

# Non-Contact Microwave Carotid Sensing for Blood Pressure Estimation with Deep Learning Method

Jianing Zhang, Mao Wang, Shigeru Shimamoto

Department of Computer Science and Communications Engineering, Waseda University, Tokyo, Japan

Emails: zhangjianing2026@akane.waseda.jp, wangmao@ruri.waseda.jp, shima@waseda.jp

**Abstract**—In this paper, we propose a contactless carotid sensing method for blood pressure (BP) estimation at 2.4 GHz. Carotid pulsations are measured as the forward transmission coefficient  $S_{21}$  using a vector network analyzer (VNA). To obtain stable heartbeat representations, we perform band-pass filtering and envelope processing to suppress slow drift, posture- and distance-induced amplitude changes, and respiration components. Recognizing the different temporal characteristics of systolic blood pressure (SBP) and diastolic blood pressure (DBP), we adopt a task-specialized dual-model design: a one-dimensional convolutional neural network (1D-CNN) regressor for SBP and a one-dimensional residual network (ResNet-1D) regressor for DBP, with optional linear post-calibration. Evaluation follows a subject-independent protocol with disjoint participants between development and test cohorts. On the test cohort, SBP achieves mean absolute error (MAE) 6.20 mmHg and root mean square error (RMSE) 7.30 mmHg, while DBP achieves MAE 3.99 mmHg and RMSE 5.18 mmHg. We also conducted a comprehensive panel of agreement and consistency analyses (e.g., Bland–Altman and related metrics). The results indicate near-zero bias and good overall consistency across subjects, suggesting that microwave carotid sensing combined with task-specialized learning enables effective cuffless BP estimation and shows practical potential for continuous monitoring.

**Index Terms**—Microwave, blood pressure, contactless sensing, machine learning

## I. INTRODUCTION

Hypertension is a major, yet preventable and controllable, global health burden. According to the World Health Organization (WHO), an estimated 1.4 billion adults aged 30–79 years were living with hypertension in 2024, while population-level control rates remain suboptimal [1]. These facts imply that intermittent, office-based measurements are insufficient. There is a pressing need for more frequent and context-appropriate blood pressure assessment.

Conventional cuff-based devices remain the clinical and home reference standard with mature accuracy. However, measurements are intermittent, and cuff inflations can cause discomfort and nocturnal disturbance, limiting long-duration monitoring. Studies have reported that nighttime ambulatory blood pressure monitoring (ABPM) may provoke arousals and degrade sleep quality, affecting adherence and data quality [2], [3]. By contrast, photoplethysmography (PPG) improves convenience but still requires tight skin contact and is highly sensitive to sensor–skin coupling and motion; moreover, sweat, skin pigmentation or tattoos, ambient light, and temperature can degrade signal quality and accuracy [4]. These limitations

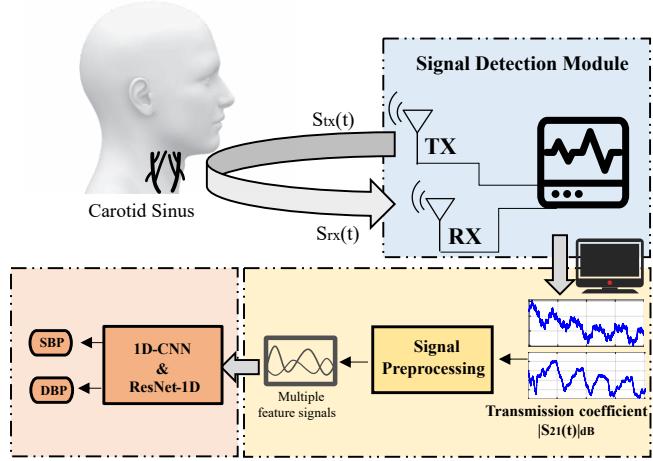


Fig. 1. Overall experimental flow.

motivate contactless approaches to improve comfort and applicability across daily-life scenarios.

Building on microwave and millimeter-wave (mmWave) sensing, Shi et al. in [5] used a single mmWave radar to form delay and Doppler representations with motion suppression and evaluated blood pressure (BP) regression under a leave-one-subject-out (LOSO) protocol; the reported mean error (ME) and standard deviation (SD) were 0.87 and 5.01 mmHg for systolic blood pressure (SBP), and 1.55 and 5.27 mmHg for diastolic blood pressure (DBP). Singh et al. in [6] synchronized two frequency-modulated continuous-wave (FMCW) radars at the chest and the wrist to compute pulse transit time (PTT) and pulse wave velocity (PWV), combining these metrics with wrist morphology for regression; on five subjects (fifty measurements each), the root mean square error (RMSE) was 3.33 mmHg (SBP) and 3.14 mmHg (DBP). Zhao et al. in [7] employed a 24 GHz linear frequency-modulated continuous-wave (LFMCW) sensor at the carotid site and applied a modified ResNet-18 for end-to-end estimation; on one hundred subjects, the record-wise mean absolute error (MAE) was 3.51 mmHg (SBP) and 2.44 mmHg (DBP). These methods are closest to ours but rely on specialized motion/time–frequency pipelines, dual-device alignment and synchronization, or band/model–specific designs with heterogeneous reporting, which complicate deployment, weaken alignment robustness, hinder comparability, and may constrain

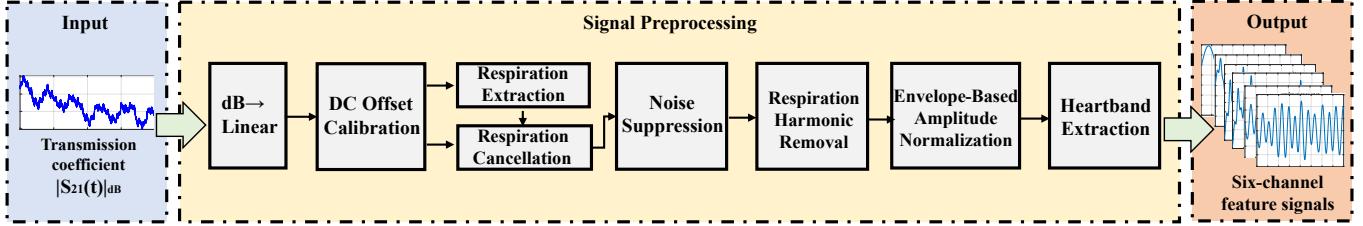


Fig. 2. Signal preprocessing pipeline.

cross-subject generalization.

In signal preprocessing, Qiu et al. in [8] presented a radar-oriented preprocessing chain with Kalman denoising and multi-scale band-pass filtering, followed by beat-wise segmentation, temporal alignment, and amplitude normalization to produce stable cardiac segments. Wang et al. in [9] obtained baseband magnitude via demodulation, applied a bidirectional fourth-order Butterworth band-pass filter and downsampling, then performed linear detrending and amplitude standardization, constructed derivative channels, and used cross-correlation for temporal alignment to emphasize the upstroke and foot while unifying timing. These processing-heavy pipelines improve input quality but introduce multi-stage dependencies and parameter sensitivities (filters, segmentation, alignment) whose cross-device and cross-scenario robustness remains to be established.

Fig. 1 presents a 2.4 GHz contactless carotid-sensing scheme. A single-tone microwave link at the carotid site is monitored to capture subtle temporal fluctuations in the received magnitude  $|S_{21}(t)|$ . To obtain stable and comparable cardiac representations, preprocessing includes: (i) respiratory reference cancellation via least squares (LS) to suppress low-frequency wander; (ii) narrowband notches at the respiratory fundamental and its low-order harmonics to attenuate harmonic leakage; and (iii) envelope-adaptive normalization (EAN), a respiration-aware automatic gain control, to stabilize gain and baseline across subjects and segments. Cardiac-band branches are then extracted for regression. Because systolic and diastolic estimation rely on different discriminative cues, a task-specialized dual model is adopted—1D-CNN for SBP and ResNet-1D for DBP—with optional linear post-calibration.

The contributions of this paper are as follows:

- We built a 2.4 GHz, single-link, single-site contactless carotid-sensing setup that can be deployed quickly on site for proof-of-concept evaluation. Using commodity RF and antenna modules together with bench-top VNA instrumentation, a reproducible prototype was realized, and the signal chain can in principle migrate to integrated, low-power wearable or edge implementations.
- We chose the carotid artery as the target site to acquire pulsations closer to central hemodynamics, reducing interference from heart sounds and structural complexity at the precordial region.
- We designed a magnitude-only signal-processing pipeline

that includes respiratory cancellation, harmonic notch filtering, and envelope-adaptive normalization (EAN) to stabilize gain and baseline and to suppress respiratory sidebands, producing consistent cardiac-band inputs that support SBP and DBP regression.

## II. MEASUREMENT METHODS

### A. Signal principle and Preprocessing

1) *Signal principle*: Measurements are performed at a single tone  $f_0 = 2.4$  GHz. The instrument records the logarithmic magnitude of the forward transmission coefficient:

$$|S_{21}(t)|_{\text{dB}} = 20 \log_{10} |S_{21}(t)| = 20 \log_{10} \left| \frac{s_{\text{rx}}(t)}{s_{\text{tx}}(t)} \right|, \quad (1)$$

where  $s_{\text{tx}}(t)$  is a single-tone excitation radiated toward the carotid region and  $s_{\text{rx}}(t)$  is the received signal after interaction with neck tissues. The transmit and receive antennas together with the subject can be modeled as a slowly time-varying two-port network at  $f_0$ :

$$S_{21}(t) = H(f_0, t) = A(t) e^{j\phi(t)}, \quad (2)$$

and only the magnitude channel is available. The linear magnitude  $A(t) = |S_{21}(t)|$  is therefore used as the observable.

Slow variations in  $A(t)$  are mainly induced by changes in near-field coupling and mutual impedance caused by minute surface displacements. When the carotid site is targeted, the observable is well approximated by the superposition of a respiratory term and a cardiac term,

$$A(t) = A_0 [1 + \alpha_r r(t) + \alpha_c c(t)] + n_A(t), \quad (3)$$

where  $r(t)$  represents respiration in 0.1 to 0.5 Hz,  $c(t)$  represents cardiac micro-motion in 0.7 to 2.0 Hz,  $A_0$  is a static baseline,  $\alpha_r$  and  $\alpha_c$  are coupling coefficients that depend on placement and subject factors, and  $n_A(t)$  aggregates instrument noise and small motion. Fig. 1 summarizes the overall sensing-to-regression concept.

2) *Preprocessing*: The objective is to extract pulse-related components from the observed magnitude sequence. Fig. 2 illustrates the overall pipeline. The raw  $A(t) = |S_{21}(t)|$  is converted into six time-domain channels that serve as inputs to the downstream models, as shown in Fig. 3. All processing is performed in the linear domain of  $A(t)$ .

Interface unification is applied first. Baseline correction removes very slow additive drift and device offsets so that

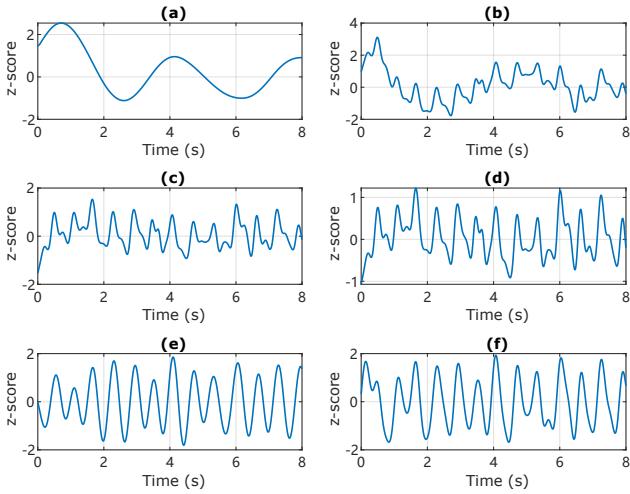


Fig. 3. Six-channel signals: (a) breathing reference (0.1–0.5 Hz); (b) respiration-removed, low-pass (5 Hz cutoff); (c) respiration-removed, notch suppressed; (d) envelope-adaptive normalization (EAN) signal; (e) heartbeat band, narrow (0.9–2.0 Hz); (f) heartbeat band, wide (0.8–3.0 Hz).

segments from different subjects share a comparable zero level and scale. Sampling-rate unification avoids step-size discrepancies across recordings and keeps filter and segmentation parameters consistent.

Respiratory cancellation follows. From (3) the observable contains a respiratory term  $r(t)$  whose low-frequency wander can obscure cardiac details in 0.7–2.0 Hz. A smoothed respiratory reference  $\tilde{r}(t)$  is obtained by band-pass filtering in 0.1–0.5 Hz and is discretized as  $\tilde{r}[n]$ . The coupling strength and a slow baseline term are estimated by ridge-regularized least squares and then subtracted:

$$\hat{\mathbf{w}} = \arg \min_{\mathbf{w}} \|\mathbf{a} - \mathbf{X}\mathbf{w}\|_2^2 + \lambda \|\mathbf{w}\|_2^2, \quad (4)$$

where  $\mathbf{w}$  is the coefficient vector,  $\mathbf{a}$  is the observation vector formed by  $A[n]$ , and  $\mathbf{X}$  stacks  $\mathbf{x}[n] = [\tilde{r}[n], \tilde{r}[n-1], \dots, \tilde{r}[n-L+1], 1, n/N]^\top$ . Here  $\lambda$  is the ridge parameter,  $L$  is the lag order, and  $N$  is the segment length.

$$s_{\text{clean}}[n] = A[n] - \mathbf{x}[n]^\top \hat{\mathbf{w}}, \quad (5)$$

where  $s_{\text{clean}}[n]$  denotes the residual after LS-based removal of the respiratory coupling. To suppress nonphysiological high-frequency noise such as quantization spikes and short transients, and to provide a stable baseline for the subsequent stages, a light low-pass is applied:

$$s_{\text{lp}}[n] = \text{LP}_{5\text{Hz}}\{s_{\text{clean}}[n]\}, \quad (6)$$

where  $s_{\text{lp}}[n]$  is the sequence after a 5 Hz cutoff low-pass, and  $\text{LP}_{5\text{Hz}}\{\cdot\}$  denotes the low-pass operator with a 5 Hz cutoff. The choice of 5 Hz is above the upper edge of the cardiac-wide band at 3.0 Hz, so distortion to the target band is negligible.

Interference near the cardiac neighborhood is then addressed in two steps. The first step targets harmonic leakage: the interaction between periodic respiration and the system response produces narrow spectral lines at the respiratory fundamental

$f_r$  and its second harmonic  $2f_r$ . Narrowband notches are placed at  $f_r$  and  $2f_r$ , yielding narrowband-suppressed signal  $s_{\text{notch}}[n]$ . This change is most evident in the spectrum because the sharp lines at  $f_r$  and  $2f_r$  are attenuated, whereas the time-domain trace may look similar to (5).

The second step handles slowly varying gain and baseline drift. Unlike the initial baseline correction that treats additive drift, real recordings often exhibit multiplicative scaling caused by posture and scattering geometry, which can be approximated as

$$s_{\text{notch}}[n] \approx g[n] s_c[n] + u[n], \quad (7)$$

where  $g[n] > 0$  is a slowly varying gain envelope,  $s_c[n]$  is the cardiac-near component, and  $u[n]$  is a small additive residual. To stabilize this multiplicative variation, envelope-adaptive normalization (EAN) is applied:

$$s_{\text{norm}}[n] = \frac{s_{\text{notch}}[n]}{e[n]}, \quad (8)$$

where  $e[n]$  denotes the smoothed envelope obtained after cardiac-near prefiltering and is used as a proxy for the slowly varying gain  $g[n]$ . Dividing  $s_{\text{notch}}[n]$  by  $e[n]$  yields the amplitude-normalized trace  $s_{\text{norm}}[n]$ , which stabilizes inter-segment gain and baseline. From  $s_{\text{norm}}[n]$ , we then extract the branches for the following regression.

Two cardiac branches are finally extracted to balance morphology sensitivity and robustness. A narrow branch of 0.9 to 2.0 Hz emphasizes the upstroke and peak details, and a wider branch of 0.8 to 3.0 Hz preserves more envelope trends. All inputs used for learning are standardized within each segment by  $z$ -scoring to prevent cross-segment statistical leakage. The envelope  $e(t)$  is only an intermediate for (8) and is not provided to the models.

By default, each training and evaluation sample is an 8 s segment resampled to 120 Hz. We stack six channels to form the input tensor of shape  $C \times T = 6 \times 960$  (channels  $\times$  time).

## B. Models

1) *Physiology-informed model selection*: Physiology and waveform analysis indicate that systolic blood pressure (SBP) is dominated by rapid, localized morphology around the early-systolic upstroke and the late-systolic peak, with augmentation influenced by the timing of forward and reflected waves [10]. In contrast, diastolic blood pressure (DBP) reflects longer time-scale behavior related to peripheral resistance and arterial compliance, manifested as a slow diastolic envelope and decay [11]. Guided by these mechanisms, SBP is modeled with a 1D-CNN that emphasizes local and fast patterns, whereas DBP is modeled with a ResNet-1D whose residual stacking enlarges the effective temporal field for integrating slow envelopes and low-frequency trends [12]. This pairing is consistent with reports that convolutional architectures, particularly ResNet variants, are effective for cuffless BP estimation from pulsatile waveforms [13], [14].

For comparability, both architectures produce two scalar outputs, SBP and DBP, and we report the intended pairings

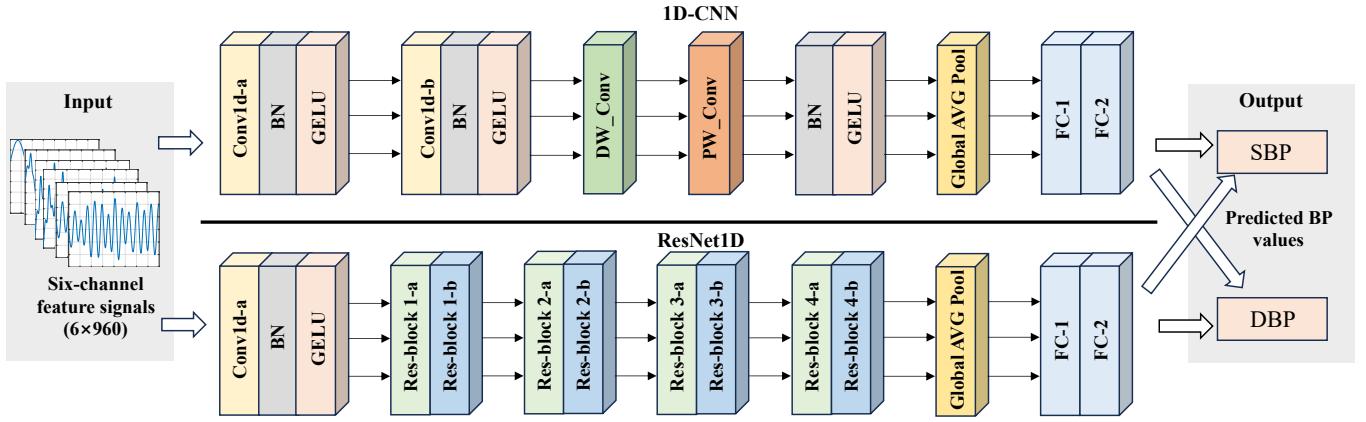


Fig. 4. Model structure for 1D-CNN and ResNet-1D.

TABLE I  
1D-CNN (SBP).

Layer	Type	$k/s/p$	OutCh	OutLen
Input	Six-channel signals	–	6	960
Conv1a	Conv1d	7/1/3	32	960
Conv1b	Conv1d	5/1/2	64	960
DW5	Depthwise Conv1d	5/1/2	64	960
PW	Pointwise $1 \times 1$	1/1/0	96	960
GAP	Global Average Pool	–	96	1
FC-1	Linear (+GELU)	–	64	1
FC-2	Linear (SBP/DBP)	–	2	1

$k/s/p$  = kernel/stride/padding.

TABLE II  
RESNET-1D (DBP).

Layer	Type / Structure	$k/s/p$ (c1; c2)	OutCh	OutLen
Input	Six-channel signals	–	6	960
Stem	Conv7 (+BN, GELU)	7/2/3	32	480
Res-block 1a	Conv(3) $\times 2$ + id skip	3/1/1; 3/1/1	32	480
Res-block 1b	Conv(3) $\times 2$ + id skip	3/1/1; 3/1/1	32	480
Res-block 2a	Conv(3) $\times 2$ + $1 \times 1$ proj skip	3/2/1; 3/1/1	64	240
Res-block 2b	Conv(3) $\times 2$ + id skip	3/1/1; 3/1/1	64	240
Res-block 3a	Conv(3) $\times 2$ + $1 \times 1$ proj skip	3/2/1; 3/1/1	96	120
Res-block 3b	Conv(3) $\times 2$ + id skip	3/1/1; 3/1/1	96	120
Res-block 4a	Conv(3) $\times 2$ + $1 \times 1$ proj skip	3/2/1; 3/1/1	128	60
Res-block 4b	Conv(3) $\times 2$ + id skip	3/1/1; 3/1/1	128	60
GAP	Global Average Pool	–	128	1
FC-1	Linear (+GELU)	–	64	1
FC-2	Linear (SBP/DBP)	–	2	1

(c1; c2) =  $k/s/p$  of conv1/conv2; “id skip” = identity shortcut; “proj skip” =  $1 \times 1$  projection shortcut.

together with cross-task baselines, where each architecture is evaluated on the non-designated target, under identical data splits and training settings. Layer-wise architectures are summarized in Table I and Table II; high-level blocks are shown in Fig. 4.

2) *1D-CNN for SBP regression*: The input  $\mathbf{X} \in \mathbb{R}^{C \times T}$  is a multi-branch stack of preprocessed signals such as the envelope-normalized trace and the cardiac narrow/wide bands. A convolutional block is summarized by

$$\text{CB}_{k,s}^c(\mathbf{x}) = \phi(\text{BN}(\text{Conv}_{k,s}^c(\mathbf{x}))), \quad (9)$$

with 1D convolution  $\text{Conv}_{k,s}^c$ , batch normalization (BN), and activation  $\phi$  GELU. The temporal encoder is the composition

$$\mathbf{H}_{\text{SBP}} = (\text{CB}_{k_L,s_L}^{c_L} \circ \dots \circ \text{CB}_{k_2,s_2}^{c_2} \circ \text{CB}_{k_1,s_1}^{c_1})(\mathbf{X}), \quad (10)$$

followed by global average pooling (GAP) over time to obtain a compact representation. A linear head then outputs scalar SBP.

Layer-level notes aligned with Fig. 4 and Table I. The first two temporal convolutions keep stride one and preserve sequence length so that upstroke slope and peak sharpness are not attenuated. Next, a depthwise convolution operates channel-by-channel to retain short, per-branch receptive fields and emphasize local morphology, and a subsequent pointwise  $1 \times 1$  convolution mixes information across channels to increase

capacity with little parameter cost. BN and GELU stabilize optimization on short segments, and GAP removes absolute timing within the 8s window, improving robustness to minor misalignment. Overall, the stack is biased toward local amplitude/gradient cues that are informative for SBP while remaining lightweight; the head produces both SBP and DBP for consistent reporting across architectures.

3) *ResNet-1D for DBP regression*: The DBP branch is a residual 1D network built from BasicBlocks with identity shortcuts;  $1 \times 1$  projection shortcuts are used when stride or channels change. A block is expressed as

$$\text{BB}_{k,s}^c(\mathbf{x}) = \mathbf{x}_{\text{skip}} \oplus \text{BN}_2 \left( \text{Conv}_{k,1}^c \left( \sigma(\text{BN}_1(\text{Conv}_{k,s}^c(\mathbf{x}))) \right) \right), \quad (11)$$

where  $\oplus$  is elementwise addition,  $\sigma$  is GELU, and  $\mathbf{x}_{\text{skip}} = \mathbf{x}$  when shapes match, otherwise  $\mathbf{x}_{\text{skip}} = \text{Conv}_{1,s}^c(\mathbf{x})$ . The backbone is the composition of  $L$  such blocks,

$$\mathbf{H}_{\text{DBP}} = (\text{BB}_{k_L,s_L}^{c_L} \circ \dots \circ \text{BB}_{k_2,s_2}^{c_2} \circ \text{BB}_{k_1,s_1}^{c_1})(\mathbf{X}), \quad (12)$$

optionally with dilation to enlarge the temporal field without a large parameter cost. Global average pooling aggregates the

sequence, and a linear head yields scalar DBP.

Layer-level notes aligned with Fig. 4 and Table II. A strided stem enlarges the temporal field at the outset. Each stage transition halves the sequence length, progressively aggregating longer-range context, suppressing high-frequency noise, and reducing sensitivity to minor timing jitter—properties desirable for modeling diastolic decay and low-frequency envelopes. Identity shortcuts preserve gradient flow when shapes match; projection shortcuts align the residual and main paths when stride or channel count changes. Within each block, two 1D convolutions with BN and GELU provide stable nonlinear modeling on short recordings. After the final stage, GAP summarizes the sequence and the head yields the two BP estimates used for reporting.

### III. EXPERIMENTS AND RESULTS ANALYSIS

The dataset comprises 270 eight-second carotid recordings from 7 adult participants, each paired with a synchronous cuff-based blood pressure reference. Unless otherwise noted, a subject-stratified 7/2/1 split (train/validation/test) is adopted. Performance on this small, controlled cohort is reported using MAE, RMSE, Pearson’s correlation coefficient  $r$ , and Bland–Altman statistics (bias  $\pm$  SD in mmHg and limits of agreement, LoA). MAE reflects the average absolute error and is interpretable while being less sensitive to outliers. RMSE penalizes large deviations more heavily and thus reflects stability. Pearson’s  $r$  quantifies linear association but not agreement, so Bland–Altman analysis supplies bias and LoA ( $\pm 1.96 \times \text{SD}$ ), which is standard practice for assessing agreement with a clinical reference.

We use 1D-CNN for SBP and ResNet-1D for DBP as the intended pairing, and we also evaluate both models on both targets as cross-task baselines for comparison.

#### A. Measurement preparation and setup

Fig. 5 sketches the measurement scene. A vector network analyzer (Keysight FieldFox N9914A) drove a TX–RX patch-antenna pair aimed at the carotid while the participant sat with head support at a 6 cm stand-off. We recorded 8s single-tone CW trials at 2.4 GHz and logged the magnitude of  $|S_{21}(t)|$ ; a home cuff monitor provided the concurrent BP reference.

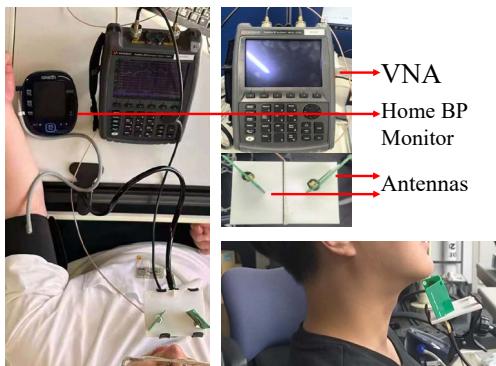


Fig. 5. Experiment scenes and devices.

TABLE III  
ACQUISITION CONFIGURATION AND PROTOCOL

Item	Setting
Instrument	Keysight FieldFox N9914A (VNA)
Measurement	$S_{21}$ magnitude, single-tone (zero-span)
Carrier $f_0$	2.4 GHz (CW; start = stop = 2.4 GHz)
IF bandwidth $B_{\text{IF}}$	3 kHz
Sweep points $N$	10001
Record length $T$	8 s
Effective rate $F_s$	1250 Hz
Sampling interval $\Delta t$	0.8 ms
TX power $P_{\text{tx}}$	3 dBm
Antennas	patch antennas (placement fixed)
Stand-off distance	6 m (to carotid sinus)
Posture	Seated, head supported
Reference Resource	Home cuff monitor, concurrent

To ensure repeatability, antenna placement/polarization and cabling were kept fixed across trials. Full instrument settings are listed in Table III.

#### B. Performance and difference comparisons

Table IV reports subject-wise metrics, with intended pairings marked within each target block. To facilitate pairwise comparison, we include a  $\Delta$  row defined as  $\Delta = \text{1D-CNN} - \text{ResNet-1D}$ . Positive values indicate the 1D-CNN entry is larger; negative values indicate the ResNet-1D entry is larger. For composite statistics, differences are computed component-wise: bias and SD separately; LoA lower and upper limits separately.

For SBP, the 1D-CNN yields lower MAE and RMSE, higher correlation  $r$ , and tighter dispersion (lower SD and narrower LoA) than ResNet-1D; the  $\Delta$  row within the SBP block reflects these consistent gains.

For DBP, the advantage shifts to ResNet-1D, which attains lower MAE/RMSE, higher  $r$ , and reduced spread relative to CNN; the corresponding  $\Delta$  indicates a small increase in error and dispersion when using CNN for DBP.

Cross-task baselines follow the expected pattern: DBP based on 1D-CNN widens LoA and increases SD, whereas SBP based on ResNet-1D attenuates localized features and lowers  $r$ . Mean bias remains close to zero for both targets, and residual error is dominated by spread.

#### C. Scatter and agreement analysis

The scatter plots in Fig. 6 cluster near the identity with least-squares slopes close to one, indicating strong linearity. Mild compression at the high-SBP end coincides with sparser coverage in that range and the larger SBP RMSE.

Bland–Altman analysis shows mean bias near zero and approximately balanced dispersion across the range, with no clear proportional bias. Outliers appear mainly near range extremes: high SBP and low DBP, which is consistent with edge-of-range scarcity and residual respiratory or motion leakage when amplitudes are small.

### IV. CONCLUSION AND FUTURE WORK

We presented a magnitude-only, single-link 2.4 GHz carotid-sensing pipeline for cuffless blood pressure estimation

TABLE IV  
SUBJECT-WISE TEST PERFORMANCE UNDER A FIXED SPLIT ( $n = 27$ ; 7 SUBJECTS).  $\Delta$ : DIFFERENCE WITHIN EACH TARGET BLOCK.

Target	Model	MAE	RMSE	$r$	bias $\pm$ SD (mmHg)	LoA
SBP	<b>1D-CNN</b>	<b>6.20</b>	<b>7.30</b>	<b>0.936</b>	<b><math>-1.18 \pm 7.34</math></b>	<b>[<math>-15.57, 13.21</math>]</b>
	ResNet-1D	6.69	8.90	0.907	$-2.32 \pm 8.75$	[ $-19.47, 14.84$ ]
	$\Delta$	-0.49	-1.60	+0.029	$+1.14 / -1.41$	[ $+3.90, -1.63$ ]
DBP	1D-CNN	4.62	5.80	0.931	$1.81 \pm 5.61$	[ $-9.19, 12.81$ ]
	<b>ResNet-1D</b>	<b>3.99</b>	<b>5.18</b>	<b>0.942</b>	<b><math>0.63 \pm 5.24</math></b>	<b>[<math>-9.65, 10.90</math>]</b>
	$\Delta$	+0.63	+0.62	-0.011	$+1.18 / +0.37$	[ $+0.46, +1.91$ ]

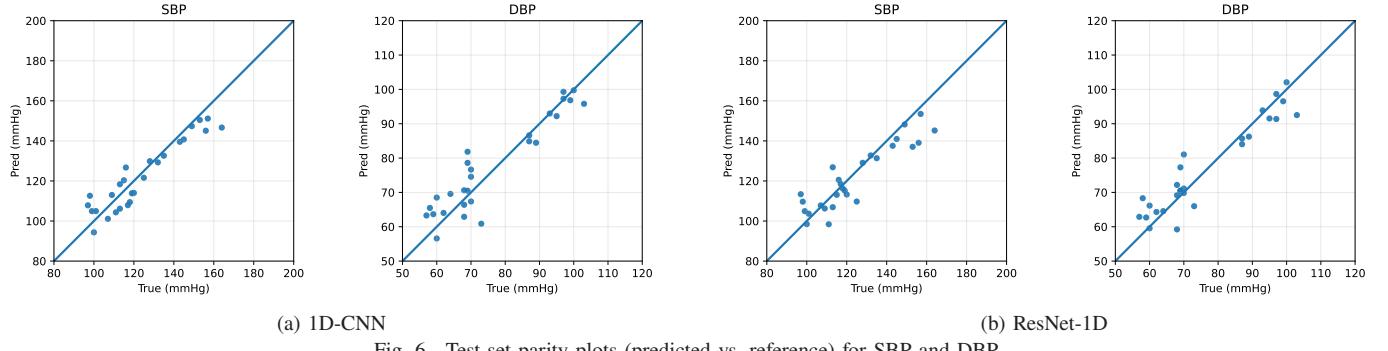


Fig. 6. Test-set parity plots (predicted vs. reference) for SBP and DBP.

based on  $|S_{21}(t)|$  in single-tone CW. With a lightweight band-pass/envelope preprocessing chain and a task-specialized pairing, the system achieves competitive subject-wise accuracy with consistent agreement and small overall bias on this small, controlled cohort, so the present results should be regarded as an initial proof of concept, which provides a basis for future studies with more compact, low-power sensing hardware.

We will (i) scale to larger and more diverse cohorts with multi-session/site validation; (ii) enrich sensing with phase or multi-frequency channels and assess robustness to posture and placement; (iii) perform systematic ablations with uncertainty quantification; and (iv) optimize hardware, firmware, and runtime toward real-time, privacy-preserving deployment.

#### ACKNOWLEDGMENT

This research was supported by the Project Research of the Waseda Research Institute for Science and Engineering, and JST ASPIRE (Grant No. JPMJAP2326), Japan.

#### REFERENCES

- [1] World Health Organization, “Global report on hypertension 2025: high stakes – turning evidence into action,” <https://www.who.int/publications/item/9789240115569>, 2025.
- [2] J. E. Dimsdale, T. V. Coy, S. Ancoli-Israel, J. Clausen, and C. C. Berry, “The effect of blood pressure cuff inflation on sleep a polysomnographic examination,” *American journal of hypertension*, vol. 6, no. 10, pp. 888–891, 1993.
- [3] C. Pena-Hernandez, K. Nugent, and M. Tuncel, “Twenty-four-hour ambulatory blood pressure monitoring,” *Journal of Primary Care & Community Health*, vol. 11, p. 2150132720940519, 2020.
- [4] D. Ray, T. Collins, S. I. Woolley, and P. V. Ponnappalli, “A review of wearable multi-wavelength photoplethysmography,” *IEEE Reviews in Biomedical Engineering*, vol. 16, pp. 136–151, 2021.
- [5] Z. Shi, T. Gu, Y. Zhang, and X. Zhang, “mmbp: Contact-free millimetre-wave radar based approach to blood pressure measurement,” in *Proceedings of the 20th ACM Conference on Embedded Networked Sensor Systems*, 2022, pp. 667–681.
- [6] L. Singh, S. You, B. J. Jeong, C. Koo, and Y. Kim, “Remote estimation of blood pressure using millimeter-wave frequency-modulated continuous-wave radar,” *Sensors*, vol. 23, no. 14, p. 6517, 2023.
- [7] H. Zhao, Z. Wang, C. Ding, B. Xue, D. Xu, Z. Wang, H. Hong, and X. Zhu, “Non-contact blood pressure estimation from the carotid pulse wave signal using a modified resnet-18 network with Ifmcw radar,” *IEEE Sensors Journal*, 2024.
- [8] Y. Qiu, X. Ma, X. Li, S. Fan, Z. Deng, and X. Huang, “Non-contact blood pressure estimation from radar signals by a stacked deformable convolution network,” *IEEE Journal of Biomedical and Health Informatics*, vol. 28, no. 8, pp. 4553–4564, 2024.
- [9] P. Wang, M. Yang, X. Zhang, J. Wang, C. Wang, and H. Jia, “Non-contact blood pressure monitoring using radar signals: A dual-stage deep learning network,” *Bioengineering*, vol. 12, no. 3, p. 252, 2025.
- [10] J. P. Mynard, A. Kondiboyina, R. Kowalski, M. M. Cheung, and J. J. Smolich, “Measurement, analysis and interpretation of pressure/flow waves in blood vessels,” *Frontiers in Physiology*, vol. 11, p. 1085, 2020.
- [11] N. Westerhof, J.-W. Lankhaar, and B. E. Westerhof, “The arterial windkessel,” *Medical & biological engineering & computing*, vol. 47, no. 2, pp. 131–141, 2009.
- [12] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [13] G. Slapničar, N. Mlakar, and M. Lustrek, “Blood pressure estimation from photoplethysmogram using a spectro-temporal deep neural network,” *Sensors*, vol. 19, no. 15, p. 3420, 2019.
- [14] C. Qin, Y. Li, C. Liu, and X. Ma, “Cuff-less blood pressure prediction based on photoplethysmography and modified resnet,” *Bioengineering*, vol. 10, no. 4, p. 400, 2023.