with a still-sizeable fluence.
2. In a layer of thickness $\delta \ll$ tumour thickness, essentially *all* of that remaining light is absorbed.
3. An enormous volumetric heat source is therefore deposited in a thin shell at the skin–tumour interface.
4. The interior of the tumour receives almost no light, contrary to what aggressive plasmonic enhancement was meant to achieve.
5. By thermal conduction, the resulting hot shell diffuses heat both into the tumour (slowly) *and* back into the skin (where conductive transport is favourable), overheating the epidermis.

This is the *self-shielding* mechanism: the tumour, once made too absorbing, blocks the light from reaching its own interior, and the misdirected energy preferentially damages the overlying skin. The same effect has been reported in the PPTT literature [46, 47] but is seldom quantified explicitly in coupled optical-thermal simulations.

5.11 Strategy 3: Sweet-Spot Search

5.11.1 Parametric Sweep

The previous section shows that doing too little ($\mu_a^{\text{tumour}} = 0.80 \text{ cm}^{-1}$) yields no selectivity, and doing too much ($\mu_a^{\text{tumour}} = 80 \text{ cm}^{-1}$) is actively harmful. A monotonic trade-off therefore connects these regimes, and an optimum should exist between them. To find it, the (ϕ, T) pair was recomputed for $\mu_a^{\text{tumour}} \in \{0.8, 2, 5, 10, 20, 40, 80, 160\} \text{ cm}^{-1}$.

5.11.2 Sweet-Spot Results

The results are reported in Table 5.8 and plotted in Figure 5.5.

Table 5.8: Sweet-spot search over μ_a^{tumour} .

μ_a^{tumour}	δ (mm)	$T_{\text{tumour}}^{\text{max}}$	$T_{\text{skin}}^{\text{max}}$	Selectivity	Status
0.8 (baseline)	1.62	55.2	55.2	1.000	no gain
2.0	0.99	59.6	59.4	1.003	sweet spot
5.0	0.55	63.7	63.7	1.001	marginal
10.0	0.37	66.4	66.6	0.998	deteriorating
20.0	0.24	68.5	71.7	0.955	worse than baseline
80.0 (PPTT)	0.07	70.1	73.9	0.948	strong self-shielding
160.0	0.04	70.0	76.3	0.918	severe self-shielding




Figure 5.5: Sweet-spot analysis: peak temperatures (left), selectivity (centre), and optical penetration depth (right) as functions of μ_a^{tumour} . The selectivity peaks at a modest enhancement ($\mu_a^{\text{tumour}} \approx 2 \text{ cm}^{-1}$) before being eroded by self-shielding for larger values. The penetration depth δ crosses the tumour thickness ($\approx 2 \text{ mm}$) near the same value, marking the transition between the deep-penetration regime and the self-shielding regime.

5.11.3 Reading the Sweet-Spot Curve

The selectivity curve in Figure 5.5 (centre) is shallow but unambiguous: it peaks at $\mu_a^{\text{tumour}} \approx 2 \text{ cm}^{-1}$, corresponding to a $\times 2.5$ enhancement over baseline. The peak selectivity (1.003) is modest in absolute terms but it identifies a regime in which:

- the tumour penetration depth $\delta \approx 1 \text{ mm}$ is comparable to the tumour thickness, so that the light reaches the interior;
- the absorbed power in the tumour is significantly higher than in the skin;

- the conductive losses to the skin are not yet catastrophic.

Beyond this sweet-spot, every additional factor of absorption produces less and less tumour heating while adding more and more thermal load to the skin. From a nanomedicine standpoint, the message is that AuNP formulations should be designed for moderate, uniform absorption gains, not for the highest achievable plasmon resonance.

5.12 Discussion

The three explorations of Part B converge on a single message: within the 1D layered geometry studied here, the surface delivery of laser energy is intrinsically selectivity-limited. Surface cooling cannot discriminate between layers because the skin thermal-diffusion time is short relative to the pulse, so any heat removed from the skin is also removed from the tumour through the same conductive path. Plasmonic enhancement cannot be pushed far before self-shielding overheats the very layer it should protect. The same validated model also points to the engineering routes that step *outside* this regime:

1. **Interstitial light delivery** — placing the source $S(z)$ inside the tumour bypasses the skin entirely.
2. **Short-pulse selective photothermolysis** — reducing the pulse below τ_{skin} preserves the optically created temperature differential, the classical paradigm of laser dermatology [4].
3. **Spectral targeting** — a wavelength for which $\mu_a^{\text{tumour}} \gg \mu_a^{\text{skin}}$ without nanoparticles, e.g. via endogenous chromophores or tumour molecular markers.
4. **Distributed nanoheaters** — formulations that spread absorption across the tumour so that δ stays comparable to the tumour thickness at moderate μ_a .

Each can be explored within the same PINN/FDM framework by altering the source term, the pulse profile, or the optical map; none requires a new methodology, only a new physical hypothesis. Two physical findings deserve emphasis. The **Q -balance effect** of Part A is a robust property of surface delivery in the diffusion regime: the lower μ_a of the skin is compensated by the higher fluence near the source, so the volumetric absorbed power — and hence the peak temperature — coincides between skin and tumour, and no perfusion contrast restores selectivity through this geometry. The **self-shielding limit** of Part B is more counter-intuitive: increasing μ_a^{tumour} to large values collapses δ below the tumour thickness and re-routes the absorbed power toward the skin, so selectivity peaks at a modest $\mu_a^{\text{tumour}} \approx 2 \text{ cm}^{-1}$ ($\times 2.5$ enhancement) and degrades beyond it. From a nanomedicine standpoint, this argues for moderate, uniformly distributed absorption gains rather than maximum plasmon resonance.

5.13 Chapter Summary

The thermal half of the thesis consolidates four findings:

1. **The Pennes bioheat PINN works.** A six-hidden-layer tanh network ($\sim 25,000$ parameters), trained with $\lambda_{\text{PDE}} = 100$ on 4,000 collocation points plus interface- and transition-aware anchors, reproduces the FDM temperature field to $L^2 \approx 0.10$ and the peak temperatures to within $\sim 0.1^\circ\text{C}$ in about five minutes on a T4 GPU.
2. **Spatially variable perfusion is captured correctly.** The 67:1 skin-to-tumour perfusion contrast drives the expected heat retention in the hypovascular tumour, even though the resulting selectivity does not exceed unity.
3. **The Q -balance effect quantifies a fundamental limit of surface hyperthermia.** For the baseline optics, $\mu_a^{\text{skin}} \phi^{\text{skin}} \approx \mu_a^{\text{tumour}} \phi^{\text{tumour}} \approx 4.8 \text{ W cm}^{-3}$, and no surface-cooling protocol breaks this balance.
4. **Plasmonic photothermal therapy has a self-shielding limit.** Pushing $\mu_a^{\text{tumour}} \rightarrow 80 \text{ cm}^{-1}$ collapses δ to $\sim 70 \mu\text{m}$, well below the 2-mm tumour thickness, redirecting absorbed power into the skin; only a modest sweet-spot near $\mu_a^{\text{tumour}} \approx 2 \text{ cm}^{-1}$ remains useful.

Beyond the specific numbers, this chapter shows that a validated PINN/FDM framework can serve as a *numerical microscope* for laser–tissue physics, exposing limitations of clinically motivated strategies that are not always explicit in the experimental literature. The General Conclusion synthesises these findings with those of Chapter 4 and outlines the corresponding research agenda.

General Conclusion

This thesis has investigated the application of Physics-Informed Neural Networks (PINNs) to a unified optical–thermal modelling framework for laser–tissue interaction. The starting point was a deliberate choice not to treat PINNs as an alternative numerical solver to be benchmarked against finite differences on a single problem, but rather as a unifying methodology capable of supporting — within one PyTorch/DeepXDE code base — the forward solution of the photon-diffusion equation, the inverse recovery of optical parameters, the time-dependent solution of the Pennes bioheat equation, and the parametric exploration of clinically motivated treatment strategies. The five progressive studies developed across Chapters 3–5 — the two forward photon-diffusion solutions, the inverse and sensitivity analyses, the coupled bioheat solution, and the photothermal exploration — together support the case for that unification.

Several methodological lessons emerge from the implementation work that are, in our view, transferable to comparable medical-physics projects. The most important is that the conservative formulation is non-negotiable in heterogeneous tissues: whenever the diffusion or thermal coefficient varies in space, the gradient term ($D'(z) \partial\phi/\partial z$ in the optical problem, $k'(z) \partial T/\partial z$ in the bioheat problem) must be included analytically, since omitting it raised the L^2 error by roughly two orders of magnitude. Closely related, replacing piecewise-constant coefficients with sigmoid-smoothed transitions (slope $k \approx 50$ – 100) keeps automatic differentiation well-behaved at the interfaces while leaving the physical jumps practically intact, and anchoring a dense subset of collocation points within ± 0.05 cm of each interface (and around the pulse-rise and pulse-fall instants) raises the effective resolution exactly where it matters. For the transient bioheat problem, the output transform $T = T_{\text{body}} + T_{\text{scale}} y_{NN}$ combined with a PDE-heavy loss reweighting ($\lambda_{\text{PDE}} \gg \lambda_{\text{IC,BC}}$) turns an otherwise hard two-dimensional problem into a tractable one within a few minutes on a single GPU. Finally, a two-phase schedule of Adam followed by L-BFGS reliably reduces the final residual by one to three orders of magnitude, whereas Adam alone plateaus too high and L-BFGS alone is too sensitive to the initial parameter cloud. These rules were not formulated in advance but emerged from systematic comparison during the implementation, and they constitute, beyond the academic content of the thesis, a practical guide for subsequent students and for the laboratory group.

Three physical phenomena were identified or quantified that go beyond the standard demonstration of PINN accuracy. The first is the identifiability gap in multi-layer inverse problems: when a coefficient is small compared with the measurement noise floor, its contri-

bution to the local fluence is masked and the inverse problem becomes weakly informative, so that the recovery error scales inversely with the magnitude of each coefficient — the small absorption coefficients of skin ($\mu_a^{\text{skin}} = 0.20$) and fat ($\mu_a^{\text{fat}} = 0.05$) reach 43% and 28% relative error, while the larger scattering coefficients are recovered to within a few percent. The PINN does not conceal this failure mode but exposes it through the noisy trajectories of the affected trainable variables, which is itself useful diagnostic information. The second is the Q -balance effect in surface hyperthermia: the volumetric optical power deposited in the skin coincides almost exactly with that deposited in the tumour despite the four-fold higher tumour absorption, because the lower μ_a of the skin is compensated by the higher local fluence near the source; the resulting thermal selectivity near unity is therefore a physical property of surface delivery in the diffusion regime rather than a numerical anomaly. The third is the self-shielding limit of plasmonic photothermal therapy: increasing μ_a^{tumour} by a factor of 100 to model gold-nanoparticle uptake collapses the optical penetration depth to $\sim 70 \mu\text{m}$, well below the tumour thickness, so that the light is absorbed in a thin shell at the skin–tumour interface and conductive transport delivers heat preferentially back into the skin; a parametric sweep locates only a modest sweet-spot near $\mu_a^{\text{tumour}} \approx 2 \text{ cm}^{-1}$, in agreement with reports advocating moderate plasmonic concentrations.

Several limitations circumscribe the present work and should be lifted in any follow-up. All computations are one-dimensional in space, whereas real tissue is three-dimensional and the conclusions on selectivity and self-shielding would have to be revisited in a geometry with more complex conductive paths. The diffusion approximation itself breaks down close to the source and at interfaces where the radiance is strongly anisotropic, so a Monte Carlo or radiative-transfer reference would tighten the optical baseline. Tissue properties are assumed constant in time, while in reality μ_a , μ'_s , k , ρc and the perfusion rate ω_b all evolve with temperature, coagulation, and damage. No experimental validation has been carried out — all measurements in the inverse experiments are synthetic — and the training was conducted on a single laptop-class GPU within a limited iteration budget, so the absolute accuracy figures are not pushed to the limits of the methodology. These restrictions do not invalidate the qualitative conclusions drawn from the work.

Four directions stand out as natural continuations. The identifiability gap could be partially closed by adding measurements from multiple source–detector configurations or wavelengths, moving toward realistic diffuse-optical-tomography reconstructions. Replacing the fixed thermal parameters by Arrhenius-driven evolutions of k , ρc , and ω_b , together with a damage integral $\Omega(t) = \int A e^{-E_a/RT} dt$, would turn the bioheat model into a quantitative dosimetry tool, a coupling the PINN framework accommodates without structural change. The one-dimensional model is a methodological testbed that generalises directly to two or three dimensions, at the price of larger networks and tighter collocation budgets, and coupling it with patient-specific anatomy from MRI or CT would bring it closer to clinical relevance. Finally, the self-shielding limit motivates quantitative experiments on interstitial light delivery, short-pulse selective photothermolysis, spectrally targeted absorbers, and uniformly distributed nano-heaters, each of which can be implemented within the same coupled framework by altering the source term, the pulse profile, or the optical map. In

all of these directions, the reproducible code, notebooks, and figures accompanying this manuscript are intended to let other students begin where this work has stopped.

References

- [1] Simon R Arridge. “Optical tomography in medical imaging.” In: *Inverse Problems* 15.2 (1999), R41.
- [2] David A Boas et al. “Imaging the body with diffuse optical tomography.” In: *IEEE Signal Processing Magazine* 18.6 (2001), pp. 57–75.
- [3] M. Ferrari and V. Quaresima. “A brief review on the history of human functional near-infrared spectroscopy (fNIRS) development and fields of application.” In: *NeuroImage* 63 (2012), pp. 921–935.
- [4] R Rox Anderson and John A Parrish. “Selective photothermolysis: precise microsurgery by selective absorption of pulsed radiation.” In: *Science* 220.4596 (1983), pp. 524–527.
- [5] Xiaohua Huang et al. “Plasmonic photothermal therapy (PPTT) using gold nanoparticles.” In: *Lasers in Medical Science* 23.3 (2008), pp. 217–228.
- [6] Leon R Hirsch et al. “Nanoshell-mediated near-infrared thermal therapy of tumors under magnetic resonance guidance.” In: *Proceedings of the National Academy of Sciences* 100.23 (2003), pp. 13549–13554.
- [7] A. Ishimaru. *Wave Propagation and Scattering in Random Media*. Academic Press, 1978.
- [8] Steven L Jacques. “Optical properties of biological tissues: a review.” In: *Physics in Medicine & Biology* 58.11 (2013), R37.
- [9] Harry H Pennes. “Analysis of tissue and arterial blood temperatures in the resting human forearm.” In: *Journal of Applied Physiology* 1.2 (1948), pp. 93–122.
- [10] Ashley J Welch and Martin J C van Gemert. *Optical-Thermal Response of Laser-Irradiated Tissue*. 2nd. Springer, 2011.
- [11] George Em Karniadakis et al. “Physics-informed machine learning.” In: *Nature Reviews Physics* 3.6 (2021), pp. 422–440.
- [12] Salvatore Cuomo et al. “Scientific machine learning through physics-informed neural networks: Where we are and what’s next.” In: *Journal of Scientific Computing* 92.3 (2022), p. 88.
- [13] I. E. Lagaris, A. Likas, and D. I. Fotiadis. “Artificial neural networks for solving ordinary and partial differential equations.” In: *IEEE Transactions on Neural Networks* 9.5 (1998), pp. 987–1000.

- [14] 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.” In: *Journal of Computational Physics* 378 (2019), pp. 686–707.
- [15] Markolf H Niemz. *Laser-Tissue Interactions: Fundamentals and Applications*. 3rd. Springer, 2007.
- [16] Francis A Duck. *Physical Properties of Tissues: A Comprehensive Reference Book*. Academic Press, 1990.
- [17] R. C. Haskell et al. “Boundary conditions for the diffusion equation in radiative transfer.” In: *JOSA A* 11 (1994), pp. 2727–2741.
- [18] L. G. Henyey and J. L. Greenstein. “Diffuse radiation in the galaxy.” In: *Astrophysical Journal* 93 (1941), pp. 70–83.
- [19] K. Furutsu and Y. Yamada. “Diffusion approximation for a dissipative random medium and the applications.” In: *Physical Review E* 50 (1994), pp. 3634–3640.
- [20] W.-F. Cheong, S. A. Prahl, and A. J. Welch. “A review of the optical properties of biological tissues.” In: *IEEE Journal of Quantum Electronics* 26.12 (1990), pp. 2166–2185.
- [21] W. S. McCulloch and W. Pitts. “A logical calculus of the ideas immanent in nervous activity.” In: *Bulletin of Mathematical Biophysics* 5 (1943), pp. 115–133.
- [22] K. Hornik, M. Stinchcombe, and H. White. “Multilayer feedforward networks are universal approximators.” In: *Neural Networks* 2 (1989), pp. 359–366.
- [23] M. Dissanayake and N. Phan-Thien. “Neural-network-based approximations for solving partial differential equations.” In: *Communications in Numerical Methods in Engineering* 10 (1994), pp. 195–201.
- [24] A. Krizhevsky, I. Sutskever, and G. E. Hinton. “ImageNet classification with deep convolutional neural networks.” In: *NIPS*. 2012, pp. 1097–1105.
- [25] Yuyao Chen et al. “Physics-informed neural networks for inverse problems in nano-optics and metamaterials.” In: *Optics Express* 28.8 (2020), pp. 11618–11633.
- [26] Shengze Cai et al. “Physics-informed neural networks for heat transfer problems.” In: *Journal of Heat Transfer* 143.6 (2021), p. 060801.
- [27] Y. Khoo, J. Lu, and L. Ying. “Solving parametric PDE problems with artificial neural networks.” In: *European Journal of Applied Mathematics* 32 (2021), pp. 421–435.
- [28] M. Raissi, A. Yazdani, and G. E. Karniadakis. “Hidden fluid mechanics: Learning velocity and pressure fields from flow visualizations.” In: *Science* 367 (2020), pp. 1026–1030.
- [29] Salah A Faroughi et al. “Physics-guided, physics-informed, and physics-encoded neural networks and operators in scientific computing: Fluid and solid mechanics.” In: *Journal of Computing and Information Science in Engineering* 24.4 (2024), p. 040802.

- [30] S. Wang, X. Yu, and P. Perdikaris. “When and why PINNs fail to train: A neural tangent kernel perspective.” In: *Journal of Computational Physics* 449 (2022), p. 110768.
- [31] A. D. Jagtap and G. E. Karniadakis. “Extended physics-informed neural networks (XPINNs).” In: *Communications in Computational Physics* 28 (2020), pp. 2002–2041.
- [32] M. Tancik et al. “Fourier features let networks learn high frequency functions in low dimensional domains.” In: *NeurIPS*. 2020.
- [33] S. Wang, S. Sankaran, and P. Perdikaris. “Respecting causality is all you need for training physics-informed neural networks.” In: *arXiv preprint arXiv:2203.07404* (2022).
- [34] S. Markidis. “The old and the new: Can physics-informed deep-learning replace traditional linear solvers?” In: *Frontiers in Big Data* 4 (2021).
- [35] Lu Lu et al. “DeepXDE: A deep learning library for solving differential equations.” In: *SIAM Review* 63.1 (2021), pp. 208–228.
- [36] X. Glorot and Y. Bengio. “Understanding the difficulty of training deep feedforward neural networks.” In: *AISTATS*. 2010, pp. 249–256.
- [37] Adam Paszke et al. *PyTorch: An imperative style, high-performance deep learning library*. Advances in Neural Information Processing Systems 32. 2019.
- [38] Charles R Harris et al. “Array programming with NumPy.” In: *Nature* 585.7825 (2020), pp. 357–362.
- [39] Pauli Virtanen et al. “SciPy 1.0: fundamental algorithms for scientific computing in Python.” In: *Nature Methods* 17.3 (2020), pp. 261–272.
- [40] Hamid Dehghani et al. “Near infrared optical tomography using NIRFAST: Algorithm for numerical model and image reconstruction.” In: *Communications in Numerical Methods in Engineering* 25.6 (2009), pp. 711–732.
- [41] Hans-Joachim Schwarzmaier et al. “MR-guided laser-induced interstitial thermotherapy of recurrent glioblastoma multiforme: preliminary results in 16 patients.” In: *European Journal of Radiology* 59.2 (2006), pp. 208–215.
- [42] F. A. Jolesz, ed. *Intraoperative Imaging and Image-Guided Therapy*. Springer, 2014.
- [43] W. C. Dewey. “Arrhenius relationships from the molecule and cell to the clinic.” In: *International Journal of Hyperthermia* 10.4 (1994), pp. 457–483.
- [44] Gerard C van Rhoon et al. “CEM43 thermal isoeffect dose threshold for cell death in tissue: A revisit.” In: *European Radiology* 23.8 (2013), pp. 2215–2227.
- [45] S. Cai et al. “Physics-informed neural networks (PINNs) for fluid mechanics: A review.” In: *Acta Mechanica Sinica* 37 (2022), pp. 1727–1738.
- [46] Moustafa R K Ali, Yue Wu, and Mostafa A El-Sayed. “Gold-nanoparticle-assisted plasmonic photothermal therapy advances toward clinical application.” In: *The Journal of Physical Chemistry C* 123.25 (2019), pp. 15375–15393.

- [47] Prashant K Jain et al. “Noble metals on the nanoscale: optical and photothermal properties and some applications in imaging, sensing, biology, and medicine.” In: *Accounts of Chemical Research* 41.12 (2008), pp. 1578–1586.

Appendix A

Selected Code Listings

This appendix collects the most representative code excerpts from the four studies. The full notebooks and Python scripts are bundled with the manuscript as supplementary material (see Section A.6). All listings use Python 3 with DeepXDE 1.x and the PyTorch backend.

A.1 Forward PINN: Single-Layer Diffusion

The core of the single-layer forward solver consists of three elements: the PDE residual, the geometry/boundary specification, and the two-phase training schedule.

```
1 import deepxde as dde
2 import numpy as np
3 import torch
4
5 # Physical parameters
6 mu_a, mu_s = 0.10, 10.0
7 D = 1.0 / (3.0 * (mu_a + mu_s))
8 L = 1.0
9
10 # Gaussian source
11 def source(z):
12     return 10.0 * dde.backend.exp(-((z - 0.1) / 0.05)**2)
13
14 # PDE: -D phi'' + mu_a phi = S(z)
15 def pde(z, phi):
16     d2 = dde.grad.hessian(phi, z)
17     return -D * d2 + mu_a * phi - source(z)
18
19 # Geometry and BCs
20 geom = dde.geometry.Interval(0, L)
21 bc_l = dde.icbc.DirichletBC(
22     geom, lambda z: 1.0,
23     lambda z, on_b: on_b and np.isclose(z[0], 0))
24 bc_r = dde.icbc.DirichletBC(
25     geom, lambda z: 0.0,
26     lambda z, on_b: on_b and np.isclose(z[0], L))
```

```

27
28 data = dde.data.PDE(geom, pde, [bc_l, bc_r],
29                     num_domain=200, num_boundary=2)
30
31 # Network and two-stage training
32 net = dde.nn.FNN([1] + [32]*4 + [1], "tanh", "Glorot uniform")
33 model = dde.Model(data, net)
34
35 model.compile("adam", lr=1e-3, loss_weights=[1, 100, 100])
36 model.train(iterations=10000)
37
38 model.compile("L-BFGS")
39 model.train()

```

Listing A.1: Forward PINN core — the single-layer forward problem.

A.2 Conservative Multi-Layer Formulation

The decisive snippet of the three-layer forward problem is the conservative PDE that includes the analytical derivative $D'(z)$ of the sigmoid-smoothed diffusion coefficient.

```

1 def D_func(z):
2     bkd = dde.backend
3     k = 100.0
4     s1 = 1.0 / (1.0 + bkd.exp(-k * (z - z1)))
5     s2 = 1.0 / (1.0 + bkd.exp(-k * (z - z2)))
6     return D_skin + (D_fat - D_skin)*s1 + (D_muscle - D_fat)*s2
7
8 def D_derivative(z):
9     # ds/dz = k * s * (1 - s) for each sigmoid
10    bkd = dde.backend
11    k = 100.0
12    s1 = 1.0 / (1.0 + bkd.exp(-k * (z - z1)))
13    s2 = 1.0 / (1.0 + bkd.exp(-k * (z - z2)))
14    ds1 = k * s1 * (1 - s1)
15    ds2 = k * s2 * (1 - s2)
16    return (D_fat - D_skin)*ds1 + (D_muscle - D_fat)*ds2
17
18 def pde_conservative(z, phi):
19     dphi = dde.grad.jacobian(phi, z)
20     d2phi = dde.grad.hessian(phi, z)
21     return (-D_func(z)*d2phi
22            - D_derivative(z)*dphi
23            + mu_a_func(z)*phi - source(z))

```

Listing A.2: Conservative multi-layer PDE with analytic $D'(z)$ — the three-layer forward problem.

A.3 Inverse PINN with Trainable Optical Variables

The inverse problem promotes μ_a and μ'_s to trainable variables and adds a measurement-data term to the loss.

```

1 mu_a_var = dde.Variable(0.5)      # poor initial guess
2 mu_s_var = dde.Variable(5.0)
3
4 def pde_inv(z, phi):
5     d2 = dde.grad.hessian(phi, z)
6     D_var = 1.0 / (3.0 * (mu_a_var + mu_s_var))
7     return -D_var * d2 + mu_a_var * phi - source(z)
8
9 # Synthetic noisy data
10 z_obs = z_grid_subsample
11 phi_obs = phi_true_subsample * (1 + 0.02*np.random.randn(...))
12 observe = dde.icbc.PointSetBC(z_obs, phi_obs.reshape(-1,1),
13                               component=0)
14
15 data = dde.data.PDE(geom, pde_inv, [bc_l, bc_r, observe],
16                    num_domain=200, num_boundary=2,
17                    anchors=z_obs)
18
19 model.compile("adam", lr=1e-3,
20              external_trainable_variables=[mu_a_var, mu_s_var],
21              loss_weights=[1, 100, 100, 100])
22 variable_logger = dde.callbacks.VariableValue(
23     [mu_a_var, mu_s_var], period=500,
24     filename="variables.dat")
25 model.train(iterations=15000, callbacks=[variable_logger])

```

Listing A.3: Trainable optical parameters and data-fit term — the inverse problem.

A.4 Bioheat PINN with Output Transform

The bioheat problem introduces the transient bioheat equation, the output transform that normalises the temperature scale, and the perfusion term.

```

1 def pde_bioheat(x, T):
2     z = x[:, 0:1]
3     t = x[:, 1:2]
4     dT_dt = dde.grad.jacobian(T, x, i=0, j=1)
5     dT_dz = dde.grad.jacobian(T, x, i=0, j=0)
6     d2T_dz2 = dde.grad.hessian(T, x, i=0, j=0)
7     Q = mu_t(z) * phi_interp(z) * pulse(t) * I_f
8     perf = perfusion_t(z) * (T - T_arterial)
9     return (rc_t(z)*dT_dt
10            - k_t(z)*d2T_dz2 - kd_t(z)*dT_dz
11            - Q + perf)

```

```

12
13 # Output transform:  $T = T_{\text{body}} + T_{\text{scale}} * y_{\text{NN}}$ 
14 net = dde.nn.FNN([2] + [64]*6 + [1], "tanh", "Glorot uniform")
15 net.apply_output_transform(lambda x, y: 37.0 + 25.0 * y)
16
17 # Loss weights emphasise PDE; IC enforced by transform
18 model = dde.Model(data, net)
19 model.compile("adam", lr=1e-3,
20               loss_weights=[100., 10., 10., 10.])
21 model.train(iterations=25000)
22 model.compile("L-BFGS", loss_weights=[100., 10., 10., 10.])
23 model.train()

```

Listing A.4: Bioheat PINN with output transform and perfusion — the bioheat problem.

A.5 Photothermal Exploration with FDM Solver

For the parametric studies of Chapter 5 the validated FDM solver is preferred for its speed. The core thermal time-stepping loop is reproduced below.

```

1 def fdm_step(T, dt, dz, kv, rcv, muv, phg, perfv, kp,
2               pulse_val, T_surface):
3     # Heat source
4     Q = muv * phg * pulse_val * I_f
5     # Perfusion sink
6     Q_p = perfv * (T - T_arterial)
7     # Conservative flux divergence
8     flux = np.zeros_like(T)
9     flux[1:-1] = (kp[1:]*(T[2:] - T[1:-1])
10                  - kp[:-1]*(T[1:-1] - T[:-2])) / dz**2
11     # Explicit update of interior
12     T[1:-1] += dt * (flux[1:-1] + Q[1:-1] - Q_p[1:-1]) \
13                 / rcv[1:-1]
14     # Boundary conditions
15     T[0] = T_surface
16     T[-1] = T_body
17     return T

```

Listing A.5: Conservative explicit FDM step for the bioheat equation — the photothermal exploration.

A.6 Reproducibility Notes

All notebooks were tested on Google Colab with a T4 GPU (15 GB VRAM) and DeepXDE 1.15.0 / PyTorch 2.x. To reproduce a result:

1. Open the corresponding notebook in Colab.

2. Activate Runtime → Change runtime type → T4 GPU.

3. Run all cells (Runtime → Run all).

Random seeds are fixed throughout (`torch.manual_seed(42)`, `np.random.seed(42)`) so that the results are bit-reproducible on the same hardware.