



**Politecnico
di Torino**

ScuDo

Scuola di Dottorato - Doctoral School
WHAT YOU ARE, TAKES YOU FAR

Doctoral Dissertation

Doctoral Program in Electronic and telecommunication Engineering (38th cycle)

From Morphology to Physiology: Limits of Data-Driven Learning and Physics-Informed Neural Networks for Cuffless Blood Pressure Estimation

By

Federico Delrio

Supervisor(s):

Prof. Eros Pasero, Supervisor

Dr. Vincenzo Randazzo, Co-Supervisor

Doctoral Examination Committee:

Prof. Pasquale Daponte, Referee, Università del Sannio

Prof. Abdelhamid Daamouche, Referee, Boumerdes University

Politecnico di Torino

2026

Declaration

I hereby declare that, the contents and organization of this dissertation constitute my own original work and does not compromise in any way the rights of third parties, including those relating to the security of personal data.

Federico Delrio
2026

* This dissertation is presented in partial fulfillment of the requirements for **Ph.D. degree** in the Graduate School of Politecnico di Torino (ScuDo).

I would like to dedicate this thesis to my loving parents and expecially to the smile of my dear mother, hoping this would make her proud. I know you would been there in the front row, almost crying with joy. Hope this won't make you too bored, tho, it's full of strange words!

Abstract

Recent advances in deep learning have demonstrated remarkable performance in the analysis of cardiovascular biosignals, particularly in tasks involving electrocardiogram (ECG) pattern recognition and intra-subject prediction. Convolutional and recurrent neural networks are able to accurately detect morphological features such as P waves, QRS complexes, and T waves, achieving state-of-the-art results when trained and evaluated on data originating from the same subjects. These capabilities have been further confirmed in real-world scenarios through the development of PulsECG, a custom wearable device designed for simultaneous ECG and photoplethysmogram (PPG) acquisition, where neural networks exhibit strong intra-subject blood pressure estimation performance under controlled conditions.

However, this thesis shows that such performance does not reliably extend to inter-subject and cross-database scenarios. When models are required to generalize across different individuals, their accuracy degrades substantially, revealing a strong dependence on subject-specific signal morphology rather than on underlying physiological mechanisms. This limitation is not confined to blood pressure estimation: similar generalization failures emerge in stress detection tasks, where subjective labels and individual autonomic responses further exacerbate model ambiguity and performance saturation.

Through extensive experimental analysis, this work argues that purely data-driven approaches implicitly learn individual signal signatures instead of causal physiological relationships, leading to over-optimistic results under improper evaluation protocols and to inherent performance ceilings under subject-wise validation. To address this fundamental issue, the thesis introduces Physics-Informed Neural Networks (PINNs) as a principled modeling framework that integrates cardiovascular hemodynamics with neural representations. By embedding one-dimensional blood flow equations, Windkessel boundary conditions, and physiologically constrained

parameters into the learning process, the proposed PINN framework aims to regularize the solution space and reduce degeneracy in the ECG–PPG–blood pressure mapping .

Experimental results on large-scale clinical databases demonstrate that, while current data availability and subject metadata are insufficient to achieve fully generalized and calibration-free blood pressure estimation, physics-informed models produce physiologically consistent arterial pressure waveforms and exhibit improved robustness compared to purely data-driven baselines. This thesis concludes that reliable inter-subject generalization in cuffless blood pressure monitoring is not yet attainable with existing datasets alone. Nevertheless, it establishes a methodological foundation upon which future systems, leveraging ECG and PPG together with easily accessible non-invasive subject information, including anthropometrics (age, sex, height, weight), short-term physiological states (e.g., hydration level, bladder filling), and lifestyle-related factors (such as recent alcohol or caffeine consumption) known to modulate blood pressure dynamics, can build. In this perspective, Physics-Informed Neural Networks emerge as a necessary step toward interpretable, personalized, and generalizable cardiovascular monitoring.

Contents

List of Figures	xv
------------------------	-----------

List of Tables	xxii
-----------------------	-------------

1 Introduction	1
1.1 Motivation	1
1.2 From Biosignals to Physiological Inference	2
1.3 Subjectivity and the Generalization Problem	5
1.4 Research Questions and Objectives	7
1.5 Contributions of the Thesis	9
1.6 Structure of the Thesis	11
2 Physiological Background	14
2.1 Cardiovascular Physiology	14
2.1.1 Electrical Activity of the Heart	14
2.1.2 Mechanical Pressure Generation	16
2.1.3 Arterial Tree and Vascular Properties	18
2.2 Hemodynamics	19
2.2.1 Blood Flow and Fluid Mechanics	19
2.2.2 Windkessel Models	21
2.2.3 Pressure–Area Relationships	23

2.3	Photoplethysmography	25
2.3.1	Optical Modeling and Blood Volume Dependence	25
2.3.2	From Blood Volume to Flow	27
2.3.3	From Flow to Pressure	27
2.4	Autonomic Nervous System and Stress Physiology	28
2.4.1	Sympathetic–Parasympathetic Balance	28
2.4.2	Electrodermal Activity	29
2.4.3	Heart Rate Variability and Respiration	30
2.5	Summary of Key Physiological Challenges	30
3	Dataset and Preprocessing	32
3.1	Overview	32
3.1.1	QT Database	32
3.1.2	MIMIC-III Waveform Database	33
3.1.3	VitalDB / PulseDB	33
3.1.4	Cuff-Less Blood Pressure Estimation Dataset	34
3.1.5	StressID Dataset	35
3.2	Dataset Biases and Structural Limitations	36
4	ECG Morphology and Feature Learning	39
4.1	Problem Formulation and Reference Annotations	39
4.2	Related Work on Automatic ECG Feature Detection	40
4.3	Methodological Overview	42
4.3.1	Dataset and Annotation Framework	42
4.3.2	Baseline Elimination	43
4.3.3	Continuous Wavelet Transform and Scalogram Derivation	44
4.3.4	Data Segmentation	47

4.4	Deep Learning Architecture	48
4.5	Results	49
4.6	Discussion and Conclusion	52
5	Data-Driven Blood Pressure Estimation from ECG and PPG	56
5.1	Cuffless Blood Pressure Estimation: Motivation and Problem Setting	56
5.2	State of the Art on Data-Driven Blood Pressure Estimation	59
5.3	Conv1D BiLSTM for Cuffless Arterial Blood Pressure Estimation from ECG and PPG	62
5.3.1	Problem Formulation and Modeling Strategy	62
5.3.2	Dataset and Signal Preprocessing	63
5.3.3	Database	63
5.3.4	Signal Filtering and Normalization	63
5.3.5	Windowing Strategy	64
5.3.6	Physiological Thresholding and Artifact Pruning	64
5.3.7	PPG Derivatives as Additional Input Channels	65
5.3.8	Network Input–Output Definition	66
5.3.9	SBP/DBP Extraction from ABP for Evaluation	67
5.3.10	Extracted Dataset Statistics and Pressure Distributions	68
5.3.11	Network Architecture	69
5.3.12	Experiments and Results	70
5.3.13	Data split and reproducibility	70
5.3.14	Training configuration and metrics	71
5.3.15	Experiment 1: Baseline configuration (MaxPool, 3 BiLSTM, batch size 64)	71
5.3.16	Experiment 2: MaxPool removal (no MaxPool, 3 BiLSTM, batch size 64)	72

5.3.17	Experiment 3: Reduced recurrent depth (MaxPool, 2 BiLSTM, batch size 64)	73
5.3.18	Experiment 4: Increased batch size (MaxPool, 3 BiLSTM, batch size 256)	74
5.3.19	Qualitative ABP waveform reconstructions	75
5.3.20	Discussion and Conclusions	76
5.4	Wearable-Oriented Cuffless Blood Pressure Estimation with PulsECG	77
5.4.1	Introduction and Motivation	78
5.4.2	State of the Art in Wearable Cardiovascular Devices	79
5.4.3	PulseECG System Architecture	79
5.4.4	ECG PCB	80
5.4.5	PPG PCB and Final Prototype	81
5.4.6	Neural Network for Blood Pressure Prediction: Context and Rationale	81
5.4.7	Dataset Description and Acquisition Protocol	82
5.4.8	Neural Network Structure	83
5.4.9	Results	84
5.4.10	Discussion and Conclusions	88
5.5	Blood Pressure Estimation via Dual-Branch Wide Space–Time Neural Networks, on a Single Subject	91
5.5.1	Database Creation	92
5.5.2	Data Pre-processing	93
5.5.3	Windowing Strategy	94
5.5.4	Blood Pressure Value Distribution	94
5.5.5	Dual-Branch Wide Space–Time Architecture	95
5.5.6	Results	98
5.5.7	Discussion and Conclusion	100

5.6	Analysis of Generalization Failure in Data-Driven BP Estimation . .	102
5.6.1	Cross-Dataset Transfer: MIMIC to PulsECG	102
5.7	Multi-Subject Generalization Study on the PulsECG Database (7 Subjects, 158 Recordings)	114
5.7.1	Dataset overview and windowing protocol	114
5.7.2	Experiment 1: within-subject modeling (leave-one-file-out) .	115
5.7.3	Experiment 2: leave-one-subject-out generalization (dual- branch ECG/PPG network)	118
5.7.4	Experiment 3: adding anthropometrics (age, height, weight)	121
5.7.5	Effect of Anthropometric Information on Cross-Subject Gen- eralization	124
5.7.6	Joint interpretation and implications for the thesis	126
5.8	Synthesis and Motivation for Physics-Informed Modeling	127
6	Physics-Informed Neural Networks for Cuffless Blood Pressure Estima- tion	130
6.1	What are PINNs?	131
6.2	From Data-Driven Models to Physics-Informed Learning	133
6.3	Why Physics Is Needed	134
6.4	Hemodynamic Model Foundations	136
6.4.1	Physical Variables and Learnable Hemodynamic Parameters	136
6.4.2	Mass and Momentum Conservation	138
6.4.3	Pressure–Area Constitutive Laws	140
6.4.4	Windkessel Models	142
6.5	State of the Art on Physics-Informed Neural Networks for Cuffless Blood Pressure Estimation	145
6.6	First Physics-Informed Approach Based on Inflow Proxy and Periph- eral PPG	146

6.6.1	Inflow Approximation from ECG and PPG	148
6.6.2	Physiological Modeling of Vessel Geometry and Mechanical Properties	151
6.7	Database	156
6.7.1	Physics Informed Neural Network Architecture	157
6.7.2	Physics Informed Loss Function	160
6.7.3	Training and Results	163
6.7.4	Discussion and Scope of the First-Generation Model	166
6.7.5	Limitations and Directions for Further Development	166
6.8	Second-Generation PINN: Physiological and Anatomical Priors	168
6.8.1	Motivation for a Revised Framework	168
6.8.2	PPG Link to Blood Flow and Blood Pressure	171
6.8.3	From Beer–Lambert Law to a Linear PPG $\sim V_b$ Model . . .	171
6.8.4	From Blood Volume to Flow: Continuity	173
6.8.5	From Flow to Pressure: Toward $P(t) \propto \text{PPG}(t)$	173
6.8.6	Comparison with Previous PPG–Pressure Relations and Practical Implications	174
6.8.7	Physiological Origin and Role of Geometric and Mechanical Priors	176
6.8.8	Physiological Modeling of Geometry and Wall Mechanics: $A_0(x)$, $h_0(x)$, and $E(x)$	177
6.8.9	Differences with respect to the first-generation architecture .	180
6.8.10	Physics-Informed Loss Function (Second Generation) . . .	181
6.8.11	Training Strategy and Experimental Protocol	186
6.8.12	Cross-database Evaluation on VitalDB	189
6.9	Results Discussion	191
6.10	Conclusion	196

7	Physiological Stress Modeling and Latent-State Learning	199
7.1	From Blood Pressure to Stress: A Common Modeling Challenge . .	199
7.2	Stress as a Physiological and Cognitive Process	200
7.3	The StressID Dataset	201
7.4	Baseline Performance and Initial Objective	201
7.5	A Structural Analogy with Blood Pressure Estimation	202
7.6	Database Statistics and Experimental Protocol	203
7.6.1	Windowing Strategy and Data Representation	203
7.6.2	Subject-Wise Cross-Validation	203
7.6.3	Class Distribution and Label Ambiguity	204
7.7	Single-Signal Modeling	204
7.7.1	Network Architecture	204
7.7.2	Results	205
7.8	Multimodal Physiological Modeling	206
7.8.1	Dual-Signal Fusion	206
7.8.2	Three-Signal Fusion	206
7.9	Label Ambiguity and the Hard Performance Limit	207
7.9.1	Stress as a Latent Physiological State	207
7.9.2	Empirical Evidence from Synthetic Labels	208
7.9.3	Dimensionality Reduction Analysis	208
7.9.4	Interpretation	209
7.10	Cross-Attention Between Physiological Modalities	214
7.10.1	Cross-Attention Mechanism	214
7.10.2	Experimental Results	215
7.10.3	Interpretation	216
7.11	From Stress Classification to Latent-State Modeling	217

7.12	Toward a Physics-Informed Model of Stress	218
7.12.1	Autonomic Nervous System as a Latent Dynamical System	219
7.12.2	Physics-Informed Objective as a Constrained Latent-State Estimation Problem	222
7.13	Discussion and Theoretical Implications	225
7.14	Concluding Remarks and Future Directions	226
8	Discussion, Conclusions, and Future Directions	228
8.1	Unifying Perspective of the Thesis	228
8.1.1	From Direct Mapping to Latent-State Inference	228
8.1.2	Partial Observability, Ill-Posedness, and Identifiability	229
8.2	Interpretation of the Data-Driven Results	231
8.2.1	Intra-Subject Learning as a Practical Upper Bound	231
8.2.2	Inter-Subject Generalization and Morphological Bias	233
8.2.3	Cross-Database Transfer and Measurement Domain Shift	234
8.2.4	Wearable Signals and Real-World Variability	235
8.3	Physics-Informed Modeling: Contributions and Limitations	236
8.3.1	What Physics Adds: Structural Regularization and Physio- logical Coherence	237
8.3.2	What Physics Cannot Resolve Alone	238
8.4	Stress Modeling as a Mirror Problem	239
8.4.1	Limits of Supervision with Subjective Labels	240
8.4.2	Task-Based Labels as a Controlled Proxy	241
8.4.3	Toward Latent Autonomic State Estimation	242
8.5	Methodological Lessons and Best Practices	243
8.5.1	Validation Protocols and Generalization Criteria	244
8.5.2	Quality Control and Failure Mode Analysis	245

8.5.3	Evaluation Metrics Beyond Pointwise Error	245
8.6	Limitations of the Present Work	246
8.6.1	Observability Constraints	246
8.6.2	Dataset and Population Bias	247
8.6.3	Modeling Assumptions and Simplifications	248
8.7	Conclusions	249
8.8	Future Research Directions	251
8.8.1	Integration of Subject-Specific Descriptors	251
8.8.2	Improved Modeling of Hemodynamic Inputs	252
8.8.3	Advances in Physics-Informed Architectures	252
8.8.4	Extensions to Stress and Multimodal Physiological Modeling	253
8.8.5	Toward Real-World Wearable Deployment	253

List of Figures

1.1	High-level pipeline for data-driven cuffless blood pressure estimation from ECG and PPG.	3
1.2	Schematic comparison of evaluation protocols: in-domain (potentially optimistic) versus strict subject-wise/cross-database generalization.	3
1.3	Roadmap of the thesis: progression from signal-level learning to physiology-informed modeling.	3
1.4	Forward (physiology→signals) versus inverse inference (signals→physiology), highlighting the ill-posed nature of cuffless estimation.	4
1.5	Why evaluation protocol matters: random/window splits can leak subject-specific signatures, while subject-wise splits better reflect deployment on unseen individuals.	6
1.6	Link between the research questions, the corresponding objectives, and the thesis chapters where they are addressed.	9
1.7	High-level map of the main thesis contributions and where they are developed in the manuscript.	10
1.8	Compact overview of the thesis structure, highlighting the progression from signal processing to physiology-informed modeling and discussion.	13
2.1	Example of a sinus-rhythm ECG trace with annotated waveform components and fiducial points (P wave, QRS complex, and T wave).	15

4.1	Example of ECG baseline elimination preserving waveform morphology.	44
4.2	Top: segmentation example of ECG with relative label for a QRS complex (red). Bottom: 2D Scalogram obtained through CWT method from the ECG.	47
4.3	Deep Neural Network Architecture	50
4.4	Training loss (blue) and validation loss (orange) for QRS (top left), P wave (top right), and T wave (bottom) networks.	52
4.5	Example of prediction of QRS complexes with the relative original label.	53
4.6	Example of prediction of P wave with the relative original label. . .	53
4.7	Example of prediction of T wave with the relative original label. . .	54
5.1	Example of input vector over time (x -axis) composed, from top to bottom (y -axis), by: PPG, ECG, VPPG, APPG; and then the relative ABP to be predicted.	67
5.2	Fraction of ABP signal (y -axis) over time (x -axis) with: (a) highlighted peaks (SBP, <i>orange stars</i>) and minimum points (DBP, <i>green stars</i>); (b) extracted SBP (mean of peaks) and DBP (mean of minimums).	68
5.3	Distribution of: SBP and DBP extracted values (left); σ of extracted values of SBP, DBP (right).	69
5.4	Structure of the first network.	70
5.5	Error (y -axis) over epochs (x -axis) for the first experiment (batch size = 64, 3 BiLSTM and MaxPool layer): (a) MSE on training (<i>blue line</i>) and validation (<i>red line</i>) sets; (b) MAE on training (<i>blue line</i>) and validation (<i>red line</i>) sets.	72
5.6	Error (y -axis) over epochs (x -axis) for the second experiment (batch size = 64, 3 BiLSTM and no MaxPool layer): (a) MSE on training (<i>blue line</i>) and validation (<i>red line</i>) sets; (b) MAE on training (<i>blue line</i>) and validation (<i>red line</i>) sets.	72

5.7	Comparison of MAE (<i>y-axis</i>) over epochs (<i>x-axis</i>) between the configurations in Fig. 5.5b (with MaxPool) and Fig. 5.6b (without MaxPool).	73
5.8	Error (<i>y-axis</i>) over epochs (<i>x-axis</i>) for the third experiment (batch size = 64, 2 BiLSTM and MaxPool layer) compared to the previous two networks: (a) MSE on training and validation sets; (b) MAE on training and validation sets.	74
5.9	Comparison of MAE (<i>y-axis</i>) over epochs (<i>x-axis</i>): (a) all network configurations (training only); (b) first and last configurations showing training and validation trends.	75
5.10	Examples of predicted ABP (<i>y-axis</i>) over time (<i>x-axis</i>) from the baseline network configuration (batch size = 64, MaxPool and 3 BiLSTM layers). Predictions are shown for random inputs from the test set to illustrate the model's ability to reproduce different waveform morphologies.	76
5.11	3D model of the PulsECG device. Electrodes are visible on the sides, while the ECG and PPG PCBs are stacked to minimize space requirements.	80
5.12	PulsECG block diagram showing the ECG PCB (left) and PPG PCB (right) connected via flat cable.	80
5.13	Final PulsECG prototype enclosed in a 3D-printed case.	81
5.14	PSD comparison of ECG signals acquired by PulsECG (blue) and GE B105 (red).	85
5.15	PSD comparison of PPG signals acquired by PulsECG (blue) and GE B105 (red).	85
5.16	CSP comparison for ECG signals acquired by PulsECG (blue) and GE B105 (red).	86
5.17	CSP comparison for PPG signals acquired by PulsECG (blue) and GE B105 (red).	87
5.18	Time-domain comparison of ECG signals acquired by PulsECG (blue) and GE B105 (red).	87

5.19	Time-domain comparison of PPG signals acquired by PulsECG (blue) and GE B105 (red).	88
5.20	Scatter plot of predicted versus measured SBP values.	88
5.21	Scatter plot of predicted versus measured DBP values.	89
5.22	Distribution of systolic (left) and diastolic (right) blood pressure values in the single-subject dataset.	94
5.23	Dual-branch wide space–time neural network architecture for single-subject blood pressure estimation. ECG and PPG signals are processed independently through convolutional encoders, fused into a wide feature representation, and temporally modeled via stacked BiLSTM layers before final regression of SBP and DBP.	96
5.24	Mean Absolute Error (MAE) for systolic (SBP) and diastolic (DBP) blood pressure predictions across the 10-fold cross-validation. Each point corresponds to the MAE computed on the test set of the respective fold.	99
5.25	Histograms of waveform-derived blood pressure quantities computed from the MIMIC-based ABP segments: (a) SBP, (b) DBP, (c) ESBP, and (d) EDBP.	104
5.26	Training and validation loss over 40 epochs for the MIMIC-trained ABP reconstruction model.	105
5.27	Training and validation MAE over 40 epochs for the MIMIC-trained ABP reconstruction model.	105
5.28	Discrepancies between mean (unweighted) model estimates and external sphygmomanometer values over the 19 PulsECG recordings: $\Delta\text{SBP} = -2.33 \pm 17.85$ mmHg and $\Delta\text{DBP} = -7.67 \pm 7.91$ mmHg. .	106
5.29	External validation discrepancies when restricting the input to ECG and PPG only: $\Delta\text{SBP} = -7.83 \pm 12.32$ mmHg and $\Delta\text{DBP} = -12.02 \pm 7.82$ mmHg.	107
5.30	Time-domain comparison between PPG acquired by PulsECG (blue) and representative PPG segments from the MIMIC-derived dataset (orange).	108

5.31	Time-domain comparison between ECG acquired by PulsECG (blue) and representative ECG segments from the MIMIC-derived dataset (orange).	108
5.32	External validation discrepancies after retraining with structured noise injection (ECG+PPG only): $\Delta\text{SBP} = 5.23 \pm 10.38$ mmHg and $\Delta\text{DBP} = -11.33 \pm 7.07$ mmHg.	109
5.33	Transfer learning attempt (frozen early layers): scatter plot for SBP predictions vs reference values.	110
5.34	Transfer learning attempt (frozen early layers): scatter plot for DBP predictions vs reference values.	111
5.35	Fine-tuning with all layers unfrozen: scatter plot for SBP predictions vs reference values.	111
5.36	Fine-tuning with all layers unfrozen: scatter plot for DBP predictions vs reference values.	112
5.37	Fine-tuning with an expanded external dataset and window-level BP assignment by linear interpolation: SBP scatter plot.	112
5.38	Fine-tuning with an expanded external dataset and window-level BP assignment by linear interpolation: DBP scatter plot.	113
5.39	Experiment 1 (leave-one-file-out): subject 1.	116
5.40	Experiment 1 (leave-one-file-out): subject 2.	116
5.41	Experiment 1 (leave-one-file-out): subject 3.	116
5.42	Experiment 1 (leave-one-file-out): subject 4.	117
5.43	Experiment 1 (leave-one-file-out): subject 5.	117
5.44	Experiment 1 (leave-one-file-out): subject 6.	117
5.45	Experiment 1 (leave-one-file-out): subject 7.	118
5.46	Experiment 2 (LOSO): training MAE example (subject 1 as test). . .	119
5.47	Experiment 2 (LOSO): training MAE example (subject 5 as test). . .	119
5.48	Experiment 2 (LOSO): test MAE trend example (subject 1 as test). .	120
5.49	Experiment 2 (LOSO): test MAE trend example (subject 5 as test). .	120

5.50	Experiment 3 (LOSO + age/height/weight): training example (subject 1 as test).	122
5.51	Experiment 3 (LOSO + age/height/weight): training example (subject 5 as test).	122
5.52	Experiment 3 (LOSO + age/height/weight): test MAE trend example (subject 1 as test).	123
5.53	Experiment 3 (LOSO + age/height/weight): test MAE trend example (subject 5 as test).	123
6.1	Example of a generated inflow proxy $Q_{\text{in}}(t)$ obtained from the logarithmic derivative of the PPG signal modulated by ECG-derived information.	150
6.2	Nominal vessel radius profile $r_0(x)$ along the normalized proximal-to-distal coordinate x , consistent with the exponential diameter tapering in Eq. (6.25).	153
6.3	Nominal spatial profile of the arterial wall Young's modulus $E(x)$ along the normalized coordinate x (cf. Eq. (6.30)).	154
6.4	Nominal spatial profile of arterial wall thickness $h_0(x)$ along the normalized coordinate x (cf. Eq. (6.30)).	155
6.5	Schematic overview of the proposed PINN architecture, highlighting the temporal encoders for $Q_{\text{in}}(t)$ and $PPG(t)$, the spatial embedding, feature fusion, and the prediction heads for pressure and flow.	159
6.6	Predicted arterial blood pressure waveforms at different spatial locations $x \in [0, 1]$ along the modeled arterial segment. The horizontal axis represents time (8 seconds). The black curve denotes the ground-truth ABP at the reference location.	165
6.7	Radius profile (dark blue) and wall-thickness (orange) used as physiological prior	178
6.8	Stergiopoulos-style exponential fit of the effective Young's modulus as a function of vessel radius, used here to construct a physiologically plausible stiffness trend (cf. the relation above).	179

6.9	Q(t) example prediction. (Remember that no ground truth was used for the Q)	191
6.10	Example of 3D plot ABP to show the prediction from $x=0$ and $x=1$ in 2 seconds of prediction (for clarity). As it can be seen, the waveform slowly shrinks in amplitude as x approaches the capillaries at $x=1$, following the real physiology	192
6.11	In blue: example of the ABP prediction at $x=0.5$, where ABP is usually measured. In black the real ground truth of the ABP.	195
6.12	Example of the SBP and DBP prediction and their respective uncertainties from aorta ($x=0$) and fingertip ($x=1$).	195
6.13	Example of the ABP prediction from aorta ($x=0$) and fingertip ($x=1$). In red the real ground truth of the ABP.	196
7.1	t-SNE projection of physiological features with subjective stress labels.	210
7.2	t-SNE projection with synthetic task-based labels.	211
7.3	PCA projection with subjective labels.	212
7.4	PCA projection with synthetic labels.	213

List of Tables

4.1	Average scores across all folds for the analyzed ECG features	51
4.2	Comparison of detection performance with different methods	52
5.1	Pruning thresholds for the standard deviation σ of peak distances and peak heights in PPG, ECG, and ABP windows.	65
5.2	Summary of the ablation experiments on the Conv1D–BiLSTM architecture for ABP waveform prediction.	75
5.3	CSP comparison for ECG signals	86
5.4	CSP comparison for PPG signals	86
5.5	Experiment 2 (LOSO): SBP/DBP MAE for the held-out subject in each fold (reported values).	121
5.6	Experiment 3 (LOSO + age/height/weight): SBP/DBP MAE for the held-out subject in each fold (reported values).	124
5.7	Change in MAE from Experiment 2 to Experiment 3 ($\Delta = \text{MAE}_{\text{Exp2}} - \text{MAE}_{\text{Exp3}}$). Positive values indicate improvement.	124
5.8	Global MAE per subject before and after adding anthropometric information.	125
6.1	Calibration-free blood pressure estimation performance reported in prior works using datasets with reliable subject separation. Results from MIMIC-II and other datasets reported in the original VitalDB paper are omitted when subject separation cannot be guaranteed. . .	191

7.1	Performance comparison of single-signal models (mean $\pm$ std over 4 folds).	205
7.2	Performance comparison of dual-signal models (mean $\pm$ std over 4 folds).	206
7.3	EDA + RR with and without cross-attention (mean $\pm$ std).	215
7.4	EDA + ECG with and without cross-attention (mean $\pm$ std).	215

Chapter 1

Introduction

1.1 Motivation

Continuous monitoring of cardiovascular and autonomic function represents a central challenge in modern healthcare, with implications ranging from early disease detection to long-term risk stratification and personalized medicine. Among the measurable physiological variables, arterial blood pressure occupies a particularly critical role, as it reflects both cardiac performance and vascular state and is directly associated with cardiovascular morbidity and mortality [1]. Despite its importance, blood pressure is still predominantly assessed through cuff-based sphygmomanometry, a method that is inherently intermittent, intrusive, and poorly suited for continuous or ambulatory monitoring [2].

In recent years, the widespread diffusion of wearable sensors capable of acquiring biosignals such as electrocardiogram (ECG) and photoplethysmogram (PPG) has opened new perspectives for non-invasive and continuous cardiovascular monitoring [3–5]. These signals can be collected unobtrusively in everyday environments and have motivated a growing body of research on cuffless blood pressure estimation and stress detection using machine learning techniques [6, 3–5, 7–19]. In parallel, advances in deep learning have demonstrated an unprecedented ability to extract complex patterns from high-dimensional time-series data [20–22], often achieving impressive predictive performance under controlled experimental conditions.

However, the apparent success of these approaches raises fundamental questions about what is actually being learned from physiological signals. Many state-of-the-

art models report low prediction errors when training and testing data are drawn from the same subjects or from the same database, yet their performance often deteriorates sharply when evaluated on unseen individuals or external datasets [23–25]. This gap between reported accuracy and real-world robustness suggests that current models may be exploiting subject-specific morphological regularities rather than learning generalizable physiological relationships (see Fig. 1.2).

This issue is not confined to blood pressure estimation. Similar limitations emerge in stress detection, where physiological responses are strongly modulated by individual differences, contextual factors, and subjective perception. In both domains, biosignals reflect complex interactions between multiple physiological subsystems, and the same observable signal patterns may correspond to different underlying states across subjects. As a result, purely data-driven mappings from signals to physiological or psychological quantities are prone to degeneracy and ambiguity.

The motivation of this thesis is therefore twofold. On one hand, it aims to critically analyze the capabilities and limitations of deep learning models applied to cardiovascular and stress-related biosignals, with particular attention to the distinction between intra-subject performance and inter-subject generalization. On the other hand, it explores whether embedding physiological knowledge into the learning process can mitigate these limitations by constraining the solution space toward physically plausible and interpretable regimes.

Rather than pursuing incremental improvements in predictive accuracy, this work is driven by the need to understand why performance saturates, why generalization fails, and which elements are fundamentally missing from current modeling approaches. By addressing these questions, the thesis seeks to contribute not only new models, but also a clearer conceptual framework for interpreting the role of machine learning in physiological inference and wearable health monitoring. The high-level pipeline is summarized in Fig. 1.1.

1.2 From Biosignals to Physiological Inference

Physiological biosignals such as electrocardiogram (ECG) and photoplethysmogram (PPG) provide indirect and partial observations of the underlying cardiovascular




Fig. 1.1 High-level pipeline for data-driven cuffless blood pressure estimation from ECG and PPG.




Fig. 1.2 Schematic comparison of evaluation protocols: in-domain (potentially optimistic) versus strict subject-wise/cross-database generalization.





Fig. 1.3 Roadmap of the thesis: progression from signal-level learning to physiology-informed modeling.



Non-uniqueness: many physiological states can explain similar waveforms.

Fig. 1.4 Forward (physiology→signals) versus inverse inference (signals→physiology), highlighting the ill-posed nature of cuffless estimation.

dynamics. The ECG reflects the electrical depolarization and repolarization processes that trigger myocardial contraction, while the PPG captures peripheral changes in blood volume driven by the propagation of pressure and flow waves along the arterial tree. Although these signals are closely linked to cardiovascular function, they do not directly encode absolute physiological quantities such as arterial blood pressure.

From a modeling perspective, the inference of physiological variables from biosignals constitutes an inverse problem. The forward process, from cardiac activation to pressure generation and peripheral waveform formation, is governed by complex, nonlinear, and subject-dependent mechanisms involving cardiac mechanics, vascular geometry, arterial stiffness, blood rheology, and peripheral resistance. Reversing this process using only surface-level observations is therefore intrinsically ill-posed: multiple physiological configurations can give rise to similar ECG and PPG morphologies, the inverse mapping is non-unique without additional priors/constraints, a limitation that has long been recognized in cardiovascular modeling literature [26–28] (Fig. 1.4).

Deep learning models have proven exceptionally effective at approximating complex mappings in high-dimensional spaces. When applied to biosignals, these models can identify subtle temporal and morphological patterns that correlate with physiological states. In practice, this capability often translates into excellent predictive performance when training and evaluation data are drawn from the same subject population or acquisition context. In such cases, neural networks implicitly exploit regularities that are stable within individuals, including electrode placement, tissue properties, and subject-specific cardiovascular characteristics [23–25].

However, the success of these models in pattern recognition does not necessarily imply successful physiological inference. A model that accurately predicts blood pressure or stress labels may do so by leveraging correlations that are valid only within a limited statistical regime, without capturing the causal relationships that govern cardiovascular dynamics. As a result, high accuracy in constrained settings can coexist with poor robustness and limited interpretability.

This distinction between signal-level learning and physiology-level inference is central to the challenges addressed in this thesis. Without explicit constraints that reflect physical and physiological laws, data-driven models are free to converge toward solutions that are numerically optimal yet physiologically implausible. Recent work on physics-informed learning has highlighted how embedding governing equations and mechanistic priors into neural networks can regularize ill-posed inverse problems and improve interpretability [29].

Recognizing this gap motivates a shift in perspective: from treating ECG and PPG as generic time-series inputs to viewing them as observations of a latent physiological system. This viewpoint naturally raises questions about how prior knowledge, mechanistic models, and subject-specific information can be integrated with learning-based approaches to reduce ambiguity and improve generalization. These questions form the basis for the analysis of subjectivity and generalization issues discussed in the following section.

1.3 Subjectivity and the Generalization Problem

One of the most persistent challenges in machine learning applied to physiological biosignals is the limited ability of models to generalize across subjects. While high predictive accuracy is often reported in controlled experimental settings, a closer inspection reveals that such performance is frequently achieved under intra-subject or intra-database evaluation protocols. When models are tested on unseen individuals or external datasets, their performance commonly degrades [30–33], sometimes dramatically. This discrepancy highlights a fundamental tension between subject-specific learning and population-level generalization.

Physiological signals are inherently shaped by individual characteristics. Factors such as age, sex, vascular anatomy, arterial stiffness, autonomic tone, medication,




Fig. 1.5 Why evaluation protocol matters: random/window splits can leak subject-specific signatures, while subject-wise splits better reflect deployment on unseen individuals.

lifestyle, and transient physiological states influence both signal morphology and the relationship between observable waveforms and latent physiological quantities. As a consequence, ECG and PPG recordings carry a strong subject-specific imprint that extends beyond noise or measurement artifacts. Even under standardized acquisition conditions, inter-individual variability remains a dominant component of the signal space.

Deep learning models trained on biosignals tend to exploit this variability when it is statistically stable within the training data. In intra-subject settings, networks can implicitly learn subject-dependent transformations that map signal morphology to target labels with high precision. However, these transformations are not guaranteed to remain valid for new subjects whose physiological parameters fall outside the learned distribution. In such cases, the model’s apparent understanding of the task does not transfer, revealing that it has learned correlations rather than invariant physiological relationships.

This issue is further exacerbated by common evaluation practices in the literature. Window-based random splits, insufficient subject separation between training and test sets, and limited reporting of cross-database performance can lead to overly optimistic conclusions. Under these conditions, models may inadvertently memorize subject-specific patterns, masking the true difficulty of the generalization problem. The resulting performance metrics, while numerically impressive, provide limited insight into the model’s ability to operate reliably in real-world scenarios (Fig. 1.5).

The same phenomenon emerges clearly in stress detection and affective computing. Physiological responses to stressors are highly individualized, and subjective labels often reflect personal perception rather than a well-defined physiological ground truth. As a result, different subjects may exhibit similar signal patterns under different psychological states, or conversely, distinct physiological responses under comparable conditions. This non-separability further challenges purely data-driven approaches and mirrors the degeneracy observed in blood pressure estimation from ECG and PPG.

Rather than interpreting these limitations as a consequence of insufficient data alone, this thesis frames subjectivity as a structural property of physiological systems. The mapping from biosignals to physiological or psychological states is not universally invariant, and models that ignore this fact are prone to failure outside narrow experimental regimes. Addressing this challenge requires rethinking both modeling strategies and evaluation protocols, placing greater emphasis on physiological plausibility, interpretability, and the explicit handling of inter-subject variability.

These considerations motivate the exploration of modeling frameworks that go beyond unconstrained data-driven learning. By incorporating physiological structure and mechanistic priors, it may be possible to reduce the degrees of freedom available to the model and guide learning toward solutions that are robust across individuals. This perspective underlies the transition from purely data-driven approaches to physics-informed and physiology-guided models developed in the subsequent chapters.

More generally, this thesis argues that several biosignal-based inference problems cannot be adequately formulated as direct regression tasks. Instead, they should be interpreted as constrained inference problems in which observable physiological signals provide only partial information about underlying latent physiological states.

1.4 Research Questions and Objectives

The considerations discussed in the previous sections naturally lead to a set of research questions that go beyond the pursuit of incremental performance improvements. Rather than focusing exclusively on prediction accuracy, this thesis aims to

investigate the underlying reasons for the observed successes and failures of machine learning models applied to physiological biosignals (Fig. 1.6).

The central research questions addressed in this work can be summarized as follows:

- What information do deep learning models actually extract from ECG and PPG signals when performing cardiovascular inference?
- Why do data-driven models achieve high performance in intra-subject or intra-database settings but fail to generalize across subjects and datasets?
- To what extent is the mapping from biosignals to physiological quantities, such as blood pressure, intrinsically degenerate?
- Can the integration of physiological and physical knowledge constrain learning and improve interpretability and robustness?
- Are similar generalization limitations observable in related domains, such as stress detection, and do they share common structural causes?

Based on these questions, the objectives of the thesis are defined as follows:

- To analyze deep learning architectures for ECG and PPG processing, with particular emphasis on feature extraction and morphological learning.
- To experimentally evaluate data-driven blood pressure estimation approaches under both intra-subject and inter-subject validation protocols, including real-world wearable scenarios.
- To critically assess the role of subject-specific variability and evaluation methodology in shaping reported model performance.
- To develop and study physics-informed neural network frameworks that embed cardiovascular hemodynamics as soft constraints within the learning process.
- To explore stress-related physiological modeling as a parallel case study affected by similar issues of subjectivity and non-separability.




Fig. 1.6 Link between the research questions, the corresponding objectives, and the thesis chapters where they are addressed.

Importantly, this thesis does not aim to deliver a clinically deployable cuffless blood pressure monitoring system. Instead, its objective is to provide a structured analysis of the limits of current data-driven approaches and to establish a methodological foundation for future models that combine biosignal processing, physiological knowledge, and subject-specific information. This perspective guides both the experimental design and the interpretation of the results presented in the following chapters.

1.5 Contributions of the Thesis

The contributions of this thesis are primarily methodological and conceptual, supported by extensive experimental evidence. Rather than proposing a single end-to-end solution for cuffless physiological monitoring, this work aims to clarify the capabilities and limitations of machine learning when applied to cardiovascular and stress-related biosignals.

The main contributions can be summarized as follows (Fig. 1.7):




Fig. 1.7 High-level map of the main thesis contributions and where they are developed in the manuscript.

- A systematic investigation of deep learning models for ECG morphology learning and feature extraction, highlighting their ability to capture stable intra-subject patterns while exposing their limitations in terms of physiological interpretability.
- An experimental analysis of data-driven blood pressure estimation from ECG and PPG signals on large-scale clinical datasets, with explicit comparison between intra-subject and inter-subject validation protocols.
- The design and validation of a custom wearable device for ECG and PPG acquisition, enabling the assessment of cuffless blood pressure estimation methods in a real-world, non-clinical setting.
- A critical examination of inter-subject generalization failure in blood pressure estimation, framed as a consequence of subject-specific variability and intrinsic degeneracy in the biosignal-to-physiology mapping.
- The development of physics-informed neural network formulations that integrate one-dimensional cardiovascular hemodynamics and Windkessel boundary conditions as soft constraints, with the goal of improving physiological consistency and interpretability.

- An exploratory study on stress modeling using multimodal physiological signals, revealing parallels with blood pressure estimation in terms of subjectivity, label noise, and non-separability.

Beyond these specific contributions, this thesis provides a unified perspective on the role of physiology-aware modeling in wearable health monitoring. By explicitly addressing why purely data-driven approaches succeed in constrained settings yet fail to generalize, the work establishes a conceptual framework that can guide future research toward more robust, interpretable, and personalized physiological inference systems.

1.6 Structure of the Thesis

The remainder of this thesis is organized into eight chapters that progressively develop the methodological and conceptual framework of the work. The overall logical organization of the thesis is summarized in Fig. 1.8. The chapters follow a progressive reasoning, moving from physiological foundations and dataset considerations to increasingly complex inference problems based on biosignals. The final chapter provides a unified discussion of the methodological insights emerging from the different studies and outlines possible future research directions.

Chapter 2 introduces the physiological foundations underlying cardiovascular biosignals. In particular, the origin and interpretation of electrocardiography (ECG), photoplethysmography (PPG), and arterial blood pressure (ABP) signals are discussed, together with the physiological mechanisms linking electrical cardiac activity to hemodynamic responses. This chapter provides the physiological context necessary to interpret the signal processing and modeling approaches introduced later in the thesis.

Chapter 3 describes the datasets used throughout the experimental studies. Particular attention is devoted to the structure, acquisition protocols, and potential sources of bias present in publicly available physiological databases. The chapter also discusses common pitfalls in dataset usage, including subject leakage, population heterogeneity, and differences in clinical acquisition environments, all of which play a crucial role in the interpretation of machine learning results.


Chapter 4 presents a study on automatic detection of significant ECG features using hybrid deep learning models. This chapter serves as an example of a physiological signal processing task in which the relationship between signal morphology and the target variables is relatively well defined. The results provide a useful methodological benchmark and highlight the conditions under which data-driven approaches can perform reliably.

Chapter 5 focuses on cuffless blood pressure estimation using purely data-driven models based on ECG and PPG signals. Different experimental protocols are analyzed, including subject-specific and cross-subject settings. The chapter critically examines the limitations of direct regression approaches and discusses the challenges related to inter-subject variability and generalization across datasets.

Chapter 6 introduces a physics-informed modeling strategy for blood pressure estimation. Building on the limitations identified in the previous chapter, this chapter presents a first Physics-Informed Neural Network (PINN) formulation that incorporates simplified hemodynamic principles into the learning process. The model integrates physiological constraints derived from blood flow dynamics in order to regularize the estimation problem.

Chapter 7 extends the discussion to the broader problem of physiological stress estimation from biosignals. The chapter analyzes the difficulties associated with stress prediction, including the subjectivity of stress labels and the indirect relationship between observable signals and psychological states. A dynamical latent-state formulation is introduced as a conceptual framework to interpret stress-related physiological modulation.

Finally, Chapter 8 provides a global discussion of the results presented throughout the thesis and synthesizes the main methodological insights emerging from the different studies. The chapter revisits the central challenges of biosignal-based physiological inference, highlighting the role of latent physiological dynamics, observability limitations, and dataset biases. The thesis concludes with a set of future research directions aimed at developing more physiologically grounded and generalizable models for health monitoring applications.



Progression:

signals → learning → physiology constraints → broader validation

Fig. 1.8 Compact overview of the thesis structure, highlighting the progression from signal processing to physiology-informed modeling and discussion.

Chapter 2

Physiological Background

2.1 Cardiovascular Physiology

2.1.1 Electrical Activity of the Heart

The electrical activity of the heart represents the initiating event of the cardiac cycle and constitutes the primary source of information captured by the electrocardiogram (ECG). From a physiological standpoint, the ECG reflects the propagation of action potentials across myocardial tissue, originating from the sinoatrial node and spreading through the atrial and ventricular conduction pathways. These electrical depolarization and repolarization processes coordinate myocardial contraction but do not directly encode mechanical or hemodynamic quantities.

The ECG waveform is conventionally described in terms of its main morphological components [12] : the P wave, associated with atrial depolarization; the QRS complex, corresponding to rapid ventricular depolarization; and the T wave, reflecting ventricular repolarization. The temporal structure of these components provides precise information about the timing and sequence of cardiac electrical events. As such, ECG morphology is highly informative with respect to cardiac rhythm, conduction properties, and beat-to-beat variability. An example of ECG with the main fiducial points and waveform components is shown in Fig. 2.1.




Fig. 2.1 Example of a sinus-rhythm ECG trace with annotated waveform components and fiducial points (P wave, QRS complex, and T wave).

Importantly, the electrical activation of the myocardium precedes mechanical contraction [34], with a non-negligible electromechanical delay separating depolarization from force generation. This delay, together with the complex coupling between electrical excitation and myocardial mechanics, implies that electrical signals alone cannot uniquely determine the resulting pressure or flow dynamics. While changes in heart rate and activation timing influence cardiac output, the amplitude and shape of arterial pressure waveforms depend on additional factors that are not directly observable in the ECG.

From the perspective of biosignal-based modeling, this distinction has important implications. The ECG provides a reliable and repeatable marker of cardiac timing and rhythm, making it a powerful reference signal for synchronization and beat segmentation. At the same time, its relationship with arterial blood pressure is indirect and highly modulated by subject-specific properties, including myocardial contractility, preload, afterload, and vascular characteristics. Consequently, similar ECG morphologies may correspond to substantially different hemodynamic states

across individuals or even within the same individual under different physiological conditions.

These considerations clarify both the strength and the limitation of ECG-based inference. On one hand, the stability and high signal-to-noise ratio of ECG recordings enable deep learning models to effectively capture morphological patterns and temporal dependencies. On the other hand, the lack of a direct and invariant mapping between ECG features and mechanical outcomes introduces intrinsic ambiguity into tasks such as blood pressure estimation. This ambiguity plays a central role in the generalization issues discussed in later chapters and motivates the integration of complementary signals and physiological constraints in subsequent modeling approaches.

2.1.2 Mechanical Pressure Generation

Arterial blood pressure is not a primary physiological signal, but rather an emergent quantity resulting from the interaction between cardiac mechanical activity and the properties of the vascular system [34]. Following electrical activation, myocardial contraction generates a time-varying ventricular pressure that drives blood ejection into the arterial tree. The resulting arterial pressure waveform is therefore shaped by both cardiac and vascular factors and cannot be attributed to a single localized mechanism.

From a simplified mechanical perspective, ventricular contraction produces a stroke volume SV over the systolic phase, generating a cardiac output

$$Q(t) = \frac{dV(t)}{dt}, \quad (2.1)$$

where $V(t)$ denotes ventricular volume. The pressure developed within the arterial system depends on how this flow interacts with arterial compliance and peripheral resistance. In its most basic form, this relationship can be approximated by lumped-parameter models in which pressure $P(t)$ satisfies

$$C \frac{dP(t)}{dt} + \frac{P(t)}{R} = Q(t), \quad (2.2)$$

where C represents arterial compliance and R the total peripheral resistance. While highly simplified, this expression highlights a key principle: pressure is jointly determined by inflow dynamics and vascular properties.

In reality, the generation of arterial pressure is further modulated by ventricular contractility, preload, and afterload. Variations in myocardial force generation, end-diastolic volume, and vascular impedance can significantly alter systolic and diastolic pressure levels even in the presence of similar electrical activation patterns. Consequently, identical ECG waveforms may correspond to different pressure responses depending on the mechanical and vascular state of the subject.

Another important aspect is the distributed nature of pressure generation and propagation. Once blood is ejected into the aorta, pressure waves travel along the arterial tree, undergoing reflection and attenuation at bifurcations and impedance mismatches. The pressure measured at a peripheral site is therefore the result of superimposed forward and reflected waves rather than a direct representation of ventricular pressure. This spatial separation between pressure generation and pressure observation further complicates the inverse inference of blood pressure from proximal electrical signals.

These considerations underscore a fundamental limitation of purely signal-based approaches to blood pressure estimation. While cardiac electrical activity sets the timing of mechanical contraction, it does not uniquely determine the amplitude or shape of the resulting pressure waveform. Mechanical pressure generation depends on a combination of cardiac output, vascular compliance, resistance, and wave propagation effects, all of which exhibit strong inter-subject variability and temporal modulation.

For modeling purposes, this implies that any attempt to infer blood pressure from biosignals must either explicitly account for these mechanical factors or accept an inherent level of ambiguity. This observation motivates the incorporation of hemodynamic structure into learning-based models, as explored in subsequent chapters, where pressure is treated not as a direct function of signal morphology but as the outcome of constrained physiological dynamics.

2.1.3 Arterial Tree and Vascular Properties

The arterial tree plays a fundamental role in shaping the pressure and flow waveforms generated by cardiac contraction. Rather than acting as a passive conduit, the arterial system behaves as a distributed, dynamic medium whose geometric and mechanical properties strongly influence cardiovascular dynamics. Once blood is ejected from the heart, pressure and flow waves propagate along the arterial network, undergoing transformation as a result of vessel geometry, elasticity, and impedance mismatches at bifurcations [26].

Arterial vessels exhibit progressive changes in radius, wall thickness, and elastic properties from central to peripheral segments. Large elastic arteries, such as the aorta, are characterized by high compliance, allowing them to store part of the stroke volume during systole and release it during diastole. This buffering effect reduces pulsatility and maintains a more continuous peripheral flow. In contrast, muscular arteries and arterioles are stiffer and contribute more significantly to wave reflection and peripheral resistance. As a result, pressure waveforms evolve along the arterial path, often exhibiting increased systolic peaks and altered waveform morphology at distal sites.

From a mechanical standpoint, arterial stiffness and compliance are key determinants of pulse wave velocity and pressure amplification. The relationship between vessel elasticity and wave propagation can be expressed qualitatively through the dependence of wave speed c on arterial wall properties and lumen geometry, commonly approximated as [26, 35]:

$$c \propto \sqrt{\frac{Eh}{\rho R}}, \quad (2.3)$$

where E denotes the effective elastic modulus of the arterial wall, h the wall thickness, R the vessel radius, and ρ the blood density. Variations in these parameters across subjects and along the arterial tree directly affect the timing and amplitude of reflected waves, contributing to inter-individual differences in observed pressure and PPG signals.

The arterial network also introduces spatial heterogeneity into cardiovascular measurements. Pressure and flow at a given location depend not only on local vessel properties but also on upstream and downstream conditions, including reflection sites and peripheral impedance. Consequently, measurements acquired at different

anatomical locations cannot be directly compared without accounting for propagation effects. This is particularly relevant for peripheral sensing modalities such as PPG, which capture signals at sites far from the heart.

These properties highlight a central challenge for biosignal-based inference: the arterial tree acts as a subject-specific and spatially distributed filter that transforms cardiac output into observable pressure and volume waveforms. Without explicit modeling of this filtering process, learning-based approaches may inadvertently associate signal morphology with blood pressure in a way that is valid only for specific anatomical or physiological configurations.

In the context of this thesis, the arterial tree is therefore treated not merely as anatomical background, but as an essential component of cardiovascular dynamics. Its geometric and mechanical characteristics motivate the use of simplified yet physiologically grounded models, such as one-dimensional flow equations and lumped-parameter terminations, to capture wave propagation effects while maintaining computational tractability. These modeling choices form the basis of the hemodynamic frameworks introduced in the following sections.

2.2 Hemodynamics

2.2.1 Blood Flow and Fluid Mechanics

Blood flow in the arterial system results from the mechanical action of the heart and is governed by the fundamental laws of fluid mechanics. At the macroscopic scale of large and medium arteries, blood can be modeled as an incompressible fluid with density ρ , while flow remains predominantly laminar and pulsatile. Although blood exhibits non-Newtonian behavior at the microcirculatory level, this approximation is well justified for modeling arterial wave propagation and pressure dynamics.

The motion of an incompressible viscous fluid can be described by the Navier–Stokes equations [27, 28]

$$\nabla \cdot \mathbf{u} = 0, \quad (2.4)$$

$$\rho \left(\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} \right) = -\nabla p + \mu \nabla^2 \mathbf{u}, \quad (2.5)$$

where $\mathbf{u}(\mathbf{x}, t)$ denotes the velocity field, $p(\mathbf{x}, t)$ the pressure field, and μ the dynamic viscosity. When coupled with appropriate vessel wall mechanics and boundary conditions, these equations provide a complete physical description of blood motion within deformable arteries.

Despite their generality, direct three-dimensional solutions of Eqs. (2.4)–(2.5) are impractical in most physiological inference scenarios. Such models require detailed subject-specific vascular geometries, material properties, and boundary conditions, and entail a computational cost incompatible with learning-based approaches and real-time wearable monitoring.

To balance physiological fidelity and computational tractability, the Navier–Stokes equations are commonly reduced by integrating them over the vessel cross-section. This approach exploits the predominantly axial nature of arterial flow and the relatively slow spatial variation of vessel geometry. Let x denote the axial coordinate along the vessel, $A(x, t)$ the cross-sectional area, $Q(x, t)$ the volumetric flow rate, and $p(x, t)$ the cross-sectionally averaged pressure.

Under standard assumptions, conservation of mass yields the one-dimensional continuity equation

$$\frac{\partial A(x, t)}{\partial t} + \frac{\partial Q(x, t)}{\partial x} = 0. \quad (2.6)$$

Similarly, conservation of momentum leads to

$$\frac{\partial Q(x, t)}{\partial t} + \frac{\partial}{\partial x} \left(\frac{Q(x, t)^2}{A(x, t)} \right) + \frac{A(x, t)}{\rho} \frac{\partial p(x, t)}{\partial x} + \mathcal{F}(Q, A) = 0, \quad (2.7)$$

where $\mathcal{F}(Q, A)$ represents viscous and frictional losses due to velocity gradients near the vessel wall. In practical arterial models, this term is often approximated as proportional to the flow rate normalized by vessel area,

$$\mathcal{F}(Q, A) = K_r \frac{Q(x, t)}{A(x, t)}, \quad (2.8)$$

where K_r is an effective viscous parameter that aggregates wall shear effects and unresolved velocity profile dynamics.

Equations (2.6)–(2.8) [36] form a hyperbolic system of partial differential equations governing the spatio-temporal evolution of flow and pressure along the arterial axis. Pressure does not appear as an independent variable but is coupled to ves-

sel deformation through a constitutive pressure–area relationship, introduced in Section 2.2.4.

A defining property of the one-dimensional hemodynamic formulation is its ability to capture wave propagation phenomena. Linearization of the governing equations reveals that pressure and flow disturbances propagate with a finite wave speed

$$c = \sqrt{\frac{A}{\rho} \frac{\partial p}{\partial A}}, \quad (2.9)$$

which depends explicitly on vessel geometry and wall elasticity. As pressure waves travel along the arterial tree, reflections at impedance mismatches and distal vascular loading alter waveform shape and amplitude.

These distributed effects imply that arterial pressure observed at peripheral sites is not a direct reflection of ventricular pressure but the result of a spatially extended dynamical system. Consequently, peripheral measurements encode information about both cardiac ejection dynamics and vascular properties along the arterial path.

From an inverse modeling perspective, the one-dimensional hemodynamic equations define a constrained dynamical system in which pressure, flow, and vessel area evolve according to known physical laws. In practice, however, only partial and indirect observations are available. Flow is rarely measured directly, pressure is unknown, and vessel properties vary substantially across subjects.

This intrinsic incompleteness renders the inference of blood pressure from biosignals an ill-posed inverse problem.

The formulation introduced in this section provides the physical foundation for the hemodynamic models developed in Chapters 6 and 7, where arterial pressure wave propagation from the aortic root to peripheral sites is modeled in a physiology-consistent and learning-compatible manner.

2.2.2 Windkessel Models

While one-dimensional hemodynamic equations describe the distributed propagation of pressure and flow waves along the arterial tree, appropriate boundary conditions are required to obtain a well-posed system. In particular, the distal arterial network, consisting of small arteries, arterioles, and the microcirculation, cannot be explicitly

resolved within a one-dimensional framework without introducing excessive complexity. Instead, its aggregate hemodynamic effect is commonly represented using lumped-parameter Windkessel models [37, 34].

From a macroscopic perspective, the distal vasculature behaves as a hydraulic load characterized by resistance and compliance. Large elastic arteries store part of the stroke volume during systole and release it during diastole, while peripheral resistance regulates mean arterial pressure and flow distribution. Windkessel models capture these effects in a compact and physiologically interpretable manner, making them well suited for coupling with one-dimensional arterial flow equations.

The two-element Windkessel model consists of a compliance C and a resistance R , representing arterial storage and peripheral resistance, respectively. Let $p(x, t)$ and $Q(x, t)$ denote pressure and flow at a terminal boundary location x . The governing equation can be written in implicit form as

$$C \frac{\partial p(x, t)}{\partial t} + \frac{p(x, t)}{R} - Q(x, t) = 0. \quad (2.10)$$

This formulation captures the exponential decay of pressure during diastole but neglects characteristic impedance effects, limiting its ability to reproduce realistic systolic dynamics and wave reflections.

To account for wave reflections and proximal impedance, the three-element Windkessel model introduces an additional resistance R_1 in series with the parallel combination of compliance C and distal resistance R_2 . This configuration provides a more accurate representation of arterial loading conditions at the interface between distributed and lumped domains.

In the notation adopted throughout this thesis, the three-element Windkessel model can be written as

$$Q(x, t) \left(1 + \frac{R_1}{R_2} \right) + CR_1 \frac{\partial Q(x, t)}{\partial t} + \frac{p(x, t)}{R_2} - C \frac{\partial p(x, t)}{\partial t} = 0. \quad (2.11)$$

This formulation is equivalent to the classical three-element Windkessel representation and explicitly highlights the coupling between pressure, flow, resistance, and compliance at the terminal boundary.

The Windkessel parameters carry clear physiological meaning. The resistance terms R_1 and R_2 represent viscous losses and peripheral vascular resistance, which are modulated by vasomotor tone, autonomic activity, and long-term vascular remodeling. The compliance C reflects the elastic storage capacity of the arterial system and is strongly influenced by arterial stiffness, age, and pathological conditions such as hypertension.

These parameters exhibit substantial inter subject variability and temporal modulation, reinforcing the need for subject specific modeling rather than fixed, population-averaged values.

Within a one-dimensional arterial flow framework, Windkessel models act as boundary constraints linking distal pressure and flow. The values of $p(x, t)$ and $Q(x, t)$ predicted by the distributed hemodynamic equations at the terminal location are required to satisfy Eq. (2.11) [34], thereby closing the system formed by the conservation of mass and momentum equations.

This coupling enables realistic pressure wave reflection and attenuation phenomena to be captured without explicitly modeling the microcirculation, while maintaining physiological consistency across the arterial domain.

In the physics-informed neural network formulations developed in this thesis, Windkessel equations are enforced as soft boundary constraints rather than as fixed boundary conditions. The residual of Eq. (2.11) is included in the loss function, allowing the network to learn spatio-temporal pressure and flow fields that satisfy both distributed hemodynamic equations and lumped vascular dynamics.

This approach enables the estimation of Windkessel parameters directly from data, providing a principled mechanism to account for inter-subject variability in peripheral resistance and compliance. By embedding physiologically grounded boundary conditions into the learning process, the model reduces ambiguity in the inverse problem of blood pressure estimation and promotes solutions that remain consistent with cardiovascular mechanics.

2.2.3 Pressure–Area Relationships

To close the one-dimensional hemodynamic system, a constitutive relationship linking arterial pressure and vessel cross-sectional area is required. This relationship

accounts for the mechanical response of the arterial wall to internal pressure and provides the coupling between fluid dynamics and vessel mechanics.

A commonly adopted approximation models the arterial wall as a thin, linearly elastic structure operating around a reference configuration [26, 28]. Under this assumption, pressure variations induce proportional changes in vessel area, leading to a pressure–area law of the form

$$p(x, t) = p_{ext} + \frac{\beta(x)}{A_0(x)} (\sqrt{A(x, t)} - \sqrt{A_0(x)}), \quad (2.12)$$

where p_0 and A_0 denote the reference pressure and cross-sectional area, respectively, and β is an effective stiffness parameter that encapsulates the elastic properties of the arterial wall.

The parameter β depends on both geometric and material properties of the vessel, including wall thickness, elastic modulus, and lumen radius. In simplified formulations, it can be expressed, up to proportionality constants depending on the chosen wall model, as

$$\beta = \frac{Eh}{R_0}, \quad (2.13)$$

where E denotes the Young's modulus of the arterial wall, h its thickness, and R_0 the reference vessel radius. This expression highlights the direct influence of arterial stiffness and geometry on pressure dynamics and wave propagation.

Although the linear pressure–area relationship represents an approximation, it captures the dominant mechanical behavior of large and medium arteries under physiological pressure ranges. In particular, it provides a direct link between changes in flow-induced vessel deformation and the resulting pressure variations, enabling the propagation of pressure waves along the arterial tree to be modeled consistently.

Importantly, the parameters A_0 and β exhibit strong inter-subject variability and may vary along the arterial path. Differences in arterial stiffness due to age, vascular remodeling, or pathological conditions directly affect the pressure–area relationship and, consequently, the observed pressure waveform. Treating these parameters as subject-specific rather than fixed constants is therefore essential for realistic hemodynamic modeling.

Within the one-dimensional framework, Eq. (2.12) closes the system formed by the conservation of mass and momentum equations. When combined with

appropriate boundary conditions, such as Windkessel models, this constitutive law enables pressure, flow, and vessel area to be described as a coupled dynamical system. This formulation provides the mechanical foundation for physics-informed neural network models, where the pressure–area relationship is enforced as a soft constraint during learning, guiding the inference toward physiologically plausible solutions.

2.3 Photoplethysmography

Photoplethysmography (PPG) is a non-invasive optical technique widely employed for monitoring cardiovascular activity at peripheral sites [13, 38]. By measuring time-varying changes in light absorption associated with pulsatile blood volume variations, PPG provides indirect information about arterial dynamics and pressure wave propagation. Due to its simplicity, low cost, and compatibility with wearable devices, PPG has become a central sensing modality in cuffless blood pressure estimation and peripheral hemodynamic monitoring.

Despite its widespread use, PPG does not measure arterial pressure directly. Instead, it reflects a cascade of physiological and optical processes that transform pressure waves generated at the heart into observable variations in peripheral blood volume and optical absorption. Understanding this transformation is essential for interpreting PPG signals and for designing models that aim to infer blood pressure from optical measurements.

2.3.1 Optical Modeling and Blood Volume Dependence

Photoplethysmography relies on the measurement of light intensity variations induced by pulsatile changes in arterial blood volume within superficial tissue layers. Let $I(t)$ denote the optical intensity detected by the photodiode and I_0 the emitted light intensity. Within a small tissue region beneath the probe, light propagation is governed by absorption and scattering processes associated with multiple tissue constituents.

A generalized Beer–Lambert formulation can be used to describe the detected signal,

$$I(t) = I_0 \exp\left(-\sum_j \varepsilon_j c_j L_j(t) - G\right), \quad (2.14)$$

where ε_j and c_j denote the extinction coefficients and concentrations of the absorbing species, $L_j(t)$ their effective optical pathlengths, and G collects geometry-dependent scattering losses.

Among these contributions, only the arterial blood compartment exhibits significant beat-to-beat variation. Denoting its optical pathlength by $L_b(t)$, and assuming a fixed measurement geometry, the pathlength contribution of arterial blood can be approximated as proportional to the enclosed blood volume $V_b(t)$,

$$L_b(t) = \alpha V_b(t), \quad \alpha > 0. \quad (2.15)$$

By grouping all non-pulsatile terms into a constant K_0 , the detected intensity can be expressed as

$$I(t) = I_0 \exp(-K_0 - k V_b(t)), \quad k = \varepsilon_{\text{Hb,eff}} \alpha, \quad (2.16)$$

where $\varepsilon_{\text{Hb,eff}}$ denotes an effective hemoglobin extinction coefficient at the illumination wavelength.

In practical PPG acquisition systems, analog and digital conditioning remove slow baseline components and rescale the signal around a steady operating point. Under small pulsatile variations of blood volume, a first-order linearization of (2.16) yields a signal that is approximately proportional to deviations in arterial blood volume,

$$\text{PPG}(t) = \beta + \kappa V_b(t), \quad \kappa > 0, \quad (2.17)$$

where β absorbs static offsets and κ accounts for optical scaling and electronic gain.

Decomposing the blood volume into slow and pulsatile components, $V_b(t) = V_{\text{DC}} + \tilde{V}(t)$, the PPG signal can be written as

$$\text{PPG}(t) = \underbrace{(\beta + \kappa V_{\text{DC}})}_{\text{DC component}} + \kappa \tilde{V}(t), \quad (2.18)$$

making explicit that the clinically relevant information is carried by the pulsatile term.

2.3.2 From Blood Volume to Flow

The temporal evolution of arterial blood volume within the optical sampling region is governed by mass conservation. Considering a small control volume, changes in blood volume arise from the imbalance between inflow and outflow. At cardiac frequencies, venous outflow and low-compliance compartments vary slowly compared to the arterial pulsation.

By isolating the pulsatile component through high-pass filtering, as an approximation used for the pulsatile component, the rate of change of blood volume can be approximated as

$$\frac{d\tilde{V}(t)}{dt} \approx Q_p(t), \quad (2.19)$$

where $Q_p(t)$ denotes the pulsatile component of arterial inflow.

Using the linearized PPG–volume relationship in (2.18), this leads to

$$\tilde{V}(t) = \frac{1}{\kappa} \left(\text{PPG}(t) - \text{PPG}_{\text{DC}}(t) \right), \quad (2.20)$$

and therefore

$$\boxed{Q_p(t) \propto \frac{d}{dt} \text{PPG}(t)} \quad (2.21)$$

where the proportionality constant absorbs optical and calibration factors [8, 9].

This result establishes a direct link between the first temporal derivative of the PPG waveform and pulsatile blood flow.

2.3.3 From Flow to Pressure

To relate pulsatile flow to arterial pressure, the local compliance of the arterial wall is introduced,

$$C = \frac{dV}{dP}. \quad (2.22)$$

For small pressure excursions within a cardiac cycle, compliance can be assumed approximately constant. Differentiating the pressure–volume relationship yields

$$\frac{dV}{dt} = C \frac{dP}{dt}. \quad (2.23)$$

Combining this expression with (2.21) gives

$$Q_p(t) \propto \frac{d}{dt} \text{PPG}(t) \approx C \frac{dP(t)}{dt}. \quad (2.24)$$

Integrating both sides over time and neglecting constant offsets removed during preprocessing yields

$$\boxed{P(t) \propto \text{PPG}(t)} \quad (2.25)$$

under local linearization, after DC removal, and assuming quasi-constant compliance within a beat. This proportionality provides a theoretical justification for the widespread empirical observation that PPG waveform morphology closely follows arterial pressure dynamics at the measurement site. Importantly, the proportionality constant depends on local arterial compliance and optical scaling factors, which vary across subjects and measurement conditions. Consequently, while PPG can reproduce the temporal structure of pressure waveforms, absolute pressure estimation requires additional physiological constraints or calibration.

2.4 Autonomic Nervous System and Stress Physiology

2.4.1 Sympathetic–Parasympathetic Balance

The autonomic nervous system (ANS) provides a fast and adaptive control layer that regulates cardiovascular function over multiple time scales. Its sympathetic branch promotes increases in heart rate and contractility and modulates vascular tone through vasoconstriction, thereby affecting peripheral resistance and the effective loading conditions of the arterial tree. Conversely, parasympathetic (vagal) activity predominantly acts on the heart, slowing sinoatrial firing and shaping beat-to-beat variability, with more limited direct influence on systemic vascular resistance [39].

This autonomic control is central to blood pressure regulation because arterial pressure is a regulated outcome rather than a passive byproduct of cardiac timing. In practical terms, the same ECG timing pattern can correspond to different mechanical and hemodynamic states depending on sympathetic tone, baroreflex-mediated adjustments, and peripheral vasomotor control [39]. This has direct implications for biosignal inference: parameters that summarize distal impedance and arterial

compliance (e.g., Windkessel resistances and compliance, and effective stiffness in pressure–area laws) are not static across subjects and can drift within-subject under changes in autonomic state.

From the modeling standpoint adopted in this thesis, ANS-mediated variability motivates subject-specific parameterization and physiologically constrained learning. In particular, inter-subject generalization is challenged by latent differences in vascular tone and compliance that are not fully observable in ECG and peripheral PPG alone, thus contributing to the ill-posedness of cuffless blood pressure estimation.

2.4.2 Electrodermal Activity

Electrodermal activity (EDA) measures variations in the electrical conductance of the skin driven by eccrine sweat gland activity. Importantly, eccrine sweat glands are innervated by sympathetic cholinergic fibers, making EDA a relatively specific marker of sympathetic arousal [40, 41]. For this reason, EDA is widely used in psychophysiology as a non-invasive proxy of autonomic activation and arousal dynamics [41].

EDA is commonly decomposed into a slow tonic component (skin conductance level) and a phasic component (skin conductance responses) linked to transient sympathetic bursts [41]. While these components are informative about arousal, they do not directly quantify stress as a unique physiological state: similar EDA responses can emerge from cognitive effort, novelty, emotional engagement, thermoregulation, or motion artifacts, and the magnitude of EDA responses depends strongly on subject physiology and measurement conditions [40, 41].

In the context of wearable cardiovascular inference, EDA is also relevant because sympathetic activation can simultaneously affect peripheral vascular tone and local perfusion, thereby altering PPG morphology and amplitude. Consequently, autonomic arousal can act as a confounder for purely morphology-based PPG→BP mappings, reinforcing the need for models that incorporate physiological structure and allow for subject- and state-dependent parameters.

2.4.3 Heart Rate Variability and Respiration

Heart rate variability (HRV) reflects fluctuations in the inter-beat interval and provides a non-invasive window into autonomic cardiac control. Standardized measurement practices and interpretive guidelines are summarized by the Task Force report [42], while subsequent reviews emphasize both the physiological origins of HRV and common caveats in interpretation [43]. In general, vagal modulation dominates fast beat-to-beat dynamics, whereas sympathetic influences tend to act more slowly, interacting with reflex loops and vascular control [43, 42].

Respiration is tightly coupled to cardiovascular control through respiratory sinus arrhythmia (RSA), i.e., the respiration-driven modulation of vagal outflow to the heart [44]. Breathing pattern therefore shapes HRV metrics and can introduce apparent changes in autonomic indices even in the absence of altered stress levels [44]. Respiration also affects hemodynamics mechanically via intrathoracic pressure changes that modulate venous return and stroke volume, producing beat-to-beat variability in pressure and peripheral pulse morphology.

These interactions are critical for stress-related modeling because HRV and respiration carry information about autonomic regulation but remain indirect, context-sensitive, and subject-dependent observables. In practice, two subjects may exhibit similar HRV changes for different underlying reasons (e.g., different baseline vagal tone, different breathing patterns), which contributes to label noise and limited separability in stress inference tasks based on physiological signals.

2.5 Summary of Key Physiological Challenges

The physiological mechanisms reviewed in this chapter highlight a common theme: biosignals provide indirect and partial observations of a highly regulated dynamical system. ECG mainly encodes electrical timing, PPG encodes peripheral vascular volume dynamics filtered by wave propagation and local tissue optics, and autonomic regulation modulates both cardiac and vascular subsystems in a state- and subject-dependent manner [39, 41, 42].

This leads to two structural challenges for learning-based inference. First, the mapping from peripheral observations to central hemodynamic quantities is intrinsically ambiguous: multiple combinations of cardiac output, vascular stiffness,

resistance, and autonomic state may produce similar observable morphologies. Second, subject variability is not an incidental nuisance but a dominant factor that changes the effective parameters of the cardiovascular system across individuals and contexts.

These considerations motivate the methodological direction adopted in this thesis: rather than learning an unconstrained direct regression from biosignals to blood pressure, the inference task should be guided by physiological structure through conservation laws, constitutive relations, and physiologically meaningful boundary conditions. In parallel, stress-related modeling faces analogous challenges, where autonomic observables (EDA, HRV, respiration) reflect regulatory dynamics but do not uniquely identify stress intensity across subjects, motivating careful evaluation protocols and models that respect physiological constraints [45, 46].

Chapter 3

Dataset and Preprocessing

3.1 Overview

This thesis relies on four complementary datasets that jointly cover (i) clean beat-level ECG delineation benchmarks, (ii) large-scale ICU waveforms with invasive arterial blood pressure (ABP), (iii) high-quality operating-room waveforms with explicit subject identifiers, and (iv) multimodal recordings for stress-related modeling. Specifically, this thesis uses the QT Database (QTDB) [47, 48], the MIMIC-III Waveform Database [49, 48], the VitalDB [50] subset distributed within PulseDB [51], the Cuff-Less Blood Pressure Estimation Dataset [52, 53], and the StressID dataset [54].

3.1.1 QT Database

The QT Database (QTDB) is a curated collection of short multi-lead electrocardiogram (ECG) recordings with expert annotations of waveform boundaries, including P waves, QRS complexes, T waves, and, when present, U waves. Unlike large-scale clinical databases, QTDB is not designed for hemodynamic inference or blood pressure estimation, but rather as a benchmark for ECG morphology analysis and delineation tasks.

In this thesis, QTDB is used as a reference dataset to validate ECG preprocessing pipelines, segmentation strategies, and morphology-sensitive neural representations. Its high-quality annotations and relatively noise-free recordings provide a controlled

environment in which signal-level behavior can be analyzed independently of confounding hemodynamic variability. This makes QTDB particularly useful for isolating what deep learning models learn from ECG morphology alone, before coupling ECG with peripheral signals or physiological constraints.

3.1.2 MIMIC-III Waveform Database

The MIMIC-III Waveform Database is a large-scale collection of physiological waveforms acquired from patients admitted to intensive care units, including electrocardiogram (ECG), photoplethysmogram (PPG), and invasive arterial blood pressure (ABP) signals. The database covers a wide range of clinical conditions, signal qualities, and acquisition durations, reflecting real-world hospital environments.

Due to the availability of invasive ABP measurements, MIMIC-III represents a key resource for supervised learning and evaluation of cuffless blood pressure estimation models. However, its clinical heterogeneity, presence of artifacts, and strong subject-specific patterns also make it a challenging benchmark for generalization. In this thesis, MIMIC-III is primarily used to study intra-database performance and to analyze how data-driven models exploit stable morphological correlations between ECG, PPG, and ABP within the same acquisition domain.

Importantly, subject overlap and recording-specific characteristics can introduce implicit leakage if not handled carefully. For this reason, MIMIC-III is treated as a reference domain against which cross-database robustness is explicitly assessed using external datasets.

3.1.3 VitalDB / PulseDB

VitalDB is a high-resolution dataset of perioperative vital signs and waveforms acquired in operating room settings, including ECG, PPG, and invasive arterial blood pressure [50]. Compared to ICU datasets, VitalDB recordings are generally characterized by higher signal quality, more stable acquisition conditions, and clearer synchronization between modalities.

In this thesis, VitalDB signals are accessed through the PulseDB distribution, which provides curated waveform segments and explicit subject identifiers to facilitate cross-subject and cross-database evaluation [51]. This makes VitalDB partic-

ularly suitable for studying distribution shift and generalization, as models trained on MIMIC-III can be evaluated on VitalDB without shared acquisition context or patient population.

VitalDB therefore plays a central role in assessing whether learned mappings between ECG, PPG, and blood pressure reflect general physiological relationships or dataset-specific correlations. Performance degradation observed when transferring models from MIMIC-III to VitalDB is interpreted as evidence of physiological degeneracy and motivates the use of physiology-informed constraints in later chapters.

3.1.4 Cuff-Less Blood Pressure Estimation Dataset

The Cuff-Less Blood Pressure Estimation Dataset was introduced in the context of learning-based blood pressure estimation from ECG and PPG signals using deep neural architectures, including convolutional and recurrent models. The dataset was derived from the MIMIC-II Waveform Database, which provides synchronized high-resolution physiological waveforms acquired in intensive-care monitoring, including arterial blood pressure (ABP) and photoplethysmography (PPG) [48]. The UCI dataset contains 12,000 instances of simultaneous ECG, PPG, and ABP signals acquired from 942 patients [52]. Importantly, the dataset was originally preprocessed by Kachuee et al. to smooth the signals, remove unacceptable values, and apply an autocorrelation-based step to assess pulse similarity in PPG [52].

This dataset has been widely used in the literature to report low prediction errors under inter-subject evaluation protocols. However, its relatively limited subject diversity, controlled acquisition setting, and implicit calibration effects make it a representative example of datasets where strong morphology-driven learning can dominate model behavior.

In this thesis, the Cuff-Less Blood Pressure Estimation Dataset is considered as a comparative benchmark to highlight the gap between apparent predictive accuracy and true physiological generalization. Its inclusion supports a broader discussion on evaluation protocols and the risks of overestimating model performance when subject variability and cross-domain transfer are not explicitly addressed.

3.1.5 StressID Dataset

StressID is a multimodal dataset designed for stress identification from synchronized video, audio, and physiological signals. Physiological recordings include ECG, EDA, and respiration acquired using a wearable device under controlled experimental conditions [54].

The dataset contains recordings from 65 participants performing 11 tasks, including:

- Relaxation and guided breathing,
- Emotionally evocative video stimuli,
- Cognitive tasks such as mathematics and comprehension exercises,
- Public speaking scenarios.

Each task is annotated using self-assessment scores for stress, relaxation, arousal, and valence. Two discrete labeling strategies are provided:

- A **binary (2-class) label**, obtained by thresholding the self-reported stress score at 5 on a 0–10 scale.
- A **3-class label**, distinguishing relaxed, neutral, and stressed states based on combined valence, arousal, and relaxation ratings.

In this thesis, StressID is used as a parallel case study to blood pressure estimation. Although the target variable differs, the underlying challenges are analogous: physiological signals are indirect, subject-dependent, and non-uniquely related to the inferred state. By analyzing stress modeling alongside blood pressure estimation, this work highlights structural similarities between the two problems, particularly in terms of label subjectivity, inter-subject variability, and limited separability of physiological patterns.

The inclusion of StressID therefore supports the broader methodological argument of the thesis, reinforcing the need for physiology-aware modeling and careful evaluation beyond purely data-driven performance metrics.

3.2 Dataset Biases and Structural Limitations

Despite the increasing availability of large-scale physiological databases, cuffless blood pressure estimation and stress modeling remain fundamentally constrained by structural limitations inherent to the datasets themselves. These limitations are not merely practical inconveniences but deeply influence the validity of evaluation protocols, the interpretation of reported performance, and the apparent generalization capability of machine learning models.

A first and central issue concerns *subject identity and data partitioning*. Many widely used clinical waveform databases were not originally designed for machine learning tasks requiring strict inter-subject generalization. As a result, subject-level identifiers are often incomplete, ambiguous, or not explicitly enforced in standard evaluation pipelines. This is particularly evident in the MIMIC-III Waveform Database, where recordings are long, heterogeneous, and collected under intensive care conditions, but where subject-level separation is not natively integrated into most downstream machine learning benchmarks.

In works that usually use MIMIC-III, physiological signals are typically segmented into short windows to increase sample size. However, without explicit subject-wise splitting, windows originating from the same patient may appear in both training and test sets. This induces a strong form of information leakage: models can implicitly exploit subject-specific morphology, sensor placement, and clinical context rather than learning invariant physiological relationships. Consequently, reported performance under such protocols may significantly overestimate true generalization to unseen individuals, even when random train–test splits appear statistically valid.

Recent studies have explicitly investigated these evaluation pitfalls, showing that random window-level splits and incomplete subject separation can lead to overly optimistic performance estimates and fundamentally flawed conclusions about inter-subject generalization [30–33].

This issue is not exclusive to MIMIC-III but reflects a broader structural bias in many physiological datasets. Window-based sampling, while convenient, transforms long recordings into large collections of highly correlated samples. If subject identity is not carefully controlled, this correlation structure violates the independence assumptions underlying standard machine learning evaluation. The resulting models

may perform well in-distribution while failing catastrophically under subject-level or cross-database evaluation.

Complementary datasets such as VitalDB and PulseDB partially address this limitation by providing clearer subject identifiers and more controlled acquisition conditions. However, these datasets introduce other forms of bias. VitalDB recordings are often cleaner, shorter, and acquired under operating room conditions, which differ substantially from both ICU environments and wearable scenarios. As a result, models trained on MIMIC-III and evaluated on VitalDB (or vice versa) experience a pronounced distributional shift that reflects not only subject variability but also differences in hardware, clinical context, and signal quality.

Additional datasets used in the thesis, such as the Cuff-Less Blood Pressure Estimation Dataset and StressID, further highlight the trade-offs between control and realism. Smaller, curated datasets often provide explicit subject separation and well-defined protocols, but at the cost of limited population diversity and reduced physiological variability. In stress-related datasets, these issues are compounded by label subjectivity: stress annotations are frequently self-reported or task-dependent, introducing label noise and non-universality that cannot be resolved through signal processing alone.

Beyond subject identity, several other structural biases affect physiological datasets. Demographic imbalance (e.g., age, sex, health status), medication effects, and unobserved confounders such as hydration, posture, or emotional state are rarely annotated in sufficient detail. For blood pressure estimation, the absence of metadata describing vascular properties, anthropometrics, or autonomic state further amplifies the ill-posedness of the inverse problem. In such conditions, different physiological configurations may generate indistinguishable ECG and PPG morphologies, leading to intrinsic degeneracy in the mapping from biosignals to pressure.

These limitations have direct methodological consequences. High accuracy reported under intra-database or intra-subject settings cannot be straightforwardly interpreted as evidence of physiological understanding. Instead, it often reflects the ability of models to interpolate within a narrow statistical regime shaped by dataset-specific biases. Conversely, poor performance under cross-subject or cross-database evaluation should not be interpreted solely as model failure, but as an indicator of missing information and insufficient constraints in the learning problem.

For these reasons, this thesis treats dataset biases not as secondary concerns but as central elements shaping both model design and evaluation. Rather than attempting to eliminate these biases through increasingly complex preprocessing, the work emphasizes transparent validation protocols, explicit discussion of subject-level leakage, and the integration of physiological structure as a means to reduce, rather than conceal, the ambiguity inherent in biosignal-based inference. This perspective motivates the physics-informed and physiology-guided approaches developed in the subsequent chapters.

Chapter 4

ECG Morphology and Feature Learning

4.1 Problem Formulation and Reference Annotations

The automatic detection of ECG fiducial points can be formulated as a supervised learning problem, where the objective is to estimate the temporal location of clinically meaningful waveform components directly from raw or minimally processed ECG signals. Given an input ECG segment $x(t)$, the task consists in identifying a set of time instants $\{t_k\}$ corresponding to predefined electrical events of the cardiac cycle, such as the onset, peak, and offset of the P wave, QRS complex, and T wave [55].

Unlike classification-based ECG tasks, which assign a discrete label to an entire heartbeat or segment, fiducial point detection requires precise temporal localization. This framing naturally leads to a regression-based or dense prediction formulation, in which the model outputs continuous-valued estimates or probability maps over time. The difficulty of the task lies not only in the accurate identification of sharp events such as the R peak, but also in the localization of smoother and more variable components, including P and T wave boundaries, whose morphology and duration exhibit significant inter- and intra-subject variability.

A further challenge arises from the fact that ECG morphology is influenced by noise, baseline wander, and acquisition artifacts, particularly in long-term or ambulatory recordings. Consequently, reference annotations provided by expert car-

diologists play a central role in defining the learning target and in constraining model evaluation. The reliability of these annotations directly affects both training stability and performance assessment, making the choice of dataset a critical methodological decision [56].

In this work, the problem is addressed using expert-annotated ECG datasets in which fiducial points have been manually identified according to standardized clinical criteria. These annotations provide a physiologically grounded reference for supervised learning and enable a quantitative evaluation of localization accuracy at the millisecond scale. Importantly, the task formulation explicitly avoids downstream clinical interpretation, focusing instead on the faithful reconstruction of ECG morphology as defined by expert labeling.

From the perspective of this thesis, this formulation serves a dual purpose. First, it allows the assessment of deep learning architectures in a setting where the physiological meaning of the target variables is well established and relatively unambiguous. Second, it provides insight into how neural networks internally represent waveform morphology when trained to solve a temporally precise inference problem, as commonly explored in recent deep learning approaches for ECG analysis [22, 57].

These insights are later leveraged when discussing more ill-posed inverse problems, such as blood pressure estimation from ECG and PPG, where ground truth definitions are less direct and subject-dependent variability is substantially higher [7, 8, 58].

The following sections describe the reference dataset, the adopted neural architecture, and the training and evaluation protocol used to address this problem, closely following the methodological framework introduced in the original study.

4.2 Related Work on Automatic ECG Feature Detection

The automatic analysis of electrocardiographic signals has been an active research topic for several decades, motivated by the central role of ECG interpretation in cardiovascular diagnosis and monitoring. Traditional approaches to ECG feature de-

tection relied on signal processing techniques and handcrafted rules, often designed to identify specific waveform characteristics such as slope changes, amplitude thresholds, or temporal constraints. While effective under controlled conditions, these methods typically require careful parameter tuning and exhibit limited robustness in the presence of noise, inter-subject variability, and morphological distortions.

The advent of machine learning, and in particular deep learning, has significantly reshaped the landscape of ECG analysis. Neural networks have demonstrated the ability to learn hierarchical representations directly from raw signals, reducing the reliance on handcrafted features and improving robustness across diverse recording conditions [22]. This paradigm shift has enabled substantial progress in a variety of medical applications, including gene expression analysis [59], cancer classification [60], cardiovascular risk stratification for conditions such as Short-QT and Brugada syndromes [61, 21], and multi-omics data integration [62–64].

Within the specific domain of ECG interpretation, deep learning has been increasingly adopted to automate tasks traditionally performed by expert cardiologists. These include arrhythmia classification, waveform segmentation, and the detection of clinically relevant features. Convolutional neural networks (CNNs), in particular, have proven effective in capturing local temporal patterns and morphological characteristics of ECG signals. A notable example is the use of one-dimensional CNN architectures for automatic arrhythmia classification, which have achieved high accuracy while operating directly on time-domain signals [5].

Beyond diagnostic classification, deep learning has also been applied to the digitization and structural interpretation of ECG records. The work in [55] demonstrated how neural models can be used to extract structured information from ECG reports, facilitating digital storage and downstream analysis. These efforts highlight the broader potential of learning-based approaches to reduce the manual burden associated with ECG interpretation and reporting.

Despite these advances, the problem of precise ECG fiducial point detection remains particularly challenging. While sharp events such as the R peak are relatively easier to localize, the identification of P and T wave boundaries is complicated by their smoother morphology, lower amplitude, and higher variability across subjects and physiological conditions. Several studies have shown that nonlinear features and temporal dynamics play a crucial role in capturing QRS-complex information and improving detection accuracy, even when using a limited number of features [20].

More recently, deep learning approaches have been extended beyond ECG interpretation alone to support related cardiovascular inference tasks, such as non-invasive blood pressure estimation from ECG and photoplethysmography signals [7, 8, 58]. These works further emphasize the importance of accurately modeling ECG morphology, as errors in feature localization can propagate and degrade performance in downstream physiological inference tasks.

Overall, the existing literature demonstrates that deep learning offers powerful tools for ECG analysis, yet also reveals persistent challenges related to interpretability, robustness, and precise temporal localization. In this context, automatic ECG feature detection serves as a well-defined and physiologically grounded benchmark task. It allows the evaluation of neural architectures in a setting where expert annotations provide a reliable reference, while also offering insights that are transferable to more ill-posed inverse problems addressed in later chapters of this thesis.

4.3 Methodological Overview

The methodological framework adopted in this chapter is designed to address the problem of automatic ECG fiducial point detection in a controlled and physiologically grounded setting. Rather than treating ECG analysis as a generic pattern recognition task, the methodology explicitly combines signal conditioning, temporal alignment, and supervised learning based on expert annotations. This choice reflects the dual objective of the study: achieving accurate temporal localization of ECG features while preserving interpretability and physiological meaning.

The overall pipeline consists of three main stages: signal selection and annotation alignment, preprocessing and transformation of ECG waveforms, and segmentation into learning-ready input windows. Each stage is tailored to minimize ambiguity in the learning targets and to ensure that the neural network is exposed to consistent and clinically meaningful signal representations.

4.3.1 Dataset and Annotation Framework

The study is conducted using the QT Database, a well-established benchmark dataset for ECG waveform analysis [47, 48]. The database includes a diverse collection of

ECG recordings acquired under different physiological and experimental conditions, such as ambulatory monitoring, exercise tests, and long-term recordings with durations ranging from approximately 14 to 22 hours. This diversity introduces realistic variability in heart rate, waveform morphology, and noise characteristics, making the dataset particularly suitable for evaluating the robustness of automatic feature detection methods.

A defining characteristic of the QT Database is the availability of expert annotated fiducial points. Trained cardiologists manually labeled the onset, peak, and offset of the main ECG waveform components, including the P wave, QRS complex, and T wave. These annotations follow standardized clinical criteria and provide a high-quality reference for supervised learning. Importantly, the presence of expert labels allows the task to be framed as a precise temporal localization problem rather than as a coarse classification task.

Signal extraction and annotation handling are performed using the WaveForm DataBase (WFDB) software package [65], developed within the PhysioNet ecosystem [48]. The WFDB tools enable consistent access to raw ECG signals and annotation files while preserving timing accuracy at the sample level. During this stage, segments without valid annotations are excluded to prevent ambiguity in the learning targets and to ensure that the training data faithfully represent clinically interpretable waveforms.

4.3.2 Baseline Elimination

Baseline wander represents a common source of low-frequency noise in electrocardiogram recordings and constitutes a significant obstacle to accurate waveform analysis. Slow variations in the ECG baseline may arise from respiration, body movements, electrode impedance changes, or measurement conditions, and can substantially interfere with the identification of clinically relevant features, particularly the P and T waves. For this reason, baseline removal is a fundamental preprocessing step prior to any automated ECG analysis.

In this work, baseline elimination is performed using a local linear regression approach. The method operates by fitting a linear model over a sliding temporal window of approximately 1.5 seconds. The estimated baseline value is taken at the midpoint of each window, after which the window is shifted forward by one sample




Fig. 4.1 Example of ECG baseline elimination preserving waveform morphology.

and the procedure is repeated across the entire signal. This local strategy allows slow baseline trends to be captured while preserving higher-frequency waveform components.

A key characteristic of the adopted algorithm lies in its robustness to transient deviations. By assigning reduced influence to outliers during the regression process, the method effectively treats sharp ECG components, such as the P wave, QRS complex, and T wave, as deviations from the local baseline rather than as part of it. As a result, these physiologically meaningful features are preserved, while slow baseline fluctuations are progressively removed.

The effectiveness of the baseline elimination procedure is illustrated in Fig. 4.1. As shown, the algorithm substantially reduces baseline instability without distorting the morphology or timing of the main ECG waveform components. This balance between noise suppression and feature preservation is essential for subsequent stages of the pipeline, particularly for tasks involving precise temporal localization of fiducial points.

4.3.3 Continuous Wavelet Transform and Scalogram Derivation

The Wavelet Transform is a mathematical tool designed to decompose a signal into components that are localized both in time and frequency [66]. In the context of

ECG analysis, this property is particularly valuable, as cardiac waveforms exhibit transient events whose spectral content evolves rapidly over time. In this work, the Continuous Wavelet Transform (CWT) [67] is employed to construct a time–frequency representation of the ECG signal in the form of a scalogram.

Prior to the application of the CWT, ECG signals are downsampled from 250 Hz to 125 Hz. This step reduces computational complexity while preserving the frequency content relevant for ECG morphology analysis. The CWT maps the one-dimensional ECG signal into a two-dimensional matrix of wavelet coefficients, indexed by scale and temporal position.

The Continuous Wavelet Transform (CWT) provides a joint time–scale representation of a signal by projecting it onto a family of functions obtained through dilation and translation of a prototype waveform, known as the mother wavelet. Let $x(t) \in L^2(\mathbb{R})$ denote a real-valued ECG signal. The CWT of $x(t)$ is defined as

$$W_x(a, b) = \frac{1}{\sqrt{|a|}} \int_{-\infty}^{+\infty} x(t) \psi^* \left(\frac{t-b}{a} \right) dt, \quad (4.1)$$

where $a \in \mathbb{R}^+$ is the scale parameter, $b \in \mathbb{R}$ is the translation parameter, $\psi(t)$ is the chosen mother wavelet, and $(\cdot)^*$ denotes complex conjugation.

The scale parameter a controls the temporal dilation of the wavelet and is inversely related to frequency content: small scales emphasize high-frequency, short-duration components, while large scales capture slower, low-frequency structures. This property is particularly relevant for ECG analysis, where sharp events such as QRS complexes coexist with smoother and more slowly varying components such as P and T waves.

The scalogram is defined as the squared modulus of the CWT coefficients,

$$S(a, b) = |W_x(a, b)|^2, \quad (4.2)$$

and represents the distribution of signal energy across time and scale. In practice, the scalogram provides a two-dimensional representation that highlights the temporal evolution of frequency-related features, offering a richer description of ECG morphology than the original one-dimensional waveform.

For a given mother wavelet with center frequency f_c , the scale parameter can be approximately related to physical frequency as

$$f \approx \frac{f_c}{a\Delta t}, \quad (4.3)$$

where Δt denotes the sampling period. This relation provides an intuitive interpretation of the scalogram axes and allows specific ECG components to be associated with characteristic scale ranges. In particular, QRS complexes primarily concentrate energy at smaller scales, while P and T waves are represented at larger scales due to their longer temporal duration.

In numerical implementations, the ECG signal is sampled as $x[n]$, $n = 0, \dots, N - 1$, and the CWT is evaluated over a discrete set of scales $\{a_m\}$ and time indices. The discrete approximation of (4.1) can be written as

$$W_x(a_m, n) = \frac{1}{\sqrt{a_m}} \sum_{k=0}^{N-1} x[k] \psi^* \left(\frac{k-n}{a_m} \right), \quad (4.4)$$

where the choice of the scale grid determines the resolution of the time–frequency representation.

Although the CWT does not constitute an orthogonal transform and therefore introduces redundancy, this redundancy is advantageous for learning-based approaches. The dense coverage of time–scale space ensures that subtle morphological variations are preserved, facilitating robust feature extraction by convolutional neural networks. In this framework, the scalogram serves as a structured intermediate representation that bridges classical time–frequency analysis and data-driven deep learning models.

Each coefficient captures the similarity between the signal and a scaled and shifted version of the selected mother wavelet, thereby encoding localized frequency information across multiple resolutions.

The resulting scalogram provides a visual representation of the signal’s time–frequency structure by displaying the absolute values of the CWT coefficients. As illustrated in Fig. 4.2 (bottom), the scalogram highlights how different frequency components evolve over time, offering an enriched description of ECG morphology compared to the original one-dimensional waveform. Color intensity in the scalogram corresponds to the magnitude of the wavelet coefficients, with darker colors indicating higher energy content and lighter colors denoting lower amplitudes.

By transforming ECG signals into two-dimensional representations, the scalogram enables convolutional neural networks to exploit both temporal and spectral features simultaneously. This representation enhances the model's ability to capture subtle morphological variations associated with different ECG components, such as P waves, QRS complexes, and T waves. In the adopted framework, scalograms serve as structured inputs that bridge classical signal processing techniques and deep learning-based feature extraction.




Fig. 4.2 Top: segmentation example of ECG with relative label for a QRS complex (red). Bottom: 2D Scalogram obtained through CWT method from the ECG.

4.3.4 Data Segmentation

Data segmentation is a fundamental step when processing long ECG recordings for deep learning applications. Convolutional neural networks operate most effectively on inputs of fixed and limited size, while excessively long sequences can degrade learning performance and hinder convergence. To address this issue, both ECG signals and their corresponding annotations are partitioned into shorter, fixed-length segments prior to model training.

In the adopted framework, each ECG recording is divided into windows of 125 samples, corresponding to one second of signal duration at a sampling frequency of 125 Hz. To preserve temporal continuity and reduce boundary effects, consecutive segments overlap by 10 samples. This overlapping strategy ensures smoother transitions between adjacent windows and mitigates the loss of relevant morphological information at segment boundaries.

Each ECG lead is segmented independently, and the resulting segments from both leads are retained as parallel input channels. In this way, the model can exploit complementary information encoded across leads during feature extraction. The segmented windows, each consisting of 125 samples per lead, constitute the direct input to the deep neural network.

Importantly, the segmentation process is designed to preserve physiological relevance. Each window is constructed so as to contain either a complete QRS complex or a partial QRS complex together with its associated label. This constraint ensures that each segment includes meaningful cardiac activity and prevents the network from being trained on windows devoid of diagnostically relevant information. As a result, the segmentation strategy supports accurate localization and characterization of ECG fiducial points.

Applying this procedure to the selected dataset yields a total of 15 089 segmented windows. An illustrative example of the segmentation process, together with the corresponding scalogram representation, is shown in Fig. 4.2 (top).

4.4 Deep Learning Architecture

The deep learning model adopted in this work is a hybrid architecture that combines convolutional and recurrent components, specifically two convolutional neural networks (CNNs) followed by two Bidirectional Long Short-Term Memory (BiLSTM) layers. The architecture is implemented in Python using the TensorFlow framework and was trained on a cloud-based environment equipped with GPU acceleration, enabling efficient optimization of the model parameters.

The convolutional layers are responsible for extracting salient features from the input ECG representations. By operating on the scalogram-derived inputs, the CNNs learn local spatial patterns associated with characteristic ECG morphologies, such as QRS complexes and P–T wave structures. These layers act as automatic feature extractors, transforming the input data into progressively higher-level representations suitable for temporal modeling.

To enable the transition from the convolutional blocks to the recurrent layers, a time-distributed flattening operation is applied. This step reshapes the two-dimensional feature maps produced by the CNNs into one-dimensional vectors while

preserving the temporal ordering of the data. The resulting sequence of feature vectors is then provided as input to the recurrent stage of the network.

Temporal dependencies within the ECG signal are modeled using Bidirectional Long Short-Term Memory networks. LSTM architectures are particularly well suited for ECG analysis due to their ability to capture long-range temporal relationships and to retain contextual information from previous time steps. The bidirectional formulation further enhances this capability by allowing the network to exploit both past and future temporal context when processing each segment, which is especially beneficial for detecting waveform components with variable duration and morphology.

The recurrent stage consists of two stacked BiLSTM layers, each followed by a dropout layer with a dropout rate of 25% to reduce overfitting. The output of the BiLSTM block is a sequence of predictions indicating the likelihood of the presence of a specific ECG feature at each time step. A time-distributed output layer is employed to generate one prediction per temporal position, ensuring alignment between the model output and the input signal structure.

Subsequently, batch normalization and an additional dropout layer (25%) are applied to improve training stability and generalization. The final classification stage is implemented using a dense layer with a softmax activation function, producing a binary output that represents the presence or absence of the target ECG feature.

The model is trained using the categorical cross-entropy loss function, while parameter optimization is performed using the Adaptive Moment Estimation (Adam) optimizer, selected for its computational efficiency and robust convergence properties. A detailed schematic of the overall network architecture is shown in Fig. 4.3.

4.5 Results

The detection of QRS complexes, P waves, and T waves follows the same modeling pipeline, network architecture, and training procedure described in the previous sections. Separate models are trained for each waveform component, while preserving an identical structural design and learning strategy across tasks.

To improve robustness and reduce bias due to data partitioning, a 10-fold cross-validation (CV) protocol is adopted. The dataset is randomly divided into ten




Fig. 4.3 Deep Neural Network Architecture

mutually exclusive subsets, with each fold acting once as a validation set while the remaining folds are used for training. This strategy enables a reliable estimation of generalization performance and mitigates overfitting effects.

For each fold, standard evaluation metrics are computed, including accuracy, categorical cross-entropy loss, sensitivity, specificity, and F1-score. Receiver Operating Characteristic (ROC) curves and Area Under the Curve (AUC) metrics are also evaluated to assess classification performance across different decision thresholds. Model training is carried out for 30 epochs using categorical cross-entropy as the loss function.

For all three detection tasks (QRS complex, P wave, and T wave), the final reported results correspond to the average performance across the ten folds. The aggregated scores are reported in Table 4.1. The proposed approach demonstrates high detection accuracy for all waveform components, achieving validation accuracies of 0.962 for P waves, 0.977 for QRS complexes, and 0.964 for T waves.

Table 4.1 Average scores across all folds for the analyzed ECG features

	Accuracy	Loss	Sensitivity	Specificity	F1-score
P wave	0.962	0.1228	0.878	0.974	0.8557
QRS complex	0.977	0.0726	0.923	0.981	0.9126
T wave	0.964	0.1174	0.934	0.974	0.9247

Figure 4.4 reports the training loss curves for the three networks corresponding to QRS complex, P wave, and T wave detection. Different axis scales are used to accommodate the distinct loss magnitudes associated with each task. For clarity, only the first 16 training epochs are displayed, as the loss function exhibits minimal variation beyond approximately the 15th epoch. This choice allows a clearer visualization of the convergence behavior during the most informative training phase.

Qualitative examples of model predictions on randomly selected test segments are shown in Fig. 4.5 for QRS complexes, Fig. 4.6 for P waves, and Fig. 4.7 for T waves. These examples illustrate the model's ability to accurately localize ECG waveform components across diverse morphologies.

To further contextualize the performance of the proposed method, a comparative analysis with established approaches in the literature is conducted. Table 4.2 summarizes the results obtained with the proposed 2D-CNN-BiLSTM model alongside ECG-SegNet [68] and Hidden Markov Model (HMM)-based methods [69]. Both

Table 4.2 Comparison of detection performance with different methods

Methods	P wave	QRS Complex	T wave
2D-CNN-BiLSTM	0.962	0.971	0.965
ECG-SegNet (Bidirectional LSTM)	0.920	0.940	0.920
HMM on raw ECG data	0.551	0.790	0.560
HMM on wavelet-encoded ECG	0.742	0.944	0.882




Fig. 4.4 Training loss (blue) and validation loss (orange) for QRS (top left), P wave (top right), and T wave (bottom) networks.

HMM implementations on raw ECG data and wavelet-encoded ECG signals are included for completeness.

The comparison highlights that the proposed architecture achieves superior sensitivity and specificity across all waveform components, significantly outperforming classical HMM-based approaches and improving upon previously reported deep learning methods.

4.6 Discussion and Conclusion

This chapter addressed the problem of automatic ECG feature detection, framing it as a supervised learning task with well-defined physiological targets and expert-




Fig. 4.5 Example of prediction of QRS complexes with the relative original label.




Fig. 4.6 Example of prediction of P wave with the relative original label.

provided annotations. The proposed hybrid 2D-CNN-BiLSTM architecture demonstrated strong performance in the localization of P waves, QRS complexes, and T waves, confirming that deep learning models can effectively learn ECG morphology when the inference problem is well posed and the reference labels are physiologically unambiguous.

From a methodological perspective, this study highlights how the combination of time–frequency representations and sequence-aware architectures enables robust feature extraction in cardiac signals. The use of Continuous Wavelet Transform scalograms enriches the representation of ECG dynamics, while convolutional layers capture localized morphological patterns and recurrent layers model temporal dependencies across successive cardiac events. The achieved results, which outperform classical signal-processing and probabilistic approaches, demonstrate that deep learning can reliably replicate expert-level ECG feature annotation under controlled conditions.




Fig. 4.7 Example of prediction of T wave with the relative original label.

Beyond the quantitative performance, the significance of this chapter lies in its role within the broader context of the thesis. ECG fiducial point detection represents a learning problem in which the physiological meaning of the target variables is clear, stable across subjects, and only weakly affected by inter-individual variability. While signal quality, noise, and morphology differ across recordings, the underlying electrical events defining P, QRS, and T waves remain invariant. As a result, the mapping from signal to label is largely non-degenerate, making this task an ideal benchmark for studying how neural networks encode physiological waveform structure.

This stands in sharp contrast to the problems addressed in the subsequent chapters, particularly blood pressure estimation from ECG and PPG signals. In those settings, the learning target is not directly observable in the signal and depends on latent physiological parameters such as vascular compliance, peripheral resistance, and autonomic regulation. By first establishing that deep learning models can successfully solve a physiologically grounded and well-defined inference task, this chapter provides a reference point against which more ill-posed inverse problems can be analyzed.

Importantly, the results presented here suggest that strong model performance alone is not sufficient to guarantee physiological inference. The success of the proposed model is closely tied to the availability of expert annotations and to the limited ambiguity of the task. This observation reinforces the central argument introduced in Chapter 1: when physiological constraints are weak or implicit, data-driven models may achieve high accuracy without learning generalizable or causal relationships.

In this sense, ECG feature detection serves as a controlled baseline that informs the interpretation of later failures and limitations observed in cuffless blood pressure estimation and stress modeling. The transition from this chapter to the following ones marks a shift from morphology reconstruction to physiological inference, where ambiguity, subjectivity, and degeneracy play a dominant role. The insights gained here motivate the integration of explicit physiological structure into learning frameworks, a direction that is explored through physics-informed and physiology-guided models in the remainder of the thesis.

Finally, while the proposed approach demonstrates strong performance on curated datasets, further validation on larger and more diverse populations remains necessary to assess its robustness in real-world settings. Nevertheless, the chapter establishes a clear methodological and conceptual foundation for the subsequent exploration of physiology-aware learning in cardiovascular and autonomic signal analysis.

Chapter 5

Data-Driven Blood Pressure Estimation from ECG and PPG

5.1 Cuffless Blood Pressure Estimation: Motivation and Problem Setting

The estimation of arterial blood pressure (BP) from non-invasive biosignals is one of the most intensively investigated problems in contemporary cardiovascular signal processing. Blood pressure is a fundamental physiological variable, tightly linked to cardiovascular risk, organ perfusion, and long-term morbidity and mortality. Continuous and unobtrusive BP monitoring would enable early detection of pathological trends, longitudinal risk stratification, and personalized health management, particularly in ambulatory and wearable scenarios.

Despite its clinical relevance, direct BP measurement remains inherently invasive or intermittent. Invasive arterial catheters provide accurate beat-to-beat pressure waveforms but are restricted to intensive care environments, while cuff-based sphygmomanometers offer only sparse measurements and are poorly suited for continuous or high-frequency monitoring. These limitations have driven a sustained research effort toward *cuffless* blood pressure estimation, aiming to infer systolic and diastolic blood pressure (SBP/DBP) from non-invasive biosignals acquired by portable or wearable devices [14, 8, 7].

Among the available biosignals, the electrocardiogram (ECG) and the photoplethysmogram (PPG) have emerged as the most widely adopted inputs for cuffless BP estimation. ECG records the electrical activity that triggers cardiac contraction, while PPG captures peripheral blood volume changes induced by the propagating pressure wave through optical sensing. Although neither signal directly measures arterial pressure, their temporal structure and morphology are influenced by cardiovascular dynamics, vascular properties, and autonomic regulation. As a result, ECG and PPG provide complementary information, and their joint use has become a standard choice in learning-based BP prediction pipelines [58, 7].

From a methodological standpoint, most data-driven approaches formulate cuffless BP estimation as a supervised learning problem, seeking to learn a mapping of the form

$$(\text{ECG}, \text{PPG}) \longrightarrow (\widehat{\text{SBP}}, \widehat{\text{DBP}}),$$

where the target quantities are derived from reference arterial pressure measurements. In practice, this mapping is typically implemented using deep neural networks, which automatically extract features from raw or lightly processed signals. The literature reports a wide variety of architectures, including convolutional neural networks (CNNs) [57], recurrent neural networks (RNNs) [8], and hybrid CNN–RNN solutions [58], sometimes combined with multimodal fusion strategies [70, 71]. This architectural diversity reflects an implicit assumption: given sufficient data, neural networks can approximate the complex and subject-dependent relationship between electrical activation, peripheral perfusion, and arterial pressure [72, 73, 20].

However, blood pressure estimation differs fundamentally from the ECG morphology learning tasks discussed in Chapter 4. While ECG fiducial point detection relies on well-defined electrical events with standardized physiological meaning, blood pressure is a *latent hemodynamic quantity*. It is not directly observable in ECG or PPG signals, but instead emerges from a complex interaction between cardiac contractility, vascular compliance, peripheral resistance, wave reflections, and autonomic modulation. Consequently, the inverse mapping from surface biosignals to absolute BP values is inherently ill-posed and potentially non-unique.

This distinction has important methodological implications. In ECG morphology learning, the learning target is physiologically unambiguous and largely invariant across subjects. In contrast, identical ECG and PPG morphologies may correspond to substantially different blood pressure levels across individuals, depending on age,

arterial stiffness, vascular tone, and autonomic state. As a result, high predictive accuracy on a given dataset does not necessarily imply physiological generalization or subject-independent validity.

Nevertheless, recent deep learning approaches have reported impressive performance in BP prediction tasks when evaluated under standard experimental protocols [8, 7]. Convolutional and recurrent architectures have been shown to exploit subtle temporal and morphological patterns in ECG and PPG signals, achieving low errors on large-scale clinical datasets. At the same time, a growing body of evidence indicates that such performance can be strongly influenced by dataset-specific biases, subject overlap between training and test sets, and implicit calibration effects. In particular, when subject-wise separation is enforced or when models are evaluated on unseen individuals or acquisition setups, predictive performance often degrades markedly.

This chapter focuses on the *pre-physics-informed* phase of the research trajectory developed in this thesis. Rather than presenting data-driven models as final solutions, the objective is threefold: (i) to formalize what can realistically be achieved using purely data-driven architectures, (ii) to clarify the role of dataset design and evaluation protocols in shaping reported performance, and (iii) to motivate, through concrete empirical evidence, the need for more structured modeling approaches.

To this end, the chapter integrates three complementary contributions that progressively expose both the potential and the limitations of deep learning for cuffless BP estimation. These include learning BP from large public clinical datasets, within-subject absolute BP prediction using dual-branch architectures, and wearable-oriented acquisition and inference through the PulsECG platform. In addition, two exploratory subject-wise analyses are discussed, highlighting a central theme of the chapter: the tension between *apparent accuracy* and *true generalization*.

By explicitly addressing these issues, this chapter establishes the conceptual and empirical foundation for the physics-informed and physiology-guided models introduced in the following chapter. The transition to PINN-based approaches is thus motivated not as an aesthetic refinement, but as a methodological response to concrete limitations observed in data-driven BP estimation: limited identifiability, fragile out-of-distribution behavior, and evaluation instability under subject-wise generalization.

5.2 State of the Art on Data-Driven Blood Pressure Estimation

The problem of cuffless blood pressure estimation from non-invasive biosignals has been widely explored over the past decade, driven by the increasing availability of large-scale physiological datasets and by advances in machine learning. Early approaches relied primarily on handcrafted features derived from ECG and PPG signals, such as pulse transit time (PTT), pulse arrival time (PAT), and morphological descriptors, combined with linear or shallow regression models. While these methods offered a degree of interpretability, their performance was often limited by sensitivity to noise, calibration requirements, and poor generalization across subjects and recording conditions.

The introduction of deep learning significantly reshaped this landscape by enabling end-to-end feature learning directly from raw or minimally processed signals. Convolutional neural networks (CNNs) have been extensively employed to extract local temporal and morphological patterns from ECG and PPG waveforms, leveraging their ability to model translation-invariant features and multi-scale structures [57, 74]. Recurrent architectures, such as Long Short-Term Memory (LSTM) networks, have been adopted to capture temporal dependencies across cardiac cycles, addressing the inherently sequential nature of cardiovascular signals [8].

Hybrid CNN–RNN architectures have emerged as a particularly popular design choice, combining convolutional layers for local feature extraction with recurrent layers for temporal integration. Several studies have demonstrated that such architectures can outperform purely convolutional or purely recurrent models in blood pressure estimation tasks [58, 7]. In these frameworks, CNNs typically operate on short windows of ECG and PPG to extract spatial or morphological features, while LSTMs or BiLSTMs aggregate information over time to produce SBP and DBP estimates.

In parallel, alternative representations of the input signals have been explored to enrich the information available to neural networks. Time–frequency transformations, including wavelet-based encodings, have been used to expose spectral dynamics associated with pulse propagation and vascular response, enabling models to exploit both temporal and frequency-domain cues [74]. More recently, spatio-temporal architectures have been proposed to jointly model signal evolution across time and

across multiple processing branches, aiming to improve absolute BP prediction accuracy [75].

A common experimental setting in the literature involves large public clinical datasets derived from intensive care unit (ICU) monitoring, most notably the MIMIC family of databases [76]. These datasets provide synchronized ECG, PPG, and invasive arterial blood pressure (ABP) waveforms, enabling supervised learning with high-quality reference targets. Their scale and accessibility have made them a de facto benchmark for cuffless BP estimation, and many reported state-of-the-art results are obtained under this setting [52, 8].

However, several works have highlighted that performance reported on ICU datasets may not translate directly to real-world or wearable scenarios. ICU data are acquired under controlled conditions using clinical-grade sensors, and often exhibit limited subject diversity in terms of physiological states and activity levels. Moreover, evaluation protocols based on random window splitting can inadvertently introduce strong within-subject correlations between training and test sets, leading to optimistic performance estimates [73, 24, 23].

To address these concerns, some studies have focused on subject-specific or within-subject modeling, showing that deep learning models can achieve very low prediction errors when trained and tested on data from the same individual [58, 7]. While these results demonstrate the expressive power of neural architectures, they also raise questions about generalizability and identifiability: strong performance in a single-subject setting does not necessarily imply that the learned mapping reflects a universal physiological relationship.

Across these diverse methodological choices, a recurring theme emerges: purely data-driven models can achieve high accuracy under favorable conditions, yet their performance is fragile with respect to subject-wise generalization, dataset shift, and changes in acquisition context. Architectural complexity alone does not resolve these issues, and improvements in benchmark metrics often coexist with limited interpretability and weak physiological grounding.

In summary, the existing literature on data-driven cuffless blood pressure estimation demonstrates both the promise and the limitations of deep learning approaches. CNNs, RNNs, and hybrid architectures are capable of extracting informative patterns from ECG and PPG signals, but the inverse mapping from surface biosignals to absolute blood pressure remains ill-posed and highly subject-dependent. These

observations motivate a transition from purely statistical learning toward models that explicitly incorporate physiological structure, a direction that is developed in the subsequent chapters of this thesis.

Within this context, the present chapter consolidates three complementary contributions that collectively explore the capabilities and limitations of data-driven approaches to cuffless blood pressure estimation. Although developed at different stages of the research trajectory, these works share a common methodological core based on deep learning applied to ECG and PPG signals, while differing in dataset characteristics, modeling choices, and evaluation focus.

The first contribution investigates non-invasive blood pressure estimation on large public clinical datasets, adopting convolutional and recurrent architectures trained on synchronized ECG, PPG, and reference blood pressure signals derived from ICU monitoring. This setting enables benchmarking against existing methods and highlights the apparent effectiveness of end-to-end learning when evaluated under standard protocols.

The second contribution extends this paradigm toward absolute blood pressure prediction through a structured space–time architecture, aiming to capture richer temporal dependencies and to reduce reliance on implicit calibration on a smaller custom dataset.

The third contribution shifts the focus toward wearable-oriented acquisition and inference, through the PulsECG platform. By operating under realistic sensing constraints and explicitly controlling the acquisition pipeline, this work emphasizes the challenges of deploying cuffless BP estimation outside clinical environments, and reveals how dataset scale and subject diversity critically shape achievable performance.

In the following sections, these three works are presented in detail, following a coherent progression from large-scale clinical data to subject-specific modeling and wearable deployment. Together, they provide a structured empirical basis for analyzing the strengths and failure modes of purely data-driven blood pressure estimation, and set the stage for the physiology-informed approaches introduced in the next chapter.

5.3 Conv1D BiLSTM for Cuffless Arterial Blood Pressure Estimation from ECG and PPG

This section presents the first data-driven blood pressure estimation framework developed within the research trajectory of this thesis. The model is based on a hybrid Conv1D–BiLSTM neural architecture and aims to reconstruct the arterial blood pressure (ABP) waveform from non-invasive electrocardiogram (ECG) and photoplethysmography (PPG) signals, together with the first- and second-order temporal derivatives of PPG. The work constitutes an early but representative example of deep learning–based cuffless blood pressure estimation, and serves as a reference baseline against which subsequent methodological refinements and limitations can be critically analyzed.

5.3.1 Problem Formulation and Modeling Strategy

The task addressed in this work is formulated as a regression problem in which the network learns a mapping from multimodal biosignals to a continuous arterial pressure waveform. Given synchronized segments of ECG, PPG, and derived PPG features, the objective is to estimate the corresponding ABP signal over a fixed temporal window. From the reconstructed waveform, systolic blood pressure (SBP) and diastolic blood pressure (DBP) can then be derived as secondary quantities.

This formulation follows one of the two dominant paradigms in cuffless BP estimation, namely waveform reconstruction rather than direct scalar prediction. Compared to direct SBP/DBP regression, waveform-based approaches preserve richer hemodynamic information and allow post hoc extraction of clinically relevant parameters. However, they also impose stricter requirements on temporal alignment, signal quality, and learning stability.

The approach is intentionally designed to remain purely data-driven, without incorporating explicit physiological constraints or mechanistic models. At this stage of the thesis, the goal is to assess what can be achieved through architectural design and signal preprocessing alone, under standard supervised learning assumptions.

5.3.2 Dataset and Signal Preprocessing

This section describes in detail the dataset adopted in this study and the complete preprocessing and selection pipeline used to construct learning-ready input–output pairs for ABP waveform prediction. The goal of this phase is twofold: (i) to standardize and condition ECG/PPG signals to improve the learning stability of the network, and (ii) to remove corrupted or physiologically implausible windows that would otherwise dominate the loss and mislead the model, especially in a waveform-regression setting.

5.3.3 Database

The experiments are conducted on the *Cuff-Less Blood Pressure Estimation Dataset* (UCI), explained in section 3.1.4, selected for its accessibility and ready-to-use organization [52, 53].

In the present work, additional preprocessing and stringent window-level selection are required to ensure that the training examples remain compatible with the objective of reconstructing ABP waveforms and extracting reliable SBP/DBP values.

5.3.4 Signal Filtering and Normalization

Signal conditioning is applied to ECG and PPG in order to attenuate low-frequency drift, high-frequency noise, and artifacts that would impair the extraction of temporal and morphological features. Following the preprocessing strategy in [77], PPG signals are filtered using a 4th-order Butterworth band-pass filter with passband $[0.5, 8]$ Hz. This range is intended to suppress baseline wandering and high-frequency components not informative for the pulsatile waveform.

ECG signals are filtered using an 8th-order Chebyshev type I band-pass filter with cut-off frequencies at 2 Hz and 59 Hz, to mitigate motion artifacts as well as powerline-related components. After filtering, each input channel is normalized *instance-wise* through min–max scaling to $[0, 1]$. This choice enforces comparable numerical ranges across subjects and windows while preserving the within-window morphology that drives the mapping to ABP.

Preservation of ABP amplitude information. Crucially, ABP is *not* filtered nor normalized in this stage. Since the learning target is the ABP waveform itself, any preprocessing that rescales or distorts amplitude would directly compromise the physical meaning of the predicted pressure signal and the accuracy of SBP/DBP extraction. For this reason, ABP is only used as-is as the regression target, and all conditioning steps are applied exclusively to the input signals (ECG/PPG and derivatives).

5.3.5 Windowing Strategy

Signals in the UCI dataset are sampled at 125 Hz. Each instance may contain up to 10 minutes of recording; therefore, each channel can reach up to 75,000 samples per instance. For supervised learning, the continuous recordings are divided into fixed-length windows of 8 seconds, corresponding to 1000 samples per window. The 8-second duration was chosen to include multiple cardiac cycles while maintaining quasi-stationarity and manageable computational complexity, and is aligned with comparable design choices in the literature (e.g., 1024-sample windows in [78]).

Each resulting window provides:

- input signals: ECG and PPG (plus PPG derivatives, when used);
- output signal: the corresponding 8-second ABP waveform segment.

5.3.6 Physiological Thresholding and Artifact Pruning

Although the dataset is curated, a substantial fraction of extracted windows still exhibits severe distortions or implausible pressure ranges that prevent the model from learning a stable mapping from ECG/PPG to ABP. Empirically, the most problematic windows fall into a small set of abnormal blood pressure regimes. Therefore, a first selection step removes any window whose ABP-derived summary values satisfy one of the following conditions:

- $SBP < 80 \text{ mmHg}$ or $SBP > 190 \text{ mmHg}$;
- $DBP < 50 \text{ mmHg}$ or $DBP > 120 \text{ mmHg}$;

- pulse pressure range $|\text{SBP} - \text{DBP}| < 20 \text{ mmHg}$ or $|\text{SBP} - \text{DBP}| > 120 \text{ mmHg}$.

All windows violating these constraints are discarded together with their corresponding ECG and PPG segments.

Beat-to-beat consistency pruning via peak statistics. A second selection step explicitly targets non-uniform or noisy acquisitions (e.g., due to movements during monitoring). Because ECG, PPG, and ABP peaks may become irregular in both amplitude and spacing under artifacts, each extracted window is further evaluated by computing peak locations and peak heights for each channel, and measuring the standard deviation of: (i) peak-to-peak distances, and (ii) peak heights.

Windows are pruned if the standard deviation exceeds fixed thresholds, summarized in Table 5.1. The thresholds were set empirically to filter out noisy windows while retaining physiologically plausible variability.

Table 5.1 Pruning thresholds for the standard deviation σ of peak distances and peak heights in PPG, ECG, and ABP windows.

	PPG	ECG	ABP
σ of peak distances	> 6	> 6	> 6
σ of peak heights	> 0.1	> 0.1	> 6

5.3.7 PPG Derivatives as Additional Input Channels

To enrich the input representation, the first and second temporal derivatives of PPG are optionally included as additional channels. Denoting

$$\text{VPPG}(t) = \frac{d\text{PPG}(t)}{dt}, \quad \text{APPG}(t) = \frac{d\text{VPPG}(t)}{dt},$$

these derivatives have been reported as informative features for blood pressure prediction in prior work [78, 79]. In the original implementation, derivatives are computed via numerical differentiation (e.g., gradient-based methods), and additional

filtering may be applied to avoid amplification of high-frequency noise introduced by differentiation (particularly for higher-order derivatives). When filtering is used in this derivative pipeline, resynchronization can be required to compensate for phase delays introduced by the filters.

In this work, the final input tensor is constructed using four aligned channels: PPG, ECG, VPPG, and APPG, each of length 1000 samples (8 seconds).

5.3.8 Network Input–Output Definition

After windowing, pruning, and derivative computation, each training example is represented as:

- **Input:** a multichannel tensor of shape $(1000, 4)$ formed by stacking, in this order, PPG, ECG, VPPG, and APPG.
- **Target:** the corresponding ABP waveform window of length 1000 samples (8 seconds).

An illustrative example of the resulting input channels and their associated ABP target waveform is shown in Fig. 5.1. The figure highlights how the derivatives emphasize rapid morphological changes in the PPG signal and provide complementary information to the raw PPG and ECG traces.




Fig. 5.1 Example of input vector over time (x -axis) composed, from top to bottom (y -axis), by: PPG, ECG, VPPG, APPG; and then the relative ABP to be predicted.

5.3.9 SBP/DBP Extraction from ABP for Evaluation

Although the model is trained to regress the full ABP waveform, systolic and diastolic blood pressure values are derived from the waveform to support standard evaluation protocols. SBP and DBP are computed through peak and trough detection on the ABP signal. Peak detection is performed via Python's `scipy.signal.find_peaks` routine, which requires a prominence threshold to identify meaningful systolic peaks. Because ABP amplitude ranges vary across windows and subjects, prominence is chosen adaptively on a per-window basis. In particular, prominence is empirically set as the difference between the median and the minimum of the ABP window, which provides a robust scale estimate under moderate noise.

Figure 5.2 illustrates an example ABP segment with detected systolic peaks (SBP, orange markers) and diastolic minima (DBP, green markers), together with the final

SBP/DBP values computed as the mean of the detected peaks and troughs within the window.




Fig. 5.2 Fraction of ABP signal (*y-axis*) over time (*x-axis*) with: (a) highlighted peaks (SBP, orange stars) and minimum points (DBP, green stars); (b) extracted SBP (mean of peaks) and DBP (mean of minimums).

5.3.10 Extracted Dataset Statistics and Pressure Distributions

After preprocessing, physiological thresholding, and pruning based on peak statistics, a total of 192,661 valid input tensors are retained. Since each tensor corresponds to an 8-second segment, the resulting training material amounts to 428.13 hours of synchronized ECG/PPG/ABP data.

To characterize the retained dataset and understand the effective supervision range, it is useful to analyze the empirical SBP and DBP distributions obtained after the selection steps. Figure 5.3 reports the SBP distribution (left) and DBP distribution (right) derived from the ABP windows. Due to the explicit pruning thresholds, the distributions are truncated at the imposed lower/upper bounds. As commonly observed, DBP values exhibit lower dispersion compared to SBP, which may influence the relative difficulty of predicting systolic versus diastolic pressure levels.




Fig. 5.3 Distribution of: SBP and DBP extracted values (**left**); σ of extracted values of SBP, DBP (**right**).

5.3.11 Network Architecture

The proposed model is designed as a hybrid deep learning architecture combining convolutional feature extraction and sequence modeling to reconstruct the arterial blood pressure (ABP) waveform from non-invasive biosignals. The design is inspired by the strong results reported in [77] and [78]. However, the architecture in [78] is substantially larger (approximately 120 million parameters), which makes a direct adoption impractical for this task when compared to lighter yet effective alternatives such as [77] (around 2 million parameters). For this reason, the final choice is a more compact network that relies on a stack of Conv1D layers (128 filters, kernel size 3) and Bidirectional LSTM (BiLSTM) layers.

Conv1D layers [57, 80] are widely used in time-series learning because they can extract local temporal patterns through shared kernels, making them effective in the presence of noise and morphological variability. In this setting, convolutional blocks act as trainable feature extractors that transform the multichannel input (PPG, ECG, VPPG, APPG) into a richer representation suitable for downstream temporal modeling.

BiLSTM networks [77] are adopted to capture temporal dependencies in both forward and backward directions. This is particularly relevant for waveform regression, where the instantaneous pressure value depends not only on local signal morphology but also on contextual information from neighboring samples within a cardiac cycle. The bidirectional structure allows the model to exploit such context symmetrically.

Architecturally, the network follows an encoder–decoder philosophy based on convolutional blocks, while using BiLSTM layers as the backbone instead of a purely feedforward regressor (as in [78]). A skip connection is introduced across the recurrent backbone to improve gradient flow and mitigate vanishing-gradient effects during training. The complete structure is illustrated in Fig. 5.4.




Fig. 5.4 Structure of the first network.

5.3.12 Experiments and Results

This section reports the experimental protocol and the results of an ablation study aimed at assessing the impact of selected architectural and training choices (MaxPool removal, number of BiLSTM layers, and batch size) on ABP waveform prediction performance.

5.3.13 Data split and reproducibility

Before training, all extracted windows are randomly shuffled using a fixed random seed to ensure comparability across experiments. The dataset is then split into:

- **Training set:** 80% of the available windows,

- **Validation set:** 10%,
- **Test set:** 10%.

This protocol is identical across experiments, such that observed differences can be attributed to model configuration changes rather than to data partitioning.

5.3.14 Training configuration and metrics

All experiments are performed for 30 epochs using the Adam optimizer with Keras default parameters [81, 82]. The loss function is Mean Squared Error (MSE), while Mean Absolute Error (MAE) is used as the primary monitored metric. This selection follows common practice in time-series regression literature and facilitates comparisons with state-of-the-art works, where MAE is typically reported as an interpretable absolute deviation measure, and MSE is used for optimization stability.

Due to the presence of a MaxPool layer in the baseline network configuration, the stride of the Conv1DTranspose layer is set to 2 to preserve input/output dimensional consistency. This ensures that the reconstructed waveform matches the original ABP window length (1000 samples), avoiding misalignment in the regression target.

5.3.15 Experiment 1: Baseline configuration (MaxPool, 3 BiLSTM, batch size 64)

The first experiment uses the architecture in Fig. 5.4, configured with batch size 64, a MaxPool layer in the convolutional encoder, and 3 stacked BiLSTM layers in the backbone. Training and validation trends for MSE and MAE are reported in Fig. 5.5.




Fig. 5.5 Error (y-axis) over epochs (x-axis) for the first experiment (batch size = 64, 3 BiLSTM and MaxPool layer): (a) MSE on training (blue line) and validation (red line) sets; (b) MAE on training (blue line) and validation (red line) sets.

5.3.16 Experiment 2: MaxPool removal (no MaxPool, 3 BiLSTM, batch size 64)

The second experiment investigates the effect of removing the MaxPool layer while preserving the rest of the configuration (3 BiLSTM layers, batch size 64). Removing MaxPool increases computational cost; in the reported setup, the training time increases from approximately 2.5 hours to 4 hours. The corresponding MSE/MAE curves are shown in Fig. 5.6.




Fig. 5.6 Error (y-axis) over epochs (x-axis) for the second experiment (batch size = 64, 3 BiLSTM and no MaxPool layer): (a) MSE on training (blue line) and validation (red line) sets; (b) MAE on training (blue line) and validation (red line) sets.

A direct comparison between the MAE curves of Experiment 1 (with MaxPool) and Experiment 2 (without MaxPool) is reported in Fig. 5.7. The two configurations show nearly identical convergence behavior, with only minor deviations at early epochs and similar asymptotic trends. Since MaxPool removal does not yield a consistent improvement while increasing training time, subsequent experiments retain the MaxPool layer to favor computational efficiency.




Fig. 5.7 Comparison of MAE (*y-axis*) over epochs (*x-axis*) between the configurations in Fig. 5.5b (with MaxPool) and Fig. 5.6b (without MaxPool).

5.3.17 Experiment 3: Reduced recurrent depth (MaxPool, 2 BiLSTM, batch size 64)

The third experiment evaluates the impact of the number of BiLSTM layers by reducing the recurrent backbone depth from 3 to 2 BiLSTM layers, while keeping batch size (64) and the MaxPool layer unchanged. Figure 5.8 compares the MSE and MAE trends of this configuration against the prior two experiments.




Fig. 5.8 Error (y -axis) over epochs (x -axis) for the third experiment (batch size = 64, 2 BiLSTM and MaxPool layer) compared to the previous two networks: (a) MSE on training and validation sets; (b) MAE on training and validation sets.

The comparison indicates that reducing the recurrent depth degrades performance, particularly in validation MAE, which does not reach the levels achieved by the 3-BiLSTM configurations. This suggests that the additional recurrent capacity contributes materially to learning the temporal structure required for accurate ABP waveform reconstruction.

5.3.18 Experiment 4: Increased batch size (MaxPool, 3 BiLSTM, batch size 256)

The last experiment investigates the effect of increasing the batch size to 256, while maintaining MaxPool and 3 BiLSTM layers. The resulting MAE trends are summarized in Fig. 5.9. Although the network continues to learn, the convergence appears slower compared to the smaller batch-size configurations, particularly in the early epochs. At later epochs, the training MAE approaches values similar to other configurations, but the validation MAE does not outperform the baseline batch size 64 setting.




Fig. 5.9 Comparison of MAE (y -axis) over epochs (x -axis): (a) all network configurations (training only); (b) first and last configurations showing training and validation trends.

Table 5.2 Summary of the ablation experiments on the Conv1D–BiLSTM architecture for ABP waveform prediction.

Exp.	Batch size	MaxPool	#BiLSTM	Notes / Outcome
1	64	Yes	3	Baseline configuration (Fig. 5.4)
2	64	No	3	Similar MAE convergence to Exp. 1, but higher computational cost
3	64	Yes	2	Worse validation MAE; reduced recurrent depth degrades performance
4	256	Yes	3	Slower convergence; validation MAE does not improve over Exp. 1

5.3.19 Qualitative ABP waveform reconstructions

For completeness, qualitative examples of ABP waveform predictions obtained on random samples from the test set are reported in Fig. 5.10. These examples illustrate the capability of the trained model to reproduce different ABP morphologies across windows, highlighting both the oscillatory structure associated with cardiac cycles and the variability in waveform shape.




Fig. 5.10 Examples of predicted ABP (*y-axis*) over time (*x-axis*) from the baseline network configuration (batch size = 64, MaxPool and 3 BiLSTM layers). Predictions are shown for random inputs from the test set to illustrate the model's ability to reproduce different waveform morphologies.

5.3.20 Discussion and Conclusions

In this section, a data-driven approach for non-invasive arterial blood pressure estimation has been presented and analyzed in detail. The proposed framework aims at predicting short segments of the arterial blood pressure (ABP) waveform directly from non-invasive biosignals, relying on a lightweight deep learning architecture designed to balance predictive accuracy and computational efficiency.

The core contribution of this work lies in the development of a hybrid Conv1D–BiLSTM neural network that combines convolutional encoder–decoder blocks with recurrent temporal modeling. Inspired by existing architectures in the literature, the proposed model deliberately avoids excessive parameterization, resulting in a compact architecture with a parameter count that is orders of magnitude lower than alternative state-of-the-art solutions. This design choice reflects a clear methodological intention: to explore how far purely data-driven models can go in predicting ABP

waveforms while remaining compatible with embedded and wearable deployment constraints.

An extensive ablation study was conducted to assess the influence of key architectural choices, including the presence of MaxPooling layers, the depth of the recurrent backbone, and the batch size used during training. Results consistently indicate that the original configuration, comprising three BiLSTM layers, MaxPooling, and a batch size of 64, provides the best trade-off between convergence speed, stability, and predictive accuracy. Under standard train–validation–test splits, this configuration achieves a mean absolute error of approximately 2.97 mmHg on unseen test data, which is competitive with, and in some cases superior to, previously reported results on the same dataset.

From a signal modeling perspective, the results demonstrate that convolutional and recurrent neural networks are capable of learning meaningful representations of cardiovascular dynamics directly from ECG and PPG signals, at least within the statistical structure of curated datasets. The network successfully reproduces a wide range of ABP morphologies, including variations in pulse shape and amplitude, suggesting that relevant temporal and morphological cues are effectively captured by the learned representations.

However, when interpreted in the broader context of this thesis, these findings also highlight important limitations. While the reported performance metrics are strong under conventional evaluation protocols, they do not, by themselves, guarantee subject-independent generalization or physiological interpretability.

5.4 Wearable-Oriented Cuffless Blood Pressure Estimation with PulsECG

The previous section analyzed data-driven blood pressure estimation approaches developed on large public clinical datasets, where ECG, PPG, and invasive arterial blood pressure signals are acquired under controlled ICU conditions. While such datasets enable large-scale training and benchmarking, they also introduce important limitations related to domain mismatch, subject leakage, and lack of control over the acquisition and preprocessing pipelines.

In this section, as the second contribution of this chapter, the focus shifts toward a complementary and more realistic scenario: wearable-oriented blood pressure estimation using a custom-designed acquisition device. Rather than relying on retrospective clinical databases, the proposed approach introduces *PulseECG*, a portable system specifically developed to acquire synchronized ECG and PPG signals under controlled yet everyday-like conditions. This transition allows the investigation of cuffless blood pressure estimation in a setting that is closer to real-world usage, while retaining full control over hardware design, signal acquisition, and preprocessing.

5.4.1 Introduction and Motivation

Medical devices play an increasingly central role in modern healthcare systems, particularly in the context of disease prevention, remote monitoring, and long-term management of chronic conditions. These devices enable continuous or near-continuous monitoring of physiological parameters, allowing clinicians to track patient status over time and intervene promptly when deviations occur. Such capabilities are especially relevant for cardiovascular diseases, which remain among the leading causes of morbidity and mortality worldwide [83].

The global medical device market encompasses millions of distinct products spanning thousands of generic device categories [84]. Within this vast landscape, portable and wearable devices represent a rapidly expanding segment, driven by their ease of use, reduced invasiveness, low power consumption, and suitability for long-term monitoring outside clinical environments. Electrocardiography (ECG) [85] and photoplethysmography (PPG) [86, 38] have emerged as two cornerstone sensing modalities in this context, due to their non-invasive nature and strong physiological relevance.

Recent years have also seen a growing integration of machine learning (ML) and deep learning (DL) techniques into biomedical signal analysis [73]. In cardiovascular applications, learning-based models have demonstrated promising performance in tasks such as arrhythmia detection, waveform segmentation, and blood pressure estimation [5, 87–89]. These approaches are particularly attractive for wearable devices, where computational models can compensate for reduced sensor complexity by extracting informative patterns directly from the acquired signals.

Within this framework, PulsECG was conceived as a portable device capable of acquiring synchronized ECG and PPG signals and enabling preliminary investigations into cuffless blood pressure estimation through neural networks. Unlike approaches based solely on public ICU datasets, PulsECG allows full control over hardware design, acquisition protocols, and signal conditioning, thereby addressing several reproducibility and generalization issues highlighted in the previous section.

5.4.2 State of the Art in Wearable Cardiovascular Devices

Several commercial wearable devices are currently capable of acquiring ECG and/or PPG signals, yet only a limited subset offers performance levels suitable for clinical-grade analysis. Prominent examples include devices developed by AliveCor, Withings, Apple, and Samsung.

AliveCor [90] provides ECG acquisition devices under the *Kardia* product line, offering single-lead and six-lead ECG recordings paired with cloud-based analysis and optional cardiologist review. Withings [91] and Apple [92] offer smartwatches capable of recording ECG and additional physiological parameters such as oxygen saturation, yet do not provide native cuffless blood pressure estimation without auxiliary devices. Samsung [93], with the Galaxy Watch series, introduced a PPG-based blood pressure estimation feature that requires frequent calibration using a traditional cuff-based sphygmomanometer.

Despite these advances, existing commercial solutions typically rely on proprietary algorithms, limited signal accessibility, and calibration-dependent procedures. Moreover, users are often restricted to PDF-based reports rather than numerical signal data, limiting clinical flexibility and research applicability. PulsECG addresses this gap by offering direct access to numerical ECG and PPG signals, enabling advanced offline analysis and customized signal processing pipelines, while simultaneously exploring cuffless blood pressure estimation using both modalities.

5.4.3 PulseECG System Architecture

PulsECG is a portable electronic device designed and assembled at the Neuronica Laboratory of the Polytechnic of Turin. The device adopts a bar-shaped form factor and is conceived to acquire electrocardiographic signals using two electrodes

positioned according to Einthoven's triangle [94, 95]. In addition to ECG, PulsECG integrates an optical sensor for photoplethysmographic signal acquisition, enabling multimodal cardiovascular monitoring.

The system architecture is composed of two stacked printed circuit boards (PCBs) connected via an 8-pin flat cable: an ECG PCB and a PPG PCB. A three-dimensional model and a block diagram of the device are shown in Figures 5.11 and 5.12.




Fig. 5.11 3D model of the PulsECG device. Electrodes are visible on the sides, while the ECG and PPG PCBs are stacked to minimize space requirements.




Fig. 5.12 PulsECG block diagram showing the ECG PCB (left) and PPG PCB (right) connected via flat cable.

5.4.4 ECG PCB

The ECG PCB integrates the following main components:

- **Analog front-end:** responsible for ECG signal conditioning [96], consisting of a TI INA333 instrumentation amplifier, followed by a high-pass filter ($f_c = 0.5$ Hz), a notch filter for power-line interference rejection, and a low-pass filter with $f_c = 100$ Hz.
- **Microcontroller (TI CC2640):** manages signal acquisition, sampling ECG signals through the internal ADC, acquiring PPG samples via I²C communication, and handling BLE communication.
- **Power management:** based on a 190 mAh Li-Po battery, a MAXIM MAX1555 charger, a MAX17048 battery gauge, and voltage regulation provided by the MAX1759 and TI REF2033.
- **BLE communication:** enabled by an integrated patch antenna for low-energy wireless data transmission.

5.4.5 PPG PCB and Final Prototype

The PPG PCB hosts a MAXIM MAXM86161 optical sensor, positioned near the device top to facilitate correct finger placement and reduce motion artifacts. The final assembled prototype, enclosed in a 3D-printed case, is shown in Figure 5.13.



Fig. 5.13 Final PulsECG prototype enclosed in a 3D-printed case.

5.4.6 Neural Network for Blood Pressure Prediction: Context and Rationale

Beyond signal acquisition, PulsECG was designed to support preliminary investigations into cuffless blood pressure estimation using machine learning techniques. Extensive literature exists on this topic [97, 98, 8], with many approaches relying on large clinical databases such as MIMIC [76, 52]. However, such datasets are

acquired in intensive care settings using specialized instrumentation and often lack transparency regarding preprocessing and filtering steps.

To address these limitations, the approach adopted here is based on a custom dataset acquired entirely using PulsECG. This strategy ensures full control over the acquisition pipeline, synchronization between ECG and PPG signals, and explicit knowledge of preprocessing steps. While the resulting dataset is necessarily smaller than public ICU datasets, it provides a valuable testbed for assessing the feasibility and limitations of cuffless blood pressure estimation in wearable-oriented scenarios.

5.4.7 Dataset Description and Acquisition Protocol

A custom dataset was collected from 20 healthy subjects (5 females and 15 males) aged between 22 and 84 years. All participants were seated in a relaxed posture during acquisition, replicating typical home-monitoring conditions. Prior to data collection, participants were screened to exclude acute cardiovascular or systemic diseases, although natural variability in cardiovascular profiles was preserved. In particular, some subjects exhibited hypertensive or hypotensive conditions, while one subject showed a consistently elevated heart rate due to pharmacological treatment.

Blood pressure measurements were obtained using a certified sphygmomanometer (OMRON X7 Smart), while ECG and PPG signals were acquired simultaneously using PulsECG. For each subject, a sequence of three BP measurements and two ECG–PPG recordings was performed. The protocol consisted of an initial BP measurement, followed by ECG and PPG acquisition, a second BP measurement, a second ECG–PPG recording, and a final BP measurement.

Each ECG–PPG recording lasted 30 seconds. After acquisition, the signals were segmented into 228 overlapping windows of 8 seconds. Blood pressure values associated with each window were computed by linearly interpolating between the preceding and subsequent cuff-based BP measurements, under the assumption that BP variations over a 30-second interval are approximately linear.

This acquisition and labeling strategy enabled the construction of a dataset containing a sufficient number of training samples while preserving numerical diversity in blood pressure values. At the same time, it highlights an intrinsic limitation: blood pressure labels are indirectly assigned and subject to interpolation uncertainty, an aspect that will be critically discussed later in this chapter.

5.4.8 Neural Network Structure

The neural network adopted for cuffless blood pressure estimation within the PulsECG framework is designed to process synchronized ECG and PPG signals as multivariate time series. Each input sample consists of two channels (ECG and PPG), each containing 1000 time points corresponding to an 8-second acquisition window.

The architecture follows a deep, multilayer structure aimed at extracting both local morphological features and long-range temporal dependencies. The network begins with an input layer followed by two consecutive convolutional blocks. Each block comprises a one-dimensional convolutional layer with 128 filters, kernel size equal to 3, and same padding, followed by batch normalization and a ReLU activation function. This configuration allows effective local feature extraction while ensuring numerical stability and non-linearity.

After the second convolutional block, a residual connection is stored to preserve high-resolution temporal information prior to downsampling. The feature maps are then passed through a max pooling layer, which reduces temporal resolution and computational complexity. The downsampled sequence is subsequently processed by three stacked bidirectional Long Short-Term Memory (BiLSTM) layers, each composed of 128 units. The bidirectional formulation enables the model to capture temporal dependencies in both forward and backward directions, which is particularly beneficial for cardiovascular signals characterized by quasi-periodic dynamics and context-dependent morphology.

Following the recurrent backbone, the architecture includes two additional convolutional layers (128 filters, kernel size 3, same padding), again followed by batch normalization and ReLU activation. Temporal upsampling is then performed using a transposed convolutional layer with 256 filters, kernel size equal to 2, stride 2, and same padding. This operation restores the original temporal resolution lost during max pooling.

The upsampled feature maps are concatenated with the previously stored residual connection, allowing the network to recover fine-grained temporal details. The final stage of the architecture consists of three convolutional layers (128 filters, kernel size 3) with batch normalization and ReLU activation, progressively refining the feature representation. The resulting feature map is flattened and fed into a dense

layer that outputs the final predictions for systolic and diastolic blood pressure (SBP and DBP).

The network is trained using the Adam optimizer, minimizing the mean squared error (MSE) loss function, while the mean absolute error (MAE) is monitored as an additional performance metric.

5.4.9 Results

The evaluation of the PulsECG system was conducted in two complementary stages. First, the quality of the ECG and PPG signals acquired by the device was assessed against a clinical-grade reference. Subsequently, the performance of the neural network for blood pressure prediction was evaluated using the collected dataset.

Device Validation

All electronic functionalities of the PulsECG device were extensively tested, with particular emphasis on the reliability and consistency of ECG and PPG signal acquisition. Hundreds of recordings were performed on multiple members of the Neuronica laboratory to verify deterministic behavior and signal repeatability.

Once signal stability was established, a comparative validation was conducted using a General Electric GE B105 medical device as the golden reference. The validation protocol involved:

- 5 subjects,
- 5 acquisitions per subject using the GE B105 device (ECG and PPG),
- 5 acquisitions per subject using the PulsECG prototype (ECG and PPG),

with all acquisitions performed in parallel to ensure comparable physiological conditions.

Signal quality was evaluated in both the time and frequency domains. In the frequency domain, the Fast Fourier Transform (FFT) was computed and used to

derive the Power Spectral Density (PSD) [99] according to:

$$PSD(f) = \frac{(\Delta t)^2}{T} \left| \sum_{n=1}^N x_n e^{-i\omega n \Delta t} \right|^2. \quad (5.1)$$

Figures 5.14 and 5.15 show the PSD comparison between PulsECG and GE B105 for ECG and PPG signals, respectively.




Fig. 5.14 PSD comparison of ECG signals acquired by PulsECG (blue) and GE B105 (red).




Fig. 5.15 PSD comparison of PPG signals acquired by PulsECG (blue) and GE B105 (red).

To provide a quantitative comparison of spectral content, the Cumulative Spectral Power (CSP) was computed as the normalized cumulative sum of the PSD [100]:

$$CSP(f) = \sum_{n=1}^f PSD(n). \quad (5.2)$$

Figures 5.16 and 5.17 show the CSP comparison for ECG and PPG signals, respectively.




Fig. 5.16 CSP comparison for ECG signals acquired by PulsECG (blue) and GE B105 (red).

Tables 5.3 and 5.4 report the frequencies at which 20%, 50%, and 80% of the total signal power is reached.

Table 5.3 CSP comparison for ECG signals

Device	$f_{20\%}$ [Hz]	$f_{50\%}$ [Hz]	$f_{80\%}$ [Hz]
GE B105	2.58	6.96	13.90
PulsECG	1.16	4.16	8.38

Table 5.4 CSP comparison for PPG signals

Device	$f_{20\%}$ [Hz]	$f_{50\%}$ [Hz]	$f_{80\%}$ [Hz]
GE B105	1.06	1.16	2.15
PulsECG	1.12	1.16	2.28

Finally, time-domain comparisons for ECG and PPG signals are shown in Figures 5.18 and 5.19.

Overall, the results show good agreement between PulsECG and the clinical reference. While minor spectral differences are observed in ECG signals, PPG




Fig. 5.17 CSP comparison for PPG signals acquired by PulsECG (blue) and GE B105 (red).




Fig. 5.18 Time-domain comparison of ECG signals acquired by PulsECG (blue) and GE B105 (red).

acquisition demonstrates near-optimal performance. ECG quality can be further improved through digital filtering in post-processing.

Neural Network Performance

Neural network performance was evaluated by comparing predicted and measured blood pressure values. Figures 5.20 and 5.21 show scatter plots of predicted versus ground-truth SBP and DBP values, respectively.

Both plots exhibit a clear linear trend, indicating a strong correlation between predicted and measured values. A slight tendency toward overestimation is observed, particularly for DBP, as data points generally lie above the identity line. Nonetheless, the results suggest that the proposed network is capable of capturing informative relationships between ECG–PPG dynamics and blood pressure values in this wearable-oriented setting.




Fig. 5.19 Time-domain comparison of PPG signals acquired by PulsECG (blue) and GE B105 (red).




Fig. 5.20 Scatter plot of predicted versus measured SBP values.

5.4.10 Discussion and Conclusions

This section presented the design, validation, and preliminary deployment of PulsECG, a portable device for the acquisition of ECG and PPG signals and for cuffless arterial blood pressure estimation through a data driven neural network. Within the broader context of telemedicine and wearable health monitoring, the integration of signal acquisition hardware with embedded machine learning represents a key technological direction for enabling continuous, unobtrusive cardiovascular assessment.

From a hardware perspective, the comparative analysis with a certified clinical-grade monitor (GE B105) demonstrated that the PulsECG device is capable of acquiring biologically meaningful ECG and PPG signals with satisfactory fidelity.




Fig. 5.21 Scatter plot of predicted versus measured DBP values.

Both time-domain and frequency-domain analyses confirmed a close correspondence with the reference system, particularly for the PPG signal, which exhibited an almost identical spectral power distribution. Minor discrepancies observed in the ECG spectrum, especially at low frequencies, highlight known challenges associated with compact analog front-end design and motivate further refinement through improved analog conditioning and digital filtering.

On the algorithmic side, the neural network trained on data acquired directly with PulsECG showed encouraging performance in predicting systolic and diastolic blood pressure values. The observed correlation between predicted and reference measurements indicates that ECG and PPG signals acquired in realistic, non-clinical environments retain sufficient information to support blood pressure inference. Importantly, this result was obtained without relying on ICU-derived datasets or invasive arterial pressure measurements, thereby addressing one of the major limitations of many existing cuffless BP studies.

At the same time, the limitations of the current approach are evident and must be interpreted carefully within a doctoral research context. The dataset, while acquired under controlled and reproducible conditions, remains limited in size and subject diversity. Although the inclusion of subjects with heterogeneous blood pressure

profiles partially mitigates this issue, the resulting model performance should be understood primarily as a feasibility demonstration rather than as a fully generalizable solution. Moreover, the interpolation-based assignment of blood pressure values within each acquisition window introduces an implicit modeling assumption that may not hold under rapid hemodynamic changes.

These considerations are particularly relevant when placing PulsECG within the broader narrative of this thesis. The results obtained here show that purely data-driven models can achieve good apparent accuracy when acquisition conditions, preprocessing, and subject populations are tightly controlled. However, as discussed in Section 5.1, such performance does not guarantee robustness under subject-wise generalization or physiological variability. The present study therefore reinforces a key theme of this chapter: strong numerical performance can coexist with limited identifiability of the underlying physiological mechanisms.

Nevertheless, PulsECG plays a crucial conceptual role in the research trajectory developed in this thesis. Unlike approaches based exclusively on large public datasets, it demonstrates the value of controlling the entire acquisition pipeline, from sensor design to signal preprocessing and labeling. This control is essential for disentangling methodological limitations from data-related artifacts and for understanding how learned models respond to realistic wearable signals.

Future developments will focus on several complementary directions. On the hardware side, improvements to the ECG analog front-end and the extension to multi-lead acquisition are expected to enhance signal quality and diagnostic relevance. On the data side, expanding the subject pool and incorporating longitudinal recordings will be essential to evaluate generalization across individuals and conditions. Finally, from a modeling perspective, the limitations observed here motivate the transition toward physiology-informed and physics-constrained learning frameworks, introduced in the subsequent chapters, which aim to reduce ambiguity and improve robustness beyond what purely data-driven approaches can offer.

In this sense, PulsECG should be viewed not as a final solution, but as a critical intermediate step: it bridges controlled laboratory modeling and real-world wearable deployment, while simultaneously exposing the structural challenges that motivate the introduction of explicit physiological priors in cuffless blood pressure estimation.

5.5 Blood Pressure Estimation via Dual-Branch Wide Space–Time Neural Networks, on a Single Subject

After the exploration of waveform-level arterial blood pressure reconstruction and wearable-oriented acquisition pipelines, this section addresses a closely related but conceptually distinct problem. Rather than focusing on generalization across heterogeneous populations or on device-level deployment constraints, the objective here is to investigate blood pressure prediction under a deliberately simplified and controlled setting, in which inter-subject variability is explicitly removed.

During the development of the research presented in this thesis, a fundamental methodological question naturally emerged: *to what extent can a data-driven model accurately estimate blood pressure once the physiological characteristics of a specific subject are known and consistently observed over time?* Before tackling population-level generalization, it is therefore reasonable to ask whether neural networks are capable of learning a stable, reproducible mapping between ECG, PPG, and blood pressure within a single individual.

Non-invasive blood pressure estimation based on ECG and PPG signals has attracted considerable attention in recent years, driven by the demand for continuous, unobtrusive monitoring solutions in both clinical and home-care scenarios [14]. Numerous studies have demonstrated the feasibility of convolutional, recurrent, or hybrid deep learning architectures for this task [8, 5, 7]. However, the majority of existing works focus on population-level datasets, implicitly assuming that sufficient data diversity will enable models to generalize across subjects.

Large public databases such as MIMIC-III [49, 48] have played a key role in advancing the field. Nevertheless, these datasets are acquired in highly controlled intensive-care environments, often using invasive arterial lines and hospital-grade instrumentation. Moreover, their preprocessing pipelines are frequently complex and not fully documented, which limits reproducibility and complicates physiological interpretation. Beyond these practical aspects, population-level training introduces a deeper methodological challenge: neural networks may exploit subject-specific morphological cues (e.g., ECG waveform shape or baseline heart rate) rather than learning the underlying physiological relationships between ECG, PPG, and blood pressure. As a consequence, models can achieve low numerical error while relying on spurious correlations, thereby limiting robustness and interpretability.

To explicitly isolate this issue, the study presented in this section adopts an intentionally simplified experimental design. By restricting the dataset to a single subject monitored longitudinally over multiple days and physiological conditions, inter-subject variability is eliminated by construction. This reframes the problem as a form of *subject-specific physiological modeling*, in which the neural network is no longer required to disentangle differences in vascular compliance, arterial stiffness, or autonomic regulation across individuals. Instead, it is allowed to implicitly learn a personalized cardiovascular response pattern and to exploit it consistently across time.

This approach is not intended to claim clinical generality, nor to suggest that such a model could be deployed without subject-specific adaptation. Rather, it should be interpreted as a controlled *proof-of-concept*. Its value lies in isolating a fundamental aspect of cuffless blood pressure estimation: *if the underlying physiological state of a subject is sufficiently constrained, data-driven models can indeed achieve low prediction errors*. While this result may appear intuitive, it provides an essential reference point against which the limitations observed in multi-subject settings can be meaningfully interpreted.

Within the broader structure of this thesis, this section plays a precise methodological role. It establishes an upper bound on achievable performance when subject-specific physiology is implicitly known, thereby clarifying what purely data-driven approaches can accomplish under idealized conditions. At the same time, it motivates the transition to the Physics-Informed Neural Network (PINN) framework introduced in the following chapter, whose objective is precisely to encode subject-specific physiological structure in an explicit, interpretable, and generalizable manner.

5.5.1 Database Creation

The database used in this study consists of 121 JSON files, each containing 30 seconds of synchronized fingertip ECG and PPG signals together with corresponding systolic and diastolic blood pressure measurements. All data were collected from a single subject over a time span of approximately two months, ensuring temporal diversity and reducing the risk of short-term physiological bias.

Signal acquisition was performed using the PulsECG device, which provides simultaneous recording of ECG through fingertip electrodes and PPG through an

optical sensor. During each acquisition, the subject was seated in a relaxed position to minimize motion artifacts and autonomic perturbations. Blood pressure measurements were obtained using a standard digital sphygmomanometer placed on the upper arm, synchronized with the ECG and PPG recordings.

To ensure ecological validity, acquisitions were performed multiple times per day (approximately five), at varying times across the 24-hour cycle. This strategy was designed to capture natural physiological fluctuations related to circadian rhythm, daily activities, and autonomic modulation, while maintaining strict subject consistency.

To the best of the author’s knowledge, few studies in the literature have explored blood pressure estimation in a single-subject, longitudinal, real-world setting. By eliminating inter-patient variability and ECG morphology heterogeneity, the proposed dataset minimizes potential degeneracies in neural network training and enables a clearer assessment of the intrinsic relationship between ECG, PPG, and blood pressure for a given individual.

5.5.2 Data Pre-processing

Several pre-processing steps were applied to the raw signals prior to model training:

1. **Resampling:** ECG and PPG signals were resampled to a common sampling frequency of 125 Hz, ensuring temporal alignment and standardized input dimensions.
2. **Filtering:** The ECG signal, already conditioned by the PulsECG analog front-end, was left unfiltered. The PPG signal was band-pass filtered between 0.5 Hz and 8 Hz to suppress motion artifacts and isolate the frequency components most relevant to blood pressure dynamics.
3. **Normalization:** Both signals were normalized using min–max scaling in the range $[0, 1]$, preserving relative amplitude relationships while improving numerical stability during training.

5.5.3 Windowing Strategy

Following preprocessing, each 30-second recording was segmented into overlapping windows of 1000 samples, corresponding to 8-second intervals at 125 Hz. A stride of 40 samples was used, resulting in 57 windows per file.

This windowing strategy enables the model to exploit short-term temporal dynamics while significantly increasing the number of training samples. Each window constitutes a single input instance for the neural network, preserving rich temporal information critical for accurate blood pressure estimation.

5.5.4 Blood Pressure Value Distribution

The distribution of systolic and diastolic blood pressure values across the dataset is shown in Fig. 5.22. Systolic values are primarily clustered around 120–125 mmHg, while diastolic values are centered near 70 mmHg, consistent with a normotensive physiological profile.




Fig. 5.22 Distribution of systolic (left) and diastolic (right) blood pressure values in the single-subject dataset.

Although the distributions are relatively narrow, mild asymmetries and occasional higher values are present, reflecting realistic day-to-day physiological variability. This variability is essential for enabling the neural network to learn meaningful blood pressure dynamics rather than converging to trivial mean-value predictions.

5.5.5 Dual-Branch Wide Space–Time Architecture

The neural network architecture adopted in this study is explicitly designed to exploit the complementary information provided by electrocardiogram (ECG) and photoplethysmogram (PPG) signals under a single-subject experimental setting. The model follows a *dual-branch wide space–time* design, in which ECG and PPG signals are processed independently in the early stages and subsequently fused to model their joint temporal dynamics.

Each input sample consists of two synchronized one-dimensional signals, ECG and PPG, each with a duration of 8 seconds and sampled at 125 Hz, resulting in input tensors of size 1000×1 per modality. No handcrafted features or fiducial points are extracted prior to network ingestion, in order to preserve a fully data-driven learning pipeline.

The architecture is composed of three main stages: (i) modality-specific convolutional encoding, (ii) wide feature fusion and temporal modeling, and (iii) decoding and regression of blood pressure values. A schematic representation of the full network structure is reported in Fig. 5.23.




Fig. 5.23 Dual-branch wide space-time neural network architecture for single-subject blood pressure estimation. ECG and PPG signals are processed independently through convolutional encoders, fused into a wide feature representation, and temporally modeled via stacked BiLSTM layers before final regression of SBP and DBP.

Modality-Specific Convolutional Encoding The ECG and PPG signals are first processed by two independent convolutional branches with identical architecture but separate learnable parameters. Each branch consists of two consecutive one-dimensional convolutional blocks. Each block includes:

- a Conv1D layer with 128 filters,
- kernel size equal to 3,
- same padding,
- batch normalization,

- ReLU activation.

As a result, each branch maps an input tensor of size 1000×1 into a feature map of size 1000×128 . This stage enables the extraction of modality-specific temporal features, such as ECG waveform morphology or PPG pulse shape characteristics, while preserving the original temporal resolution.

Wide Feature Fusion and Temporal Modeling The outputs of the ECG and PPG convolutional encoders are concatenated along the feature dimension, producing a fused representation of size 1000×256 . This operation constitutes the *wide* component of the architecture, allowing the network to jointly analyze electrical and hemodynamic information without enforcing early alignment constraints.

The fused feature map is then processed by an additional convolutional layer with:

- 128 filters,
- kernel size 3,
- same padding,
- batch normalization and ReLU activation,

reducing the feature dimension to 1000×128 and acting as a shared encoder prior to temporal modeling.

Temporal dependencies are subsequently modeled using a stack of three Bidirectional Long Short-Term Memory (BiLSTM) layers. Each BiLSTM layer is composed of 128 hidden units in each direction, resulting in an output dimensionality of 1000×256 per layer. The use of bidirectional recurrence allows the network to exploit both past and future contextual information within the 8-second observation window, which is particularly relevant in a single-subject longitudinal setting.

Decoding and Output Regression Following temporal modeling, the network applies two additional Conv1D layers with:

- 128 filters,

- kernel size 3,
- same padding,
- batch normalization and ReLU activation,

to refine the latent representation.

The final feature map is flattened and passed through a fully connected dense layer with two output neurons, corresponding to systolic blood pressure (SBP) and diastolic blood pressure (DBP). The network therefore performs a direct regression of absolute SBP and DBP values for each input window.

The model is trained using the Adam optimizer with default parameters. Mean squared error (MSE) is employed as the loss function, while mean absolute error (MAE) is monitored as an additional performance metric.

Architectural Rationale It is important to emphasize that this architecture is not designed to address inter-subject generalization. Instead, its expressiveness is intentionally leveraged to learn a stable, subject-specific mapping between ECG, PPG, and blood pressure. By removing inter-individual variability, the model is allowed to implicitly encode the physiological characteristics of a single cardiovascular system.

Within the context of this thesis, the proposed architecture serves as a controlled reference point: it demonstrates what can be achieved by purely data-driven models when subject physiology is fixed and repeatedly observed. The insights gained from this setting directly motivate the transition toward physics-informed and physiology-constrained approaches introduced in the following chapter.

5.5.6 Results

After preprocessing, synchronization, and windowing, the final dataset consisted of 6897 input windows extracted from 121 source files. Each window corresponds to an 8-second segment of synchronized ECG and PPG signals, sampled at 125 Hz and represented as an input tensor of size $(1000, 2)$. The corresponding target tensor contains systolic and diastolic blood pressure values and has dimensions $(6897, 2, 1)$.

Given the limited size of the dataset and its single-subject nature, particular care was devoted to the evaluation protocol in order to avoid overly optimistic

performance estimates. For this reason, a 10-fold cross-validation strategy was adopted. In each fold, 10 complete acquisition files were randomly excluded from training and used exclusively for testing, while the remaining files were used for training and validation. This file-wise split ensures that each evaluation is performed on temporally unseen data, thereby providing a more reliable assessment of the model's predictive capability.

The results of the 10-fold cross-validation are summarized in Fig. 5.24, which reports the mean absolute error (MAE) for systolic (SBP) and diastolic (DBP) blood pressure across the folds, computed on the respective test sets. The MAE for SBP (blue curve) varies across folds from values slightly above 5 mmHg down to approximately 2.5 mmHg, while the MAE for DBP (orange curve) exhibits a comparable range of variation.




Fig. 5.24 Mean Absolute Error (MAE) for systolic (SBP) and diastolic (DBP) blood pressure predictions across the 10-fold cross-validation. Each point corresponds to the MAE computed on the test set of the respective fold.

This variability across folds reflects the heterogeneous distribution of blood pressure values over time and highlights the sensitivity of the MAE metric to specific acquisition segments. In particular, a small number of windows associated with extreme SBP or DBP values may disproportionately influence the error in certain folds, a phenomenon exacerbated by the limited dataset size. Nevertheless, several

folds exhibit consistently low MAE values for both SBP and DBP, indicating that once the subject-specific physiological dynamics are sufficiently captured, the model can achieve accurate blood pressure estimation.

Beyond raw performance metrics, these results provide valuable methodological insight. They suggest that the model is capable of learning a stable mapping between ECG, PPG, and blood pressure within a constrained physiological regime, while also revealing the limitations imposed by data scarcity. The observed variability therefore serves not only as a measure of accuracy but also as an indicator of where future improvements, such as increased data volume or more structured modeling, could reduce uncertainty and improve robustness.

5.5.7 Discussion and Conclusion

This section investigated cuffless blood pressure estimation under a deliberately simplified and controlled experimental setting, in which inter-subject variability was explicitly removed. By focusing on a single subject monitored over time, the objective was not to demonstrate population-level generalization, but rather to probe a more fundamental question: whether a data-driven model can reliably learn the cardiovascular dynamics of a specific individual when the underlying physiological state remains relatively stable.

From a purely numerical perspective, the proposed dual-branch wide space–time neural network achieved encouraging performance, with mean absolute errors for systolic and diastolic blood pressure that, in several cross-validation folds, approached values commonly reported in multi-subject studies. These results confirm that, once subject-specific physiological characteristics are implicitly known to the model, ECG and PPG signals indeed contain sufficient information to support accurate blood pressure estimation.

However, the variability observed across cross-validation folds provides a crucial interpretative key. Even in the absence of inter-subject heterogeneity, performance was not uniform across time segments, suggesting that intra-subject variability, measurement noise, and uneven blood pressure distributions still play a significant role. This highlights that low prediction error does not arise solely from architectural sophistication, but also depends critically on data structure, sampling density, and physiological stationarity.

Importantly, the success of this single-subject experiment should not be interpreted as evidence that purely data-driven models can resolve the general cuffless blood pressure estimation problem. Rather, it establishes an upper bound on achievable performance under idealized conditions. In this setting, the neural network effectively learns a personalized cardiovascular response function, implicitly encoding subject-specific vascular compliance, autonomic tone, and hemodynamic timing. While this capability is useful as a methodological benchmark, it does not scale naturally to multi-subject scenarios, where such physiological parameters vary widely and are not directly observable in surface biosignals.

This observation directly connects with the broader narrative developed throughout the thesis. In Chapter 4, ECG morphology detection was shown to be a well-posed learning problem, characterized by clear physiological targets and expert-defined ground truth. In contrast, blood pressure estimation is inherently ill-posed: the mapping from ECG and PPG to arterial pressure is indirect, non-unique, and mediated by latent physiological variables. The present study demonstrates that when these latent variables are effectively “frozen” by restricting attention to a single subject, data-driven models perform well. Conversely, when this constraint is lifted, performance degradation observed in multi-subject settings is not a mere implementation issue, but a structural limitation of the learning problem.

Another important implication concerns model interpretability. Although the network achieves low errors, its internal representations remain opaque, and it is not possible to disentangle whether the learned features correspond to physiologically meaningful mechanisms or to surrogate correlations specific to the dataset. This lack of interpretability further limits the clinical relevance of purely data-driven solutions, especially in safety-critical applications such as blood pressure monitoring.

From a methodological standpoint, this section serves as a controlled proof-of-concept that clarifies *what* neural networks can learn when sufficient physiological consistency is present, and, more importantly, *what they cannot reliably infer* in the absence of explicit constraints. The findings reinforce the idea that architectural complexity alone is insufficient to guarantee generalization or physiological plausibility.

In the context of the thesis, the results presented here play a pivotal role. They justify the need for models that do not merely fit data, but explicitly incorporate cardiovascular dynamics and hemodynamic principles. This insight motivates the

transition to the Physics-Informed Neural Network (PINN) framework introduced in the following chapter, where subject-specific physiology is not implicitly memorized from data, but explicitly modeled through governing equations and structured constraints.

In summary, while single-subject data-driven blood pressure estimation represents a useful benchmark and an informative diagnostic tool, it also exposes the fundamental limitations of black-box learning in this domain. Addressing these limitations requires moving beyond correlation-based inference toward models that embed physiological knowledge, thereby enabling robust, interpretable, and generalizable cuffless blood pressure estimation.

5.6 Analysis of Generalization Failure in Data-Driven BP Estimation

5.6.1 Cross-Dataset Transfer: MIMIC to PulsECG

The first study presented in this chapter (Section 5.3) demonstrated that a purely data-driven model can achieve low numerical errors when trained and evaluated on large ICU-derived datasets, such as those curated from the MIMIC ecosystem [76, 48]. While these results are valuable to establish feasibility under controlled conditions, they also raise an unavoidable question that naturally emerged immediately after the first paper: *does the learned mapping generalize to real-world wearable acquisitions, acquired with different hardware, different preprocessing, and different noise/motion characteristics?*

This question is not merely practical but methodological. In cuffless blood pressure estimation, high accuracy on a benchmark dataset does not necessarily imply that the model has learned a physiologically meaningful inverse mapping. The risk is that the network may exploit dataset-specific regularities (instrumentation, acquisition context, preprocessing, subject overlap, or implicit calibration cues) that are not preserved outside the training domain. For this reason, after completing the initial ABP waveform reconstruction study, a set of exploratory experiments was performed to *stress-test generalization*, by applying the MIMIC-trained model to ECG/PPG recordings acquired with the PulsECG device in everyday-like conditions.

The following report, originally drafted as an internal investigation, was therefore not conceived as a “new model proposal”, but as a *failure analysis* step in the research trajectory: it quantifies how and why a high-performing ICU-trained model may fail when transferred to a wearable-oriented domain without adaptation. Although the original document was informal, the core methodological steps and the numerical outcomes are clear and informative, and they are summarized here in a thesis-grade form.

Training domain: MIMIC-based waveform dataset and ABP reconstruction model

The training dataset was built from high-quality recordings extracted from MIMIC-III via WFDB [76], with ECG, PPG (PLETH), and invasive ABP channels available. Signals were segmented into fixed-length windows of 1500 samples (corresponding to the window duration used throughout the early phase of this work). Segments were generated in large batches (up to 82,000 windows per HDF5 chunk) and subjected to a multi-stage quality control procedure to exclude artifacts and uninformative segments, the same quality control that was used in Section 5.3 together with the same filtering process. After quality filtering and aggregation across HDF5 chunks, the resulting dataset consisted of

$$(151,211, 1500, 3),$$

i.e., 151,211 windows, each of length 1500, with 3 channels (PPG, ECG, ABP). From the ABP waveform in each window, systolic and diastolic values were extracted using a dedicated routine operating on local maxima/minima of the ABP signal, the same reported in Section 5.3, yielding distributions of SBP/DBP as well. Figure 5.25 reports the histograms of these extracted quantities.




Fig. 5.25 Histograms of waveform-derived blood pressure quantities computed from the MIMIC-based ABP segments: (a) SBP, (b) DBP, (c) ESBP, and (d) EDBP.

The reference model used in this analysis is the Conv1D–BiLSTM architecture described in detail in Section 5.3). Inputs included synchronized and normalized ECG and PPG, and in one configuration also the first and second derivatives of PPG (VPPG and APPG) as auxiliary channels. The dataset was partitioned using random window shuffling into 80% training, 10% validation, and 10% test subsets, and the network was trained with Adam optimizer and MSE loss, tracking MAE.

In this new experiment, the model was trained over 40 epochs with batch size 64. Training curves are reported in Figures 5.26 and 5.27. Numerically, the training started from $\text{MAE} \approx 21.17$ ($\text{loss} \approx 1045.73$) and reached final training $\text{MAE} \approx 2.62$ ($\text{loss} \approx 15.70$) by epoch 40. Validation metrics converged rapidly and remained substantially lower than the initial training values, indicating stable learning within the training domain.




Fig. 5.26 Training and validation loss over 40 epochs for the MIMIC-trained ABP reconstruction model.




Fig. 5.27 Training and validation MAE over 40 epochs for the MIMIC-trained ABP reconstruction model.

Target domain: PulsECG external recordings and evaluation protocol

To test out-of-distribution transfer, the trained model was frozen and used for inference on external recordings acquired with PulsECG. This constituted a *true cross-dataset evaluation*: no fine-tuning or recalibration was applied, and the network was not exposed to PulsECG data during training.

The external dataset comprised 19 recordings from 6 subjects. Each recording contained ECG and PPG sampled by the device, with reference SBP/DBP values measured by an external sphygmomanometer.

For each recording, signals were segmented into overlapping windows of 1500 samples with stride 100. Windows were normalized consistently with the training preprocessing, and (in the initial configuration) auxiliary derivative channels (VPPG/APPG) were computed as filtered first and second derivatives of PPG. The model predicted an ABP waveform (or equivalently a waveform-level output later reduced to SBP/DBP estimates), and the resulting SBP/DBP per window were summarized at the file level.

Out-of-distribution results: stable intra-record predictions, poor absolute agreement

The *absolute* agreement with external BP was poor and highly variable across files. The mean error (prediction minus external) across the 19 recordings was:

$$\Delta\text{SBP} = -2.33 \pm 17.85 \text{ mmHg}, \quad \Delta\text{DBP} = -7.67 \pm 7.91 \text{ mmHg},$$

as summarized in Figure 5.28. This indicates a large dispersion in SBP error (standard deviation ≈ 18 mmHg) and a systematic underestimation trend in DBP (negative mean), despite the model having achieved low MAE within the MIMIC domain.



Fig. 5.28 Discrepancies between mean (unweighted) model estimates and external sphygmomanometer values over the 19 PulsECG recordings: $\Delta\text{SBP} = -2.33 \pm 17.85$ mmHg and $\Delta\text{DBP} = -7.67 \pm 7.91$ mmHg.

Interpretation. This pattern is highly informative for the thesis narrative: it suggests that the mapping is *not anchored* to the absolute BP scale under domain shift. In other words, the network appears to preserve internal consistency while losing calibration across acquisition setups and subject/device conditions. This is compatible with the broader concern that deep models can exploit dataset-specific correlations and implicit calibration artifacts, and that the inverse mapping from ECG/PPG to absolute BP is underdetermined without additional physiological structure.

Ablation of auxiliary channels: ECG+PPG only

To test whether derivative-based channels (VPPG/APPG) contributed positively to transfer, a second experiment was performed using only ECG and PPG as network inputs. Under this configuration, the discrepancy statistics became:

$$\Delta\text{SBP} = -7.83 \pm 12.32 \text{ mmHg}, \quad \Delta\text{DBP} = -12.02 \pm 7.82 \text{ mmHg},$$

as reported in Figure 5.29. While SBP dispersion decreased (standard deviation ≈ 12 mmHg), the mean bias worsened, and DBP underestimation increased in magnitude. Overall, adding VPPG/APPG did not provide a systematic advantage for cross-dataset transfer in this setting, and the subsequent work therefore focused primarily on ECG+PPG as the core signal pair.




Fig. 5.29 External validation discrepancies when restricting the input to ECG and PPG only: $\Delta\text{SBP} = -7.83 \pm 12.32$ mmHg and $\Delta\text{DBP} = -12.02 \pm 7.82$ mmHg.

Signal-domain investigation: ECG/PPG mismatch across databases

To better contextualize the failure, the report compared representative ECG and PPG segments from PulsECG recordings and MIMIC-derived segments in both time and frequency domains. The comparison (Figures 5.30 and 5.31) indicated that PPG morphology was qualitatively more comparable across domains, whereas ECG exhibited more pronounced differences, including different spectral patterns and apparent noise/artefact components.




Fig. 5.30 Time-domain comparison between PPG acquired by PulsECG (blue) and representative PPG segments from the MIMIC-derived dataset (orange).




Fig. 5.31 Time-domain comparison between ECG acquired by PulsECG (blue) and representative ECG segments from the MIMIC-derived dataset (orange).

Although these comparisons were exploratory (and not a complete domain adaptation study), they support a plausible mechanism: *the learned representation is sensitive to acquisition-dependent waveform characteristics*, especially for ECG, which can vary substantially with electrode placement, analog front-end response, baseline wander, and motion artefacts. In such conditions, a model trained on ICU-grade instrumentation may retain intra-domain accuracy while failing to produce calibrated outputs on wearable signals.

Noise-matching attempt: targeted noise injection into training

Motivated by the observed signal mismatch, a further exploratory experiment attempted to “bridge” the domains by injecting synthetic, frequency-targeted noise into the training pipeline. The procedure consisted of:

- Aggregating external ECG/PPG recordings into a CSV for estimating dominant frequency components via FFT and peak detection.
- Generating structured periodic noise as sums of sinusoids at detected peak frequencies, with randomized phase, small frequency perturbations, and amplitude scaled approximately as $1/f^2$ to mimic the prominence of low-frequency components.
- Injecting this noise into ECG and PPG channels during training data generation, then retraining the model (50 epochs, ECG+PPG inputs).

As expected, training on noisier inputs increased in-domain training error (minimum MAE reported ≈ 4.01 vs ≈ 2.9 in the cleaner baseline). Interestingly, when re-evaluated on the PulsECG external set, the discrepancy statistics changed to:

$$\Delta\text{SBP} = 5.23 \pm 10.38 \text{ mmHg}, \quad \Delta\text{DBP} = -11.33 \pm 7.07 \text{ mmHg},$$

reported in Figure 5.32. This indicates that SBP bias flipped sign (from underestimation to overestimation) and dispersion decreased, whereas DBP bias remained strongly negative. While not a satisfactory solution, this outcome is conceptually important: it shows that *apparent transfer can be modulated by training-domain perturbations*, but in an unstable and non-universal manner, reinforcing that the problem is not simply “add noise” but address identifiability and physiological anchoring.




Fig. 5.32 External validation discrepancies after retraining with structured noise injection (ECG+PPG only): $\Delta\text{SBP} = 5.23 \pm 10.38 \text{ mmHg}$ and $\Delta\text{DBP} = -11.33 \pm 7.07 \text{ mmHg}$.

Transfer learning and fine-tuning: qualitative evidence of domain dependence

The report also explored a transfer learning strategy in which the MIMIC-trained network was modified to predict SBP/DBP directly (rather than reconstructing the full ABP waveform) using a small PulseECG dataset labeled by sphygmomanometer measurements. In a first phase, most layers were frozen and only late layers were adapted; this approach produced poor scatter agreement (Figures 5.33 and 5.34). When the entire network was unfrozen and fine-tuned, the scatterplots improved substantially (Figures 5.35 and 5.36), and further improvements were obtained when additional external subjects were included and BP values were assigned to 8-second windows by linear interpolation between before/after cuff measurements (Figures 5.37 and 5.38).




Fig. 5.33 Transfer learning attempt (frozen early layers): scatter plot for SBP predictions vs reference values.




Fig. 5.34 Transfer learning attempt (frozen early layers): scatter plot for DBP predictions vs reference values.




Fig. 5.35 Fine-tuning with all layers unfrozen: scatter plot for SBP predictions vs reference values.




Fig. 5.36 Fine-tuning with all layers unfrozen: scatter plot for DBP predictions vs reference values.




Fig. 5.37 Fine-tuning with an expanded external dataset and window-level BP assignment by linear interpolation: SBP scatter plot.




Fig. 5.38 Fine-tuning with an expanded external dataset and window-level BP assignment by linear interpolation: DBP scatter plot.

While these steps were exploratory and not developed into a standalone publication, they provide a clear message that is fully aligned with the thesis narrative: *the model can adapt to PulsECG data only when it is allowed to substantially reshape its internal representation*, i.e., when domain-specific information is injected through fine-tuning. This supports the interpretation that the original MIMIC-trained mapping is not invariant, and that cross-dataset BP estimation is strongly constrained by acquisition conditions and subject/device physiology.

Synthesis: what this report establishes within the thesis

This cross-dataset study establishes three thesis-level conclusions that motivate the subsequent trajectory (and, ultimately, the physics-informed framework in Chapter 6):

1. **In-domain accuracy does not imply transferable calibration.** The MIMIC-trained model achieved low training-domain MAE, yet exhibited large dispersion and systematic biases on wearable acquisitions, with SBP errors

showing very high variability (up to ~ 18 mmHg standard deviation in the mean-discrepancy analysis).

2. **Domain adaptation requires subject/device-specific information.** Improvements observed via full fine-tuning suggest that the mapping must be reshaped using external data, consistent with the idea that subject-specific physiology and acquisition transfer functions are implicitly learned rather than structurally constrained.

These points directly inform the motivation for the next steps developed in this thesis: (i) building controlled acquisition pipelines and datasets (as addressed by PulsECG in Section 5.4), (ii) studying subject-specific learning limits under reduced variability (single-subject setting in Section 5.5), and (iii) moving toward models that encode hemodynamic structure so that subject-specific physiology is not merely memorized, but *represented explicitly* through interpretable constraints (Chapter 6).

5.7 Multi-Subject Generalization Study on the PulsECG Database (7 Subjects, 158 Recordings)

Following the first evidence of cross-dataset transfer failure reported in Section 5.6 (MIMIC $\rightarrow$ PulsECG), this section summarizes a second, more structured internal study aimed at probing *generalization within PulsECG itself*, i.e., whether a purely data-driven network can learn a mapping from ECG/PPG to blood pressure that transfers across different individuals when trained on a small, heterogeneous cohort. The content reported here is a thesis-oriented rewrite of an internal experimental report.

5.7.1 Dataset overview and windowing protocol

The PulsECG dataset considered in this report consists of ECG and PPG recordings acquired from 7 subjects, for a total of 158 files, in the time span of 1 or 2 months. Each file contains approximately 30 s of synchronized ECG and PPG, together with the corresponding reference cuff measurements (SBP/DBP) acquired using a digital sphygmomanometer. In addition, basic anthropometric information was available per

subject (age, height, weight; and sex, as collected during the acquisition campaign), spanning an age range of approximately 24–70 years.

For learning-based experiments, each 30 s file was segmented into 65 fixed windows, yielding on the order of $\sim 10^4$ windows overall (approximately $158 \times 65 = 10,270$). Consistently with the previous sections of this chapter, each window was treated as a short paired segment of (ECG, PPG) used to regress SBP/DBP at the window level (rather than reconstructing a full ABP waveform, which is not available in PulsECG).

5.7.2 Experiment 1: within-subject modeling (leave-one-file-out)

A first experiment explored a *subject-specific* setting: a separate model was trained for each subject using only that subject's files. The evaluation followed a leave-one-file-out protocol: at each fold, one full file was held out for testing, while the remaining files were split into training and validation (70%/30% within the non-test partition).

Figures 5.39–5.45 report representative prediction outcomes across subjects. A key observation of this experiment is that, even when inter-subject variability is removed by construction, the prediction error can remain large on some subjects, with deviations reaching ~ 20 mmHg in the worst cases. This indicates that (i) the per-subject data volume is still limited (roughly ~ 20 –25 files per subject in this dataset), and (ii) the signal-to-label mapping may be weakly constrained under the adopted measurement protocol (short acquisition, cuff reference, and potential motion/contact variability).




Fig. 5.39 Experiment 1 (leave-one-file-out): subject 1.




Fig. 5.40 Experiment 1 (leave-one-file-out): subject 2.




Fig. 5.41 Experiment 1 (leave-one-file-out): subject 3.




Fig. 5.42 Experiment 1 (leave-one-file-out): subject 4.




Fig. 5.43 Experiment 1 (leave-one-file-out): subject 5.




Fig. 5.44 Experiment 1 (leave-one-file-out): subject 6.




Fig. 5.45 Experiment 1 (leave-one-file-out): subject 7.

5.7.3 Experiment 2: leave-one-subject-out generalization (dual-branch ECG/PPG network)

The second experiment directly tested *cross-subject generalization* in a leave-one-subject-out (LOSO) setting. In each fold, one subject was held out as test set, while the remaining six subjects were used for training/validation (80%/20% split within the non-test partition).

Model family. The adopted architecture was a dual-branch network, where ECG and PPG are processed by two separate feature-extraction streams (convolutional and recurrent components) before fusion for SBP/DBP regression. Training was allowed up to 1000 epochs with early stopping, and learning curves were inspected together with test performance trends over epochs.

Representative training MAE and test MAE curves are shown in Fig. 5.46–5.49. A recurrent pattern observed in this report is that *apparent optimization success on the pooled training distribution does not translate into stable performance on the held-out subject*. In several folds, as training proceeds and training/validation MAE decreases, the test MAE may plateau or even degrade, consistent with rapid subject-specific overfitting.




Fig. 5.46 Experiment 2 (LOSO): training MAE example (subject 1 as test).




Fig. 5.47 Experiment 2 (LOSO): training MAE example (subject 5 as test).




Fig. 5.48 Experiment 2 (LOSO): test MAE trend example (subject 1 as test).




Fig. 5.49 Experiment 2 (LOSO): test MAE trend example (subject 5 as test).

Table 5.5 reports the reported LOSO MAE values. The results show large fold-to-fold variability: while some subjects yield relatively low MAE (e.g., subject 3), multiple folds exhibit errors that are incompatible with practical BP monitoring (e.g., SBP MAE $\approx 23\text{--}27.5$ mmHg).

Table 5.5 Experiment 2 (LOSO): SBP/DBP MAE for the held-out subject in each fold (reported values).

Held-out subject	SBP MAE [mmHg]	DBP MAE [mmHg]
1	9.44	10.81
2	23.0	8.2
3	4.8	3.7
4	8.4	4.2
5	7.8	6.3
6	27.5	7.5
7	22.8	5.1

5.7.4 Experiment 3: adding anthropometrics (age, height, weight)

To partially mitigate the ambiguity of inferring vascular state from ECG/PPG alone, a third experiment augmented the model input with subject-level covariates (age, height, weight). The goal was not to “solve” generalization with demographics, but rather to test whether the network benefits from coarse, physically meaningful proxies of arterial stiffness and body size that may help align subjects in representation space.

Figures 5.50–5.53 show representative training curves and test trends. The resulting LOSO MAEs are listed in Table 5.6. A comparative view (Table 5.7) highlights that improvements are not uniform: in some folds SBP improves (notably in previously extreme cases), while in others performance worsens, especially for DBP. This behavior is coherent with the small- N regime: adding covariates increases model capacity and can either regularize or destabilize learning depending on the fold composition and label noise.




Fig. 5.50 Experiment 3 (LOSO + age/height/weight): training example (subject 1 as test).




Fig. 5.51 Experiment 3 (LOSO + age/height/weight): training example (subject 5 as test).




Fig. 5.52 Experiment 3 (LOSO + age/height/weight): test MAE trend example (subject 1 as test).




Fig. 5.53 Experiment 3 (LOSO + age/height/weight): test MAE trend example (subject 5 as test).

Table 5.6 Experiment 3 (LOSO + age/height/weight): SBP/DBP MAE for the held-out subject in each fold (reported values).

Held-out subject	SBP MAE [mmHg]	DBP MAE [mmHg]
1	8.25	6.48
2	10.0	17.0
3	13.0	8.5
4	8.6	4.6
5	8.0	4.8
6	20.3	5.7
7	16.0	6.0

Table 5.7 Change in MAE from Experiment 2 to Experiment 3 ($\Delta = \text{MAE}_{\text{Exp2}} - \text{MAE}_{\text{Exp3}}$). Positive values indicate improvement.

Subject	ΔSBP [mmHg]	ΔDBP [mmHg]
1	1.19	4.33
2	10.0	-8.8
3	-8.2	-4.8
4	-0.2	-0.4
5	-0.2	1.5
6	7.2	1.8
7	6.8	-0.9

5.7.5 Effect of Anthropometric Information on Cross-Subject Generalization

To quantitatively assess the impact of introducing subject-level anthropometric information (age, height, weight, and sex), the leave-one-subject-out results obtained with ECG–PPG only (Experiment 2) are compared against those obtained after augmenting the input with anthropometrics (Experiment 3).

For each held-out subject, a global error metric is defined as:

$$\text{MAE}_{\text{tot}} = \frac{\text{MAE}_{\text{SBP}} + \text{MAE}_{\text{DBP}}}{2}, \quad (5.3)$$

which provides a compact measure of overall blood pressure prediction accuracy, balancing systolic and diastolic components.

Table 5.8 reports the resulting total MAE before and after the inclusion of anthropometric variables, together with the corresponding improvement $\Delta\text{MAE}_{\text{tot}}$ (positive values indicate performance gains).

Table 5.8 Global MAE per subject before and after adding anthropometric information.

Subject	MAE _{tot} (ECG+PPG)	MAE _{tot} (+Anthro)	$\Delta\text{MAE}_{\text{tot}}$
1	10.13	7.37	+2.76
2	15.60	13.50	+2.10
3	4.25	10.75	−6.50
4	6.30	6.60	−0.30
5	7.05	6.40	+0.65
6	17.50	13.00	+4.50
7	13.95	11.00	+2.95

Overall, the introduction of anthropometric covariates leads to a reduction of the total MAE in 5 out of 7 subjects. The improvement is particularly pronounced for subjects that exhibited the largest baseline errors, with gains exceeding 4 mmHg in the most challenging case (subject 6). This behavior suggests that age, height, weight, and sex provide coarse but informative priors that partially compensate for inter-subject variability in vascular properties.

However, the effect is not uniform. For subject 3, the total MAE substantially increases after adding anthropometrics, and for subject 4 the change is negligible. This divergence highlights a critical limitation of small-cohort learning: when the number of subjects is limited, additional covariates can either regularize or destabilize the learned mapping, depending on how representative the held-out subject is with respect to the training distribution.

From a methodological perspective, these results indicate that anthropometric variables act as weak physiological constraints rather than universal correctors. They can mitigate extreme generalization failures but do not guarantee consistent improvements across all subjects. This observation reinforces the central argument of this chapter: while subject metadata may improve robustness in some cases, reliable cross-subject blood pressure estimation ultimately requires models that

encode cardiovascular dynamics more explicitly, motivating the physics-informed approaches developed in the subsequent sections.

5.7.6 Joint interpretation and implications for the thesis

Overall, this report provides a second, independent confirmation of a central thesis message: *in the small-data, multi-subject regime, purely data-driven BP estimation from ECG/PPG is dominated by inter-subject variability and overfitting dynamics*. Concretely:

- **Within-subject learning is not trivially solved:** even leave-one-file-out training can exhibit large errors, suggesting that measurement noise, contact variability, and limited physiological excitation reduce identifiability of SBP/DBP from short ECG/PPG segments.
- **Cross-subject generalization is unstable:** LOSO results show that training progress on pooled subjects does not guarantee improvements on the held-out subject; in several folds test MAE worsens as the network fits the training subjects more tightly.
- **Demographics help inconsistently:** adding age/height/weight changes the test dynamics and improves some extreme SBP cases, but does not robustly solve generalization and can degrade DBP on specific folds. This suggests that coarse covariates alone are insufficient as a “physiology encoder”.

A noteworthy empirical clue is the presence of a transient regime (reported here around 0–120 epochs) where test MAE does not systematically degrade. This is compatible with a *bias–variance* trade-off typical of low- N : early in training, the model behaves closer to a population-average predictor, whereas prolonged training encourages memorization of subject-specific morphology and label idiosyncrasies, harming LOSO transfer.

Consequences for subsequent developments. These findings motivate, in a directly operational way, the transition from “black-box” regression toward structured modeling of subject physiology:

- **Need for explicit subject state:** rather than relying on demographics, the model should incorporate a compact representation of vascular state (latent parameters or identifiable physical quantities).
- **Need for mechanistic regularization:** the instability and sensitivity to folds indicate that the learning problem is underconstrained; physics-based priors (as introduced in the subsequent PINN chapter) provide additional constraints that can reduce hypothesis space and improve robustness.
- **Need for careful evaluation protocols:** the strong LOSO variability emphasizes that window-level random splits can be misleading; subject-wise and acquisition-wise splits are essential whenever the objective is real-world deployment.

In summary, the 7-subject PulsECG study acts as a bridge between the initial success on curated public datasets (Section 5.3), the explicit cross-dataset failure analysis (Section 5.6), and the physics-informed methodology developed next: the results clarify *why* learning BP from ECG/PPG without a structured physiological scaffold becomes fragile in realistic acquisition conditions and small cohorts.

5.8 Synthesis and Motivation for Physics-Informed Modeling

This chapter has systematically explored the capabilities and limitations of purely data-driven approaches for non-invasive, cuffless blood pressure estimation from ECG and PPG signals. Through a sequence of progressively more controlled experiments, ranging from large public datasets to custom wearable acquisitions, several recurring patterns have emerged that ultimately motivate the transition toward physics-informed neural modeling.

From waveform reconstruction to generalization limits. The first contribution presented in this chapter demonstrated that deep neural networks, when trained on curated datasets, are capable of accurately reconstructing short segments of arterial blood pressure waveforms. The Conv1D–BiLSTM architecture achieved low mean absolute error under standard train–validation–test protocols, confirming that ECG

and PPG indeed encode sufficient information to recover arterial pressure dynamics under favorable conditions.

However, subsequent analyses revealed that this apparent success is tightly coupled to the statistical properties of the training data. When the same architecture was transferred across datasets, specifically from MIMIC-III to PulsECG acquisitions, the predictive accuracy deteriorated substantially, despite stable and smooth predictions within each individual recording. This behavior indicates that the learned mapping is highly sensitive to acquisition conditions, preprocessing pipelines, sensor characteristics, and latent subject-specific physiology.

Subject specificity as a double-edged sword. The single-subject study further clarified this point. By removing inter-subject variability, the neural network was shown to reliably learn a stable mapping between ECG, PPG, and blood pressure for a given individual, achieving consistently low errors across time. While this result is expected, it plays a crucial conceptual role: it confirms that the learning problem itself is not ill-posed, but that its difficulty arises primarily from uncontrolled physiological heterogeneity across subjects.

In this sense, purely data-driven models implicitly learn a *subject-specific cardiovascular state*, even when trained in a population setting. When such a state is not shared between training and testing distributions, generalization fails.

Limits of metadata-based compensation. The introduction of anthropometric variables (age, height, weight, and sex) partially mitigated these failures. For several subjects, the total MAE decreased markedly, especially in cases characterized by large baseline errors. Nonetheless, the improvements were inconsistent and, in some cases, detrimental. This confirms that anthropometric information provides only coarse, static priors and cannot adequately represent the dynamic and continuous nature of cardiovascular physiology. But let's keep in mind that the dataset was incredibly small, and maybe the introduction of anthropometric information in a bigger database could help a lot. Future experiments will surely point in this way.

Thus, while subject metadata can regularize learning in small datasets, it does not resolve the fundamental ambiguity of the inverse problem: multiple physiological states may produce similar ECG and PPG morphologies while corresponding to different pressure levels.

A structural limitation of purely data-driven learning. Taken together, the results of this chapter expose a structural limitation of black-box learning in cuffless blood pressure estimation. Data-driven models are highly effective at interpolating within the physiological manifold spanned by the training data, but they lack an explicit representation of the mechanisms governing blood flow, arterial compliance, wave propagation, and pressure–area relationships. As a consequence, they struggle to extrapolate beyond the empirical distribution, particularly when confronted with new subjects, sensors, or acquisition protocols.

This limitation is not merely a matter of dataset size or architectural complexity; rather, it reflects the absence of explicit physiological structure in the learning process.

Toward physics-informed subject modeling. These observations naturally motivate the introduction of Physics-Informed Neural Networks (PINNs) in the following chapter. Instead of relying solely on statistical correlations, PINNs allow embedding governing equations of hemodynamics directly into the learning objective, thereby constraining the solution space to physiologically admissible states.

From this perspective, the data-driven models explored in this chapter serve a dual purpose. First, they establish a strong empirical baseline and clarify what can be achieved without explicit physics. Second, they provide a clear diagnostic of *why* generalization fails, identifying subject-specific physiology as the missing latent variable.

The PINN framework developed next can therefore be interpreted as an attempt to explicitly model this latent physiological state, bridging the gap between subject-specific learning and population, level generalization through physically grounded constraints.

In summary, Chapter 5 demonstrates both the promise and the inherent limitations of data-driven cuffless blood pressure estimation. These findings set the stage for Chapter 6, where physics-informed learning is introduced as a principled strategy to overcome the generalization barriers identified here.

Chapter 6

Physics-Informed Neural Networks for Cuffless Blood Pressure Estimation

The investigations presented in the previous chapters converge toward a common methodological conclusion. While data-driven neural networks can achieve low prediction errors in controlled experimental settings, their reliability rapidly deteriorates when confronted with inter-subject variability, heterogeneous acquisition protocols, or cross-database evaluation scenarios.

This behavior does not arise from a lack of model capacity or architectural sophistication. On the contrary, increasingly complex deep learning models have demonstrated strong interpolation capabilities within homogeneous datasets. The limitation instead lies at a conceptual level: purely data-driven approaches lack explicit constraints that reflect the physical and physiological mechanisms governing blood pressure generation and propagation. As a consequence, learned representations tend to encode dataset-specific or subject-specific correlations rather than the underlying cardiovascular dynamics.

These observations motivate a shift in perspective. Rather than asking whether neural networks can implicitly discover cardiovascular physics from data alone, this chapter addresses a complementary question: can known physical and physiological principles be explicitly embedded into the learning process to guide the model toward solutions that are more robust, interpretable, and physiologically plausible?

Physics-Informed Neural Networks (PINNs) provide a natural framework to pursue this objective. By incorporating governing equations of blood flow, arterial mechanics, and boundary conditions directly into the training objective, PINNs restrict the admissible solution space to regimes consistent with established hemodynamic theory, while preserving the expressive power of deep learning.

In the context of cuffless blood pressure estimation, this approach represents a conceptual evolution from signal-based regression toward a hybrid paradigm in which data and physics cooperate. This chapter presents the development of such a framework throughout this research, progressing from an initial physics-informed formulation to a second-generation architecture enriched with anatomical and physiological priors.

The chapter is organized as follows. First, the physiological and mathematical foundations of one-dimensional arterial hemodynamics are introduced. Next, the first-generation PINN model is presented and critically analyzed. Building upon its limitations, an enhanced second-generation framework is then described. Finally, experimental results on large-scale and external datasets are discussed, with particular attention to generalization, interpretability, and implications for future cuffless blood pressure monitoring systems.

6.1 What are PINNs?

Physics-Informed Neural Networks (PINNs) are a class of *scientific machine learning* methods in which a neural network is not trained solely to fit observed data, but also to (approximately) satisfy the physical laws governing the system of interest. In practice, the neural network acts as a continuous *surrogate model* for the solution of a set of differential equations (ODEs/PDEs), while physical constraints are enforced through loss terms built from equation residuals and boundary/initial conditions.

Core idea

Let $\mathbf{u}(\mathbf{x})$ denote the physical state (e.g., pressure and flow) defined over space–time variables $\mathbf{x} = (t, \mathbf{s})$, and let $\mathcal{N}[\cdot]$ be a differential operator describing the system

dynamics. A physics-based model can be written compactly as

$$\mathcal{N}[\mathbf{u}(\mathbf{x})] = \mathbf{0}, \quad \mathbf{x} \in \Omega, \quad (6.1)$$

with boundary and/or initial conditions (e.g., $\mathbf{u} = \mathbf{u}_b$ on $\partial\Omega$). In a PINN, one introduces a neural approximation $\mathbf{u}_\theta(\mathbf{x})$ (parameterized by weights θ) and defines the physics residual

$$\mathbf{r}_\theta(\mathbf{x}) = \mathcal{N}[\mathbf{u}_\theta(\mathbf{x})], \quad (6.2)$$

which can be evaluated at a set of interior *collocation points* $\{\mathbf{x}_f\}$ within the domain. The derivatives required to compute $\mathbf{r}_\theta$ are obtained via *automatic differentiation*, avoiding explicit numerical discretization of derivatives.

Training objective

Training is performed by minimizing a composite loss, typically of the form

$$\mathcal{L}(\theta) = \lambda_d \mathcal{L}_{\text{data}} + \lambda_f \mathcal{L}_{\text{phys}} + \lambda_b \mathcal{L}_{\text{bc/ic}}, \quad (6.3)$$

where:

- $\mathcal{L}_{\text{data}}$ penalizes mismatch between predictions and available measurements (e.g., pressure samples or peripheral signals);
- $\mathcal{L}_{\text{phys}}$ penalizes violations of the governing equations, e.g., $\mathcal{L}_{\text{phys}} = \frac{1}{N_f} \sum_{\mathbf{x}_f} \|\mathbf{r}_\theta(\mathbf{x}_f)\|_2^2$;
- $\mathcal{L}_{\text{bc/ic}}$ enforces boundary and initial conditions (or compatibility constraints) where they are known.

The weights $\lambda_d, \lambda_f, \lambda_b$ control the relative importance of these terms.

Why they matter in biomedical applications

In biomedical settings, and particularly in hemodynamics, data are often noisy, incomplete, and unevenly sampled across space and time. PINNs are therefore attractive because they:

- naturally incorporate physiological knowledge (mass and momentum conservation, constitutive relations, and boundary-condition models);
- reduce the space of admissible solutions, promoting physically plausible predictions even under *data scarcity*;
- enable estimation of unobserved physiological parameters (e.g., arterial stiffness or peripheral resistances) by treating them as learnable variables alongside θ .

In summary, PINNs combine the expressiveness of neural models with the structure imposed by physical equations, providing a bridge between purely signal-driven regression and physiological modeling.

6.2 From Data-Driven Models to Physics-Informed Learning

The experimental studies presented in the previous chapter reveal a recurring pattern in cuffless blood pressure estimation. Models trained and evaluated under controlled conditions, such as subject-specific learning or within-database validation, consistently achieve low prediction errors. However, when these same models are tested on unseen subjects or external datasets, their performance degrades substantially.

In the early stages of this research, purely data-driven architectures were intentionally adopted to establish a reference baseline. By progressively increasing model capacity and refining signal preprocessing, it was possible to reconstruct arterial blood pressure waveforms with high accuracy when training and testing were conducted on the same population. These results initially suggested that ECG and PPG signals contained sufficient information to support accurate cuffless blood pressure estimation.

This assumption proved fragile under more realistic evaluation protocols. Cross-database experiments and external validation revealed systematic prediction biases, despite stable temporal behavior within individual recordings. Importantly, the networks did not fail by producing noisy or unstable outputs. Instead, they converged toward incorrect physiological operating points, indicating that the learned mappings were not governed by true cardiovascular dynamics.

This behavior was observed across different architectures and training configurations, suggesting that the limitation is intrinsic to black-box regression approaches applied to physiological systems. In the absence of explicit physical constraints, neural networks are free to exploit dataset-specific correlations, sensor artifacts, and population biases that do not generalize beyond the training domain. As a result, the learned relationship between ECG, PPG, and blood pressure becomes highly dependent on the acquisition context and difficult to interpret physiologically.

These findings indicate that architectural refinement alone is insufficient to resolve the generalization problem. Instead, it becomes necessary to guide the learning process using prior knowledge derived from cardiovascular physiology and hemodynamics. In this framework, the role of the neural network shifts from discovering physical laws to approximating unknown or subject-specific components within a mathematically constrained model.

Physics-Informed Neural Networks offer a principled mechanism to implement this transition. By embedding the governing equations of blood flow and arterial mechanics directly into the loss function, PINNs restrict the solution space to pressure and flow fields that satisfy known conservation laws and boundary conditions. Observational data, such as ECG- and PPG-derived signals, thus act in conjunction with physical constraints rather than in isolation.

The remainder of this chapter develops this approach progressively, starting from a minimal physics-informed formulation and advancing toward a more comprehensive and physiologically grounded model.

6.3 Why Physics Is Needed

The need for physics-informed modeling in cuffless blood pressure estimation does not arise from a lack of expressive power in modern neural networks, but from the intrinsic structure of the cardiovascular system. Blood pressure is not a purely signal-level quantity; it is the macroscopic manifestation of underlying physical processes governed by fluid dynamics, vessel mechanics, and boundary interactions between the heart and the arterial tree.

Electrocardiogram and photoplethysmogram signals provide indirect and partial observations of these processes. ECG reflects the electrical activation of the my-

ocardium, while PPG encodes local volumetric changes associated with pulse wave propagation and peripheral vascular tone. Although both signals are causally related to arterial pressure, their relationship with pressure is mediated by latent variables such as blood flow, arterial compliance, wave reflections, and peripheral resistance quantities that are not directly observed but follow well-established physical laws.

In purely data-driven approaches, these latent mechanisms are implicitly absorbed into the network parameters. As shown in the previous chapter, this implicit encoding is sufficient to achieve low errors under constrained conditions, but it lacks robustness when the physiological operating point changes. Without explicit constraints, the network has no incentive to preserve conservation laws, respect realistic pressure–flow relationships, or maintain physically meaningful wave dynamics.

Introducing physics into the learning process serves two complementary purposes. First, it reduces the solution space by enforcing known governing equations, thereby eliminating spurious mappings that are statistically plausible but physiologically invalid. Second, it provides a natural inductive bias that guides the network toward representations aligned with cardiovascular dynamics, improving stability across subjects, acquisition protocols, and time.

Importantly, the objective of physics-informed learning in this context is not to replace data with equations, but to combine them. Physiological models of the arterial system are necessarily simplified and rely on parameters that vary across individuals and conditions. Neural networks, in turn, are well suited to approximate unknown functions, infer subject-specific parameters, and compensate for model mismatches. Physics-Informed Neural Networks offer a principled framework in which these two components coexist: physical laws define the structure of admissible solutions, while data drive personalization and adaptation.

For these reasons, the following sections introduce the hemodynamic foundations adopted in this work and describe how they are embedded into the training objective of the proposed PINN architectures. The focus is on one-dimensional arterial blood flow models, which strike a balance between physiological fidelity and computational tractability, and are particularly well suited for integration with sparse, non-invasive measurements.

6.4 Hemodynamic Model Foundations

6.4.1 Physical Variables and Learnable Hemodynamic Parameters

This chapter adopts a one-dimensional (1D) hemodynamic formulation to describe pressure and flow wave propagation along a simplified arterial path from the aortic root to the fingertip. The state of the system is described by the space–time fields $\{A(x, t), Q(x, t), p(x, t)\}$, i.e., cross-sectional area, volumetric flow rate, and pressure, respectively. Before introducing the governing equations, this section summarizes the set of *learnable* physiological and calibration parameters used in the proposed PINN, consistently with the formulation adopted in the Wiley/Sensors reports.

Learnable parameter vector. For each input window, the network predicts a compact parameter vector

$$\theta \in \mathbb{R}^{12} = (\rho, K_r, \beta, A_0, R_1, R_2, C, p_0, p_{ext}, \alpha, Q_s, Q_{off}), \quad (6.4)$$

which governs the coefficients of the 1D PDE system, the Windkessel termination, and the distal PPG-based anchoring.

Physical parameters (hemodynamics).

- ρ [kg m^{-3}]: blood density, typically treated as weakly learnable around physiological ranges.
- K_r [–]: effective distributed viscous/friction coefficient appearing in the momentum balance (lumped representation of dissipative losses).
- A_0 [m^2]: reference lumen area (baseline geometry) used in the pressure–area constitutive closure.
- β [mmHg m^{-1}]: stiffness/elasticity parameter controlling the pressure–area coupling (arterial compliance surrogate).

Terminal (Windkessel) parameters. At the distal end, the downstream vasculature is represented via a three-element Windkessel model:

- R_1 [Pa s m^{-3}]: proximal characteristic resistance (impedance-related term);
- R_2 [Pa s m^{-3}]: distal peripheral resistance (microvascular load);
- C [$\text{m}^3 \text{Pa}^{-1}$]: compliance of the distal bed.

Pressure reference and external pressure.

- p_0 [mmHg]: baseline pressure offset used to anchor the distal pressure level (calibration/intercept term).
- p_{ext} [mmHg]: external (reference) pressure entering the transmural pressure definition in the constitutive law; typically close to zero but kept learnable with weak regularization.

PPG-to-pressure anchoring. At the fingertip boundary, the measured PPG is used as a soft anchor for the pressure waveform:

$$p(L, t) \approx p_0 + \alpha \text{PPG}(t), \quad (6.5)$$

where

- α [mmHg]: linear gain mapping the (normalized) PPG morphology to pressure amplitude.

Flow calibration. Since true inflow/flow measurements are not available in the second-generation PINN that will be discussed later, the model optionally couples the predicted flow to a surrogate derived from the PPG upstroke (e.g., normalized $d\text{PPG}/dt$), using two simple calibration scalars:

- Q_s [dimensionless]: multiplicative scale for the surrogate flow amplitude;
- Q_{off} [$\text{m}^3 \text{s}^{-1}$]: additive offset to adjust the mean flow level.

Modeling assumptions (summary). The 1D model relies on standard hemodynamic assumptions: incompressible Newtonian blood, axisymmetric flow, negligible radial velocity components, and an elastic wall described through a pressure–area relation. Under these hypotheses, mass and momentum conservation yield the governing PDE system introduced next.

6.4.2 Mass and Momentum Conservation

The macroscopic dynamics of arterial blood flow are described, under appropriate assumptions, by the one-dimensional conservation laws of mass and momentum. These equations represent a reduced-order formulation of the Navier–Stokes equations, obtained by integrating the three-dimensional fluid equations over the vessel cross-section and retaining the dominant axial dynamics.

Let $x \in [0, L]$ denote the spatial coordinate along the arterial centerline and t the time variable. As introduced in Section 6.4.1, the primary state variables are:

- the cross-sectional area $A(x, t)$,
- the volumetric flow rate $Q(x, t)$,
- the internal blood pressure $p(x, t)$.

Blood is modeled as an incompressible Newtonian fluid with (approximately) constant density ρ , and the vessel is assumed to be axisymmetric with slowly varying geometry.

Conservation of mass. For an incompressible fluid, the integral form of mass conservation over an infinitesimal arterial segment yields

$$\frac{\partial A(x, t)}{\partial t} + \frac{\partial Q(x, t)}{\partial x} = 0, \quad (6.6)$$

which expresses the balance between temporal variations of the vessel cross-section and spatial variations of volumetric flow. Physically, Eq. (6.6) enforces that any local increase in arterial area must be accompanied by a corresponding inflow of blood, and vice versa.

Conservation of momentum. The one-dimensional momentum balance is obtained by integrating the axial component of the Navier–Stokes equations over the vessel cross-section and applying standard closure assumptions on the velocity profile and wall shear stress. The resulting equation reads

$$\frac{\partial Q(x,t)}{\partial t} + \frac{\partial}{\partial x} \left(\frac{Q^2(x,t)}{A(x,t)} \right) + \frac{A(x,t)}{\rho} \frac{\partial p(x,t)}{\partial x} = - \frac{2\pi R(x,t) \tau_w(x,t)}{\rho}, \quad (6.7)$$

where $R(x,t)$ is the vessel radius and τ_w denotes the wall shear stress.

In practical reduced-order hemodynamic models, the viscous term on the right-hand side is commonly expressed through a phenomenological friction representation. Assuming laminar flow and a prescribed velocity profile, the wall shear contribution can be approximated as

$$- \frac{2\pi R \tau_w}{\rho} \approx - \frac{K_r}{\rho} \frac{Q(x,t)}{A(x,t)}, \quad (6.8)$$

where K_r is the effective friction coefficient introduced in Section 6.4.1, lumping viscous and profile-dependent effects.

This leads to the commonly adopted one-dimensional momentum equation

$$\frac{\partial Q}{\partial t} + \frac{\partial}{\partial x} \left(\frac{Q^2}{A} \right) + \frac{A}{\rho} \frac{\partial p}{\partial x} + \frac{K_r}{\rho} \frac{Q}{A} = 0. \quad (6.9)$$

Hyperbolic structure and wave propagation. Equations (6.6) and (6.9) form a nonlinear hyperbolic system of partial differential equations. In the absence of viscous dissipation, the system admits wave-like solutions corresponding to the propagation of pressure and flow disturbances along the arterial tree. This property underlies the physiological interpretation of arterial pulse waves and establishes a direct link between pressure dynamics and peripheral optical measurements such as PPG.

Role within the PINN framework. Within the Physics-Informed Neural Network formulation, Eqs. (6.6) and (6.9) are enforced as soft constraints through dedicated residual loss terms. The neural network approximates the continuous fields $p(x,t)$ and $Q(x,t)$, while the cross-sectional area $A(x,t)$ is obtained through the constitutive pressure–area relation introduced in the following subsection.

Although derived under simplifying assumptions, these equations encode essential physiological principles—mass conservation, momentum balance, and wave propagation—that remain valid across a wide range of arterial conditions. Their inclusion therefore acts as a structural regularizer, anchoring the learned solutions to physically admissible cardiovascular dynamics.

6.4.3 Pressure–Area Constitutive Laws

The one-dimensional conservation equations introduced above describe the coupled evolution of cross-sectional area $A(x, t)$, volumetric flow rate $Q(x, t)$, and pressure $p(x, t)$. However, the system remains underdetermined unless a constitutive relationship linking arterial deformation to internal pressure is specified. This mechanical closure reflects the elastic properties of the arterial wall and plays a central role in pulse-wave propagation, impedance, and pressure amplification phenomena.

In large and medium-sized arteries, vessels can be reasonably modeled as thin-walled elastic tubes undergoing small-to-moderate radial deformations around a reference configuration. Under these assumptions, the arterial cross-sectional area is expressed as a function of the local transmural pressure.

General pressure–area relationship. A generic constitutive law can be written as

$$p(x, t) - p_0 = f(A(x, t), A_0(x)), \quad (6.10)$$

where p_0 denotes a reference pressure (introduced in Section 6.4.1), and $A_0(x)$ is the reference cross-sectional area at $p = p_0$. The function $f(\cdot)$ encodes the elastic response of the arterial wall.

Linearized pressure–area law. In this work, arterial mechanics are modeled using a linear elastic approximation,

$$A(x, t) = A_0(x) \left(1 + \frac{p(x, t) - p_0}{\beta(x)} \right), \quad (6.11)$$

where $\beta(x)$ is the stiffness parameter introduced earlier, aggregating wall elasticity, thickness, and material properties. This formulation assumes small relative

deformations,

$$\left| \frac{p(x,t) - p_0}{\beta(x)} \right| \ll 1,$$

which is generally satisfied in normotensive and mildly hypertensive conditions.

Equation (6.11) yields

$$\frac{\partial A}{\partial t} = \frac{A_0(x)}{\beta(x)} \frac{\partial p}{\partial t}, \quad \frac{\partial A}{\partial x} = \frac{\partial A_0}{\partial x} \left(1 + \frac{p - p_0}{\beta} \right) + \frac{A_0(x)}{\beta(x)} \frac{\partial p}{\partial x}, \quad (6.12)$$

allowing consistent computation of the derivatives required by the mass and momentum equations.

Physiological interpretation. The parameter $A_0(x)$ represents the baseline arterial geometry, while $\beta(x)$ governs local arterial compliance. Lower values of β correspond to more compliant vessels, whereas higher values reflect increased stiffness, as observed with aging or pathological conditions.

In classical hemodynamic models, $A_0(x)$ and $\beta(x)$ are prescribed from anatomical measurements or empirical laws. Within the PINN framework adopted here, these parameters may be treated as fixed priors or as learnable (possibly spatially varying) quantities constrained by the governing equations, enabling subject-specific physiological adaptation.

Role in wave propagation. The pressure–area relationship determines the local pulse-wave velocity

$$c(x) = \sqrt{\frac{A_0(x)}{\rho \beta(x)}}, \quad (6.13)$$

which directly influences pulse transit time, pressure amplification, and reflection phenomena, linking arterial mechanics to observable features in ECG and PPG signals.

Consistency with PINN formulation. Embedding Eq. (6.11) into the physics-based loss function provides the mechanical closure required to express all PDE residuals solely in terms of neural network outputs and their derivatives. Together with the conservation laws, this constitutive model forms the physical backbone of the proposed PINN architectures. Peripheral effects and boundary conditions are

introduced separately through lumped-parameter models, discussed in the following section.

6.4.4 Windkessel Models

The one-dimensional equations in Section 6.4.2–6.4.3 describe wave propagation along an arterial segment, but they do not specify how the segment is *loaded* at its downstream end. In practice, the distal vasculature (small arteries, arterioles, capillaries, venules) is not explicitly resolved in 1D models, yet it strongly shapes the pressure and flow waveforms through resistive dissipation and compliant storage.

To account for the net impedance of the unresolved microcirculation, it is standard to terminate the 1D domain at $x = L$ with a *lumped-parameter* model. Among the available choices, Windkessel models provide a compact and physiologically interpretable approximation of the peripheral load.

In the following, both the two-element and three-element Windkessel models are described.

Notation at the outlet. Let

$$p_L(t) := p(L, t), \quad Q_L(t) := Q(L, t)$$

the pressure and flow at the distal end of the 1D segment. All pressures are expressed in mmHg and all flows in consistent volumetric units, as in the remainder of this chapter.

Two-element Windkessel (RC): distal resistance and compliance

The two-element Windkessel (often called RC or “WK2”) represents the downstream vasculature as a *distal resistance* R in parallel with a *compliance* C (capacitor). Physiologically, R captures the steady dissipation mainly due to arterioles, while C approximates the ability of the microvascular bed to store volume and smooth pulsatility.

The total outflow from the terminal node splits into:

- a resistive component $Q_R(t) = \frac{p_L(t)}{R}$,
- a capacitive component $Q_C(t) = C \frac{dp_L(t)}{dt}$,

so that the conservation of flow at the terminal node yields

$$Q_L(t) = \frac{p_L(t)}{R} + C \frac{dp_L(t)}{dt}. \quad (6.14)$$

Equation (6.14) is the WK2 boundary condition: it links the distal pressure waveform to the distal flow by enforcing that part of the incoming flow is stored elastically ($C dp/dt$) and the remainder is dissipated through the peripheral resistance (p/R).

Remarks. WK2 is simple and stable, but it does not explicitly represent the impedance mismatch between the last resolved artery and the downstream load. In wave propagation problems, this may lead to unrealistic high-frequency reflections, because the characteristic impedance of the 1D segment is not matched by the terminal model.

Three-element Windkessel (RCR): characteristic impedance and distal bed

The three-element Windkessel (RCR or “WK3”) introduces an additional *proximal resistance* R_1 in series with the parallel RC branch, yielding a more realistic representation of wave transmission and reflection at the outlet. The parameters are:

- R_1 : proximal resistance, often interpreted as the *characteristic impedance-like* component of the terminal artery;
- C : compliance of the downstream vascular bed;
- R_2 : distal resistance representing microvascular dissipation (arteriolar/capillary network).

Let $p_C(t)$ be the pressure across the parallel RC branch (i.e., the pressure at the node after R_1). Then:

$$\text{series drop: } p_L(t) - p_C(t) = R_1 Q_L(t), \quad (6.15)$$

$$\text{parallel branch: } Q_L(t) = \frac{p_C(t)}{R_2} + C \frac{dp_C(t)}{dt}. \quad (6.16)$$

By eliminating $p_C(t)$, one obtains a single differential relation between (p_L, Q_L) . From (6.15), $p_C = p_L - R_1 Q_L$ and thus

$$\frac{dp_C}{dt} = \frac{dp_L}{dt} - R_1 \frac{dQ_L}{dt}.$$

Substituting into (6.16) gives

$$Q_L(t) = \frac{p_L(t) - R_1 Q_L(t)}{R_2} + C \left(\frac{dp_L(t)}{dt} - R_1 \frac{dQ_L(t)}{dt} \right). \quad (6.17)$$

Rearranging terms yields the common WK3 boundary form

$$Q_L(t) \left(1 + \frac{R_1}{R_2} \right) + CR_1 \frac{dQ_L(t)}{dt} - \frac{p_L(t)}{R_2} - C \frac{dp_L(t)}{dt} = 0 \quad (6.18)$$

Interpretation. Compared to WK2, the WK3 model improves the representation of the *high-frequency* behavior and wave reflections, because the series element R_1 partially captures the impedance seen by traveling waves at the terminal artery. At low frequencies (slow variations), the WK3 behaves approximately like a single resistance $R_1 + R_2$ with compliance C shaping the transient.

Boundary conditions. In the Physics-Informed Neural Network formulations developed in this thesis, the one-dimensional arterial segment is terminated at the distal end by means of a three-element Windkessel (RCR) model, which is adopted as an outlet boundary condition to account for the effective impedance of the unresolved peripheral vasculature. The formal definition and implementation of the boundary conditions within the PINN framework will be introduced later in this chapter.

6.5 State of the Art on Physics-Informed Neural Networks for Cuffless Blood Pressure Estimation

The application of Physics-Informed Neural Networks (PINNs) to cuffless blood pressure estimation represents a relatively recent development within the broader field of physiological signal modeling. Early approaches to non-invasive BP estimation were predominantly data-driven, relying on direct mappings between ECG and PPG signal morphology and blood pressure values through convolutional or recurrent neural architectures [8, 7, 72]. While these models demonstrated encouraging accuracy under intra-database or subject-dependent evaluation protocols, their limited physiological grounding motivated the exploration of alternative formulations.

PINNs, originally introduced by Raissi et al. for solving partial differential equations through constrained learning [29], provided a natural framework for incorporating cardiovascular physics into neural models. In the cardiovascular domain, these methods exploit the availability of well-established hemodynamic descriptions, such as Windkessel-type lumped parameter models [34] and one-dimensional formulations of blood flow derived from mass and momentum conservation laws.

Initial PINN-based studies for blood pressure estimation focused on enforcing simplified hemodynamic relationships within the loss function, using pressure or flow measurements as sparse supervision signals. These early formulations demonstrated that embedding physical constraints can regularize the learning process and promote physiologically consistent pressure waveforms, even when training data are limited [101]. In particular, the inclusion of differential constraints reduced the tendency of purely data-driven models to generate non-physical solutions when extrapolating beyond the training distribution.

Parallel research lines explored hybrid architectures combining deep feature extraction with classical regression or reduced-order models. CNN-based encoders coupled with support vector regression or linear models achieved low mean absolute errors on small cohorts, highlighting the potential of combining learned representations with structured inference [102]. Although not strictly physics-informed in the PDE sense, these approaches contributed to shaping the conceptual space in which physiological priors and machine learning coexist.

More recent contributions extended first-generation PINNs by incorporating richer representations of vascular dynamics. Operator-learning frameworks such as BP-DeepONet explicitly modeled the mapping between boundary conditions and spatio-temporal pressure fields, embedding Navier–Stokes-inspired dynamics and Windkessel boundary conditions into neural operators [36]. Similarly, physics-informed temporal architectures augmented with adversarial or contrastive learning mechanisms were proposed to capture long-term periodicity and temporal coherence in arterial pressure signals [101, 103].

Within this evolving landscape, first-generation PINN models for cuffless blood pressure estimation can be characterized by three defining features. First, they rely on non-invasive peripheral measurements, primarily ECG and PPG, as the sole inputs available at inference time. Second, they embed simplified but physiologically meaningful hemodynamic models as soft constraints during training. Third, arterial blood pressure measurements are typically employed only during training, as reference supervision to anchor the pressure scale, while remaining absent during deployment [104].

These early physics-informed formulations established the feasibility of coupling cardiovascular modeling with neural learning for cuffless BP estimation, and laid the methodological foundation for more advanced architectures that integrate spatial, temporal, and physiological representations within a unified framework.

6.6 First Physics-Informed Approach Based on Inflow Proxy and Peripheral PPG

Following the theoretical and methodological framework introduced in the previous sections, the first concrete attempt to integrate hemodynamic modeling and data-driven learning within a unified Physics-Informed Neural Network was developed in the works presented at WIRN 2025 and ICO. These contributions represent the initial instantiation of the proposed research direction and can be regarded as the first-generation PINN formulation investigated in this thesis.

The core idea underlying this first approach is to reformulate cuffless arterial blood pressure estimation as a constrained spatio-temporal inference problem, in which the evolution of pressure and flow along an abstract arterial pathway is guided

by known physical laws. Rather than relying directly on ECG morphology, the model introduces an explicit approximation of the cardiac inflow waveform $Q_{\text{in}}(t)$, interpreted as the primary mechanical driver of arterial pressure generation. This choice reflects the hemodynamic view that arterial pressure waves originate from ventricular ejection and subsequently propagate through the arterial tree according to conservation laws and vascular loading.

In this formulation, two non-invasive signals are assumed to be available at inference time: (i) a fingertip photoplethysmogram (PPG), providing a local measurement related to peripheral blood volume changes, and (ii) a proxy of the cardiac inflow $Q_{\text{in}}(t)$, heuristically reconstructed from ECG and PPG-derived information. Arterial blood pressure measurements are not used as inputs at deployment, but only as supervisory signals during training at a single reference location, in order to anchor the pressure scale.

The Physics-Informed Neural Network is trained to predict continuous pressure and flow fields $\{p(x,t), Q(x,t)\}$ over a normalized spatial coordinate $x \in [0, 1]$, representing an abstract proximal-to-distal arterial segment. The spatial dimension is not introduced to recover true multi-site pressure measurements, but rather to enable the enforcement of physical consistency through one-dimensional hemodynamic equations and terminal Windkessel boundary conditions. In this way, the learned solution is encouraged to remain coherent with pulse-wave propagation phenomena, even though only single-site invasive pressure measurements are available in the training data.

From a modeling perspective, this first-generation PINN combines three key elements. First, the governing equations of one-dimensional blood flow are embedded directly into the loss function, enforcing mass and momentum conservation in a soft, differentiable manner. Second, a lumped-parameter Windkessel model is used as a terminal constraint, capturing the effect of downstream vascular compliance and resistance. Third, additional supervision is provided through a phenomenological relationship between distal pressure and PPG, linking volumetric changes at the measurement site to the predicted pressure waveform.

6.6.1 Inflow Approximation from ECG and PPG

The first-generation PINN formulation requires an estimate of the cardiac inflow waveform $Q_{\text{in}}(t)$, intended to represent the mechanical excitation generated by ventricular ejection and driving arterial pressure propagation. Since direct measurements of aortic flow are not available in standard non-invasive or ICU waveform datasets, a surrogate inflow signal is constructed from peripheral photoplethysmography (PPG) and electrocardiogram (ECG).

Optical model of PPG and logarithmic transformation

Photoplethysmography can be interpreted, under a Beer–Lambert approximation, as a measurement of light attenuation due to blood volume variations in the illuminated tissue. The recorded PPG signal can be written as

$$\text{PPG}(t) = I_0 e^{-\mu(V_b(t) + V_t)}, \quad (6.19)$$

where I_0 denotes the incident optical intensity, μ is an effective attenuation coefficient, $V_b(t)$ represents the pulsatile blood volume, and V_t is a slowly varying or approximately constant background volume. Taking the natural logarithm yields

$$\log \text{PPG}(t) = \log I_0 - \mu (V_b(t) + V_t). \quad (6.20)$$

Differentiating with respect to time and neglecting the derivative of the non-pulsatile component V_t over the considered time scale gives

$$\frac{d}{dt} \log \text{PPG}(t) = -\mu \frac{dV_b(t)}{dt}. \quad (6.21)$$

Equation (6.21) shows that the time derivative of the logarithm of the PPG signal is directly proportional to the instantaneous rate of pulsatile blood volume change. This quantity constitutes the core element used to construct a surrogate of the cardiac inflow waveform.

Heuristic link between pulsatile volume change and inflow

The arterial inflow generated by ventricular ejection is characterized by a sharp systolic upstroke and a predominantly forward-flow profile, particularly near the proximal arterial segments. Based on this observation, the inflow waveform is assumed to be proportional to the rate of pulsatile blood volume change, modulated by a factor accounting for beat-to-beat ejection strength and cardiac timing.

Formally, the inflow proxy is defined as

$$Q_{\text{in}}(t) \propto \text{SV} \frac{d}{dt} \log \text{PPG}(t) \mathcal{F}(\text{ECG})(t), \quad (6.22)$$

where SV is a stroke-volume-related scaling factor and $\mathcal{F}(\text{ECG})(t)$ is a modulation function derived from the ECG signal. The role of $\mathcal{F}(\text{ECG})(t)$ is to emphasize systolic ejection phases and impose the correct temporal alignment between electrical activation and mechanical inflow.

In practice, $\mathcal{F}(\text{ECG})(t)$ is obtained by applying a low-pass filter with cutoff frequency at 8 Hz to the ECG waveform, producing a smooth, non-oscillatory modulation that suppresses high-frequency components while preserving the dominant cardiac rhythm. This operation transforms the ECG into a slowly varying envelope that gates the pulsatile component extracted from the PPG derivative.

Stroke-volume scaling from PPG amplitude

To account for variations in ejection magnitude across beats, a stroke-volume-related factor is estimated directly from PPG morphology. Under the Beer–Lambert model, the pulsatile blood volume variation over a cardiac cycle corresponds to the integral of the volume change rate. Consistently, the stroke volume proxy is approximated as

$$\text{SV} \propto \int_{\text{systole}}^{\text{diastole}} \frac{d}{dt} \log \text{PPG}(t) dt \propto \max(\text{PPG}) - \min(\text{PPG}), \quad (6.23)$$

computed within a cardiac cycle or within the considered analysis window. This formulation exploits the peak-to-peak PPG amplitude as a surrogate of pulsatile blood volume displacement, providing a simple and computationally efficient estimate of relative stroke volume. The proportionality constant is not explicitly required, as the inflow proxy is used in relative terms within the learning framework.




Fig. 6.1 Example of a generated inflow proxy $Q_{in}(t)$ obtained from the logarithmic derivative of the PPG signal modulated by ECG-derived information.

Final inflow proxy expression

Combining the above components, the inflow approximation can be written explicitly as

$$Q_{in}(t) = k_Q [\max(\text{PPG}) - \min(\text{PPG})] \frac{d}{dt} \log \text{PPG}(t) \mathcal{F}(\text{ECG})(t), \quad (6.24)$$

where k_Q is a global scaling coefficient, for now arbitral, so that the amplitude of Q_{in} is around 5 L/s. In discrete time, the derivative of $\log \text{PPG}(t)$ is computed numerically using finite differences, and the resulting waveform is used directly as input to the PINN.

Figure 6.1 illustrates an example of the resulting $Q_{in}(t)$ waveform generated from a representative PPG and ECG segment.

Interpretation of the inflow proxy

The constructed inflow proxy $Q_{in}(t)$ is not intended as a direct quantitative estimate of the true aortic volumetric flow. Rather, it is designed to encode the essential mechanical information required to excite the arterial system within the physics-informed formulation, namely the timing, relative magnitude, and pulsatile structure of ventricular ejection.

From a physiological perspective, the electrocardiogram provides information about the electrical activation of the heart, which precedes and governs the mechanical contraction of the ventricles. Although ECG does not measure blood flow directly, the depolarization and repolarization dynamics captured by the QRS complex are tightly coupled to the onset and temporal structure of systolic ejection. For this reason, the ECG-derived modulation term $\mathcal{F}(\text{ECG})(t)$ is introduced to encode the timing and phase of cardiac mechanical activity, shaping the inflow proxy so that its dominant features are aligned with ventricular contraction.

The photoplethysmogram, on the other hand, reflects blood volume changes in the peripheral microvasculature and is therefore causally linked to the propagation of the arterial pressure and flow waves generated at the heart. Even though PPG is measured at a distal site, its pulsatile component arises from the same inflow-driven wave that travels through the arterial tree. The logarithmic derivative of the PPG signal captures the local rate of blood volume change, which is assumed to be proportional, up to scaling and distortion, to the underlying flow dynamics that originated from the cardiac inflow.

By combining ECG and PPG, the inflow proxy leverages complementary information. The ECG contributes temporal structure and ejection timing associated with cardiac mechanics, while the PPG contributes a flow-related pulsatile signal that carries the imprint of arterial propagation. Their product yields a waveform that is mechanically plausible as a driving inflow for the one-dimensional hemodynamic equations, without requiring invasive measurements or explicit estimation of ventricular pressure or valve dynamics.

Within the present framework, $Q_{\text{in}}(t)$ should therefore be interpreted as a physically motivated excitation signal. Its role is to provide the PINN with a coherent temporal and morphological driver consistent with cardiovascular physiology, enabling the learning of pressure and flow evolution under the imposed hemodynamic constraints.

6.6.2 Physiological Modeling of Vessel Geometry and Mechanical Properties

In order to enforce physiologically meaningful solutions within the hemodynamic formulation, the one-dimensional blood flow equations are complemented by explicit

models for the nominal vessel geometry and mechanical properties along the arterial segment. In particular, the spatially varying reference cross-sectional area $A_0(x)$, the pressure–area coupling parameter $\beta(x)$, and the distributed viscous resistance $K_r(x)$ are introduced as structured priors guiding the solution space of the physics-informed model.

The spatial coordinate $x \in [0, 1]$ denotes a normalized proximal-to-distal arterial path, abstractly representing the propagation of the pressure wave from central elastic arteries toward peripheral muscular vessels. All parameter profiles are therefore expressed as smooth functions of x , consistent with known large-scale trends in arterial physiology.

Nominal cross-sectional area $A_0(x)$

The reference cross-sectional area $A_0(x)$ defines the vessel geometry in the absence of pulsatile pressure variations. Physiologically, arterial diameter decreases progressively from the aorta to the peripheral arteries, leading to an increase in characteristic impedance and wave speed.

This tapering is modeled through an exponential decay of the vessel diameter,

$$d(x) = d_0 e^{-\lambda x}, \quad (6.25)$$

where d_0 denotes the effective proximal diameter and λ controls the rate of tapering. The corresponding cross-sectional area is then

$$A_0(x) = \frac{\pi}{4} d^2(x) = \frac{\pi}{4} d_0^2 e^{-2\lambda x}. \quad (6.26)$$

In the present formulation, nominal values are chosen in agreement with large-artery physiology:

$$d_0 = 25 \text{ mm}, \quad \lambda \in [0.8, 1.2], \quad (6.27)$$

which yield distal diameters in the range of 6 to 10 mm, consistent with brachial and radial arteries. Rather than being fixed, these parameters define a reference profile around which the model is softly regularized, allowing adaptation to individual recordings. The corresponding nominal radius tapering $r_0(x) = \frac{1}{2}d(x)$ is shown in Fig. 6.2.




Fig. 6.2 Nominal vessel radius profile $r_0(x)$ along the normalized proximal-to-distal coordinate x , consistent with the exponential diameter tapering in Eq. (6.25).

Pressure–area coupling and elasticity parameter $\beta(x)$

The relationship between pressure and vessel deformation is closed through a linearized tube law,

$$A(x, t) = A_0(x) \left(1 + \frac{p(x, t) - p_0}{\beta(x)} \right), \quad (6.28)$$

where p_0 is a reference pressure and $\beta(x)$ quantifies local arterial stiffness. Smaller values of β correspond to more compliant vessels, while larger values indicate increased stiffness.

From classical arterial mechanics, $\beta(x)$ can be expressed as

$$\beta(x) = \sqrt{\pi} E(x) h_0(x), \quad (6.29)$$

where $E(x)$ is the Young modulus of the arterial wall and $h_0(x)$ is the wall thickness. Both quantities are known to vary along the arterial tree, with proximal arteries being more elastic and distal arteries progressively stiffer.

These trends are captured using exponential profiles,

$$E(x) = E_0 e^{\gamma x}, \quad h_0(x) = h_0^* e^{-\nu x}, \quad (6.30)$$




Fig. 6.3 Nominal spatial profile of the arterial wall Young's modulus $E(x)$ along the normalized coordinate x (cf. Eq. (6.30)).

with nominal values

$$E_0 = 0.7\text{--}0.9 \text{ MPa}, \quad h_0^* = 2.5\text{--}3.0 \text{ mm}, \quad \gamma \approx 1.5 \text{ m}^{-1}, \quad \nu \approx 2.0 \text{ m}^{-1}. \quad (6.31)$$

The corresponding nominal profiles for the Young's modulus and wall thickness are shown in Fig. 6.3 and Fig. 6.4, respectively.

This formulation yields an overall increase of $\beta(x)$ toward the periphery, reflecting the transition from elastic to muscular arteries. In the PINN framework, $\beta(x)$ is not enforced deterministically but constrained to remain close to this physiologically plausible profile.

Distributed viscous resistance $K_r(x)$

Viscous energy losses along the arterial segment are modeled through a distributed resistance term $K_r(x)$, appearing in the momentum equation as

$$K_r(x) \frac{Q(x,t)}{A(x,t)}. \quad (6.32)$$

This term accounts for frictional dissipation due to blood viscosity and wall shear stress.




Fig. 6.4 Nominal spatial profile of arterial wall thickness $h_0(x)$ along the normalized coordinate x (cf. Eq. (6.30)).

Blood viscosity is known to increase toward smaller vessels due to plasma skimming and hematocrit redistribution. To reflect this effect, viscosity is modeled as a linearly increasing function of the normalized spatial coordinate,

$$\mu(x) = \mu_0 + (\mu_{\max} - \mu_0)x, \quad (6.33)$$

with

$$\mu_0 = 0.004 \text{ Pa} \cdot \text{s}, \quad \mu_{\max} = 0.006 \text{ Pa} \cdot \text{s}. \quad (6.34)$$

Under a Poiseuille-like approximation, the resistance coefficient is then expressed as

$$K_r(x) = \frac{8\pi\mu(x)}{\rho}, \quad (6.35)$$

where $\rho \approx 1060 \text{ kg m}^{-3}$ is the blood density. This formulation induces progressively stronger damping of flow toward distal locations, consistent with observed attenuation of pulsatile energy in peripheral arteries.

Role as soft physiological priors

The profiles $A_0(x)$, $\beta(x)$, and $K_r(x)$ should be interpreted as structured physiological priors rather than fixed anatomical truths. Their role is to restrict the admissible solution space of the physics-informed model, preventing convergence toward non-

physical regimes while preserving sufficient flexibility to accommodate inter-subject variability.

In the subsequent neural formulation, these quantities are either directly parameterized or weakly regularized toward the profiles described above, enabling the model to learn patient-specific deviations while remaining grounded in cardiovascular physiology.

6.7 Database

The dataset employed in this study is derived from the MIMIC-III Waveform Database, which provides high-resolution, time-synchronized physiological signals recorded in intensive care settings. The signals of interest include electrocardiogram (ECG), photoplethysmogram (PPG), and invasive arterial blood pressure (ABP), which together form a suitable basis for the investigation of hemodynamic relationships.

Since the structure and characteristics of the MIMIC database have already been discussed in detail in previous chapters, only the aspects directly relevant to the present formulation are summarized here. A dedicated preprocessing pipeline is applied to ensure signal integrity, discarding segments affected by artifacts, missing samples, or physiologically implausible values, similar to previous works already discussed in earlier chapters. This step is essential to guarantee stable computation of temporal derivatives and reliable enforcement of physics-based constraints.

The cardiac inflow proxy $Q_{\text{in}}(t)$ is constructed offline from ECG and PPG signals according to the methodology described in the previous subsection. The resulting multichannel data are segmented into overlapping windows of 1000 samples, corresponding to 8-second intervals at the original sampling frequency of 125 Hz. This choice represents a compromise between temporal resolution and computational tractability, while allowing each window to capture multiple cardiac cycles and their associated hemodynamic variability.

The sampling frequency of 125 Hz exceeds by more than a factor of three the highest physiologically relevant frequency content of the signals, which is typically below 40 Hz. This margin satisfies the Nyquist criterion and, at the same time,

improves the numerical stability and accuracy of temporal derivatives appearing in the PINN loss formulation.

After preprocessing and windowing, approximately 156,000 valid segments are retained and used for training and validation. These windows constitute the dataset on which the first-generation physics-informed neural network is trained and evaluated.

6.7.1 Physics Informed Neural Network Architecture

The Physics Informed Neural Network is designed to jointly learn arterial pressure and flow dynamics by combining non invasive temporal measurements with spatial regularization induced by hemodynamic equations. The network approximates continuous pressure and flow fields $\{p(x, t), Q(x, t)\}$ over a normalized spatial coordinate $x \in [0, 1]$, representing an abstract proximal to distal arterial pathway.

It is important to stress that, in typical clinical recordings, invasive arterial blood pressure is measured at a single catheter location. As a consequence, ground truth pressure information is available only at one spatial position, here denoted as a reference coordinate x_{ref} . Predictions at other spatial locations are not directly supervised and must not be interpreted as measured quantities. Instead, the spatial dimension is introduced to enable physics based regularization. The governing equations and terminal boundary conditions constrain the predicted fields to remain consistent with pulse wave propagation and vascular loading along the arterial segment.

Input representation and spatial sampling The input space consists of both temporal and spatial components. For each sample, the temporal inputs include 8 seconds of the inflow proxy $Q_{\text{in}}(t)$ and the photoplethysmogram signal $PPG(t)$, corresponding to 1000 time samples at a sampling frequency of 125 Hz. In addition, a spatial coordinate $x \in [0, 1]$ is provided as input to the network.

For each pair of temporal signals, a set of spatial points $\{x_n\}_{n=1}^N$ is sampled from the interval $[0, 1]$. The points are initially drawn from a uniform distribution and subsequently sorted in ascending order to facilitate numerical differentiation along the spatial dimension. To ensure stability of the learned dynamics and proper enforcement of boundary conditions, the endpoints $x = 0$ and $x = 1$ are always included in the sampled set. As a result, the network processes N inputs of the

form $(x_n, Q_{\text{in}}(t), PPG(t))$, where the same temporal signals are paired with different spatial locations.

Temporal feature extraction The inflow proxy $Q_{\text{in}}(t)$ and the PPG signal are processed by two parallel temporal branches with identical structure. Each branch consists of two consecutive residual convolutional blocks with skip connections. Each block includes 32 one dimensional convolutional filters with kernel size 5, followed by batch normalization and ReLU activation. Dropout with rate 0.1 is applied to improve regularization and reduce overfitting.

The outputs of the convolutional blocks are flattened and projected into a 32 dimensional latent space. These embeddings encode the temporal morphology and dynamic patterns of the input signals and serve as compact representations for subsequent processing.

Spatial embedding The spatial coordinate x is processed by a dedicated dense subnetwork. This branch consists of two fully connected layers with 32 neurons each, interleaved with dropout layers. The purpose of this branch is to embed the spatial position into a latent representation that can interact with the temporal features and modulate the predicted dynamics along the arterial segment.

Feature fusion and learnable physiological parameters The latent features extracted from the temporal and spatial branches are concatenated into a single joint representation. This fused feature vector is passed to a learnable parameter module that outputs physiologically constrained values for a set of arterial properties, including blood density ρ , pressure area coupling parameter β , nominal cross sectional area A_0 , and distributed resistance K_r .

These quantities are not treated as fixed constants, but as sample dependent parameters learned by the network under implicit physiological constraints. Their role is to guide the solution toward physically meaningful regimes while allowing adaptation to inter subject variability.

Prediction layers The final stage of the network combines the fused latent representation and the learned physiological parameters and propagates them through




Fig. 6.5 Schematic overview of the proposed PINN architecture, highlighting the temporal encoders for $Q_{in}(t)$ and $PPG(t)$, the spatial embedding, feature fusion, and the prediction heads for pressure and flow.

fully connected layers with 128 and 256 neurons. Dropout with rate 0.1 is applied to improve generalization. The network outputs continuous estimates of pressure and flow for each spatial location and time sample according to

$$p(x, t) = f_{\theta}(x, Q_{in}, PPG), \quad Q(x, t) = g_{\theta}(x, Q_{in}, PPG), \quad (6.36)$$

where f_{θ} and g_{θ} denote the learned mappings parameterized by the network weights.

Output activations are selected to reflect basic physical sign conventions. Pressure is constrained to be non negative, and flow is predominantly forward along the modeled segment. These constraints are imposed softly, allowing the model to retain sufficient flexibility in the presence of wave reflections and oscillatory components.

A schematic overview of the complete network architecture is reported in Fig. 6.5.

6.7.2 Physics Informed Loss Function

Training is formulated as the minimization of a composite objective that combines data fidelity terms with physics based constraints. The network outputs continuous fields $\{p(x, t), Q(x, t)\}$ over $x \in [0, 1]$ and over the temporal support of each window. The loss is designed so that the predicted fields are simultaneously consistent with the available measurements at a reference location and with the governing hemodynamic model along the entire spatial domain.

Global objective

Let θ denote the trainable weights of the neural network. For each input window, the model is evaluated at a set of spatial points $\{x_n\}_{n=1}^N \subset [0, 1]$ and at the discrete time samples $\{t_m\}_{m=1}^T$, with $T = 1000$. The total loss is expressed as a weighted sum of complementary components,

$$\mathcal{L}(\theta) = \lambda_{\text{PDE}} \mathcal{L}_{\text{PDE}} + \lambda_W \mathcal{L}_W + \lambda_{\text{ABP}} \mathcal{L}_{\text{ABP}} + \lambda_{\text{PPG}} \mathcal{L}_{\text{PPG}} + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}}, \quad (6.37)$$

where the $\lambda_{\bullet} > 0$ coefficients control the relative strength of each constraint. In this first generation formulation, they are treated as fixed hyperparameters.

PDE residual loss $\mathcal{L}_{\text{PDE}}$

The core physics term enforces compliance with the one dimensional conservation laws of mass and momentum. Since the PINN outputs $\{p(x, t), Q(x, t)\}$, the vessel area $A(x, t)$ required by the PDE system is obtained through the linearized tube law introduced previously,

$$A(x, t) = A_0(x) \left(1 + \frac{p(x, t) - p_0}{\beta(x)} \right). \quad (6.38)$$

The PDE residuals are then defined as

$$\begin{aligned} r_1(x, t) &= \frac{\partial A}{\partial t}(x, t) + \frac{\partial Q}{\partial x}(x, t), \\ r_2(x, t) &= \frac{\partial Q}{\partial t}(x, t) + \frac{\partial}{\partial x} \left(\frac{Q^2}{A} \right)(x, t) + \frac{A}{\rho} \frac{\partial p}{\partial x}(x, t) + K_r(x) \frac{Q}{A}(x, t). \end{aligned} \quad (6.39)$$

The PDE residual loss penalizes the squared magnitude of these residuals over the sampled space time grid,

$$\mathcal{L}_{\text{PDE}} = \frac{1}{NT} \sum_{n=1}^N \sum_{m=1}^T \left(r_1(x_n, t_m)^2 + r_2(x_n, t_m)^2 \right). \quad (6.40)$$

All spatial and temporal derivatives in Eq. (6.39) are computed through automatic differentiation of the neural outputs with respect to the continuous inputs x and t , which avoids explicit finite difference schemes and keeps the PDE constraints fully differentiable.

Terminal Windkessel constraint $\mathcal{L}_W$

To represent the downstream vascular load beyond the modeled segment, a three element Windkessel boundary condition is enforced at the distal boundary $x = 1$ (equivalently $x = L$ in dimensional notation). Let $p_L(t) \equiv p(1, t)$ and $Q_L(t) \equiv Q(1, t)$. The Windkessel dynamics introduce the terminal resistances R_1 and R_2 and the compliance C , yielding the constraint residual

$$r_W(t) = Q_L(t) R_1 + C R_1 \frac{dQ_L(t)}{dt} - \frac{p_L(t)}{R_2} - C \frac{dp_L(t)}{dt}. \quad (6.41)$$

The corresponding loss penalizes the squared residual over time,

$$\mathcal{L}_W = \frac{1}{T} \sum_{m=1}^T r_W(t_m)^2. \quad (6.42)$$

This term couples the learned pressure and flow at the distal boundary, encouraging physiologically consistent reflections and damping through the lumped peripheral impedance.

Reference ABP supervision $\mathcal{L}_{\text{ABP}}$

Invasive ABP is available only at a single catheter location in the recordings. Accordingly, direct supervision is imposed at a single reference coordinate x_{ref} within the normalized domain. Denoting by $\text{ABP}(t)$ the measured arterial pressure waveform

for the current window, the data fidelity term is

$$\mathcal{L}_{\text{ABP}} = \frac{1}{T} \sum_{m=1}^T (p(x_{\text{ref}}, t_m) - \text{ABP}(t_m))^2, \quad (6.43)$$

where x_{ref} is fixed during training and represents the approximately consistent measurement site across windows. This loss anchors the absolute pressure scale of the predicted field, while the spatial evolution away from x_{ref} is primarily shaped by the physics terms.

PPG based distal constraint $\mathcal{L}_{\text{PPG}}$

To provide an additional non invasive constraint at the distal site, the fingertip PPG waveform is linked to the predicted distal pressure through an empirical logarithmic relation. Let $\text{PPG}(t)$ denote the measured PPG signal in the considered window. The distal pressure is constrained by

$$p(1, t) \approx p_0 + \alpha \log(\text{PPG}(t) + \varepsilon), \quad (6.44)$$

where p_0 represents a baseline pressure, α is a scaling coefficient, and $\varepsilon > 0$ prevents numerical issues when $\text{PPG}(t)$ approaches zero. The corresponding loss is

$$\mathcal{L}_{\text{PPG}} = \frac{1}{T} \sum_{m=1}^T (p(1, t_m) - p_0 - \alpha \log(\text{PPG}(t_m) + \varepsilon))^2. \quad (6.45)$$

In this formulation, α and p_0 are treated as learnable quantities within the model, enabling the constraint to adapt to inter window variability while maintaining a physiologically motivated monotonic link between distal perfusion changes and pressure.

Physiological regularization $\mathcal{L}_{\text{reg}}$

Finally, the learnable physiological parameters are softly regularized around nominal values in order to discourage non physical solutions. Let ρ denote the predicted blood density (or a learnable effective density term) and let $A_0(x)$ and $K_r(x)$ denote the geometry and resistance profiles used by the model. The regularization loss

penalizes deviations from reference priors,

$$\mathcal{L}_{\text{reg}} = \|\rho - \rho_0\|^2 + \|A_0 - A_0^{\text{nom}}\|^2 + \|K_r - K_r^{\text{nom}}\|^2, \quad (6.46)$$

where $\|\cdot\|$ denotes an L^2 norm over the discretized spatial grid (or over the corresponding parameter vector representation). In practice, this term acts as a soft prior that stabilizes training by preventing the network from compensating modeling errors through extreme parameter values.

Remarks on identifiability and enforcement

The loss in Eq. (6.37) is constructed so that different terms constrain complementary aspects of the solution. The ABP term anchors the pressure scale at a reference point, the PPG term injects distal information from non invasive sensing, and the Windkessel and PDE terms enforce global hemodynamic consistency. In this first generation setting, the primary objective is feasibility and stable convergence of a physics guided learning formulation, rather than strict parameter identifiability. Accordingly, the physiological parameters should be interpreted as effective, sample dependent quantities that regularize the predicted fields in a plausible regime.

Optimization is performed with Adam using learning rate 10^{-3} . At each iteration, the spatial points $\{x_n\}$ are resampled for each waveform window, so that the PDE residual is enforced over diverse collocation sets across epochs, improving coverage of the spatial domain while keeping the computational burden tractable.

6.7.3 Training and Results

After defining the network architecture and the composite physics-informed loss, the model is trained using mini-batch optimization. In this context, it is important to clarify the structure of a training batch, which differs from standard supervised learning due to the presence of an explicit spatial dimension.

Each batch is composed of 16 distinct waveform pairs, where each pair corresponds to one physiological recording of the inflow proxy $Q_{\text{in}}(t)$ and the photoplethysmogram $PPG(t)$. For each waveform pair, the spatial domain $x \in [0, 1]$ is discretized by sampling N spatial points. Consequently, a single batch contains

$16 \times N$ input triplets of the form $(x, Q_{\text{in}}(t), PPG(t))$, all sharing the same temporal signals but evaluated at different spatial locations.

In the present implementation, the batch size is fixed to 16 waveform pairs and the number of spatial points is set to $N = 11$. Among these points, the midpoint $x = 0.5$ is always explicitly included, together with the boundary locations $x = 0$ and $x = 1$. This configuration yields a total of 176 input instances per batch. The relatively small value of N is chosen to mitigate numerical instabilities and gradient explosion that may arise when enforcing spatial derivatives in the PDE residual, particularly in the presence of sharp pressure upstrokes. Moreover, limiting the number of spatial collocation points prevents losses applied at specific locations—such as the ABP supervision at x_{ref} and the Windkessel and PPG constraints at the boundaries—from being dominated by the globally enforced PDE residual.

The dataset is randomly partitioned into 80% training, 10% validation, and 10% test sets. Training is performed for 50 epochs using the Adam optimizer with learning rate 10^{-3} , requiring approximately two hours of computation on an NVIDIA GeForce RTX 3050 GPU. Inference consists of a single forward pass of the network evaluated at the sampled (x, t) points and is therefore expected to be compatible with near real-time operation. A detailed latency analysis is left for future work.

Hemodynamic closure and computation of derivatives

The PDE residual loss introduced in Section 6.7.2 is formulated in terms of cross-sectional area $A(x, t)$, flow $Q(x, t)$, and pressure $p(x, t)$. Since the neural network directly predicts $\{p(x, t), Q(x, t)\}$, the vessel area is recovered through the linearized pressure–area relation

$$A(x, t) = A_0(x) \left(1 + \frac{p(x, t) - p_0}{\beta(x)} \right), \quad (6.47)$$

where $A_0(x)$ and $\beta(x)$ are the nominal geometry and elasticity profiles introduced previously. This closure ensures consistency between the predicted variables and allows computing $\partial A / \partial t$ directly from $\partial p / \partial t$.

All spatial and temporal derivatives required by the PDE residuals and boundary constraints are computed through automatic differentiation of the network outputs with respect to the continuous inputs x and t . Time is treated as a continuous

coordinate associated with the sampled waveform points, enabling direct evaluation of $\partial(\cdot)/\partial t$ and $\partial(\cdot)/\partial x$ without resorting to finite-difference schemes. This approach avoids numerical artifacts related to discrete differencing across sparse or irregular spatial samples and preserves full differentiability of the loss.

To improve numerical stability during training, several stabilization strategies are adopted. First, all input signals and the spatial coordinate x are normalized to the interval $[0, 1]$. Second, boundary points at $x = 0$ and $x = 1$, together with a fixed reference location x_{ref} , are always included in each spatial set. Third, small ε terms are introduced in denominators to prevent numerical singularities. Finally, the number of spatial collocation points per waveform is intentionally limited, preventing the PDE residual from dominating the optimization or inducing unstable gradients. More advanced nondimensionalization strategies and adaptive loss balancing are deferred to future developments.

Results and qualitative observations

Figure 6.6 reports representative examples of predicted arterial blood pressure waveforms along the normalized arterial coordinate, from the proximal end ($x = 0$) to the distal site ($x = 1$). The predicted pressure fields exhibit physiologically consistent trends. Near the proximal location, the systolic peak appears broader and wave reflections are shifted further from the primary upstroke. Moving toward the distal end, the waveform progressively sharpens and reflected components become more pronounced, in agreement with established hemodynamic behavior.




Fig. 6.6 Predicted arterial blood pressure waveforms at different spatial locations $x \in [0, 1]$ along the modeled arterial segment. The horizontal axis represents time (8 seconds). The black curve denotes the ground-truth ABP at the reference location.

6.7.4 Discussion and Scope of the First-Generation Model

The results presented in this section demonstrate the feasibility of a physics informed neural network formulation for arterial blood pressure estimation from non invasive signals. By coupling a lightweight neural architecture with first principle hemodynamic constraints, the proposed model is able to converge stably and produce pressure and flow fields that are internally consistent and physiologically plausible along a normalized arterial coordinate.

The predicted pressure waveforms exhibit coherent qualitative trends associated with pulse wave propagation and vascular loading, including progressive waveform sharpening toward distal locations and attenuation of pulsatile energy. These behaviors emerge despite the limited supervision available during training, which is restricted to a single pressure measurement site, and highlight the regularizing role of the embedded physical constraints. At the same time, deviations remain in the accurate reconstruction of systolic peaks and in the fine scale morphology of the waveform, indicating that the present formulation should be regarded as a proof of feasibility rather than a fully optimized estimator.

The primary contribution of this first generation model lies in establishing that a physics guided learning strategy can successfully integrate non invasive inputs, heuristic inflow information, and reduced order hemodynamic models within a single differentiable framework. Rather than aiming at state of the art numerical accuracy, this phase is intended to validate the conceptual soundness of the approach and to identify the critical modeling components that govern stability, interpretability, and generalization.

6.7.5 Limitations and Directions for Further Development

The present formulation intentionally prioritizes methodological exploration over exhaustive performance evaluation. Several aspects required for deployment oriented validation and quantitative benchmarking are therefore left to subsequent developments.

Evaluation protocol and generalization In the current implementation, waveform windows are randomly split into training, validation, and test sets. Given the large

number of overlapping windows extracted from the same subjects and recordings, this strategy may introduce information leakage and inflate apparent performance. A strict subject wise partitioning will be adopted in future studies to properly assess generalization to unseen individuals and to quantify robustness across heterogeneous physiological conditions.

Quantitative metrics and clinical reporting The results discussed in this chapter primarily emphasize convergence behavior and qualitative waveform plausibility. Future work will incorporate standard cuffless blood pressure evaluation metrics, including mean absolute and root mean squared errors for systolic, diastolic, and mean arterial pressure, waveform level error measures, and correlation analyses. Where appropriate, Bland Altman analysis and stratification across clinically relevant pressure ranges will be employed to facilitate comparison with established reporting standards.

Role and robustness of the inflow proxy The inflow signal $Q_{in}(t)$ used in this work is a heuristic proxy derived from electrocardiogram and photoplethysmogram information. While sufficient for a feasibility study, this approximation is inherently sensitive to amplitude scaling, signal drift, contact pressure, and vasomotor effects, particularly in the photoplethysmographic signal. A key methodological limitation of the current generation therefore lies in the absence of a learnable and data adaptive inflow representation.

A central direction for further development is the replacement of the heuristic inflow construction with a dedicated neural module that infers cardiac inflow directly from raw non invasive signals. Such a module is expected to improve robustness while preserving the temporal structure required by the hemodynamic model.

Spatial interpretation of predicted fields Because invasive arterial pressure is available only at a single measurement location, predicted values at other spatial coordinates cannot be validated directly and should not be interpreted as ground truth. In the present work, the spatially resolved pressure and flow fields serve as an internal consistency check enforced by the governing equations and boundary conditions. A rigorous experimental validation of spatial evolution would require

multisite pressure measurements or datasets with simultaneous proximal and distal recordings, which remain an explicit target for future investigation.

Toward a generalized and personalized framework The next phase of this research aims to extend the conditioning of the physics informed neural network beyond waveform inputs and spatial coordinate. Incorporating readily available subject information such as age, sex, height, and weight is expected to enable subject adaptive estimation of key physiological parameters, including vascular compliance and Windkessel elements. This evolution would support improved cross subject and cross database robustness while preserving the calibration free nature of the inference stage.

Overall, the first generation model presented in this chapter establishes a principled and interpretable foundation for physics guided blood pressure estimation. It clarifies the role of hemodynamic priors, highlights the limitations of heuristic inflow modeling, and motivates the architectural and methodological extensions that will be introduced in the subsequent generation.

6.8 Second-Generation PINN: Physiological and Anatomical Priors

6.8.1 Motivation for a Revised Framework

The first-generation Physics-Informed Neural Network introduced in the previous section demonstrated the feasibility of integrating non-invasive signals with reduced-order hemodynamic modeling. By embedding conservation laws, Windkessel boundary conditions, and phenomenological constraints, the model converged stably and produced internally consistent pressure and flow fields along a normalized arterial coordinate.

At the same time, the analysis of this first formulation revealed a number of structural limitations that motivate a revised modeling strategy. These limitations do not undermine the validity of the physics-guided approach itself, but rather expose the need for a more explicit and physiologically grounded representation of anatomical and mechanical variability.

Limitations of weakly structured physiological parameters. In the first generation model, key physiological quantities, such as vessel geometry, elasticity, and viscous dissipation, were introduced either as global parameters or as weakly regularized spatial profiles. While sufficient to stabilize training, this representation remained overly coarse. In particular, the reference lumen area A_0 , wall thickness h , and elastic modulus E were not modeled as coherent anatomical descriptors, but rather absorbed implicitly into effective parameters (e.g., β) or learned in a loosely constrained fashion.

This simplification limited both interpretability and generalization. From a physiological standpoint, arterial tapering, wall thickening, and stiffening toward the periphery are well-established features of the arterial tree. From a modeling standpoint, failing to encode these trends explicitly increases the burden placed on the neural network, which must compensate for missing anatomical structure by adjusting waveform-level mappings. As a result, part of the inter-subject and inter-database variability is inadvertently absorbed into the learned signal-to-pressure relationship, rather than being attributed to meaningful changes in vascular properties.

Revisiting anatomical and mechanical priors. The second-generation framework directly addresses this issue by explicitly revising the modeling of arterial geometry and mechanics. Rather than treating A_0 , h , and E as secondary or implicit quantities, they are introduced as structured, spatially varying priors along the proximal-to-distal coordinate. This choice reflects both anatomical knowledge and modeling necessity.

The nominal lumen area $A_0(x)$ is constrained to follow a physiologically plausible tapering profile, while wall thickness $h(x)$ and Young's modulus $E(x)$ are modeled with opposing spatial trends, capturing the transition from elastic central arteries to stiffer, muscular peripheral vessels. These profiles do not enforce subject-specific anatomy, but define a realistic baseline around which the model can adapt. In this way, anatomical variability is elevated from a nuisance factor to a first-class component of the inference process.

Arbitrariness of the inflow proxy. A second major limitation of the first generation formulation concerns the treatment of cardiac inflow. The inflow signal $Q_{\text{in}}(t)$ was constructed through a heuristic combination of ECG and PPG-derived features. While this proxy proved sufficient to excite the hemodynamic model and enable

feasibility studies, it is fundamentally arbitrary. Its amplitude scaling, temporal shaping, and physiological interpretation are not uniquely determined, nor grounded in a validated mechanical model of ventricular ejection.

More importantly, the use of a fixed, externally defined inflow decouples the most uncertain component of the system from the physics-informed learning process. Any mismatch between the proxy and the true aortic flow must be compensated downstream by adjustments in pressure, compliance, or resistance, thereby reducing identifiability and interpretability of the learned parameters. This arbitrariness becomes particularly problematic when generalizing across datasets with different sensor characteristics or population statistics.

Clarifying the distal optical constraint. The revised framework also reconsiders the role of the photoplethysmogram as a distal observation. Rather than assuming a generic proportionality between PPG morphology and pressure, the second-generation model explicitly acknowledges the Beer–Lambert origin of the signal and the approximations required to obtain a usable working model. Under small-signal conditions and after removal of the DC component by the acquisition frontend, the pulsatile PPG can be interpreted as a proxy for local blood volume variations. Through a linearized compliance assumption, this leads to a distal pressure anchoring of the form

$$p(L, t) \approx p_0 + \alpha \text{PPG}(t), \quad (6.48)$$

where both p_0 and α are treated as learnable, window-dependent quantities. This formulation makes explicit what is assumed and what is learned, and integrates the optical constraint coherently within the physics-informed loss.

Toward a more identifiable and interpretable model. Taken together, these considerations motivate a second-generation Physics-Informed Neural Network in which anatomical descriptors, mechanical properties, and distal observations are modeled explicitly rather than absorbed implicitly. The objective is not merely improved numerical accuracy, but a restructuring of the learning problem that reduces degeneracy, improves robustness under cross-database shifts, and enables meaningful interpretation of learned physiological parameters. The following sections detail how these revised priors are integrated into the model architecture and loss formulation.

6.8.2 PPG Link to Blood Flow and Blood Pressure

The previous sections established that purely data-driven mappings between ECG, PPG, and arterial blood pressure are intrinsically ill-posed unless guided by physiological structure. This guidance is made explicit by formalizing the relationship between the optical plethysmogram, local blood volume, flow, and pressure.

The underlying idea is conceptually simple but physically grounded: (i) the PPG sensor measures pulsatile changes in arterial blood volume within a small tissue region; (ii) conservation of mass links volume changes to flow; (iii) under small beat-to-beat excursions, local vascular compliance links flow to pressure dynamics.

In the following, this chain

$$\text{optics} \rightarrow \text{volume} \rightarrow \text{flow} \rightarrow \text{pressure}$$

is translated into a minimal, linearized working model. The objective is not to derive a complete optical–hemodynamic model, but to justify the distal pressure anchoring used in the proposed PINN while keeping all assumptions explicit and interpretable.

6.8.3 From Beer–Lambert Law to a Linear PPG $\sim V_b$ Model

Let $I(t)$ denote the photodetector intensity measured by the PPG sensor, and let $L(t)$ represent the effective optical pathlength through the illuminated vascular bed. Denote by $V_b(t)$ the arterial blood volume within the small tissue region probed by the sensor.

Under general conditions, light attenuation in biological tissue obeys a generalized Beer–Lambert law:

$$I(t) = I_0 \exp\left(-\sum_j \varepsilon_j c_j L_j(t) - G\right), \quad (6.49)$$

where ε_j and c_j are the extinction coefficients and concentrations of optical absorbers, $L_j(t)$ their effective pathlengths, and G accounts for scattering and geometric losses.

Among all absorbing components, only the arterial blood compartment exhibits appreciable variations at cardiac frequencies. All remaining contributions can there-

fore be grouped into a constant term. Let $L_b(t)$ denote the optical pathlength associated with pulsatile arterial blood.

For a fixed sensor geometry and small deformations, the arterial optical pathlength can be approximated as proportional to the enclosed blood volume:

$$L_b(t) = \alpha V_b(t), \quad \alpha > 0, \quad (6.50)$$

where α depends on probe geometry and photon scattering statistics.

Substituting into (6.49) and collecting constant terms yields

$$I(t) = I_0 \exp(-K_0 - k V_b(t)), \quad k := \epsilon_{\text{Hb,eff}} \alpha, \quad (6.51)$$

with K_0 absorbing all static contributions.

Linearization under small pulsatile excursions. In practice, PPG acquisition systems remove the DC component and operate around a steady operating point. Let $V_b(t) = \bar{V}_b + \tilde{V}(t)$, with $|\tilde{V}(t)| \ll \bar{V}_b$. A first-order Taylor expansion of (6.51) around $\bar{V}_b$ gives

$$I(t) \approx \bar{I}(1 - k \tilde{V}(t)), \quad \bar{I} := I_0 e^{-K_0 - k \bar{V}_b}. \quad (6.52)$$

Hence, the pulsatile component of the detected signal satisfies

$$\Delta I(t) \propto \tilde{V}(t). \quad (6.53)$$

After conditioning, normalization, and sign convention, the displayed plethysmogram can therefore be modeled as

$$\boxed{\text{PPG}(t) = \beta + \kappa V_b(t), \quad \kappa > 0,} \quad (6.54)$$

where β absorbs the DC offset. Separating slow and pulsatile components,

$$\text{PPG}(t) = \text{PPG}_{\text{DC}}(t) + \kappa \tilde{V}(t). \quad (6.55)$$

This linearized Beer–Lambert model constitutes the optical foundation of the subsequent derivations.

6.8.4 From Blood Volume to Flow: Continuity

Consider the illuminated tissue region as a small control volume. Mass conservation yields

$$\frac{dV_b(t)}{dt} = Q_{\text{in}}(t) - Q_{\text{out}}(t). \quad (6.56)$$

At cardiac frequencies, venous outflow and slow redistribution processes vary negligibly compared to arterial inflow. After high-pass filtering (as implicitly performed by PPG front-ends), the pulsatile component satisfies

$$\frac{d\tilde{V}(t)}{dt} \approx Q_p(t), \quad (6.57)$$

where $Q_p(t)$ denotes the pulsatile arterial inflow into the sensing region.

Using (6.55), this gives

$$\tilde{V}(t) = \frac{1}{\kappa} (\text{PPG}(t) - \text{PPG}_{\text{DC}}(t)), \quad (6.58)$$

and therefore

$$\boxed{Q_p(t) \propto \frac{d}{dt} \text{PPG}(t)}, \quad (6.59)$$

up to a proportionality constant that absorbs optical scaling and geometric factors.

This result formalizes the widely used observation that the first derivative of the PPG waveform reflects pulsatile blood flow.

6.8.5 From Flow to Pressure: Toward $P(t) \propto \text{PPG}(t)$

To relate flow to arterial pressure, the local compliance definition is invoked

$$C = \frac{dV}{dP}. \quad (6.60)$$

For small beat-to-beat excursions around the operating point, compliance can be assumed approximately constant. Differentiating $V(t) = CP(t)$ yields

$$\frac{dV}{dt} = C \frac{dP}{dt}. \quad (6.61)$$

Combining (6.57) and (6.61) gives

$$Q_p(t) \approx C \frac{dP(t)}{dt}. \quad (6.62)$$

Using (6.59),

$$\frac{d}{dt} \text{PPG}(t) \propto C \frac{dP(t)}{dt}. \quad (6.63)$$

Integrating in time,

$$\text{PPG}(t) \propto C P(t) + \text{const.} \quad (6.64)$$

Since the integration constant corresponds to the DC baseline already removed in practice, this yields

$$\boxed{P(t) \propto \text{PPG}(t)}. \quad (6.65)$$

Equation (6.65) provides a theoretical justification for the use of PPG as a distal pressure surrogate. The proportionality constant embeds both local vascular compliance and optical scaling, while the waveform morphology directly reflects arterial pressure dynamics. In the proposed PINN, this relationship is exploited as a soft distal constraint, with scale and offset treated as learnable quantities rather than fixed calibration parameters.

6.8.6 Comparison with Previous PPG–Pressure Relations and Practical Implications

In earlier formulations of the present framework and in parts of the cuffless BP literature, the relationship between distal pressure and PPG was often expressed using a logarithmic mapping, typically of the form

$$p(t) \approx p_0 + \alpha \log(\text{PPG}(t) + \epsilon), \quad (6.66)$$

motivated by the exponential dependence of detected light intensity on optical pathlength in the Beer–Lambert law.

This choice was physically consistent at the level of raw photodetector intensity, where light attenuation depends exponentially on absorber concentration and pathlength. However, this formulation implicitly overlooks an important practical aspect: the PPG signal available to the model is *not* the raw optical intensity, but the output

of a signal acquisition pipeline that already performs several nonlinear operations, including DC removal, normalization, gain control, and often band-pass filtering.

As shown in Sec. 6.8.3, once the Beer–Lambert law is linearized around a steady operating point and the signal is conditioned as in standard PPG front-ends, the pulsatile component becomes approximately linear with respect to local blood volume variations. Under these conditions, the exponential structure that motivates the logarithmic mapping in (6.66) is effectively removed at the hardware and preprocessing level.

This observation is further supported empirically. When distal arterial pressure waveforms $p(t)$ and PPG signals are temporally aligned and overlaid, their morphologies already exhibit a striking similarity in shape, including systolic upstroke, peak timing, and diastolic decay. In such cases, no systematic nonlinear distortion is observed that would necessitate a logarithmic transformation to reconcile the two signals.

From a modeling perspective, retaining a logarithmic constraint in this setting introduces an unnecessary nonlinearity that is no longer justified by either physiology or signal acquisition physics. For this reason, the present second-generation formulation adopts a linear distal calibration,

$$p(t) \approx p_0 + \alpha \text{PPG}(t), \quad (6.67)$$

where both p_0 and α are treated as learnable, window-specific parameters.

This choice aligns more closely with the effective signal processing chain of real-world PPG devices, preserves the observed waveform similarity between pressure and PPG, and avoids introducing artificial nonlinearities that could bias the physics-informed optimization. Importantly, this linear relation should not be interpreted as a claim of universal proportionality between pressure and PPG, but rather as a locally valid approximation under small pulsatile excursions and conditioned signals, precisely the regime in which the PINN operates.

6.8.7 Physiological Origin and Role of Geometric and Mechanical Priors

To confine the PINN to physiologically credible regimes while preserving subject specificity, *soft* priors are imposed on a compact set of parameters appearing in the 1D hemodynamic model and in the Windkessel termination. These include blood density ρ , reference cross-section $A_0(x)$, wall mechanics through $\beta(x)$, distributed viscous losses $K_r(x)$, and terminal elements (R_1, R_2, C) , all defined with respect to a baseline pressure p_0 and an external pressure P_{ext} . In all cases, $x \in [0, 1]$ denotes the normalized arc length from the aortic root ($x = 0$) to the fingertip ($x = 1$).

A crucial point is that the physiological priors adopted in this work are *not* derived from a single coherent anatomical dataset, nor are they intended to represent subject-specific ground truth. Instead, they are constructed by aggregating nominal values and empirical trends reported across multiple heterogeneous sources in the cardiovascular and hemorheological literature, including classical works by Ku [105], Pries et al. [106], and Olufsen [26], as well as additional anatomical tables and review material commonly used in large-scale hemodynamic modeling.

As a consequence, these priors should be interpreted as providing *orders of magnitude and monotonic trends*, rather than millimeter-accurate morphometry. While the absolute values of vessel radius, wall thickness, or elastic modulus may vary substantially across individuals, measurement protocols, and physiological conditions, the qualitative relationships they encode are robust and well established:

- arterial radius decreases progressively from the aorta to the periphery,
- wall thickness decreases distally,
- effective wall stiffness increases as vessel radius decreases,
- apparent blood viscosity increases in small vessels due to the Fåhræus Lindqvist effect.

In the proposed framework, these trends are implemented as *soft anchors* through quadratic regularization terms around nominal spatial profiles. This design choice deliberately avoids enforcing rigid anatomical constraints, which would be unjustified given the lack of subject-specific geometric information in wearable sensing

scenarios. Instead, the priors act as weak structural guidance that prevents the PINN from drifting toward non-physiological solutions while still allowing adaptation driven by the ECG and PPG inputs.

The intention is therefore not to prescribe exact arterial geometry, but to restrict the solution space to physiologically plausible regimes and to stabilize the inverse problem of reconstructing pressure and flow from distal observations. Within this context, deviations from the nominal profiles are not only allowed but expected, as they implicitly capture subject-specific morphology, vascular tone, and aging effects encoded in the input signals.

6.8.8 Physiological Modeling of Geometry and Wall Mechanics:

$A_0(x)$, $h_0(x)$, and $E(x)$

Following the general rationale introduced in Section 6.8.7, this section details how geometric and mechanical priors are constructed and incorporated into the PINN. The objective is not to reproduce subject-specific arterial anatomy, but to provide physiologically grounded spatial profiles that regularize the hemodynamic equations and constrain the inverse problem to realistic regimes.

Reference cross-sectional area $A_0(x)$ and wall thickness $h_0(x)$. The reference lumen area $A_0(x)$ and arterial wall thickness $h_0(x)$ are derived from nominal anatomical values reported for large and medium arteries across the heart-to-fingertip pathway, including aorta, subclavian, brachial, radial/ulnar, and digital segments [26]. These values were collected from heterogeneous literature sources, anatomical tables, and modeling studies, and subsequently merged into smooth spatial profiles defined over the normalized coordinate $x \in [0, 1]$.

Rather than enforcing piecewise-constant vessel properties or explicit branching geometry, continuous fitted functions $R_0(x)$ and $h_0(x)$ are adopted to capture the dominant monotonic trends observed in vivo: progressive tapering of the lumen and thinning of the arterial wall toward the periphery. The resulting reference profiles are illustrated in Fig. 6.7.

The reference area is then defined as

$$A_0(x) = \pi R_0(x)^2, \quad (6.68)$$




Fig. 6.7 Radius profile (dark blue) and wall-thickness (orange) used as physiological prior

and enters the pressure–area constitutive relation as the baseline geometry. Importantly, deviations from these nominal profiles are penalized softly rather than enforced strictly, allowing the network to adapt geometry implicitly through the data while preventing unphysical configurations (e.g., distal dilation or wall thickening).

Effective wall elasticity and pressure–area coupling. The elastic response of the arterial wall is encoded through the pressure–area coupling coefficient

$$\beta(x) = \sqrt{\pi} E(x) h_0(x), \quad (6.69)$$

where $E(x)$ denotes an *effective* Young’s modulus of the wall. It should be emphasized that $E(x)$ should not be interpreted as a local material constant in the classical sense: arterial walls are anisotropic, viscoelastic, layered structures, and reported moduli vary by nearly an order of magnitude depending on vessel type, frequency, and experimental protocol [107, 108].

To obtain physiologically reasonable spatial trends, the widely used empirical relationship proposed by Stergiopoulos et al.[19] is adopted, which fits effective elastic




Fig. 6.8 Stergiopoulos-style exponential fit of the effective Young's modulus as a function of vessel radius, used here to construct a physiologically plausible stiffness trend (cf. the relation above).

moduli as a function of vessel radius (Fig. 6.8):

$$E(r_0) = k_1 \exp(k_2 r_0) + k_3, \quad (6.70)$$

with constants

$$(k_1, k_2, k_3) = (1.89 \times 10^7, -18.34, 3.53 \times 10^6). \quad (6.71)$$

This formulation captures a key physiological fact: large elastic arteries (e.g., aorta) are substantially more compliant than distal muscular arteries, whose effective stiffness increases as radius decreases. When combined with the wall thickness profile $h_0(x)$, this relation yields a spatially varying $\beta(x)$ that governs local wave speed and pressure–area sensitivity, consistent with classical 1D hemodynamic modeling [26].

Role within the PINN framework. Within the PINN, the triplet $\{A_0(x), h_0(x), E(x)\}$ does not act as ground truth anatomy, but as a structured physiological prior that regularizes the PDE residuals and stabilizes the learning of pressure and flow fields.

By enforcing realistic monotonic trends in geometry and stiffness, these priors reduce the degeneracy of the inverse ECG–PPG–ABP problem, prevent non-physical compensations between pressure and flow, and promote internally consistent solutions across the spatial domain.

Crucially, the soft nature of these constraints allows subject-specific deviations to emerge naturally from the data, while preserving the global hemodynamic structure required for meaningful wave propagation modeling.

6.8.9 Differences with respect to the first-generation architecture

While the overall architectural backbone of the PINN remains intentionally lightweight and structurally similar to the first generation formulation, several key design choices distinguish the present second-generation model and reflect a shift toward a more explicit physiological interpretation.

Input representation. In the first-generation PINN, the temporal inputs consisted of a heuristic inflow proxy $Q_{\text{in}}(t)$ and the PPG signal. In the revised architecture, this proxy is replaced by the raw ECG signal, paired with PPG. This change reflects a conceptual reorganization: rather than prescribing an explicit surrogate for aortic inflow, the network is now encouraged to infer hemodynamically relevant timing and excitation information directly from ECG morphology, while leaving the absolute scaling of flow to learnable calibration parameters. As a result, the inflow is no longer treated as a partially known quantity but as a latent boundary condition constrained by physics.

Learnable parameterization. The first-generation model inferred a limited subset of physiological parameters, primarily ρ , β , A_0 , and K_r , which acted as soft regulators of the PDE residuals. In contrast, the second-generation PINN introduces an explicit window-wise parameter vector that also includes Windkessel elements (R_1, R_2, C) , pressure offsets (p_0, p_{ext}) , a linear PPG–pressure gain α , and inflow calibration terms (Q_s, Q_{off}) . This expanded parameter layer unifies geometric, mechanical, boundary, and calibration quantities within a single differentiable module, allowing them to be learned jointly with the pressure and flow fields rather than treated as auxiliary constants.

Distal pressure calibration. In the first-generation formulation, the distal pressure–PPG relationship was implicitly encoded through logarithmic or heuristic constraints motivated by optical considerations. In the revised model, this relation is made explicit and linear,

$$p(L, t) = p_0 + \alpha \text{PPG}(t),$$

consistent with the linearized Beer–Lambert derivation presented in the second-generation framework and with the fact that clinical PPG signals are already pre-processed by acquisition hardware. This choice removes unnecessary nonlinearities while preserving physiological interpretability.

Spatial sampling and differentiation. While the first-generation PINN relied on sorted spatial samples to facilitate numerical differentiation, the second-generation architecture abandons spatial ordering entirely. Collocation points are sampled randomly along x , and spatial derivatives are computed exclusively via automatic differentiation. This modification decouples the physics enforcement from any implicit finite-difference structure and aligns the implementation more closely with the continuous PINN paradigm.

Conceptual shift. Overall, the first-generation PINN can be interpreted as a proof-of-feasibility model, demonstrating that a physics-guided network can converge stably when combining non-invasive signals with hemodynamic constraints. The second-generation architecture moves beyond this stage by restructuring the network around interpretable latent parameters, reducing heuristic assumptions (notably on inflow), and embedding optical, mechanical, and boundary modeling more explicitly into the learning process.

6.8.10 Physics-Informed Loss Function (Second Generation)

Training is formulated as the minimization of a composite objective that blends (i) waveform supervision at the virtual observation site, (ii) a distal PPG–pressure calibration consistent with the linearized optical model, (iii) physics residuals of the 1D hemodynamic equations, (iv) a Windkessel terminal constraint, and (v) soft physiological priors that prevent drift toward non-credible parameter regimes. For each ECG/PPG window of duration 8 s ($T=1000$ samples at 125 Hz), a set of spatial

collocation points $x_n \in [0, 1]$ is sampled uniformly at random, always including the boundaries $x=0$ and $x=1$ as well as the reference point $x=0.5$.

Global objective

Let θ denote the trainable network weights (including the parameter layer producing the window-wise physiological parameters). The total loss is

$$\begin{aligned} \mathcal{L}_{\text{total}} = & \lambda_{\text{ABP}} \mathcal{L}_{\text{ABP}} + \lambda_{\text{PPG}} \mathcal{L}_{\text{PPG}} + \lambda_{\text{PDE}} \mathcal{L}_{\text{PDE}} + \lambda_{\text{WK}} \mathcal{L}_{\text{WK}} + \lambda_Q \mathcal{L}_Q \\ & + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}} + \lambda_\rho \mathcal{L}_\rho + \lambda_{K_r} \mathcal{L}_{K_r} + \lambda_{p_0} \mathcal{L}_{p_0} + \lambda_R \mathcal{L}_R, \end{aligned} \quad (6.72)$$

where each $\lambda_\bullet > 0$ controls the strength of a complementary constraint.

Data-fidelity terms

ABP supervision at the virtual observation point ($x=0.5$). Invasive ABP is available only at a single catheter site; this is represented by supervising the predicted pressure at a fixed reference coordinate $x_{\text{ref}}=0.5$:

$$\mathcal{L}_{\text{ABP}} = \frac{1}{T} \sum_{k=1}^T \left(p(0.5, t_k) - \text{ABP}(t_k) \right)^2 = \left\langle \left(p(0.5, t) - \text{ABP}(t) \right)^2 \right\rangle_t. \quad (6.73)$$

This term anchors the absolute pressure scale, while the spatial evolution away from x_{ref} is shaped primarily by physics and boundary constraints.

PPG–pressure calibration at the fingertip ($x=1$). At the distal boundary, the predicted pressure is softly tied to the measured PPG via the linear calibration derived from the linearized Beer–Lambert model and from the fact that clinical PPG waveforms are typically already preprocessed (DC removal and gain normalization):

$$\mathcal{L}_{\text{PPG}} = \frac{1}{T} \sum_{k=1}^T \left[p(1, t_k) - (p_0 + \alpha \text{PPG}(t_k)) \right]^2 = \left\langle \left(p(1, t) - (p_0 + \alpha \text{PPG}(t)) \right)^2 \right\rangle_t, \quad (6.74)$$

where (p_0, α) are learned *per window*. Compared to the first-generation logarithmic anchoring, this formulation removes an unnecessary nonlinearity and better matches

the empirical observation that $p(t)$ and $\text{PPG}(t)$ already exhibit similar beat-to-beat morphology after acquisition conditioning.

Physics residuals

Pressure–area coupling and area reconstruction. The PINN directly predicts $(p(x, t), Q(x, t))$. The cross-sectional area needed by the PDE system is obtained by inverting the pressure–area law around the reference configuration:

$$p(x, t) - P_{ext} = \frac{\beta(x)}{A_0(x)} \left(\sqrt{A(x, t)} - \sqrt{A_0(x)} \right), \quad (6.75)$$

$$\Rightarrow \sqrt{A(x, t)} = \sqrt{A_0(x)} + \frac{A_0(x)}{\beta(x)} (p(x, t) - P_{ext}), \quad (6.76)$$

which defines $A(x, t)$ deterministically given $p(x, t)$ and the learned fields (A_0, β, P_{ext}) .

PDE residuals (continuity and momentum). At each collocation sample (x_n, t_k) , the 1D conservation laws are enforced through residuals

$$r_1(x, t) = \frac{\partial A}{\partial t}(x, t) + \frac{\partial Q}{\partial x}(x, t), \quad (6.77)$$

$$r_2(x, t) = \frac{\partial Q}{\partial t}(x, t) + \frac{\partial}{\partial x} \left(\frac{Q^2}{A} \right)(x, t) + \frac{A}{\rho} \frac{\partial p}{\partial x}(x, t) + K_r(x) \frac{Q}{A}(x, t), \quad (6.78)$$

where the convective term can be expanded as $\frac{\partial}{\partial x} (Q^2/A) = \frac{2Q}{A} \frac{\partial Q}{\partial x} - \frac{Q^2}{A^2} \frac{\partial A}{\partial x}$. To reduce sensitivity to occasional large residual spikes in early epochs, a high clipping threshold is applied before squaring, $\tilde{r}_i = \text{clip}(r_i; -r_{\max}, r_{\max})$, and define

$$\mathcal{L}_{\text{PDE}} = \left\langle \tilde{r}_1(x, t)^2 + \tilde{r}_2(x, t)^2 \right\rangle_{x, t}. \quad (6.79)$$

All derivatives are computed through automatic differentiation of the network outputs with respect to continuous inputs (x, t) .

Windkessel terminal constraint at $x=1$. To model the downstream vascular load beyond the modeled segment, a three-element Windkessel boundary condition is

enforced at the outlet:

$$r_{\text{WK}}(t) = Q(1, t) \left(1 + \frac{R_1}{R_2}\right) + CR_1 \frac{dQ}{dt}(1, t) - \frac{1}{R_2} p(1, t) - C \frac{dp}{dt}(1, t), \quad (6.80)$$

and penalize its mean-square value over time,

$$\mathcal{L}_{\text{WK}} = \frac{1}{T} \sum_{k=1}^T r_{\text{WK}}(t_k)^2 = \langle r_{\text{WK}}(t)^2 \rangle_t. \quad (6.81)$$

Flow surrogate coupling and admissibility

Because the true inlet flow waveform is unobserved, the second-generation framework avoids prescribing an arbitrary $Q_{\text{in}}(t)$ (as done in the first generation) and instead introduces a minimal, learnable surrogate that only constrains timing and sign. Motivated by the continuity argument linking pulsatile flow to the PPG upstroke, the predicted flow is aligned to a normalized PPG derivative:

$$q_{\text{ppg}}(t) = Q_{\text{off}} + Q_s \frac{d\text{PPG}/dt}{\sigma(d\text{PPG}/dt) + 10^{-6}}, \quad \mathcal{L}_{Q, \text{ppg}} = \langle (Q(0, t) - q_{\text{ppg}}(t))^2 \rangle_t. \quad (6.82)$$

Here Q_s (dimensionless) and Q_{off} ($\text{m}^3 \text{s}^{-1}$) are learned per window, and $\sigma(\cdot)$ denotes the standard deviation computed within the window.

To discourage physically implausible solutions, the flow sign and offset are further regularized:

$$\mathcal{L}_{Q, \text{neg}} = \langle \max(0, -Q(0, t)) \rangle_t, \quad \mathcal{L}_{Q, \text{off}} = Q_{\text{off}}^2, \quad (6.83)$$

and softly keep the inlet scale near unity in normalized units via a relative deviation penalty, $\mathcal{L}_{Q, \text{scale}} = \text{penalty}(Q_s; 1, \pm 50\%)$. The resulting flow-related contribution is

$$\mathcal{L}_Q = \mathcal{L}_{Q, \text{ppg}} + \mathcal{L}_{Q, \text{neg}} + \mathcal{L}_{Q, \text{off}} + \mathcal{L}_{Q, \text{scale}}. \quad (6.84)$$

Physiological priors and parameter regularization

Let $A_0^{\text{nom}}(x)$, $\beta^{\text{nom}}(x)$, and $K_r^{\text{nom}}(x)$ denote the nominal spatial templates introduced in Sec. 6.8.8. The learned fields are softly anchored toward these priors:

$$\mathcal{L}_{A_0} = \langle (A_0(x) - A_0^{\text{nom}}(x))^2 \rangle_x, \quad (6.85)$$

$$\mathcal{L}_\beta = \langle (\beta(x) - \beta^{\text{nom}}(x))^2 \rangle_x, \quad (6.86)$$

$$\mathcal{L}_{K_r} = \langle (K_r(x) - K_r^{\text{nom}}(x))^2 \rangle_x. \quad (6.87)$$

A tight prior on blood density and a weak prior on external pressure are used:

$$\mathcal{L}_\rho = (\rho - \rho_{\text{ref}})^2, \quad \rho_{\text{ref}} \approx 1050\text{--}1060 \text{ kg m}^{-3}, \quad \mathcal{L}_{\text{pext}} = P_{\text{ext}}^2, \quad (6.88)$$

and define the grouped regularizer

$$\mathcal{L}_{\text{reg}} = \mathcal{L}_{A_0} + \mathcal{L}_\beta + \mathcal{L}_{\text{pext}}. \quad (6.89)$$

Baseline pressure prior. A gentle calibration prior is applied to p_0 :

$$\mathcal{L}_{p_0} = (p_0 - p_{0,\text{ref}})^2, \quad p_{0,\text{ref}} \approx 80 \text{ mmHg}. \quad (6.90)$$

Windkessel parameter priors. When enabled ($\lambda_R > 0$), the proximal resistance is softly tied to the characteristic impedance and the load split is constrained:

$$\mathcal{L}_{R_1} = \langle (R_1 - \rho c_0/A_0)^2 \rangle, \quad c_0(x) = \sqrt{\frac{A_0}{\rho} \frac{\partial p}{\partial A} \Big|_{A_0}} = \sqrt{\frac{133.322 \beta(x) \sqrt{A_0(x)}}{2\rho}}, \quad (6.91)$$

$$\mathcal{L}_{\text{ratio}} = \langle \max(0, R_2/R_1 - 100) \rangle + \langle \max(0, 10 - R_2/R_1) \rangle,$$

and set $\mathcal{L}_R = \mathcal{L}_{R_1} + \mathcal{L}_{\text{ratio}}$.

Remarks on enforcement and differences with the first generation

The supervision term $\mathcal{L}_{\text{ABP}}$ anchors the solution to invasive ABP at a single virtual observation point ($x=0.5$), while $\mathcal{L}_{\text{PPG}}$ calibrates the distal pressure waveform to the

optical plethysmogram using a *linear* relation consistent with the revised PPG model. The hemodynamic consistency is imposed by $\mathcal{L}_{\text{PDE}}$ and $\mathcal{L}_{\text{WK}}$. Compared to the first-generation loss, the main structural changes are: (i) replacing the logarithmic PPG constraint with the linear calibration, (ii) removing the need for an explicit inflow proxy by introducing a learnable timing-only flow surrogate through $\mathcal{L}_Q$, and (iii) strengthening identifiability and stability via explicit priors on (p_0, ρ, K_r) and optional impedance-based regularization of (R_1, R_2) . Overall, the loss is designed to preserve subject adaptability while preventing convergence to non-physical regimes.

6.8.11 Training Strategy and Experimental Protocol

This section describes the training protocol adopted for the second-generation Physics-Informed Neural Network, including dataset construction, batching and collocation strategy, optimization details, and the interpretation of the resulting performance metrics.

Datasets and window extraction

To construct a reliable and sufficiently diverse training corpus, data from two complementary sources were combined: the MIMIC-III Waveform Database and the PulseDB VitalDB subset.

From MIMIC-III, recordings containing synchronous ECG (lead II), peripheral PPG (PLETH), and invasive arterial blood pressure (ABP) were selected, all sampled at their native frequency of 125 Hz. Signals were segmented into fixed 8 s windows (1000 samples), a duration chosen to include multiple cardiac cycles while preserving quasi-stationary waveform morphology.

Each window was subjected to a compact but rigorous quality-control (QC) pipeline adapted from prior work. Windows were discarded if they contained non-interpolable missing samples, saturation or flat-line segments, abrupt baseline drifts, or physiologically implausible amplitudes (ABP outside 15–300 mmHg). Beat-level consistency was also enforced: both PPG and ABP were required to exhibit at least four systolic peaks and four diastolic troughs without excessive plateauing or peak–valley mismatch. After QC, a total of 36 080 valid 8 s windows were retained from MIMIC-III.

To reduce bias from repeated patterns originating from the same patient, extraction was limited to a maximum of ten windows per recording. This heuristic is particularly relevant because MIMIC-III does not provide explicit subject identifiers, making strict subject-wise partitioning impossible. As a consequence, standard train/validation/test splits performed at the window level may inadvertently mix segments from the same subject, potentially inflating apparent generalization performance. This limitation is explicitly acknowledged and addressed in later cross-database evaluations.

To probe robustness beyond the training distribution, the VitalDB test subset provided by PulseDB was additionally employed for external evaluation. The same ECG/PPG QC pipeline was applied, and windows with severe ECG artifacts, flat lines, or inconsistent beat structure were excluded. This step removed approximately 60% of low-quality windows, yielding a cleaner evaluation set. In total, around 170 000 windows of ECG, PPG, and ABP (each 8 s long) were extracted across 1 216 unique subjects. Unlike MIMIC-III, VitalDB includes subject identifiers, enabling stricter subject-level assessment of generalization across patients and recording setups.

Batching and spatial collocation

Training batches are constructed by first sampling a set of ECG/PPG windows and then augmenting each window with a set of spatial collocation points $x_n \in [0, 1]$. These points are drawn uniformly at random, with mandatory inclusion of the boundary locations $x = 0$ (aortic root), $x = 1$ (fingertip), and the virtual observation point $x = 0.5$.

Each batch therefore consists of triads $(x_n, \text{ECG}, \text{PPG})$, with supervised ABP labels applied only at $x = 0.5$, the distal PPG–pressure calibration enforced at $x = 1$, and the physics-based residuals evaluated at all sampled collocation points. This construction ensures that conservation laws and boundary conditions constrain the solution globally along the vessel, while direct data supervision remains localized.

Optimization details

All models were implemented in TensorFlow/Keras and trained using the Adam optimizer with an initial learning rate of 10^{-3} and default exponential decay parameters ($\beta_1 = 0.9$, $\beta_2 = 0.999$). Mixed-precision arithmetic was enabled to reduce memory usage and accelerate training, which is particularly beneficial given the convolutional front-ends used for ECG and PPG feature extraction.

Batch size was defined in terms of collocation triads, with 64 triads per update, corresponding to multiple waveform windows depending on the number of spatial samples N_x . The full composite loss described in Sec. 6.7.2 was backpropagated through both the data-driven encoders and the physics-informed parameter layer. All spatial and temporal derivatives required by the PDE residuals and Windkessel constraint were obtained via automatic differentiation.

Gradient clipping was not applied. Training stability was instead achieved through careful loss balancing, residual clipping in the PDE terms, and the use of soft physiological priors.

Regularization and stopping criteria

Architectural regularization was provided by dropout (rate 0.1) applied consistently across convolutional and dense layers. No explicit weight decay was used. Instead, regularization was enforced primarily through the physiological priors embedded in the loss function, which constrain geometry, elasticity, viscosity, density, and Windkessel parameters to remain within plausible ranges.

Training was performed for up to 100 epochs with early stopping based on the validation ABP error at the virtual observation point $x = 0.5$. This criterion was chosen because it directly reflects the primary supervised objective while remaining sensitive to overfitting.

Training on MIMIC-III and interpretation of results

The first series of experiments focused exclusively on the MIMIC-III dataset. The available windows were randomly divided into training (80%), validation (10%), and test (10%) subsets at the window level. Although this split cannot guarantee

subject independence, the large size of MIMIC-III makes it a suitable benchmark for assessing intra-record consistency and overall feasibility of the proposed framework.

On the held-out MIMIC-III test set, the model achieved a mean absolute error of 5.77 ± 3.33 mmHg when comparing the predicted pressure waveform at $x = 0.5$ to the invasive ABP. When decomposed into systolic and diastolic components, the errors were 7.32 mmHg and 5.12 mmHg, respectively. These values indicate that the model can reconstruct arterial pressure dynamics from non-invasive ECG and PPG inputs with clinically reasonable accuracy.

It is important to emphasize that these results should not be interpreted as peak performance benchmarks. Previous purely data-driven models applied to the same database have reported lower errors under similar intra-database conditions. However, the objective of the present work is fundamentally different: rather than maximizing raw accuracy under potentially leakage-prone splits, the present work deliberately introduces physics-based structure and physiological priors that constrain the solution space.

This choice trades some apparent accuracy for interpretability, stability, and the potential for improved generalization. The results obtained on MIMIC-III therefore serve primarily as a proof-of-concept, demonstrating that a heavily constrained, physics-informed formulation can converge stably and achieve reasonable performance on large-scale data. They establish a necessary intermediate step toward more meaningful evaluations on subject-separated cohorts and external databases such as VitalDB, which are essential for assessing real-world robustness and clinical viability.

6.8.12 Cross-database Evaluation on VitalDB

To assess cross-database generalization, the PINN trained exclusively on MIMIC-III was first evaluated directly on the VitalDB subset of PulseDB, without any adaptation. On this external test set, performance degraded substantially. The waveform mean absolute error at the virtual mid-point observation reached 17.02 ± 10.01 mmHg, with corresponding systolic and diastolic errors of 19.19 mmHg and 13.95 mmHg, respectively.

This degradation is not unexpected. MIMIC-III and VitalDB differ markedly in acquisition hardware, sensor characteristics, patient population, and recording

conditions. These differences induce non-trivial shifts in the distributions of ECG, PPG, and ABP waveforms, which are known to severely impact generalization in cuffless blood pressure estimation. The observed performance drop therefore highlights the intrinsic limits of models trained in a database-specific manner and reinforces the need for explicit cross-database validation.

To investigate whether the proposed physics-informed formulation can adapt to a new physiological and instrumental domain, a transfer-learning experiment was performed on VitalDB. Subject-level separation was enforced: 20% of the available subjects were used for fine-tuning, with an internal split of 80% for training and 20% for validation, while the remaining 70% of subjects were held out exclusively for testing. Importantly, no layers of the network were frozen. The entire PINN, including the signal encoders, the physiological parameter inference block, and the pressure/flow prediction heads, was allowed to adapt to the new data distribution.

Prior to fine-tuning and evaluation, the same ECG/PPG quality-control pipeline described earlier was applied. In particular, a conservative ECG filter was used to remove windows with severe artifacts or inconsistent beat structure, eliminating approximately 60% of low-quality segments. This step was essential to avoid contaminating the learning process with corrupted physiological signals.

After transfer learning, performance on the held-out VitalDB test subjects improved markedly. The waveform MAE at the mid-point observation decreased to 13.10 ± 7.38 mmHg, with systolic and diastolic errors of 15.84 mmHg and 10.42 mmHg, respectively. While these values remain higher than those obtained in the intra-database MIMIC-III setting and exceed current clinical accuracy requirements, they demonstrate that the proposed PINN can meaningfully adapt to a new population when provided with clean signals and limited subject-specific data.

To contextualize these results, Table 6.1 reports calibration-free BP estimation performance from prior works that explicitly enforce subject separation. The obtained errors are in line with those reported in the original VitalDB study and related literature, underscoring the difficulty of calibration-free, cross-subject BP estimation in realistic settings.

Beyond numerical metrics, qualitative inspection of the predicted pressure and flow fields confirms that the model preserves physiologically coherent dynamics after transfer. Figures 6.9 and 6.10 illustrate representative examples of predicted arterial pressure and volumetric flow along the normalized arterial coordinate and over time,

Table 6.1 Calibration-free blood pressure estimation performance reported in prior works using datasets with reliable subject separation. Results from MIMIC-II and other datasets reported in the original VitalDB paper are omitted when subject separation cannot be guaranteed.

Citation	SBP (MAE $\pm$ SD)	DBP (MAE $\pm$ SD)
Aguirre et al.	14.39 $\pm$ 17.87	6.57 $\pm$ 8.43
Slapničar et al.	15.41 $\pm$ –	12.38 $\pm$ –
This work	15.84 $\pm$ 13.46	10.42 $\pm$ 7.75




Fig. 6.9 $Q(t)$ example prediction. (Remember that no ground truth was used for the Q)

highlighting the internal consistency enforced by the physics-based constraints even under distribution shift.

Overall, the VitalDB experiments underline both the severity of the cross-database generalization challenge and the adaptability of the proposed framework. While the results do not yet meet clinical standards, they support the central hypothesis of this work: embedding hemodynamic structure and physiological priors enables meaningful adaptation across populations and recording conditions, providing a stronger foundation for future calibration-free and subject-robust blood pressure estimation.

6.9 Results Discussion

The results obtained in this work reinforce a broader conclusion that has emerged consistently across the cuffless blood pressure literature: ECG and PPG signals alone do not provide sufficient information for accurate, calibration-free blood pressure estimation under strict subject-disjoint evaluation [23–25]. Even when embedded within




Fig. 6.10 Example of 3D plot ABP to show the prediction from $x=0$ and $x=1$ in 2 seconds of prediction (for clarity). As it can be seen, the waveform slowly shrinks in amplitude as x approaches the capillaries at $x=1$, following the real physiology

a Physics-Informed Neural Network, the intrinsic degeneracy of the ECG–PPG–ABP mapping persists. Physics-based constraints improve interpretability, numerical stability, and internal consistency, but they cannot fully resolve the ambiguity inherent in reconstructing central or peripheral pressure from distal optical and electrical measurements alone.

This limitation reflects a fundamental identifiability issue rather than a shortcoming of a specific model architecture. ECG encodes cardiac electrical activation, while PPG reflects local volumetric changes modulated by vascular compliance, peripheral resistance, and measurement conditions. Multiple physiological configurations may therefore produce highly similar ECG and PPG waveforms while corresponding to substantially different arterial pressure profiles. Physics-informed modeling constrains this solution space, but does not collapse it to a unique mapping in the absence of additional subject-specific information.

A natural and necessary extension is therefore to enrich the input space with readily available personal descriptors such as age, sex, height, and weight. Age is among the strongest determinants of arterial compliance: progressive stiffening of the aorta and large arteries increases pulse wave velocity, alters wave reflection timing, and reshapes systolic–diastolic relationships [109]. Sex further modulates vascular aging trajectories; in particular, post-menopausal women experience accelerated arterial stiffening due to the loss of estrogen-mediated vascular protection [110, 111]. Height influences pulse transit and arrival times by determining the effective arterial path length [112], while weight and body mass index affect peripheral resistance and arterial compliance [113]. These factors jointly determine the mechanical context in which ECG and PPG are generated and observed.

Taken together, these considerations indicate that ECG and PPG cannot be interpreted in isolation [114]. Incorporating low-dimensional anthropometric metadata within a physics-informed framework—where such variables explicitly modulate compliance, resistance, and wave speed—appears to be a necessary step toward reducing degeneracy and achieving meaningful personalization in wearable blood pressure monitoring [115].

From a quantitative standpoint, comparison with the two calibration-free baselines reported in [50] (Aguirre et al. and Slapničar et al.) places the proposed approach well within the current state of the art. Across these studies and the present work, systolic blood pressure errors under strict subject separation typically fall in the

range of approximately 9–16 mmHg, while diastolic errors span roughly 5–12 mmHg. Differences across studies can be attributed to variations in dataset composition, sensor characteristics, preprocessing pipelines, and evaluation protocols. Within this context, the performance achieved by the proposed PINN is consistent with what is currently attainable for calibration-free methods under realistic cross-subject conditions.

Beyond numerical accuracy, the qualitative results highlight a key strength of the proposed framework: its ability to generate physiologically coherent pressure and flow fields along the arterial axis. Figure 6.12 shows that both systolic and diastolic pressures decrease progressively from the aorta ($x = 0$) to the fingertip ($x = 1$), with increasing uncertainty toward distal locations. This behavior aligns with known hemodynamic principles, including wave damping and enhanced reflection in peripheral vessels.

The spatiotemporal pressure projection in Figure 6.10 reproduces a forward-traveling pulse wave with gradual amplitude attenuation toward the periphery, consistent with physiological wave propagation. Similarly, the predicted volumetric flow waveform in Figure 6.9, although not supervised by any ground-truth flow measurement, aligns with cardiac phasing and exhibits plausible amplitude and damping behavior along the vessel. This suggests that the physics constraints effectively regularize the latent flow dynamics without imposing an explicit inflow waveform.

At the waveform level, Figure 6.11 illustrates that systolic and diastolic phases are captured with good fidelity at the virtual observation point, with residual discrepancies mainly localized around fine morphological features such as the dicrotic notch. The three-dimensional visualization in Figure 6.10 further emphasizes that the model encodes both temporal periodicity and spatial gradients in a coherent and physiologically meaningful manner.

Overall, these results indicate that while the proposed PINN does not yet overcome the fundamental limitations of calibration-free blood pressure estimation from ECG and PPG alone, it succeeds in embedding cardiovascular structure into the learning process. The model produces internally consistent, physiologically plausible solutions and demonstrates adaptability across datasets, thereby providing a solid methodological foundation for future extensions incorporating personalized physiological information.




Fig. 6.11 In blue: example of the ABP prediction at $x=0.5$, where ABP is usually measured. In black the real ground truth of the ABP.




Fig. 6.12 Example of the SBP and DBP prediction and their respective uncertainties from aorta ($x=0$) and fingertip ($x=1$).




Fig. 6.13 Example of the ABP prediction from aorta ($x=0$) and fingertip ($x=1$). In red the real ground truth of the ABP.

6.10 Conclusion

This chapter investigated the feasibility of reconstructing arterial blood pressure waveforms from non-invasive ECG and PPG signals by explicitly embedding cardiovascular physiology into a Physics-Informed Neural Network framework. Rather than aiming at maximal predictive accuracy alone, the central objective was to assess whether a constrained, hemodynamics-driven model could learn pressure and flow fields that remain consistent with known arterial mechanics, despite the intrinsic limitations of the available measurements.

Starting from the recognition that ECG–PPG–ABP mappings are fundamentally degenerate, the proposed Heart-to-Finger PINN was designed to reduce this ambiguity by introducing explicit priors on vessel geometry, wall mechanics, viscous losses, and distal vascular loading. In this sense, the contribution of this work is not limited to waveform prediction at a single observation point, but extends to the construction of a coherent spatiotemporal representation of arterial pressure and flow along a simplified arterial pathway. The resulting model produces pressure, flow, and auxiliary physiological quantities that evolve consistently in space and time according to the governing equations, even though only sparse supervision is available.

The experiments on MIMIC-III confirm that, within a single database, the framework can achieve reasonable reconstruction accuracy while maintaining stable and interpretable internal dynamics. However, the results on VitalDB highlight a critical and unavoidable aspect of cuffless blood pressure estimation: models trained on one population and acquisition setup do not transfer seamlessly to another. The observed

degradation under direct cross-database evaluation is therefore not a failure of the method, but rather an explicit manifestation of the physiological and instrumental variability that any practical system must confront. The subsequent improvement obtained through subject-level fine-tuning demonstrates that the proposed PINN can adapt to new distributions, provided that sufficient representative data are available.

Importantly, this chapter should not be interpreted as claiming calibration-free clinical readiness. On the contrary, the results reinforce a broader conclusion emerging from recent literature: ECG and PPG alone, even when combined with physics-informed modeling, are insufficient to uniquely determine arterial blood pressure across heterogeneous populations. What physics-informed learning offers instead is a principled way to structure this uncertainty, preventing non-physiological solutions and enabling the estimation of latent hemodynamic variables in a consistent manner.

Within the broader context of this thesis, this work serves as a foundational step toward more expressive and personalized cardiovascular models. By demonstrating that physiological priors can be embedded into neural architectures without sacrificing flexibility, it opens the door to extensions that incorporate richer and more structured sources of information.

A particularly relevant direction concerns the definition of the inflow $Q(t)$. In the present framework, Q is introduced through a heuristic surrogate derived from the PPG upstroke and calibrated by learnable scaling parameters. While sufficient for feasibility and stability, this choice remains inherently arbitrary and does not fully exploit the substantial body of physiological knowledge relating cardiac output and inflow characteristics to subject-specific factors. Future developments should therefore aim at replacing the current surrogate with more principled formulations of $Q(t)$, grounded in established relationships between cardiac output, heart rate, stroke volume, and anthropometric variables such as age, sex, height, and body composition, as reported in other correlated works in the literature.

In this perspective, subject-level metadata could be used to constrain or parameterize the inflow model itself, reducing arbitrariness and improving physiological consistency. More generally, integrating personalized inflow modeling together with additional contextual variables affecting short-term blood pressure variability, or with supplementary sensing modalities providing partial observations of flow or pressure at multiple sites, represents a key step toward reducing the intrinsic degeneracy of ECG–PPG–ABP mappings.

In summary, this chapter establishes a methodological framework rather than a final solution. It shows that enforcing hemodynamic structure within a learning-based model is both feasible and beneficial, and that such an approach provides a meaningful bridge between purely data-driven prediction and mechanistic cardiovascular modeling. The insights gained here motivate the subsequent directions explored in this thesis and lay the groundwork for future developments toward more personalized, interpretable, and physiologically grounded blood pressure estimation systems.

Chapter 7

Physiological Stress Modeling and Latent-State Learning

7.1 From Blood Pressure to Stress: A Common Modeling Challenge

The previous chapters have shown that estimating arterial blood pressure from ECG and PPG signals is not merely a regression problem. Although deep neural networks can achieve low prediction errors under controlled settings, their performance degrades when the physiological variability of subjects is not explicitly modeled. This limitation ultimately led to the introduction of Physics-Informed Neural Networks, where the missing element between input signals and output pressure values was identified as the underlying hemodynamic state governed by conservation laws and vascular mechanics.

A remarkably similar situation arises in the context of physiological stress detection.

At first glance, stress prediction appears to be a standard classification task: given physiological signals such as electrocardiogram (ECG), electrodermal activity (EDA), and respiration (RR), predict whether a subject is stressed or not. However, as in the blood pressure problem, a deeper analysis reveals that the mapping between input signals and output labels is not direct. There exists an intermediate latent process, the autonomic nervous system (ANS) state and the cognitive appraisal of

stressors, that mediates the relationship between external stimuli and measurable physiological responses.

In other words, stress is not directly encoded in the signals. Rather, physiological signals reflect the activity of sympathetic and parasympathetic pathways, which themselves are influenced by cognitive evaluation, emotional context, prior experiences, and individual sensitivity thresholds. The observable signals are therefore indirect projections of an unobserved dynamical state.

This chapter investigates stress detection from physiological signals under this perspective. Similarly to the blood pressure case, purely data-driven models are first analyzed, and their limitations are critically examined. The results suggest the presence of a structural bottleneck: performance saturates around a certain threshold regardless of architectural complexity. This phenomenon motivates a transition toward latent-state modeling approaches, where the autonomic dynamics are explicitly represented and constrained.

7.2 Stress as a Physiological and Cognitive Process

Stress can be defined as the physiological and psychological response to internal or external demands that challenge an individual's adaptive capacity. Acute stress responses are often beneficial, enabling rapid mobilization of energy and enhanced alertness. However, chronic exposure to stress is associated with cardiovascular disease, cognitive impairment, and psychiatric disorders.

Traditional assessment methods rely on self-reported questionnaires and subjective ratings. While useful, these approaches suffer from recall bias, low temporal resolution, and inter-individual variability in perception. Physiological signals, on the other hand, provide continuous and objective measurements reflecting autonomic nervous system activity.

Among these signals:

- **Electrodermal Activity (EDA)** reflects sympathetic activation through changes in skin conductance.
- **Electrocardiogram (ECG)** encodes cardiac electrical activity and indirectly heart rate variability.

- **Respiration (RR)** captures breathing dynamics, which are modulated by both sympathetic and parasympathetic pathways.

These modalities are closely linked to autonomic regulation and have therefore been widely studied as biomarkers of stress. However, physiological activation does not necessarily coincide with subjective perception. Different individuals may exhibit similar physiological responses while reporting different stress levels, or conversely, report stress without strong autonomic changes.

This discrepancy introduces a fundamental modeling challenge: the stress level is partially subjective, while the inputs are objective physiological signals.

7.3 The StressID Dataset

To investigate this problem, the StressID dataset, already introduced in Chapter 3.1.5 is adopted as experimental framework.

In this chapter, the binary labeling scheme is initially adopted in order to maintain comparability with baseline approaches and simplify the interpretation of classification performance.

7.4 Baseline Performance and Initial Objective

In the original StressID study, handcrafted feature extraction was performed prior to classification. Specifically, 35 ECG features, 23 EDA features, and 40 respiration features were computed for each task window. Classical machine learning algorithms such as Random Forest and Support Vector Machines were then trained on these features.

Using only physiological signals, the best reported performance achieved an F1-score of approximately 0.73 and an accuracy around 0.72. Multimodal fusion including video and audio did not substantially improve performance, suggesting that physiological information already captured a significant portion of the discriminative signal.

The primary objective of this work is to investigate whether comparable or improved performance can be achieved directly from raw physiological signals, without relying on handcrafted feature engineering. In particular, convolutional neural networks are employed to learn temporal representations directly from windowed signal segments.

As will be shown in the following sections, the results reveal a consistent performance plateau around 70–73% accuracy, accompanied by strong overfitting tendencies. This behavior strongly suggests that the limitation does not stem from insufficient network capacity, but rather from intrinsic ambiguities in the labeling process.

7.5 A Structural Analogy with Blood Pressure Estimation

The stress detection problem mirrors the modeling issues encountered in cuffless blood pressure estimation:

- In blood pressure estimation, ECG and PPG signals do not directly encode pressure; they reflect an underlying hemodynamic state governed by conservation laws.
- In stress detection, physiological signals do not directly encode stress perception; they reflect an underlying autonomic and cognitive state.

In both cases, purely data-driven approaches attempt to approximate a mapping that implicitly depends on a hidden dynamical process. Without modeling this intermediate state, the network is forced to rely on statistical correlations that may be dataset-specific or partially spurious.

This analogy motivates the development of latent-state models for stress, in which sympathetic and parasympathetic dynamics are explicitly represented and constrained. The remainder of this chapter progressively builds toward such a formulation, starting from data-driven baselines and moving toward state-space and physics-informed approaches.

7.6 Database Statistics and Experimental Protocol

7.6.1 Windowing Strategy and Data Representation

All physiological signals (ECG, EDA, and respiration) are sampled at 125 Hz. In order to construct supervised learning examples, continuous recordings are segmented into fixed-length windows of 60 seconds. Each window therefore contains:

$$N = 60 \times 125 = 7500 \text{ samples per signal.}$$

The choice of a 60-second window reflects a compromise between temporal resolution and physiological stability. Shorter windows (e.g., 8–10 seconds) exhibited higher variance and reduced classification stability, whereas longer windows would significantly reduce the number of available training examples. Empirically, 60 seconds provided the most stable performance across cross-validation folds.

Each window is z-normalized *subject-wise*, ensuring that inter-subject amplitude differences do not dominate the learning process. This normalization step removes baseline offsets while preserving within-window temporal dynamics.

7.6.2 Subject-Wise Cross-Validation

To ensure that results reflect inter-subject generalization rather than intra-subject memorization, a subject-wise split is adopted.

A 4-fold cross-validation scheme is used:

- 48 subjects are used for training,
- 15 subjects are used for testing.

No window from a test subject is ever seen during training. This protocol is intentionally stringent, as it reflects the realistic scenario of deploying a stress detection model on unseen individuals.

7.6.3 Class Distribution and Label Ambiguity

Binary labels are derived from self-reported stress values using the threshold defined in the dataset documentation. However, inspection of the label distribution reveals substantial inter-subject variability:

- Some participants report stress in nearly all tasks,
- Others report no stress even during cognitively demanding tasks,
- Several tasks exhibit inconsistent labeling across participants.

This heterogeneity introduces label noise, which directly impacts the achievable classification performance. Unlike blood pressure values, which are physically measurable quantities, stress labels are subjective constructs influenced by perception, expectation, and personality traits.

This distinction is central to interpreting the experimental results presented below.

7.7 Single-Signal Modeling

7.7.1 Network Architecture

For single-signal experiments, a lightweight one-dimensional convolutional neural network is adopted. Each input window has shape:

$$(7500, 1),$$

corresponding to 60 seconds of a single physiological modality.

The architecture consists of:

- 3 stacked Conv1D blocks (kernel size 5, ReLU activation),
- Batch normalization after each convolution,
- Global average pooling,

- A fully connected layer (32 units),
- A final sigmoid output layer.

The total number of parameters ranges between approximately 4×10^3 and 1.5×10^6 , depending on filter width and depth explored during hyperparameter tuning.

Binary cross-entropy is used as loss function. Early stopping with patience of 10 epochs prevents severe overfitting.

7.7.2 Results

Table 7.1 reports the mean and standard deviation across the 4 subject-wise folds.

Table 7.1 Performance comparison of single-signal models (mean $\pm$ std over 4 folds).

Signal	Loss	Accuracy	AUC	F1-Weighted
EDA	0.70 ± 0.15	0.72 ± 0.016	0.75 ± 0.020	0.73 ± 0.016
RR	0.68 ± 0.033	0.71 ± 0.034	0.71 ± 0.040	0.71 ± 0.038
ECG	0.98 ± 0.22	0.63 ± 0.046	0.63 ± 0.066	0.61 ± 0.056

Several observations emerge:

- EDA achieves performance comparable to the handcrafted-feature baseline of the original study.
- Respiration (RR) yields statistically similar but worse performance.
- ECG alone performs substantially worse, indicating that raw ECG morphology is insufficient to robustly infer stress state under subject-wise generalization.

The relatively poor ECG performance is consistent with physiological expectations. While stress modulates heart rate variability, the ECG waveform morphology itself is highly subject-specific. A neural network may therefore learn subject identity cues rather than stress-related modulation, leading to poor cross-subject transfer.

Notably, increasing network capacity beyond moderate size does not improve performance; instead, overfitting becomes more pronounced. This suggests that model capacity is not the primary bottleneck.

7.8 Multimodal Physiological Modeling

7.8.1 Dual-Signal Fusion

To evaluate whether combining physiological modalities improves performance, dual-branch architectures are implemented. Each signal is processed by an independent convolutional encoder, and the resulting feature embeddings are concatenated before classification through a shared multilayer perceptron (MLP) of size 32.

Table 7.2 summarizes the results.

Table 7.2 Performance comparison of dual-signal models (mean $\pm$ std over 4 folds).

Signals	Loss	Accuracy	AUC	F1-Weighted
EDA + RR	0.642 ± 0.017	0.711 ± 0.007	0.72 ± 0.03	0.710 ± 0.005
EDA + ECG	0.647 ± 0.029	0.700 ± 0.010	0.701 ± 0.018	0.702 ± 0.011
RR + ECG	0.663 ± 0.028	0.653 ± 0.029	0.66 ± 0.05	0.638 ± 0.027

Contrary to expectations, combining modalities does not significantly improve performance compared to EDA alone. The marginal gains remain within the variance across folds.

7.8.2 Three-Signal Fusion

Using all three physiological signals together yields:

- Loss: 0.6273 ± 0.0221
- Accuracy: 0.7000 ± 0.0203
- AUC: 0.7217 ± 0.0394
- F1-weighted: 0.7008 ± 0.0218

The absence of significant improvement suggests that the limiting factor is not insufficient modality coverage, but rather the intrinsic structure of the labeling process.

Across all architectures and hyperparameter configurations explored, performance consistently saturates around 72–73% accuracy. Increasing depth or parameter count systematically increases overfitting without improving validation metrics.

7.9 Label Ambiguity and the Hard Performance Limit

The experimental results presented in the previous sections reveal a striking and consistent phenomenon: across architectures, modalities, and fusion strategies, performance stabilizes around 70–73% accuracy, without significant improvement from increasing model capacity or combining additional physiological signals.

This behavior suggests the presence of a *hard performance limit* that is not attributable to architectural underfitting. Instead, it points toward a deeper structural issue in the learning problem itself.

7.9.1 Stress as a Latent Physiological State

Unlike arterial blood pressure, which is a measurable physical quantity governed by conservation laws and mechanical constraints, stress is not directly observable. What is available to the model is a mapping:

$$\text{Physiological signals} \longrightarrow \text{Subjective label.}$$

However, stress should more rigorously be interpreted as a latent internal state $S(t)$, which influences autonomic nervous system activity and, consequently, measurable signals such as EDA, ECG, and respiration. At the same time, the subjective self-assessment label Y is a separate observation channel:

$$S(t) \rightarrow \begin{cases} \text{Physiological modulation} \\ \text{Subjective perception} \end{cases}$$

The crucial point is that these two channels are not guaranteed to be aligned.

Physiological activation does not necessarily imply conscious stress perception, and conversely, perceived stress may not always manifest as strong autonomic modu-

lation. Individual thresholds, personality traits, coping mechanisms, and reporting biases introduce variability in the subjective channel.

In statistical terms, this creates a noisy supervision problem:

$$P(Y \mid \text{signals}) \text{ is intrinsically uncertain.}$$

7.9.2 Empirical Evidence from Synthetic Labels

To probe this hypothesis, alternative task-based (synthetic) labels were introduced. In this configuration, tasks objectively designed to induce stress were assigned label 1, and relaxation-oriented tasks were assigned label 0.

Under this labeling scheme, preliminary experiments yielded:

- Accuracy $\approx 94\%$
- AUC $\approx 98\%$
- F1-weighted $\approx 94\%$

Although not obtained through full cross-validation, these values demonstrate that the physiological signals contain sufficient information to discriminate between strongly stress-inducing and clearly relaxing tasks.

The dramatic performance increase under synthetic labeling strongly supports the hypothesis that the limitation observed under subjective labeling arises primarily from label inconsistency rather than insufficient model expressiveness.

7.9.3 Dimensionality Reduction Analysis

To further investigate the structure of the data independently of the classifier, dimensionality reduction techniques were applied to handcrafted summary features extracted from each window.

For each signal (ECG, EDA, RR), the following features were computed per window:

- Mean

- Standard deviation
- Maximum
- Minimum
- 25th percentile
- 75th percentile
- Mean of first differences
- Standard deviation of first differences

This yields 8 features per modality, for a total of 24 features per window. The concatenated feature vectors were projected using PCA and t-SNE.

Subjective Labels

Figures 7.1 and 7.3 show the projections under subjective labels.

The resulting clusters exhibit substantial overlap between stress and non-stress classes. No clear decision boundary appears in low-dimensional space, indicating high entropy in the label assignment relative to physiological structure.

Synthetic Labels

In contrast, Figures 7.2 and 7.4 show projections under synthetic task-based labeling.

Here, two dominant clusters emerge with significantly improved separability. The alignment between physiological structure and label assignment is markedly stronger.

7.9.4 Interpretation

The dimensionality reduction analysis reinforces the existence of a structural mismatch between physiological states and subjective labels.

The network does not appear to be limited by insufficient representational capacity. Instead, it is confronted with an inherently ambiguous supervision signal.




Fig. 7.1 t-SNE projection of physiological features with subjective stress labels.




Fig. 7.2 t-SNE projection with synthetic task-based labels.




Fig. 7.3 PCA projection with subjective labels.




Fig. 7.4 PCA projection with synthetic labels.

This observation parallels, in a different domain, the challenges encountered in cuffless blood pressure estimation: when the mapping between input signals and target values is not governed by consistent physical constraints, purely data-driven approaches risk learning unstable or biased associations.

In the stress detection problem, the missing element is not a physical conservation law, but rather a formal model linking autonomic dynamics, latent stress state, and subjective reporting.

This suggests that future progress may require:

- Modeling stress as a latent variable,
- Incorporating uncertainty-aware or probabilistic supervision,
- Using semi-supervised or contrastive learning,
- Or embedding physiological priors related to autonomic regulation.

In other words, the limitation is epistemic rather than architectural.

7.10 Cross-Attention Between Physiological Modalities

Given the limited improvement observed when combining multiple physiological signals through simple feature concatenation, an alternative fusion strategy was explored based on cross-attention mechanisms.

The intuition behind cross-attention is that different physiological modalities may encode complementary aspects of the same latent stress state. Instead of independently extracting features and merging them through a multilayer perceptron (MLP), cross-attention allows one signal to dynamically reweight or attend to informative regions of another signal.

7.10.1 Cross-Attention Mechanism

In the dual-signal configuration, each modality (e.g., EDA and respiration) is first processed through an identical convolutional encoder, producing intermediate feature

maps:

$$\mathbf{F}_1 \in \mathbb{R}^{T \times d}, \quad \mathbf{F}_2 \in \mathbb{R}^{T \times d}.$$

A cross-attention block is then introduced such that:

$$\mathbf{Q} = \mathbf{F}_1 W_Q, \quad \mathbf{K} = \mathbf{F}_2 W_K, \quad \mathbf{V} = \mathbf{F}_2 W_V,$$

where W_Q, W_K, W_V are learnable projection matrices.

The attention output is computed as:

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d}}\right) \mathbf{V}.$$

This allows the first modality to selectively focus on temporally aligned or correlated patterns in the second modality. The attended representation is then concatenated with the original features and passed to the final classifier.

Architecturally, all other components (encoder depth, kernel size, dropout, optimizer, early stopping criteria) were kept identical to the non-attention baseline in order to isolate the effect of the cross-attention module.

7.10.2 Experimental Results

Tables 7.3 and 7.4 report the results for EDA+RR and EDA+ECG configurations, respectively, with and without cross-attention.

Table 7.3 EDA + RR with and without cross-attention (mean $\pm$ std).

Configuration	Loss	Accuracy	AUC	F1-Weighted
EDA+RR+CA	0.656 ± 0.010	0.702 ± 0.019	0.718 ± 0.019	0.703 ± 0.018
EDA+RR	0.642 ± 0.017	0.711 ± 0.007	0.72 ± 0.03	0.710 ± 0.005

Table 7.4 EDA + ECG with and without cross-attention (mean $\pm$ std).

Configuration	Loss	Accuracy	AUC	F1-Weighted
EDA+ECG+CA	0.663 ± 0.017	0.690 ± 0.029	0.69 ± 0.05	0.69 ± 0.03
EDA+ECG	0.647 ± 0.029	0.700 ± 0.010	0.701 ± 0.018	0.702 ± 0.011

Across configurations, cross-attention does not improve performance. In both cases, accuracy and F1-weighted scores are slightly lower compared to the simpler fusion strategy.

7.10.3 Interpretation

The absence of performance gains suggests that the limitation is not due to insufficient interaction modeling between modalities. If a strong and consistent multimodal coupling were present in the data, cross-attention should have provided measurable improvements by aligning complementary stress-related patterns across signals.

Instead, the results reinforce the hypothesis introduced in Section 7.9: the dominant source of uncertainty lies in the supervision signal rather than in feature interaction capacity.

When the target label is noisy or weakly aligned with the physiological structure of the data, increasing model complexity tends to amplify overfitting without improving generalization. Indeed, the cross-attention module increases the number of learnable parameters and the flexibility of the representation, yet the validation performance remains bounded by the same empirical ceiling observed in simpler architectures.

This behavior further supports the interpretation of stress detection, under subjective labeling, as an inherently ambiguous learning problem. The model is not failing to integrate modalities; rather, it is constrained by the inconsistency between latent physiological states and reported stress levels.

In this sense, the negative result obtained with cross-attention is informative: it indicates that architectural sophistication alone cannot resolve structural label noise. Future improvements are therefore more likely to emerge from better modeling of stress as a latent process or from rethinking the supervision paradigm, rather than from increasingly complex fusion mechanisms.

7.11 From Stress Classification to Latent-State Modeling

The experiments presented so far reveal a recurring pattern that echoes the challenges encountered in cuffless blood pressure estimation.

In both problems, the mapping from physiological signals to the target variable is not direct. There exists an intermediate, unobserved process that governs the transformation between input and output:

- In blood pressure estimation, ECG and PPG do not directly encode arterial pressure; rather, pressure emerges from cardiovascular dynamics governed by hemodynamics.
- In stress detection, physiological signals do not directly encode the subjective stress label; instead, they reflect autonomic nervous system activity, which itself only partially correlates with perceived stress.

In the blood pressure case, the missing link was addressed by explicitly modeling wave propagation, vessel elasticity, and boundary conditions through a physics-informed neural network. In the stress case, the missing component appears to be the latent autonomic state that mediates between physiological measurements and psychological perception.

The empirical ceiling observed around 70–73% accuracy across multiple architectures, including convolutional networks and cross-attention models, suggests that the current formulation of the problem may be incomplete. When synthetic, task-based labels are used, performance rises dramatically, indicating that the physiological signals themselves contain structured and separable information. However, when subjective labels are employed, separability degrades and performance saturates.

This observation leads to a central hypothesis:

Stress, as labeled in the dataset, is not a directly observable physiological state, but rather a subjective projection of an underlying autonomic and cognitive process.

From a modeling perspective, the task should therefore not be framed as a direct mapping:

$$(EDA, ECG, RR) \longrightarrow \text{Stress label},$$

but rather as a two-stage latent process:

$$(EDA, ECG, RR) \longrightarrow \text{Latent autonomic state} \longrightarrow \text{Subjective stress perception}.$$

In the absence of explicit modeling of this intermediate state, purely discriminative classifiers are forced to approximate a noisy projection of a hidden process. As the experiments suggest, increasing architectural complexity alone does not overcome this structural limitation.

This realization motivates a shift in perspective similar to the one adopted for blood pressure modeling: instead of treating stress detection as a black-box classification problem, this chapter explores whether it can be reformulated as a constrained dynamical system, in which autonomic regulation is modeled explicitly and neural networks are used to estimate hidden states under physiological priors.

The following section introduces a physics-informed formulation of stress modeling, grounded in simplified autonomic dynamics. The goal is not to claim a complete mechanistic theory of stress, but to test whether embedding regulatory structure into the learning process can reduce degeneracy and improve generalization.

7.12 Toward a Physics-Informed Model of Stress

The previous section argued that stress classification, when treated as a purely discriminative mapping from physiological signals to subjective labels, encounters an intrinsic performance ceiling. This limitation suggests that stress perception is not directly encoded in the observable signals, but rather emerges from an intermediate latent process.

In physiological terms, this intermediate process corresponds to the dynamics of the autonomic nervous system (ANS), which regulates cardiovascular, electrodermal, and respiratory activity through interacting sympathetic and parasympathetic branches.

Rather than predicting stress labels directly, the proposed formulation models stress as a latent dynamical state governed by simplified autonomic regulation equations.

7.12.1 Autonomic Nervous System as a Latent Dynamical System

In analogy with the hemodynamic formulation introduced in the previous chapters, stress perception is here interpreted as the macroscopic manifestation of an underlying physiological process evolving in time. The following formulation does not aim to provide an experimentally validated autonomic model, but rather to introduce a theoretically consistent dynamical hypothesis that can serve as a structural prior in a physics-informed framework.

Stress-related physiological modulation is postulated to be summarized by two latent variables:

$$\mathbf{x}(t) = \begin{bmatrix} S(t) \\ P(t) \end{bmatrix}, \quad (7.1)$$

where $S(t)$ represents an effective sympathetic activation level and $P(t)$ represents an effective parasympathetic (vagal) activation level. These quantities should not be interpreted as directly measurable physiological signals; instead, they constitute abstract state variables summarizing distributed regulatory activity across central and peripheral pathways.

Energetic and circuital intuition The proposed formulation is conceptually inspired by an energetic interpretation of autonomic regulation. Sympathetic activation may be viewed as a process of energy mobilization, whereas parasympathetic activation promotes restoration and recovery. From a systems-theoretic perspective, this suggests an analogy with elementary electrical circuits, in which activation and recovery resemble charge and discharge processes governed by first-order dynamics.

In particular, the exponential relaxation terms $-\alpha_S S$ and $-\alpha_P P$ are formally analogous to the discharge of a capacitor through a resistive element in an RC circuit. External drives act as input currents that temporarily shift the system away from equilibrium, while intrinsic decay restores baseline tone once stimulation vanishes.

This analogy does not imply a literal electrical equivalence between neural regulation and passive components; rather, it provides a structural rationale for modeling autonomic regulation as a low-order dissipative dynamical system.

Minimal antagonistic dynamics As a first-order approximation, the interaction between the two branches is assumed to be described through a coupled linear dynamical system:

$$\dot{\mathbf{x}}(t) = \mathbf{A} \mathbf{x}(t) + \mathbf{B} \mathbf{u}(t), \quad (7.2)$$

with

$$\mathbf{A} = \begin{bmatrix} -\alpha_S & -\gamma_S \\ -\gamma_P & -\alpha_P \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} \beta_S & 0 \\ 0 & \beta_P \end{bmatrix}, \quad \mathbf{u}(t) = \begin{bmatrix} u(t) \\ v(t) \end{bmatrix}. \quad (7.3)$$

The coefficients $\alpha_S, \alpha_P > 0$ represent intrinsic relaxation rates toward baseline autonomic tone. The cross-coupling terms $\gamma_S, \gamma_P \geq 0$ encode reciprocal inhibition between sympathetic and parasympathetic activity, reflecting their well-established antagonistic behavior.

The inputs $u(t)$ and $v(t)$ are not directly measurable physiological signals, nor experimentally controlled stimuli in the present work. Rather, they represent abstract modulatory drives acting at the central regulatory level of the autonomic nervous system.

More specifically, $u(t)$ models a sympathetic-oriented central drive, which may conceptually summarize the combined effect of cognitive load, emotional appraisal, perceived threat, or task-related arousal. It should therefore be interpreted as a low-dimensional representation of central processes that increase sympathetic tone.

Conversely, $v(t)$ models a parasympathetic-oriented regulatory drive, associated with recovery mechanisms, safety perception, vagal activation, or respiration-mediated modulation. It represents the aggregate influence of processes promoting autonomic stabilization and relaxation.

In the absence of explicit experimental control over stressors or recovery protocols, both $u(t)$ and $v(t)$ remain latent forcing terms. Their effect is not observed

directly, but is indirectly reflected in the temporal evolution of the hidden autonomic states $S(t)$ and $P(t)$.

In the absence of external inputs, the system reduces to

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x}. \quad (7.4)$$

For physiological plausibility, asymptotic stability of the resting equilibrium $\mathbf{x} = 0$ is required. In the two-dimensional case, this corresponds to the conditions

$$\text{tr}(\mathbf{A}) < 0, \quad \det(\mathbf{A}) > 0, \quad (7.5)$$

which translate into

$$\alpha_S + \alpha_P > 0, \quad \alpha_S \alpha_P > \gamma_S \gamma_P. \quad (7.6)$$

These inequalities ensure that mutual inhibition does not dominate intrinsic decay, preventing unrealistic divergence or self-sustained oscillations in the absence of stimulation. The model therefore remains dynamically well-posed while preserving antagonistic interaction.

Interpretative role of the model It is important to emphasize that Eq. (7.2) is not introduced as a quantitatively accurate physiological simulator. Rather, it provides a structurally coherent representation of three qualitative principles:

(i) autonomic activation exhibits exponential recovery dynamics, (ii) sympathetic and parasympathetic branches interact antagonistically, (iii) external perturbations modulate autonomic tone through additive drives.

Within this perspective, $S(t)$ and $P(t)$ act as low-dimensional latent descriptors that compress complex neurophysiological processes into a mathematically tractable form.

More elaborate nonlinear interactions could be introduced by replacing the linear cross-terms with bounded monotonic functions, e.g.,

$$\frac{dS}{dt} = -\alpha_S S - \gamma_S \phi_P(P) + \beta_S u(t), \quad (7.7)$$

$$\frac{dP}{dt} = -\alpha_P P - \gamma_P \phi_S(S) + \beta_P v(t), \quad (7.8)$$

where ϕ_S and ϕ_P are saturating functions. However, for the theoretical development pursued here, the linear antagonistic formulation is sufficient to encode the desired dynamical structure while maintaining analytical transparency.

In this framework, stress is not directly equated with $S(t)$ or $P(t)$ individually, but emerges from their temporal imbalance.

One may formally define a stress-related observable as a projection

$$\sigma(t) = h(S(t), P(t)),$$

where $h(\cdot)$ represents an unknown, possibly nonlinear mapping that transforms autonomic imbalance into perceived stress.

The autonomic state therefore constitutes an intermediate latent process that mediates between observable physiological signals and subjective stress perception.

The introduction of this dynamical hypothesis allows stress prediction to be reformulated, at least conceptually, as a constrained latent-state estimation problem rather than as a purely discriminative classification task.

7.12.2 Physics-Informed Objective as a Constrained Latent-State Estimation Problem

The autonomic dynamical hypothesis introduced in Section 7.12.1 implies a fundamental shift in the learning objective. If observable physiological signals are indirect manifestations of a hidden autonomic trajectory, then stress prediction should not be framed as a direct mapping $\mathbf{z}(t) \rightarrow y$, but rather as a constrained inference problem:

$$\mathbf{z}(t) \implies \hat{\mathbf{x}}(t) \implies \hat{y}.$$

In this perspective, the neural network is interpreted as a differentiable estimator of the latent state $\hat{\mathbf{x}}(t) = [\hat{S}(t), \hat{P}(t)]^\top$, whose trajectory must simultaneously:

(i) reconstruct observable physiological signals, (ii) remain consistent with the assumed autonomic dynamics, (iii) retain predictive information about stress labels.

Observation consistency. Let $\mathbf{z}(t)$ denote the vector of measured physiological quantities (e.g., EDA, HR, RR). The formulation assumes the existence of an observation operator

$$\hat{\mathbf{z}}(t) = h(\hat{\mathbf{x}}(t)), \quad (7.9)$$

where $h(\cdot)$ is a learnable nonlinear mapping acting as a physiological decoder.

The first component of the objective enforces reconstruction consistency:

$$\mathcal{L}_{\text{rec}} = \frac{1}{T} \int_0^T \|\hat{\mathbf{z}}(t) - \mathbf{z}(t)\|^2 dt. \quad (7.10)$$

This term ensures that the latent trajectory encodes information genuinely explanatory of physiological measurements, preventing the network from collapsing into representations that merely optimize label prediction.

Dynamical regularization. The second component enforces consistency with the theoretical autonomic model. Given the assumed dynamics

$$\dot{\mathbf{x}}(t) = \mathbf{A} \mathbf{x}(t) + \mathbf{B} \mathbf{u}(t), \quad (7.11)$$

the dynamical residual is defined

$$\mathcal{R}_{\text{ANS}}(t) = \dot{\hat{\mathbf{x}}}(t) - \mathbf{A} \hat{\mathbf{x}}(t) - \mathbf{B} \hat{\mathbf{u}}(t), \quad (7.12)$$

and the corresponding penalty

$$\mathcal{L}_{\text{dyn}} = \frac{1}{T} \int_0^T \|\mathcal{R}_{\text{ANS}}(t)\|^2 dt. \quad (7.13)$$

This term does not assert that the autonomic model is exact; rather, it biases the learned latent trajectory toward a physiologically plausible dynamical structure. In absence of this constraint, $\hat{\mathbf{x}}(t)$ would be underdetermined, since infinitely many trajectories could reconstruct $\mathbf{z}(t)$ equally well.

Stress projection term. Stress labels, act as weak supervisory signals. Instead of operating directly on instantaneous latent values, a summary functional of the trajectory is considered:

$$\hat{y} = g(\Phi(\hat{\mathbf{x}})), \quad (7.14)$$

where $\Phi(\hat{\mathbf{x}})$ denotes a low-dimensional descriptor of the autonomic imbalance (for example, time-averaged sympathetic–parasympathetic difference), and $g(\cdot)$ is the class output of the PINN.

The associated supervised loss is defined as the binary cross-entropy between the predicted probability $\hat{y} \in (0, 1)$ and the ground-truth label $y \in \{0, 1\}$:

$$\mathcal{L}_{\text{cls}} = -[y \log(\hat{y}) + (1 - y) \log(1 - \hat{y})]. \quad (7.15)$$

Importantly, this term does not drive the entire representation learning process; it only constrains the latent space to preserve discriminative structure.

Total objective. The resulting theoretical objective can therefore be written as

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{rec}} + \lambda \mathcal{L}_{\text{dyn}} + \eta \mathcal{L}_{\text{cls}} + \mu \mathcal{L}_{\text{reg}}, \quad (7.16)$$

where the hyperparameters λ, η, μ balance reconstruction fidelity, dynamical plausibility, supervised guidance, and parameter regularization.

Conceptually, the dominant learning pressure is exerted by reconstruction and dynamical consistency, while stress labels serve as secondary constraints. This reflects the theoretical stance of the present chapter: perceived stress emerges from autonomic dynamics and should not be treated as a primary observable.

7.13 Discussion and Theoretical Implications

The developments presented in this chapter extend the methodological philosophy introduced in the previous parts of the thesis to a fundamentally different physiological domain. While blood pressure estimation required embedding hemodynamic laws to constrain waveform reconstruction, stress prediction demands confronting an even more ill-posed problem: the absence of a direct, measurable physical quantity linking input signals to the target variable.

In the case of arterial blood pressure, the intermediate process—blood flow dynamics—can be described through well-established conservation laws. In contrast, stress perception is not itself a physical field but a subjective cognitive construct emerging from complex neurophysiological regulation. The autonomic nervous system therefore plays the role of a hidden mediator, analogous to pressure and flow fields in the cardiovascular setting.

The central insight of this chapter is that stress prediction should not be approached as a purely discriminative classification problem. Instead, it can be re-framed as a constrained latent-state estimation task in which observable physiological signals (EDA, HR, RR) are treated as indirect projections of an underlying autonomic dynamical system. Within this perspective, subjective stress labels are not primary drivers of learning but weak supervisory signals acting on a latent physiological representation.

This reinterpretation provides a possible explanation for the empirical limitations observed in purely data-driven experiments. When models attempt to map physiological signals directly to subjective labels, they are forced to approximate an inherently noisy and partially inconsistent mapping. The “hard performance limit” observed across architectures may therefore reflect label entropy rather than insufficient model capacity.

By introducing a minimal antagonistic dynamical prior for sympathetic and parasympathetic interaction, the proposed framework reduces the hypothesis space of admissible latent trajectories. Even in the absence of conclusive experimental validation, this structural constraint offers three conceptual advantages:

- It separates physiological consistency from subjective perception.
- It transforms stress detection into an interpretable state-space problem.

- It provides a principled bridge between signal processing and regulatory physiology.

Importantly, this formulation remains intentionally minimal. The linear antagonistic model does not attempt to reproduce the full complexity of autonomic regulation, nor does it claim clinical validity in its present form. Rather, it establishes a mathematically coherent scaffold upon which future empirical studies can build.

From a broader perspective, this chapter reinforces a recurrent theme of the thesis: whenever a direct mapping between biosignals and a target variable appears unstable, underdetermined, or dataset-dependent, the missing component is often an intermediate dynamical process. Making that process explicit—whether hemodynamic or autonomic—transforms a fragile regression problem into a structured inference task.

7.14 Concluding Remarks and Future Directions

The investigation conducted in this chapter highlights both the promise and the intrinsic difficulty of stress detection from physiological signals. While neural networks can extract discriminative patterns under controlled labeling schemes, subjective self-assessments introduce variability that cannot be eliminated by architectural tuning alone.

The proposed physics-informed latent-state formulation suggests several natural research directions:

- Empirical validation of the autonomic dynamical hypothesis through controlled experimental protocols.
- Integration of respiration-mediated vagal modulation explicitly into the dynamical prior.
- Exploration of semi-supervised and weakly supervised training strategies within the constrained framework.
- Investigation of identifiability and observability properties of the latent state under limited sensing.

Ultimately, the objective is not merely to improve classification accuracy, but to move toward a physiologically interpretable representation of stress-related regulation. If successful, such an approach would align stress modeling with the same principle that guided the evolution of the blood pressure PINN: complex biosignal interpretation becomes more robust when learning is anchored to dynamical structure.

In this sense, the present chapter does not conclude the exploration of stress modeling; rather, it opens a theoretical pathway for future work in which physiological priors and latent dynamical systems play a central role in understanding human regulatory processes.

Chapter 8

Discussion, Conclusions, and Future Directions

8.1 Unifying Perspective of the Thesis

8.1.1 From Direct Mapping to Latent-State Inference

A central conceptual outcome of this thesis is the recognition that biosignal-based estimation cannot, in general, be interpreted as a simple direct mapping from observable signals to target variables. A purely supervised formulation of the form

$$\mathbf{z}(t) \longrightarrow y,$$

where $\mathbf{z}(t)$ denotes the measured biosignals and y a clinically or psychologically relevant target, is often sufficient to obtain good predictive accuracy under controlled conditions. However, the results discussed throughout this work show that such a formulation becomes fragile as soon as the evaluation setting moves beyond constrained within-subject scenarios.

This limitation emerges consistently in both application domains addressed in the thesis. In cuffless blood pressure estimation, ECG and PPG signals provide only indirect information about the underlying cardiovascular dynamics. They reflect aspects of cardiac timing, pulse propagation, and peripheral vascular response, but they do not uniquely determine the internal hemodynamic configuration from which

pressure ultimately arises. In stress modeling, physiological signals such as electrodermal activity, heart rate variability, and respiration reflect autonomic modulation, yet they remain only partial manifestations of a more complex psychophysiological process that also depends on cognitive and contextual factors.

These observations suggest that the problem should be formulated more appropriately as a constrained inference task involving an unobserved physiological state:

$$\mathbf{z}(t) \implies \hat{\mathbf{x}}(t) \implies \hat{y},$$

where $\hat{\mathbf{x}}(t)$ denotes a latent state trajectory inferred from the observations under structural constraints. Within this perspective, the neural network is no longer interpreted merely as a flexible regressor, but as a differentiable estimator embedded in a physiological model. The objective is therefore not only to reproduce target values, but also to ensure that the inferred internal trajectory remains consistent with known dynamical principles.

This shift provides the unifying perspective of the thesis. The common difficulty across blood pressure estimation and stress inference is not simply limited model capacity, but the fact that the target variable depends on latent processes that are only partially observed through wearable signals. Direct regression compresses observation and inference into a single step. Constrained latent-state modeling restores the distinction between them and provides a more coherent formulation of the problem.

8.1.2 Partial Observability, Ill-Posedness, and Identifiability

The latent-state perspective clarifies why biosignal-based estimation is intrinsically difficult under realistic sensing conditions. The problem is not only statistical, but structural. Peripheral biosignals provide incomplete projections of an internal physiological process, and the target of interest is itself a function of that hidden process. A general formulation may be written as

$$\mathbf{z}(t) = \mathcal{H}(\mathbf{x}(t)), \quad y = \mathcal{G}(\mathbf{x}(t)),$$

where $\mathbf{x}(t)$ is the latent state, $\mathcal{H}$ is the observation operator, and $\mathcal{G}$ links the same state to the quantity to be estimated.

Under this formulation, the inverse problem is generally ill-posed. If different latent trajectories produce similar observable signals, then the mapping from $\mathbf{z}(t)$ to y is not uniquely determined. This is precisely the situation encountered in the domains studied in this thesis. In blood pressure estimation, different combinations of vascular stiffness, arterial geometry, peripheral resistance, and flow conditions may generate similar peripheral waveform morphologies while corresponding to different pressure values. In stress modeling, similar physiological activations may correspond to different subjective reports because self-perceived stress also depends on latent cognitive and contextual variables that are not directly measurable.

The empirical findings discussed in the previous chapters are consistent with this interpretation. Strong performance in intra-subject settings can be understood as the consequence of an implicitly fixed latent configuration, in which unobserved subject-specific factors remain approximately constant. By contrast, the degradation observed in inter-subject and cross-database evaluation reflects the failure of a learned direct mapping to remain valid when the latent physiological configuration or the measurement conditions change.

This point is crucial for the interpretation of generalization. The observed performance ceiling of purely data-driven models should not be attributed only to insufficient architectural sophistication. Rather, it reflects limited identifiability under partial observability. When essential determinants of the target variable remain hidden, different predictors may achieve similar empirical performance on a given dataset while relying on correlations that do not transfer beyond it.

Within this context, the role of structural priors becomes clear. Physics-informed constraints, dynamical models, and subject-level physiological descriptors do not remove ambiguity completely, but they reduce the space of admissible solutions by excluding trajectories that are inconsistent with known system structure. Their value lies precisely in transforming an unconstrained regression problem into a more stable and physiologically meaningful inference task.

For this reason, the main methodological contribution of the thesis is not an architectural claim, but a conceptual one. The results collectively support the view that reliable biosignal-based estimation requires explicit acknowledgment of latent dynamics, limited observability, and constrained inference. This perspective provides the common interpretative framework through which both the blood pressure and the stress investigations can be understood.

8.2 Interpretation of the Data-Driven Results

The experimental investigations presented in the previous chapters explored the behavior of purely data-driven models in two domains: cuffless blood pressure estimation and stress classification. Although the specific architectures and datasets differed, the results revealed consistent patterns regarding when learning remains reliable and when it deteriorates.

Rather than interpreting these outcomes solely through numerical performance metrics, this section examines what the models effectively learned under different evaluation regimes. In particular, the comparison between intra-subject, inter-subject, and cross-database validation provides insight into how latent physiological variability and measurement-domain differences influence model behavior.

The evidence collected throughout the thesis is remarkably consistent. In Chapter 5, blood pressure models achieved low error in constrained settings, including waveform reconstruction on curated clinical data and subject-specific estimation on PulsECG, where several folds reached SBP/DBP MAE values in the low single-digit range. However, once evaluation was shifted to leave-one-subject-out or cross-dataset transfer, performance degraded substantially, with some LOSO folds exhibiting SBP errors above 20 mmHg and with direct transfer from MIMIC-trained models to wearable PulsECG recordings showing large dispersion and systematic mismatch. In Chapter 7, a parallel phenomenon emerged in stress modeling: accuracy under subjective labels repeatedly saturated around 70–73% across different architectures, whereas task-based synthetic labels produced a much higher separability, with preliminary performance reaching about 94% accuracy and 98% AUC.

The results should therefore be read not only as empirical benchmarks, but also as evidence supporting the structural interpretation introduced in the previous section. If observable biosignals represent only partial projections of a latent physiological system, then the difficulty of generalization naturally depends on how stable that latent configuration remains across the evaluation setting.

8.2.1 Intra-Subject Learning as a Practical Upper Bound

Intra-subject evaluation represents the most controlled learning scenario considered in this thesis. In this setting, training and testing are performed on data originating

from the same individual, typically separated only in time. Under such conditions, subject-specific physiological parameters can be assumed approximately constant over the observation window.

The results obtained across the blood pressure experiments show that neural networks can achieve low prediction errors when trained and evaluated within a single subject. A similar behavior is observed in stress modeling when training and testing occur within the same participant or under tightly controlled experimental conditions.

This point is supported by concrete evidence from Chapter 5. In the single-subject PulsECG study, the dual-branch ECG/PPG model achieved consistently low MAE across the 10-fold file-wise validation, with several folds yielding SBP and DBP errors close to only a few mmHg. Likewise, the ABP waveform reconstruction experiments on curated clinical data showed that compact data-driven architectures can reproduce a broad range of pressure morphologies with low average error under standard in-domain splits. These findings confirm that ECG and PPG are not devoid of hemodynamic information. On the contrary, when the hidden physiological configuration is effectively fixed, the signal contains sufficient information to support accurate prediction.

This outcome indicates that the signals indeed contain information related to the target variables. If the physiological measurements were entirely uninformative, performance would remain poor even when subject variability is removed. The strong intra-subject results therefore provide an empirical upper bound on what can be achieved when latent parameters such as vascular stiffness, baseline tone, or autonomic responsiveness remain approximately fixed.

However, this regime should not be interpreted as evidence that the mapping from biosignals to physiological targets is universally identifiable. Instead, intra-subject learning implicitly performs a form of subject-specific calibration. The network learns a conditional relationship of the form

$$\mathbf{z}(t) \rightarrow y \mid \theta_{\text{subject}},$$

where θ_{subject} represents latent characteristics of the individual that remain constant within the dataset.

The practical interpretation is therefore not that the inverse problem has been solved, but that it has been simplified by holding the latent subject-specific configuration approximately constant. In this sense, intra-subject performance should be viewed as a best-case operational regime rather than as evidence of a universal biosignal-to-physiology mapping.

From this perspective, strong intra-subject performance reflects the stability of hidden subject-specific parameters rather than the existence of a universal mapping from signals to targets.

8.2.2 Inter-Subject Generalization and Morphological Bias

When evaluation moves from intra-subject to inter-subject validation, the latent physiological configuration is no longer fixed. Each individual introduces variations in vascular properties, cardiovascular regulation, baseline autonomic tone, and demographic characteristics. These factors alter the relationship between observed signal morphology and the target variable.

Across the experiments reported in this thesis, predictive performance decreases when models are trained and evaluated on different subjects. This degradation is observed consistently in both blood pressure estimation and stress classification.

The blood pressure experiments provide a particularly clear illustration of this point. In the PulsECG leave-one-subject-out study, the model could optimize successfully on the pooled training distribution while still failing on the held-out individual, and several folds reached SBP MAE values in the range of roughly 23–27.5 mmHg. This means that training loss reduction and even reasonable validation behavior within the seen population do not guarantee subject-independent inference. Rather, they may simply indicate that the model has learned population-specific correlations that cease to hold when the test subject departs from the morphology represented in training.

One explanation for this behavior lies in the presence of morphological bias. Physiological waveforms encode not only dynamic physiological states, but also relatively stable subject-specific traits. Neural networks trained on such signals may therefore exploit correlations between waveform morphology and target distributions present in the training population.

In practice, the learned predictor may implicitly approximate

$$\hat{y} \approx f(\mathbf{z}(t); \theta_{\text{population}}),$$

where $\theta_{\text{population}}$ captures statistical regularities specific to the training cohort. When a new subject exhibits different physiological characteristics, the learned correlations may no longer hold, leading to systematic prediction errors.

A comparable phenomenon appears in stress classification. Under subject-wise evaluation, increasing the number of modalities or the architectural sophistication did not substantially improve performance, which remained concentrated around 70–73% accuracy. This suggests that the bottleneck is not merely representational capacity, but the instability of the relationship between measurable physiology and the target label across individuals.

8.2.3 Cross-Database Transfer and Measurement Domain Shift

Cross-database evaluation introduces an even more demanding scenario. In this setting, training and testing data originate not only from different individuals, but also from distinct acquisition environments, sensor configurations, and preprocessing pipelines.

The experiments performed across different databases reveal that this regime produces the most pronounced performance degradation. Models trained on one dataset often exhibit systematic bias and increased prediction variance when evaluated on another.

This effect can be understood as a combination of physiological variability and measurement-domain shift. Different datasets may involve distinct patient populations, clinical contexts, or recording conditions. At the same time, sensor hardware, filtering procedures, and annotation protocols may differ across acquisition setups.

From a modeling perspective, the observation process itself becomes domain dependent. The measurement equation may therefore be expressed as

$$\mathbf{z}(t) = \mathcal{H}_d(\mathbf{x}(t)),$$

where the operator $\mathcal{H}_d$ depends on the acquisition domain d .

A model trained under one observation operator implicitly learns correlations specific to that measurement configuration. When applied to data generated by a different operator, those correlations may no longer reflect the same physiological relationships.

The cross-database experiments thus highlight an important limitation of purely data-driven regression. Even if the underlying physiological dynamics remain similar, variations in the measurement channel can significantly affect the learned representation.

8.2.4 Wearable Signals and Real-World Variability

The transition from clinical-grade recording systems to wearable devices introduces additional challenges that extend beyond subject variability and dataset differences. Wearable signals are affected by motion artifacts, variable sensor contact, environmental conditions, and simplified hardware design.

Analyses involving wearable-derived signals show that model performance typically decreases compared to controlled clinical acquisition settings. Even when preprocessing pipelines attempt to standardize filtering and normalization, residual differences in noise structure and waveform morphology remain.

This effect was observed concretely in the transfer experiments from MIMIC-derived data to PulsECG recordings. A model that achieved low error in the original clinical domain did not preserve the same behavior when applied directly to wearable acquisitions. Instead, the predictions displayed large dispersion and poor absolute calibration, and the error analysis showed substantial variability, especially for systolic pressure. Importantly, the subsequent improvement obtained through full fine-tuning indicates that the problem is not random instability, but domain misalignment: the network can adapt, but only after its internal representation is reshaped using subject-specific or device-specific information.

These effects can be interpreted as further variability in the observation process. A more general measurement model can be written as

$$\mathbf{z}(t) = \mathcal{H}_d(\mathbf{x}(t)) + \varepsilon_d(t),$$

where both the observation operator and the noise characteristics depend on the device and acquisition conditions.

In purely data-driven models, the network must implicitly disentangle physiological information from measurement artifacts. When the training data originate from a limited set of devices or recording conditions, the learned representation may encode sensor-specific patterns that do not generalize to other environments.

The wearable experiments therefore provide a realistic stress test for biosignal-based estimation. They illustrate how latent physiological variability and measurement-domain differences jointly affect model reliability, reinforcing the broader argument that direct regression from signals to targets remains fragile under real-world deployment conditions.

8.3 Physics-Informed Modeling: Contributions and Limitations

The limitations observed in purely data-driven experiments motivate the introduction of physics-informed modeling as a structural extension of the learning framework. Within this thesis, physics-informed neural networks are not presented as universally superior predictors, but rather as a principled mechanism to constrain inference in an intrinsically ill-posed problem.

The central idea consists in embedding simplified physiological principles directly into the learning objective. Instead of optimizing a model solely with respect to agreement between predicted and reference signals, the network is also required to produce internal trajectories that remain consistent with known dynamical relations. In the case of blood pressure estimation, these relations arise from simplified one-dimensional hemodynamic equations and boundary conditions inspired by Windkessel-type models.

The role of physics-informed modeling in this context is therefore not primarily to increase predictive flexibility, but to introduce structural constraints that restrict the space of admissible solutions. This section discusses both the contributions and the limitations of this approach in light of the results presented in the previous chapters.

8.3.1 What Physics Adds: Structural Regularization and Physiological Coherence

Embedding hemodynamic constraints into the training objective modifies the nature of the learning problem. In purely data-driven regression, the network is free to approximate any function that minimizes the empirical loss. When conservation laws and constitutive relations are incorporated, the model is additionally required to generate trajectories that remain dynamically consistent.

In the one-dimensional formulation adopted in this thesis, conservation of mass and momentum describe the evolution of cross-sectional area $A(x, t)$ and flow $Q(x, t)$ along the vessel. Pressure $p(x, t)$ is linked to area through a constitutive relation. Within the physics-informed framework, the residuals of these equations are evaluated at collocation points in space and time and included as additional loss terms during training.

The resulting objective function may be expressed schematically as

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{data}} + \lambda_{\text{phys}} \mathcal{L}_{\text{phys}},$$

where $\mathcal{L}_{\text{data}}$ measures agreement with observed signals and $\mathcal{L}_{\text{phys}}$ penalizes violations of the governing equations.

From an optimization perspective, this structure acts as a form of functional regularization. Among the many input-output mappings capable of fitting the data, the model is encouraged to select those that also satisfy simplified hemodynamic dynamics. Solutions that violate conservation principles or produce internally inconsistent trajectories are explicitly discouraged.

The results of Chapter 6 support this interpretation in a concrete way. The proposed Heart-to-Finger PINN did not simply output blood pressure values at a single observation point, but reconstructed pressure, flow, and auxiliary variables along a simplified arterial path. Even where absolute prediction accuracy remained limited, the model produced internally consistent spatiotemporal trajectories and physiologically interpretable quantities, thereby changing the nature of the solution from a black-box regression output to a constrained hemodynamic estimate.

This constraint introduces an additional notion of coherence into the learning process. Predicted quantities such as pressure waveforms and flow profiles can

be evaluated not only in terms of numerical accuracy, but also in terms of their compatibility with the governing equations. Even when the numerical error remains comparable to that of purely data-driven baselines, the resulting predictions are less likely to exhibit non-physical behavior or implausible parameter combinations.

Moreover, the experimental behavior across datasets reinforces the specific role of physics. On MIMIC-III, the PINN achieved reasonable reconstruction accuracy together with stable internal dynamics, while on VitalDB direct transfer still degraded, confirming that physics does not remove domain shift. However, the improvement observed after subject-level fine-tuning suggests that the physics-informed formulation provides a more structured basis for adaptation, because the model is forced to reorganize its predictions around a coherent physiological representation rather than around unconstrained correlations alone.

The main contribution of physics-informed modeling in this thesis should therefore be interpreted as structural regularization. By embedding domain knowledge into the optimization objective, the model is guided toward physiologically admissible solutions rather than relying solely on statistical correlations present in the training data.

8.3.2 What Physics Cannot Resolve Alone

Despite these advantages, physics-informed modeling does not eliminate the fundamental sources of ambiguity identified earlier. The approach constrains the solution space, but it does not fully resolve the identifiability challenges inherent to wearable physiological sensing.

The first limitation arises from partial observability. Peripheral signals such as ECG and PPG provide only indirect information about the internal cardiovascular state. Even if the inferred trajectories satisfy simplified hemodynamic equations, multiple latent configurations may remain compatible with the same observable signals.

A second limitation concerns subject-specific physiological parameters. Quantities such as arterial stiffness, vessel geometry, baseline vascular tone, and peripheral resistance vary substantially across individuals and are typically not measured in wearable monitoring scenarios. Without explicit subject descriptors or calibration

procedures, these parameters remain implicit and cannot be uniquely identified through the available signals alone.

A further limitation derives from the simplified nature of the governing equations employed. The one-dimensional flow model and Windkessel-type boundary conditions provide tractable approximations of cardiovascular dynamics, but they inevitably abstract away complex phenomena such as three-dimensional flow patterns, nonlinear wall mechanics, and distributed microvascular regulation. The imposed constraints therefore restrict the solution space to trajectories consistent with the chosen model class, which may only approximate real physiological behavior.

Finally, physics-informed training introduces additional practical challenges. Balancing the influence of data fidelity and physical consistency requires careful weighting of the respective loss terms. If the physical constraints are emphasized excessively, empirical accuracy may degrade; if they are weighted too weakly, the model may revert to shortcut learning driven primarily by statistical correlations.

For these reasons, physics-informed modeling should not be interpreted as a complete solution to the problem of calibration-free physiological inference. Its primary value lies in introducing structural guidance that improves plausibility and stability, while the remaining ambiguity must be addressed through complementary strategies. These may include the integration of subject-specific descriptors, improved observation models, or minimal calibration procedures capable of anchoring individual physiological parameters.

8.4 Stress Modeling as a Mirror Problem

The investigation of stress modeling in this thesis complements the study of cuffless blood pressure estimation by revealing a closely related structural difficulty. Although the physiological mechanisms and targets differ, both domains illustrate the limitations of direct signal-to-target regression when the quantity of interest depends on latent processes that are only partially observable.

In blood pressure estimation, the challenge lies in inferring internal hemodynamic quantities from peripheral waveforms without explicit knowledge of subject-specific cardiovascular parameters. In stress modeling, the situation is even more complex:

the supervised target is not a purely physiological variable but a subjective construct shaped by autonomic activity, cognitive appraisal, and contextual factors.

The experimental results obtained in the stress classification studies reinforce the interpretation developed throughout this thesis. Increasing architectural complexity alone did not substantially improve subject-wise performance, suggesting that the limitation arises from the structure of the problem rather than from insufficient model capacity. In this sense, stress modeling provides a complementary example of how latent physiological processes mediate the relationship between observable biosignals and the target variable.

8.4.1 Limits of Supervision with Subjective Labels

One of the most informative observations emerging from the stress experiments is the performance ceiling encountered when training classifiers using subjective self-reported labels. Despite experimenting with different architectures and multimodal combinations of physiological signals, classification accuracy remained relatively stable under subject-wise validation.

This ceiling is not a vague impression, but an empirical regularity observed across the experiments of Chapter 7. Whether using individual modalities, multimodal fusion, or more sophisticated mechanisms such as cross-attention, validation performance remained concentrated around 70–73% accuracy, with no substantial gain from additional complexity. The negative result is therefore informative: it suggests that the main limitation lies in the target formulation rather than in architectural expressiveness.

This behavior can be explained by the intrinsic ambiguity of the supervision signal. Self-reported stress levels reflect a complex psychological construct that depends not only on physiological activation but also on individual perception, coping mechanisms, and contextual interpretation.

From a modeling perspective, the supervised task can be written as

$$\mathbf{z}(t) \rightarrow y_{\text{subjective}},$$

where $\mathbf{z}(t)$ represents physiological measurements and $y_{\text{subjective}}$ denotes the reported stress level. However, if the mapping between autonomic activity and subjective

perception varies across individuals, the learning problem becomes intrinsically noisy.

In such conditions, even an ideal representation of physiological signals cannot uniquely determine the target label. The irreducible error observed in the experiments therefore reflects not only measurement noise, but also the limited alignment between the supervised target and the measurable physiological processes.

This interpretation explains why increasing model complexity did not substantially improve subject-wise performance. Architectural refinements can enhance representation learning, but they cannot resolve inconsistencies inherent to the supervision signal itself.

8.4.2 Task-Based Labels as a Controlled Proxy

A different pattern emerges when supervision is derived from task-based experimental conditions rather than from subjective ratings. When labels correspond to controlled stress-inducing or relaxation tasks, classification accuracy increases and the previously observed performance ceiling is partially alleviated.

In Chapter 7, the contrast was particularly pronounced: when synthetic task-based labels were introduced, preliminary experiments reached about 94% accuracy, 98% AUC, and 94% weighted F1. Although these results were not obtained under the full evaluation pipeline used for the subjective-label experiments, the magnitude of the improvement is too large to be dismissed as a minor methodological fluctuation. It strongly indicates that the physiological signals contain structured information about autonomic activation, but that this structure is only partially aligned with self-reported stress.

Task-based supervision introduces a more direct relationship between physiological activation and label assignment. Experimental protocols designed to induce stress produce structured autonomic responses that are temporally aligned with the annotated condition. Although such labels remain imperfect proxies of internal experience, they are generally more tightly coupled to measurable physiological dynamics.

The supervised mapping therefore becomes

$$\mathbf{z}(t) \rightarrow y_{\text{task}},$$

where y_{task} reflects the experimental condition rather than subjective perception.

The improvement observed under this regime confirms that physiological signals contain meaningful information about autonomic activation. The earlier performance ceiling was therefore not due to the absence of informative features in the signals, but rather to the weak coupling between physiological measurements and subjective supervision.

Nevertheless, task-based labels remain an approximation. Experimental conditions cannot fully reproduce the complexity of real-world stress perception, and individuals may still respond differently to the same stimulus. Thus, while task-based supervision improves learnability, it does not completely eliminate latent variability.

8.4.3 Toward Latent Autonomic State Estimation

The contrast between subjective and task-based supervision suggests that stress modeling should be interpreted as an inference problem involving latent autonomic dynamics. Physiological signals provide partial observations of an underlying regulatory process rather than direct encodings of perceived stress intensity.

Let $\mathbf{x}_{\text{ANS}}(t)$ denote a latent autonomic state describing the balance between sympathetic and parasympathetic activity. Physiological signals may be interpreted as projections of this internal state

$$\mathbf{z}(t) = \mathcal{H}_{\text{ANS}}(\mathbf{x}_{\text{ANS}}(t)),$$

while perceived stress depends on both the autonomic trajectory and additional psychological or contextual variables

$$y_{\text{stress}} = \mathcal{G}(\mathbf{x}_{\text{ANS}}(t), \mathbf{c}).$$

Within a direct classification framework, the intermediate state $\mathbf{x}_{\text{ANS}}(t)$ is never explicitly modeled. The network instead attempts to learn a direct mapping from

physiological signals to labels, implicitly absorbing both physiological dynamics and cognitive variability into a single function.

A latent-state formulation provides a more coherent alternative. By explicitly modeling autonomic dynamics through simplified differential relations, the learning problem can be reframed as the estimation of an internal physiological trajectory rather than the direct prediction of subjective labels.

Such an approach does not assume that perceived stress can be perfectly reconstructed from physiological signals alone. Instead, it aims to isolate the component of stress perception that is mediated by autonomic activity, providing a representation that may be more stable and transferable across individuals.

From this perspective, stress modeling reinforces the broader methodological argument of the thesis. Both blood pressure estimation and stress inference involve hidden physiological systems that are only partially observed through wearable signals. When the target variable depends on such latent dynamics, constrained inference frameworks provide a more consistent modeling strategy than direct regression alone.

8.5 Methodological Lessons and Best Practices

Beyond the specific architectures and application domains investigated, the experiments conducted in this thesis highlight several methodological considerations relevant to biosignal-based machine learning. These considerations arise from the systematic comparison of validation protocols, supervision strategies, and signal preprocessing pipelines across both blood pressure and stress modeling tasks.

Rather than presenting universal guidelines, the observations summarized here should be interpreted as practical lessons emerging from the empirical analysis carried out in the previous chapters. Together, they emphasize the importance of aligning evaluation protocols, data quality control, and performance metrics with the structural properties of physiological inference problems.

8.5.1 Validation Protocols and Generalization Criteria

One of the most significant methodological insights concerns the role of validation protocol design. In biosignal-based estimation, the choice of data partitioning strategy strongly influences the interpretation of reported performance.

Random window-level splits may artificially inflate predictive accuracy by allowing subject-specific waveform characteristics to appear simultaneously in both training and testing sets. Such protocols effectively measure interpolation within individuals rather than true generalization.

Subject-wise validation provides a more realistic assessment. By ensuring that signals from a given individual appear exclusively in either the training or testing partition, this approach exposes the impact of latent subject-specific variability. The consistent performance degradation observed in this thesis when moving from intra-subject to subject-wise evaluation highlights how strongly biosignal-based prediction can depend on hidden individual characteristics.

An even more stringent assessment is obtained through cross-database evaluation, where models are trained and tested on datasets collected under different acquisition conditions. This regime simultaneously introduces physiological variability and measurement-domain shift, providing a more realistic approximation of real-world deployment scenarios.

These observations suggest that different notions of generalization should be clearly distinguished in biosignal modeling:

- **Intra-subject generalization**, corresponding to temporal extrapolation within a single individual.
- **Inter-subject generalization**, corresponding to transfer across unseen individuals within the same acquisition domain.
- **Cross-domain generalization**, corresponding to transfer across datasets with different populations and measurement setups.

Explicitly reporting the validation regime is therefore essential for the meaningful interpretation of performance metrics in physiological inference tasks.

8.5.2 Quality Control and Failure Mode Analysis

Physiological signals collected in clinical or wearable environments are often affected by artifacts, sensor displacement, and synchronization errors. Without adequate quality control, such disturbances may introduce spurious correlations that machine learning models can inadvertently exploit.

The analyses conducted in this thesis indicate that model performance is highly sensitive to signal quality and preprocessing decisions. Poor-quality segments increase variability in the training objective and may distort the relationship between waveform morphology and target variables.

Systematic quality control therefore serves two complementary purposes. First, it improves signal consistency and stabilizes model training. Second, it helps distinguish between limitations arising from intrinsic physiological ambiguity and those caused by measurement artifacts.

Failure mode analysis provides an additional diagnostic tool. Examining prediction errors across subjects, experimental conditions, or signal quality levels often reveals structured patterns that remain hidden in aggregate metrics. For example, systematic bias in specific individuals may indicate unmodeled subject-specific parameters, whereas performance degradation under motion conditions may reflect sensor-related variability rather than physiological complexity.

In this sense, quality assessment should be viewed not merely as preprocessing, but as a methodological safeguard ensuring that empirical conclusions reflect genuine modeling limitations rather than avoidable measurement artifacts.

8.5.3 Evaluation Metrics Beyond Pointwise Error

A further methodological consideration concerns the interpretation of evaluation metrics. In many biosignal-based studies, model performance is summarized through pointwise error measures such as mean absolute error, root mean square error, or classification accuracy. While informative, these metrics capture only numerical agreement between predictions and reference values.

In physiological inference tasks, additional aspects of model behavior may be equally important. For blood pressure waveform estimation, relevant criteria

include preservation of waveform morphology, consistency of systolic and diastolic extrema, and stability of derived indices. A predictor may achieve low mean absolute error while still producing temporally inconsistent or physiologically implausible trajectories.

Similarly, in stress classification problems, aggregate accuracy may conceal systematic bias across subjects or experimental conditions. Analyzing confusion matrices, subject-wise performance, and distributional shifts often provides deeper insight into how the model interprets physiological patterns.

The introduction of physics-informed constraints further expands the range of meaningful evaluation criteria. Residuals of governing equations and the stability of inferred physiological quantities across domains become relevant indicators of model reliability.

Taken together, these considerations suggest that evaluation in biosignal-based machine learning should extend beyond pointwise error alone. Combining numerical accuracy with assessments of physiological coherence and cross-domain robustness provides a more informative picture of model behavior.

8.6 Limitations of the Present Work

While the results presented in this thesis support the proposed shift toward constrained latent-state modeling, several limitations must be acknowledged. These limitations arise from the sensing configuration, the characteristics of the available datasets, and the modeling assumptions adopted in the proposed frameworks.

Recognizing these constraints is essential for correctly interpreting the results and for guiding future developments in biosignal-based physiological inference.

8.6.1 Observability Constraints

The most fundamental limitation of the present work arises from partial observability. Both blood pressure estimation and stress modeling rely on peripheral biosignals that provide indirect and incomplete information about the underlying physiological processes.

In the blood pressure domain, only surface ECG and distal PPG signals are available as inputs. These measurements capture aspects of cardiac timing and peripheral blood volume changes, but they do not directly encode distributed pressure, flow, arterial geometry, or vascular compliance. Even when hemodynamic constraints are incorporated through physics-informed losses, the inverse problem remains underdetermined because multiple latent cardiovascular configurations may produce similar peripheral waveform projections.

A similar limitation appears in stress modeling. Electrodermal activity, heart rate variability, and respiration reflect autonomic modulation, yet they do not capture cognitive appraisal, emotional context, or individual psychological traits that influence perceived stress. Consequently, the supervised target, particularly in the case of subjective labels, depends on processes that are only partially reflected in measurable physiology.

These observability constraints imply that the estimation problem cannot be fully resolved through signal processing alone. The models developed in this thesis therefore aim to infer physiologically plausible trajectories given limited observations, rather than to reconstruct the complete internal state of the underlying systems.

8.6.2 Dataset and Population Bias

A second limitation concerns the characteristics of the datasets used for training and evaluation. Although widely adopted in the literature, the available databases do not fully represent the diversity of real-world populations and sensing conditions.

Clinical waveform databases often consist of recordings collected from hospitalized patients under specific medical conditions, medications, or monitoring environments. Such populations may exhibit physiological patterns that differ from those of healthy ambulatory individuals, potentially introducing bias into the learned representations.

Differences between datasets also extend to sensor hardware, acquisition protocols, preprocessing pipelines, and annotation procedures. These factors contribute to measurement-domain variability, which can affect model behavior when transferring across datasets.

In stress modeling, dataset limitations are further amplified by the nature of experimental protocols. Laboratory tasks designed to induce stress provide controlled conditions but may not capture the complexity of stress responses encountered in everyday life. As a result, learned associations between physiological patterns and stress labels may not generalize beyond the experimental context.

Consequently, while the methodological conclusions of this thesis are intended to be general, the empirical validation remains bounded by the characteristics of the available datasets.

8.6.3 Modeling Assumptions and Simplifications

A further limitation arises from the modeling assumptions adopted in the physics-informed formulations developed in this work.

In the blood pressure framework, the cardiovascular system is approximated through a one-dimensional description of arterial blood flow combined with simplified constitutive relations and reduced-order boundary conditions. While these models capture essential features of pulse wave propagation and pressure–flow coupling, they neglect phenomena such as three-dimensional flow structures, complex wall viscoelasticity, and detailed microvascular dynamics.

Similarly, Windkessel-type boundary conditions provide an effective abstraction of peripheral circulation but cannot fully represent dynamic autoregulation or subject-specific microvascular variability. The imposed constraints therefore restrict the admissible solution space to trajectories compatible with the chosen model class, which may only approximate real physiological behavior.

In stress modeling, the latent autonomic framework represents sympathetic and parasympathetic dynamics through simplified differential relations. Such abstraction does not account for the full neuroendocrine cascade, including hormonal regulation and higher-level cognitive processes influencing stress perception.

These simplifications were necessary to maintain tractability and identifiability within the available sensing configuration. However, they also define the scope of the conclusions that can be drawn from the proposed models.

Future developments may incorporate richer physiological representations, but such extensions will inevitably introduce additional parameters and identifiability

challenges. Balancing physiological realism with estimability therefore remains a central issue in wearable biosignal modeling.

8.7 Conclusions

This thesis investigated the problem of estimating physiological and psychophysiological quantities from peripheral biosignals under realistic sensing constraints. Through the study of cuffless blood pressure estimation and stress modeling, a common structural challenge emerges: the difficulty of inferring internal physiological states from signals that provide only partial and indirect observations of the underlying system.

The experimental evidence supporting this conclusion is consistent across the different chapters. In blood pressure estimation, purely data-driven models achieved low error when trained and evaluated in constrained regimes, including curated in-domain datasets and subject-specific settings. However, their behavior deteriorated markedly under leave-one-subject-out and cross-dataset transfer, with some folds reaching clinically unacceptable systolic errors and with wearable transfer experiments revealing poor absolute calibration. In stress modeling, a parallel limitation emerged: subjective-label classification repeatedly saturated around 70–73% accuracy despite changes in architecture and signal fusion strategy, whereas task-based labels produced much higher separability and substantially better predictive performance.

The experimental results demonstrate that modern neural network architectures can achieve strong predictive performance under controlled conditions, particularly in intra-subject settings where latent physiological parameters remain approximately constant. However, the transition to more realistic validation regimes, including inter-subject and cross-database evaluation, reveals substantial degradation in performance. These results indicate that purely data-driven models often rely on statistical correlations specific to a given population or acquisition setup rather than capturing invariant physiological relationships.

Within this context, the main contribution of the thesis lies in reframing biosignal-based estimation as a constrained inference problem rather than a direct signal-to-target regression task. Observable signals are interpreted as partial projections of an

evolving latent physiological state, while the quantities of interest arise as functions of this hidden process. This perspective motivates the integration of structural priors, such as simplified hemodynamic equations or dynamical models of autonomic regulation, into the learning framework.

Physics-informed modeling was therefore explored as a mechanism for incorporating physiological structure into neural networks. Rather than aiming to replace data-driven learning, the proposed approach constrains the space of admissible solutions by enforcing consistency with simplified dynamical principles. This form of structural regularization improves physiological plausibility and reduces degeneracy in the inferred trajectories, although it does not eliminate the intrinsic ambiguity associated with limited observability.

Importantly, the results of the PINN chapter show that the value of physics is not captured only by pointwise prediction error. The Heart-to-Finger framework enabled the reconstruction of coherent pressure and flow trajectories along a simplified arterial pathway and provided interpretable latent quantities that remain inaccessible to standard black-box regressors. At the same time, the persistence of degradation under cross-dataset evaluation confirms that structural priors reduce ambiguity but do not fully solve identifiability when the sensing configuration remains limited.

The parallel investigation of stress modeling further reinforces this interpretation. The performance ceiling observed when training on subjective stress labels highlights the limitations of direct supervised learning when the target variable depends on latent physiological and psychological processes that are not fully measurable. Reframing the problem in terms of latent autonomic dynamics provides a conceptual bridge between physiological inference and affective computing.

Taken together, the results presented in this thesis support three main conclusions:

- Accurate prediction within a constrained domain does not necessarily imply transferable physiological inference, as shown by the contrast between strong intra-subject or in-domain performance and the marked degradation observed in leave-one-subject-out and cross-dataset evaluations.
- Generalization failures in biosignal-based models often reflect limited observability and latent variability rather than insufficient model complexity, as evidenced by both the subject-dependent collapse of blood pressure predic-

tors and the hard performance ceiling obtained in stress classification under subjective supervision.

- Incorporating structural physiological priors provides a principled strategy to constrain inference and improve robustness under realistic sensing conditions, not by removing ambiguity completely, but by replacing unconstrained regression with physiologically admissible latent-state estimation.

More broadly, this work argues that progress in wearable physiological monitoring requires a shift in perspective. Reliable inference cannot rely solely on increasingly complex neural architectures, but must explicitly account for the latent dynamical systems that generate observable biosignals. Integrating machine learning with physiological modeling therefore represents a promising path toward more interpretable and robust biosignal-based estimation.

8.8 Future Research Directions

The results presented in this thesis highlight both the potential and the current limitations of biosignal-based physiological inference under wearable sensing constraints. Future research should therefore aim to extend the constrained latent-state perspective developed here, combining improved physiological modeling with more robust data-driven learning strategies.

Several directions appear particularly promising for advancing this framework.

8.8.1 Integration of Subject-Specific Descriptors

A central challenge identified throughout this work is the influence of subject-specific physiological variability. Parameters such as vascular stiffness, arterial geometry, and autonomic responsiveness vary substantially across individuals and strongly affect the relationship between observable signals and the target variables.

Future work should therefore investigate the integration of subject-level descriptors into the inference framework. Demographic and anthropometric variables including age, sex, height, weight, and body composition may provide valuable

priors on cardiovascular and autonomic characteristics. Incorporating such information as conditioning variables or hierarchical parameters could help reduce latent ambiguity and improve inter-subject generalization.

In addition to static descriptors, minimal calibration procedures or reference measurements may offer a practical compromise between fully calibration-free approaches and personalized physiological estimation.

8.8.2 Improved Modeling of Hemodynamic Inputs

The physics-informed framework for blood pressure estimation relies on simplified assumptions regarding hemodynamic inputs and boundary conditions. While these approximations enable tractable modeling, further refinement of these components represents an important research direction.

In particular, more principled estimation of flow-related quantities, cardiac output surrogates, and subject-dependent vascular profiles could reduce arbitrariness in the current formulation. Linking hemodynamic inputs to anthropometric variables or population-level physiological models may allow the construction of more realistic boundary conditions without requiring invasive measurements.

Exploring hybrid strategies in which certain physiological parameters are inferred from data while remaining constrained within plausible ranges may further improve both realism and identifiability.

8.8.3 Advances in Physics-Informed Architectures

From an architectural perspective, several extensions of physics-informed learning frameworks deserve further investigation. Multi-scale models combining simplified mechanistic representations with flexible neural components may provide a more balanced trade-off between physiological interpretability and predictive flexibility.

Adaptive strategies for balancing data-driven and physics-based loss components could also improve training stability. In addition, probabilistic formulations capable of quantifying uncertainty in the inferred latent states would provide valuable information for clinical interpretation and decision support.

Identifiability-aware training procedures, including sensitivity analysis and physiological parameter regularization, may further enhance the robustness of physics-informed models when operating under partial observability.

8.8.4 Extensions to Stress and Multimodal Physiological Modeling

The latent autonomic framework proposed for stress modeling suggests several possible extensions. Future studies may incorporate additional sensing modalities, such as temperature, accelerometry, or hormonal proxies, in order to better constrain the hidden autonomic state.

Combining task-based supervision with subjective labels through weakly supervised or semi-supervised learning may also help disentangle physiological activation from cognitive appraisal processes. Furthermore, integrating slower regulatory mechanisms, such as hypothalamic–pituitary–adrenal axis dynamics, could provide a richer representation of stress physiology.

More broadly, constrained latent-state modeling may be applicable to other psychophysiological constructs, including fatigue, emotional arousal, or cognitive workload, potentially providing a unified framework for wearable affective computing.

8.8.5 Toward Real-World Wearable Deployment

Finally, translating these approaches to real-world wearable monitoring remains an essential objective. Future research should evaluate model stability over longer monitoring periods and across diverse activity conditions, environmental contexts, and device configurations.

Continuous or adaptive calibration strategies may help bridge the gap between theoretical calibration-free inference and practical reliability in everyday use. Standardized cross-device evaluation protocols would also facilitate fair comparison between methods and improve reproducibility across studies.

Ultimately, progress in wearable physiological monitoring will depend on the integration of structural physiological modeling, robust machine learning methodolo-

gies, and realistic deployment constraints. The framework developed in this thesis provides a foundation for such integration, but further interdisciplinary research will be necessary to achieve reliable and interpretable physiological inference in real-world settings.

References

- [1] Flávio D Fuchs and Paul K Whelton. High blood pressure and cardiovascular disease. *Hypertension*, 75(2):285–292, 2020.
- [2] N Pilz, DS Picone, A Patzak, OS Opatz, T Lindner, L Fessler, V Heinz, and TL Bothe. Cuff-based blood pressure measurement: challenges and solutions. *Blood Pressure*, 33(1):2402368, 2024.
- [3] Federico Delrio, Valentina De Vitis, Silvia Caligari, Vincenzo Randazzo, Giansalvo Cirrincione, and Eros Pasero. Enhancing ecg analysis with a hybrid deep learning approach: Automatic detection of significant features. In *Advanced Neural Artificial Intelligence: Theories and Applications*, pages 23–33. Springer, 2025.
- [4] Vincenzo Randazzo, Jacopo Ferretti, and Eros Pasero. Anytime ECG monitoring through the use of a low-cost, user-friendly, wearable device. 21(18):6036.
- [5] Jacopo Ferretti, Vincenzo Randazzo, Giansalvo Cirrincione, and Eros Pasero. 1-D convolutional neural network for ECG arrhythmia classification. In Anna Esposito, Marcos Faundez-Zanuy, Francesco Carlo Morabito, and Eros Pasero, editors, *Progresses in Artificial Intelligence and Neural Systems*, Smart Innovation, Systems and Technologies, pages 269–279. Springer.
- [6] Majid Nour, Kemal Polat, Ümit Şentürk, and Murat Arıcan. A novel cuffless blood pressure prediction: uncovering new features and new hybrid ml models. *Diagnostics*, 13(7):1278, 2023.
- [7] Federico Delrio, Vincenzo Randazzo, Giansalvo Cirrincione, and Eros Pasero. Non-invasive arterial blood pressure estimation from electrocardiogram and photoplethysmography signals using a conv1d-bilstm neural network. *Engineering Proceedings*, 39(1):78, 2023.
- [8] Annunziata Paviglianiti, Vincenzo Randazzo, Giansalvo Cirrincione, and Eros Pasero. Neural recurrent approaches to noninvasive blood pressure estimation. In *2020 international joint conference on neural networks (IJCNN)*, pages 1–7. IEEE, 2020.
- [9] ZengDing Liu, Bin Zhou, Ye Li, Min Tang, and Fen Miao. Continuous blood pressure estimation from electrocardiogram and photoplethysmogram during arrhythmias. *Frontiers in physiology*, 11:575407, 2020.

- [10] Eoin Finnegan, Shaun Davidson, Mirae Harford, Peter Watkinson, Lionel Tarassenko, and Mauricio Villarroel. Features from the photoplethysmogram and the electrocardiogram for estimating changes in blood pressure. *Scientific Reports*, 13(1):986, 2023.
- [11] Vincenzo Randazzo, Pietro Buccellato, Jacopo Ferretti, Federico Delrio, and Eros Pasero. Pulsecg-a cuffless non-invasive blood pressure monitoring device through neural network analysis of ecg and ppg signals. In *2024 IEEE 22nd Mediterranean Electrotechnical Conference (MELECON)*, pages 1030–1035. IEEE, 2024.
- [12] National Library of Medicine. What is an electrocardiogram (ecg).
- [13] M.Sc. Sussha Cheriyeedath. Photoplethysmography (ppg).
- [14] Siti Nor Ashikin Ismail, Nazrul Anuar Nayan, Rosmina Jaafar, and Zazilah May. Recent advances in non-invasive blood pressure monitoring and prediction using a machine learning approach. 22(16):6195.
- [15] Nikoleta Sinou, Natalia Sinou, Stamatios Koutroulakis, and Dimitrios Filippou. The role of wearable devices in blood pressure monitoring and hypertension management: A systematic review. *Cureus*, 16(12), 2024.
- [16] Zirui Lou, Jun Tao, Binbin Wei, Xinyu Jiang, Simin Cheng, Zehao Wang, Chao Qin, Rong Liang, Haotian Guo, Liping Zhu, et al. Near-infrared organic photodetectors toward skin-integrated photoplethysmography-electrocardiography multimodal sensing system. *Advanced Science*, 10(36):2304174, 2023.
- [17] Jung Jae Park, Sangwoo Hong, Yeongju Jung, Phillip Won, Carmel Majidi, Minwoo Kim, and Seung Hwan Ko. Highly sensitive cuffless blood pressure monitoring with selective laser-sintered liquid metal conductors. *Advanced Functional Materials*, page 2505089, 2025.
- [18] Myeong Namkoong, Balaji Baskar, Lakhvir Singh, Heng Guo, Justin McMurray, Kimberly Branan, Md Saifur Rahman, Chin-To Hsiao, Josh Kuriakose, Joanna Hernandez, et al. Add-on soft electronic interfaces for continuous cuffless blood pressure monitoring. *Advanced materials technologies*, 8(15):2300158, 2023.
- [19] Shuwen Chen, Shicheng Fan, Zheng Qiao, Zixiong Wu, Baobao Lin, Zhijie Li, Michael A Riegler, Matthew Yu Heng Wong, Arve Opheim, Olga Korostynska, et al. Transforming healthcare: Intelligent wearable sensors empowered by smart materials and artificial intelligence. *Advanced Materials*, page 2500412, 2025.
- [20] Giansalvo Cirrincione, Vincenzo Randazzo, and Eros Pasero. A neural based comparative analysis for feature extraction from ECG signals. In Anna Esposito, Marcos Faundez-Zanuy, Francesco Carlo Morabito, and Eros Pasero,

- editors, *Neural Approaches to Dynamics of Signal Exchanges*, Smart Innovation, Systems and Technologies, pages 247–256. Springer.
- [21] Vincenzo Randazzo, Gaia Marchetti, Carla Giustetto, Erica Gugliermine, Rahul Kumar, Giansalvo Cirrincione, Fiorenzo Gaita, and Eros Pasero. Learning-based approach to predict fatal events in brugada syndrome. In *Applications of Artificial Intelligence and Neural Systems to Data Science*, pages 63–72. Springer, 2023.
- [22] Giansalvo Cirrincione, Vincenzo Randazzo, and Eros Pasero. The growing curvilinear component analysis (gccca) neural network. *Neural Networks*, 103:108–117, 2018.
- [23] Rikuto Yoshizawa, Kohei Yamamoto, and Tomoaki Ohtsuki. Investigation of data leakage in deep-learning-based blood pressure estimation using photoplethysmogram/electrocardiogram. *IEEE Sensors Journal*, 23(12):13311–13318, 2023.
- [24] Suril Mehta, Nipun Kwatra, Mohit Jain, and Daniel McDuff. Examining the challenges of blood pressure estimation via photoplethysmogram. *Scientific Reports*, 14(1):18318, 2024.
- [25] Ramakrishna Mukkamala, George S Stergiou, and Alberto P Avolio. Cuffless blood pressure measurement. *Annual Review of Biomedical Engineering*, 24(1):203–230, 2022.
- [26] Annemette Sofie Olufsen. Modeling the arterial system with reference to an anesthesia simulator. 1998.
- [27] Alberto Zingaro, Michele Bucelli, Ivan Fumagalli, Luca Dede’, and Alfio Quarteroni. Modeling isovolumetric phases in cardiac flows by an augmented resistive immersed implicit surface method. *International Journal for Numerical Methods in Biomedical Engineering*, 39(12):e3767, 2023.
- [28] Luca Formaggia, Daniele Lamponi, and Alfio Quarteroni. One-dimensional models for blood flow in arteries. 47(3):251–276.
- [29] Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019.
- [30] Ramakrishna Mukkamala, Mohammad Yavarimanesh, Keerthana Natarajan, Jin-Oh Hahn, Konstantinos G Kyriakoulis, Alberto P Avolio, and George S Stergiou. Evaluation of the accuracy of cuffless blood pressure measurement devices: challenges and proposals. *Hypertension*, 78(5):1161–1167, 2021.
- [31] Guillaume Weber-Boisvert, Benoit Gosselin, and Frida Sandberg. Intensive care photoplethysmogram datasets and machine-learning for blood pressure estimation: Generalization not guaranteed. *Frontiers in Physiology*, 14:1126957, 2023.

- [32] Mohammad Moulaeifard, Peter H Charlton, and Nils Strodthoff. Generalizable deep learning for photoplethysmography-based blood pressure estimation—a benchmarking study. *arXiv preprint arXiv:2502.19167*, 2025.
- [33] Ana Cisnal, Yanke Li, Bertram Fuchs, Mehdi Ejtehad, Robert Riener, and Diego Paez-Granados. Robust feature selection for bp estimation in multiple populations: Towards cuffless ambulatory bp monitoring. *IEEE Journal of Biomedical and Health Informatics*, 28(10):5768–5779, 2024.
- [34] Nico Westerhof, Jan-Willem Lankhaar, and Berend E Westerhof. The arterial windkessel. *Medical & biological engineering & computing*, 47(2):131–141, 2009.
- [35] Wen Wen, Rong Luo, Xiaojing Tang, Lan Tang, Hunter X Huang, Xiaoyan Wen, Shan Hu, and Bin Peng. Age-related progression of arterial stiffness and its elevated positive association with blood pressure in healthy people. *Atherosclerosis*, 238(1):147–152, 2015.
- [36] Lingfeng Li, Xue-Cheng Tai, and Raymond Hon-Fu Chan. A new method to compute the blood flow equations using the physics-informed neural operator. *Journal of Computational Physics*, 519:113380, 2024.
- [37] Shuran Zhou, Lisheng Xu, Liling Hao, Hanguang Xiao, Yang Yao, Lin Qi, and Yudong Yao. A review on low-dimensional physics-based models of systemic arteries: application to estimation of central aortic pressure. *Biomedical engineering online*, 18(1):41, 2019.
- [38] Peter Charlton. Ppg waveform.
- [39] Michael J. Joyner, Nisha Charkoudian, and B. Gunnar Wallin. The sympathetic nervous system and blood pressure in humans. *Experimental Physiology*, 95(1):2–5, 2010.
- [40] Hugo D. Critchley. Electrodermal responses: what happens in the brain. *The Neuroscientist*, 8(2):132–142, 2002.
- [41] Wolfram Boucsein. *Electrodermal Activity*. Springer, 2012.
- [42] Marek Malik, J. Thomas Bigger, A. John Camm, Robert E. Kleiger, Alberto Malliani, Arthur J. Moss, and Peter J. Schwartz. Heart rate variability: Standards of measurement, physiological interpretation, and clinical use. 17(3):354–381.
- [43] Heart rate variability: Origins, methods, and interpretive caveats. *Psychophysiology*, 34(6):623–648, 1997.
- [44] Toward understanding respiratory sinus arrhythmia: Relations to cardiac vagal tone, evolution and biobehavioral functions. *Biological Psychology*, 74(2):263–285, 2007.

- [45] Bruce S. McEwen. Stress, adaptation, and disease. allostasis and allostatic load. *Annals of the New York Academy of Sciences*, 840:33–44, 1998.
- [46] Robert-Paul Juster, Bruce S. McEwen, and Sonia J. Lupien. Allostatic load biomarkers of chronic stress and impact on health and cognition. *Neuroscience and Biobehavioral Reviews*, 35(1):2–16, 2010.
- [47] Pablo Laguna, Roger G Mark, A Goldberg, and George B Moody. A database for evaluation of algorithms for measurement of QT and other waveform intervals in the ECG. In *Computers in cardiology 1997*, pages 673–676. IEEE, 1997.
- [48] Ary L Goldberger, Luis AN Amaral, Leon Glass, Jeffrey M Hausdorff, Plamen Ch Ivanov, Roger G Mark, Joseph E Mietus, George B Moody, Chung-Kang Peng, and H Eugene Stanley. Physiobank, physiotoolkit, and physionet: components of a new research resource for complex physiologic signals. *circulation*, 101(23):e215–e220, 2000.
- [49] Alistair Johnson, Tom Pollard, and Roger Mark. MIMIC-III clinical database.
- [50] Hyung-Chul Lee, Yoonsang Park, Soo Bin Yoon, Seong Mi Yang, Dongnyeok Park, and Chul-Woo Jung. Vitaldb, a high-fidelity multi-parameter vital signs database in surgical patients. *Scientific Data*, 9(1):279, 2022.
- [51] Weinan Wang, Pedram Mohseni, Kevin L Kilgore, and Laleh Najafizadeh. Pulsedb: A large, cleaned dataset based on mimic-iii and vitaldb for benchmarking cuff-less blood pressure estimation methods. *Frontiers in Digital Health*, 4:1090854, 2023.
- [52] Mohammad Kachuee, Mohammad Mahdi Kiani, Hoda Mohammadzade, and Mahdi Shabany. Cuffless blood pressure estimation algorithms for continuous health-care monitoring. *IEEE Transactions on Biomedical Engineering*, 64:859–869, 04 2017.
- [53] Dheeru Dua and Casey Graff. UCI machine learning repository, 2017.
- [54] Hava Chaptoukaev, Valeriya Strizhkova, Michele Panariello, Bianca Dalpaos, Aglind Reka, Valeria Manera, Susanne Thümmel, Esma Ismailova, Massimiliano Todisco, Maria A Zuluaga, et al. Stressid: a multimodal dataset for stress identification. *Advances in Neural Information Processing Systems*, 36:29798–29811, 2023.
- [55] Vincenzo Randazzo, Edoardo Puleo, Annunziata Paviglianiti, Alberto Vallan, and Eros Pasero. Development and validation of an algorithm for the digitization of ecg paper images. *Sensors*, 22(19):7138, 2022.
- [56] Francesca Alfieri, Andrea Ancona, Giovanni Tripepi, Vincenzo Randazzo, Annunziata Paviglianiti, Eros Pasero, Luigi Vecchi, Cristina Politi, Valentina Cauda, and Riccardo Maria Fagugli. External validation of a deep-learning model to predict severe acute kidney injury based on urine output changes in critically ill patients. *Journal of Nephrology*, 35(8):2047–2056, 2022.

- [57] Jacopo Ferretti, Pietro Barbiero, Vincenzo Randazzo, Giansalvo Cirrincione, and Eros Pasero. Towards uncovering feature extraction from temporal signals in deep cnn: the ecg case study. In *2020 international joint conference on neural networks (IJCNN)*, pages 1–7. IEEE, 2020.
- [58] Annunziata Paviglianiti, Vincenzo Randazzo, Eros Pasero, and Alberto Vallan. Noninvasive arterial blood pressure estimation using abpnet and vital-ecg. In *2020 IEEE International Instrumentation and Measurement Technology Conference (I2MTC)*, pages 1–5. IEEE, 2020.
- [59] Matteo Stefanini, Marta Lovino, Rita Cucchiara, and Elisa Ficarra. Predicting gene and protein expression levels from dna and protein sequences with perceiver. *Computer Methods and Programs in Biomedicine*, 234:107504, 2023.
- [60] Giansalvo Cirrincione, Sergio Cannata, Giovanni Cicceri, Francesco Prinzi, Tiziana Currieri, Marta Lovino, Carmelo Militello, Eros Pasero, and Salvatore Vitabile. Transformer-based approach to melanoma detection. *Sensors*, 23(12):5677, 2023.
- [61] Eros Pasero, Fiorenzo Gaita, Vincenzo Randazzo, Pierre Meynet, Sergio Cannata, Philippe Maury, and Carla Giustetto. Artificial intelligence ecg analysis in patients with short qt syndrome to predict life-threatening arrhythmic events. *Sensors*, 23(21):8900, 2023.
- [62] Pietro Barbiero, Marta Lovino, Mattia Siviero, Gabriele Ciravegna, Vincenzo Randazzo, Elisa Ficarra, and Giansalvo Cirrincione. Unsupervised multi-omic data fusion: The neural graph learning network. In *Intelligent Computing Theories and Application: 16th International Conference, ICIC 2020, Bari, Italy, October 2–5, 2020, Proceedings, Part I*, pages 172–182. Springer, 2020.
- [63] Marta Lovino, Vincenzo Randazzo, Gabriele Ciravegna, Pietro Barbiero, Elisa Ficarra, and Giansalvo Cirrincione. A survey on data integration for multi-omics sample clustering. *Neurocomputing*, 488:494–508, 2022.
- [64] Marta Lovino, Gianpaolo Bontempo, Giansalvo Cirrincione, and Elisa Ficarra. Multi-omics classification on kidney samples exploiting uncertainty-aware models. In *International Conference on Intelligent Computing*, pages 32–42. Springer, 2020.
- [65] The WFDB software package.
- [66] Lokenath Debnath and Jean-Pierre Antoine. Wavelet transforms and their applications. *Physics Today - PHYS TODAY*, 56:68–68, 04 2003.
- [67] Luís Aguiar-Conraria and Maria Joana Soares. The continuous wavelet transform: Moving beyond uni- and bivariate analysis. *Journal of Economic Surveys*, 28(2):344–375, 2014.

- [68] Xiaohong Liang, Liping Li, Yuanyuan Liu, Dan Chen, Xinpei Wang, Shunbo Hu, Jikuo Wang, Huan Zhang, Chengfa Sun, and Changchun Liu. ECG_segnet: an ECG delineation model based on the encoder-decoder structure. *Computers in Biology and Medicine*, 145:105445, 2022.
- [69] Nicholas Hughes, Lionel Tarassenko, and Stephen J Roberts. Markov models for automated ecg interval analysis. *Advances in Neural Information Processing Systems*, 16, 2003.
- [70] Qin Tang, Jing Liang, and Fangqi Zhu. A comparative review on multi-modal sensors fusion based on deep learning. *Signal Processing*, page 109165, 2023.
- [71] Amedeo Buonanno, Giovanni Di Gennaro, Giorgio Graditi, Antonio Nogarotto, Francesco AN Palmieri, and Maria Valenti. Fusion of energy sensors with missing values. *Applied Intelligence*, 53(20):23613–23627, 2023.
- [72] Annunziata Paviglianiti, Vincenzo Randazzo, Stefano Villata, Giansalvo Cirrincione, and Eros Pasero. A comparison of deep learning techniques for arterial blood pressure prediction. 14(5):1689–1710.
- [73] Iqbal H. Sarker. Machine learning: Algorithms, real-world applications and research directions. 2(3):160.
- [74] Nabil Ibtehaz, Sakib Mahmud, Muhammad E. H. Chowdhury, Amith Khandakar, Muhammad Salman Khan, Mohamed Arselene Ayari, Anas M. Tahir, and M. Sohel Rahman. PPG2abp: Translating photoplethysmogram (PPG) signals to arterial blood pressure (ABP) waveforms. 9(11):692. Number: 11 Publisher: Multidisciplinary Digital Publishing Institute.
- [75] Federico Delrio, Vincenzo Randazzo, Giansalvo Cirrincione, and Eros Pasero. Dual-branch fusion neural network model: A preliminary demonstration of absolute, calibration-free blood pressure estimation from ecg and ppg signals. In Anna Esposito, Marcos Faundez-Zanuy, Francesco Carlo Morabito, Eros Pasero, and Gennaro Cordasco, editors, *Neural Networks: Overview of Current Theories and Applications*, volume 459 of *Smart Innovation, Systems and Technologies*, pages 41–50. Springer, Singapore, 2026.
- [76] Alistair E. W. Johnson, Tom J. Pollard, Lu Shen, Li-wei H. Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G. Mark. MIMIC-III, a freely accessible critical care database. 3(1):160035. Number: 1 Publisher: Nature Publishing Group.
- [77] Annunziata Paviglianiti, Vincenzo Randazzo, Stefano Villata, Giansalvo Cirrincione, and Eros Pasero. A comparison of deep learning techniques for arterial blood pressure prediction. *Cognitive Computation*, 14:1689–1710, 08 2021.
- [78] Sakib Mahmud, Nabil Ibtehaz, Amith Khandakar, Anas M. Tahir, Tawsifur Rahman, Khandaker Reajul Islam, Md Shafayet Hossain, M. Sohel Rahman, Farayi Musharavati, Mohamed Arselene Ayari, Mohammad Tariqul Islam,

- and Muhammad E. H. Chowdhury. A shallow u-net architecture for reliably predicting blood pressure (bp) from photoplethysmogram (ppg) and electrocardiogram (ecg) signals. *Sensors*, 22:919, 01 2022.
- [79] Sakib Mahmud, Nabil Ibtehaz, Amith Khandakar, M. Sohel Rahman, Antonio JR. Gonzales, Tawsifur Rahman, Md Shafayet Hossain, Md. Sakib Abrar Hossain, Md. Ahasan Atick Faisal, Farhan Fuad Abir, Farayi Musharavati, and Muhammad E. H. Chowdhury. Nabnet: A nested attention-guided biconv-lstm network for a robust prediction of blood pressure components from reconstructed arterial blood pressure waveforms using ppg and ecg signals. *Biomedical Signal Processing and Control*, 79:104247, 01 2023.
- [80] Jacopo Ferretti, Vincenzo Randazzo, Giansalvo Cirrincione, and Eros Pasero. 1-d convolutional neural network for ecg arrhythmia classification. *Progresses in Artificial Intelligence and Neural Systems*, pages 269–279, 2021.
- [81] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [82] tf.keras.optimizers.adam | tensorflow v2.11.0.
- [83] World Health Organization. Cardiovascular diseases, 2019.
- [84] World Health Organization. Facts about medical devices | who | regional office for africa, 2023.
- [85] National Library of Medicine. What is an electrocardiogram (ecg).
- [86] M.Sc. Susha Cheriyaedath. Photoplethysmography (ppg).
- [87] Andrew S. Tseng and Peter A. Noseworthy. Prediction of atrial fibrillation using machine learning: A review. 12.
- [88] Eros Pasero, Fiorenzo Gaita, Vincenzo Randazzo, Pierre Meynet, Sergio Cannata, Philippe Maury, and Carla Giustetto. Artificial intelligence ECG analysis in patients with short QT syndrome to predict life-threatening arrhythmic events. 23(21):8900. Number: 21 Publisher: Multidisciplinary Digital Publishing Institute.
- [89] Vincenzo Randazzo, Gaia Marchetti, Carla Giustetto, Erica Gugliemina, Rahul Kumar, Giansalvo Cirrincione, Fiorenzo Gaita, and Eros Pasero. Learning-based approach to predict fatal events in brugada syndrome. In

- Anna Esposito, Marcos Faundez-Zanuy, Francesco Carlo Morabito, and Eros Pasero, editors, *Applications of Artificial Intelligence and Neural Systems to Data Science*, Smart Innovation, Systems and Technologies, pages 63–72. Springer Nature.
- [90] AliveCor Company. AliveCor company.
- [91] Withings Company. Withings company.
- [92] Apple Company. Apple watch.
- [93] Samsung Company. Samsung galaxy watch.
- [94] The University of Nottingham. Cardiology teaching package - einthoven's triangle.
- [95] N. V. Thakor and J. G. Webster. Ground-free ecg recording with two electrodes. *IEEE Transactions on Biomedical Engineering*, BME-27(12):699–704, 1980.
- [96] J. Surda, S. Lovas, J. Pucik, and M. Jus. Spectral properties of ecg signal. In *Proceedings of the 17th International Conference Radioelektronika*, pages 1–5, Brno, Czech Republic, 2007. IEEE Xplore.
- [97] Giansalvo Cirrincione, Vincenzo Randazzo, and Eros Pasero. A neural based comparative analysis for feature extraction from ecg signals. *Neural approaches to dynamics of signal exchanges*, pages 247–256, 2020.
- [98] Tasbiraha Athaya and Sunwoong Choi. An estimation method of continuous non-invasive arterial blood pressure waveform using photoplethysmography: A u-net architecture-based approach. *Sensors*, 21:1867, 03 2021.
- [99] Vincenzo Randazzo, Jacopo Ferretti, and Eros Pasero. Anytime ecg monitoring through the use of a low-cost, user-friendly, wearable device. *Sensors*, 21(18):6036, 2021.
- [100] Vincenzo Randazzo, Jacopo Ferretti, and Eros Pasero. A wearable smart device to monitor multiple vital parameters: Vital ecg. *Electronics*, pages 9–10, 2020.
- [101] Kaan Sel, Amirmohammad Mohammadi, Roderic I Pettigrew, and Roozbeh Jafari. Physics-informed neural networks for modeling physiological time series for cuffless blood pressure estimation. *npj Digital Medicine*, 6(1):110, 2023.
- [102] Solmaz Rastegar, Hamid Gholam Hosseini, and Andrew Lowe. Hybrid cnn-svr blood pressure estimation model using ecg and ppg signals. *Sensors*, 23(3):1259, 2023.
- [103] Rui Wang et al. Adversarial contrastive learning based physics-informed temporal networks for cuffless blood pressure estimation. *arXiv preprint arXiv:2408.08488*, 2024.

- [104] Federico Delrio, Vincenzo Randazzo, Giansalvo Cirrincione, and Eros Pasero. Physics informed neural network for continuous and cuffless arterial blood pressure. In *2025 IEEE International Conference on Artificial Intelligence for Learning and Optimization (ICoAILO)*, pages 143–148. IEEE, 2025.
- [105] David N Ku. Blood flow in arteries. *Annual review of fluid mechanics*, 29(1):399–434, 1997.
- [106] Axel R Pries, D Neuhaus, and P Gaehtgens. Blood viscosity in tube flow: dependence on diameter and hematocrit. *American Journal of Physiology-Heart and Circulatory Physiology*, 263(6):H1770–H1778, 1992.
- [107] DEM Taylor. Blood flow in arteries. by da mcdonald. edward arnold, london, 1974. pp. xviii+ 496.£ 12. *Quarterly Journal of Experimental Physiology and Cognate Medical Sciences: Translation and Integration*, 60(1):65–65, 1975.
- [108] GJ Langewouters, KH Wesseling, and WJA Goedhard. The static elastic properties of 45 human thoracic and 20 abdominal aortas in vitro and the parameters of a new model. *Journal of biomechanics*, 17(6):425–435, 1984.
- [109] Luana de Rezende Mikael, Anelise Machado Gomes de Paiva, Marco Mota Gomes, Ana Luiza Lima Sousa, Paulo César Brandão Veiga Jardim, Priscila Valverde de Oliveira Vitorino, Maicon Borges Euzébio, Wátila de Moura Sousa, and Weimar Kunz Sebba Barroso. Vascular aging and arterial stiffness. *Arquivos brasileiros de cardiologia*, 109(3):253–258, 2017.
- [110] Susana Novella, Ana Paula Dantas, Gloria Segarra, Pascual Medina, and Carlos Hermenegildo. Vascular aging in women: is estrogen the fountain of youth? *Frontiers in physiology*, 3:165, 2012.
- [111] Mackenzie N Kehmeier and Ashley E Walker. Sex differences in large artery stiffness: implications for cerebrovascular dysfunction and alzheimer’s disease. *Frontiers in aging*, 2:791208, 2021.
- [112] Fatemeh Heydari, Malikeh P. Ebrahim, Jean-Michel Redoute, Keith Joe, Katie Walker, Alberto Avolio, and Mehmet R. Yuce. Clinical study of a chest-based cuffless blood pressure monitoring system. *Medical Devices & Sensors*, 3(5):e10091, 2020.
- [113] Jesse Fine, Kimberly L Branan, Andres J Rodriguez, Tananant Boonya-Ananta, Ajmal, Jessica C Ramella-Roman, Michael J McShane, and Gerard L Cote. Sources of inaccuracy in photoplethysmography for continuous cardiovascular monitoring. *Biosensors*, 11(4):126, 2021.
- [114] Liangqi Wang, Shuo Tian, and Rong Zhu. A new method of continuous blood pressure monitoring using multichannel sensing signals on the wrist. *Microsystems & Nanoengineering*, 9(1):117, 2023.
- [115] Jiao Zheng and Zhengyu Yu. A novel machine learning-based systolic blood pressure predicting model. *Journal of Nanomaterials*, 2021(1):9934998, 2021.