

# Contactless Seismocardiography via Deep Learning Radars

Unsoo Ha, Salah Assana, Fadel Adib

Massachusetts Institute of Technology

{unsoo,salah,fadel}@mit.edu

## ABSTRACT

The seismocardiogram (SCG) is a recording of a human heart's mechanical activity. It captures fine-grained cardiovascular events such as the opening and closing of heart valves and the contraction and relaxation of heart chambers. Today, SCG recordings are obtained by strapping an accelerometer at the apex of the heart to measure chest wall vibrations. These recordings can be used to diagnose and monitor various cardiovascular conditions including myocardial infarction (heart attack), coronary heart disease, and ischemia.

This paper introduces RF-SCG, a system that can capture SCG recordings without requiring any contact with the human body. The system operates by analyzing the reflections of millimeter-wave radar signals off the human body. RF-SCG can reconstruct the SCG waveform, and it can time 5 cardiovascular events within individual heartbeats with high accuracy. Our design is based on a hybrid architecture that combines signal processing with deep learning. The pipeline includes a *4D Cardiac Beamformer* that can focus on the reflections of the human heart and a deep learning pipeline (*RF-to-SCG Translator*) that can transform these reflections into SCG waveforms. Empirical evaluation with 40,000 heartbeats from 21 healthy subjects demonstrates RF-SCG's ability to robustly time five key cardiovascular events (aortic valve opening, aortic valve closing, mitral valve opening, mitral valve closing, and isovolumetric contraction) with a median error between 0.26%-1.29%.

## CCS CONCEPTS

- Applied computing → Life and medical sciences;
- Human-centered computing → Ubiquitous and mobile computing systems and tools;

## KEYWORDS

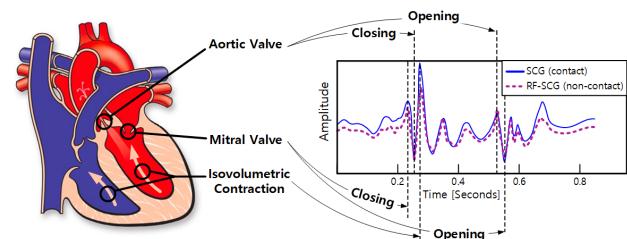
Seismocardiogram; Contactless Sensing; Wireless; Millimeter Waves; Radar; Healthcare; Well-being

### ACM Reference Format:

Unsoo Ha, Salah Assana, Fadel Adib. 2020. Contactless Seismocardiography via Deep Learning Radars . In *The 26th Annual International Conference on Mobile Computing and Networking (MobiCom '20), September 21–25, 2020, London, United Kingdom*. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3372224.3419982>

## 1 INTRODUCTION

The past decade has witnessed significant advances in using RF signals to sense people and their vital signs. Novel algorithms and software-hardware systems have enabled capturing human breathing and heart rates based on the RF signals that bounce off the human body and without requiring anybody contact [6, 58, 66]. While this research has demonstrated that RF signals carry impressive information about human vitals, it still cannot capture the level of



**Figure 1—Contactless SCG Recording.** The plot to the right shows SCG recordings of a single heartbeat obtained from an accelerometer strapped to the apex of the heart (in blue) and from RF-SCG without any body contact (in purple). RF-SCG can also precisely time five micro-cardiac movements: opening and closing of the aortic valve, opening and closing of the mitral valve, and isovolumetric contraction (of the ventricles).

detail in typical gold standard heart recordings, which are needed to understand and monitor cardiovascular conditions.

In this paper, we ask the following question: Can we use RF signals to wirelessly capture a person's seismocardiogram (SCG)? The SCG is a heart recording that is analogous to the more-commonly known electrocardiogram (ECG). In contrast to the ECG which measures the heart's electrical activity (i.e., voltage), the SCG measures the heart's mechanical activity (i.e., vibrations). Medical literature has shown that the SCG can be used to precisely time fine-grained heart activities including the opening and closing of valves, which allow blood to flow between the heart chambers and into the blood vessels [7, 12, 44, 48]. These measurements are useful in the detection and diagnosis of several cardiovascular conditions like myocardial infarction (heart attack), coronary heart disease, ischemia, and hemorrhage [38, 45, 54].

The standard approach for measuring SCG signals today relies on accelerometers that capture micro-vibrations of the chest wall [69]. The process typically requires users to take off their shirts, lie in a supine position (i.e., on their back), and affix an accelerometer near the apex of the heart using a chest strap. As a result, measuring SCG today remains intrusive and inconvenient, and it typically needs to be administered by medical practitioners in calibrated medical settings or controlled environments.

Recognizing these limitations, prior work has taken initial steps toward contactless cardiac monitoring. Early research in this space focused on measuring the heart rate or the heartbeat period [6, 40, 59, 71], but these systems could not extract micro-cardiac events *within* an individual heartbeat. More recent research [15, 32, 61] has shown how to recover micro-cardiac events, by using contactless millimeter waves. However, these prior systems still lack the granularity and/or the robustness to accurately time the majority of micro-cardiac events. Our work shares the motivation of these prior systems and aims to push the boundary of cardiac monitoring by achieving an accuracy high enough to recover micro-cardiac events that were not possible to sense before (in a contactless fashion).

We present RF-SCG, a deep-learning-based, contactless approach for measuring SCG signals that enables passive, long-term monitoring of users in everyday environments. Our approach relies on low-cost millimeter-wave radars (similar to those used in the Google Pixel 4 [19]). The radar transmits a signal and captures its reflection off the user's chest to sense cardiac micro-vibrations. Because millimeter-wave signals can traverse clothes, this approach neither requires users to take off their clothes nor requires affixing a sensor to their chest. Such an unobtrusive approach would enable lay users to routinely monitor their SCG signals (e.g., on a daily basis), and may provide early warnings of cardiovascular conditions. This capability could be particularly helpful for monitoring high-risk populations – like the elderly, neonates, or patients with arrhythmia – in their everyday environments. It may also enable *on-the-spot* heart recordings in the event of a cardiovascular emergency. For example, if someone suspects they may be suffering from a heart attack, they could use such a system to immediately measure their SCG.<sup>1</sup> Obtaining an early recording of the heart during an attack leads to better diagnosis of the type of attack and results in better treatment [27, 55]. Additionally, RF-SCG may be used by health workers to remotely monitor infectious disease patients (or burn patients) similar to how infrared thermometers enable remote monitoring of temperature while minimizing the risk of contagion.

Designing a system that wirelessly captures a person's SCG is very challenging. The system must model the relationship between the observed radio waves and cardiovascular events inside the human heart. The measured reflections change with slight changes in the user's posture or location. Even if a user slightly slouches, the location of the apex of his heart would change with respect to the millimeter-wave radar, altering the relationship between the captured reflections and the chest vibrations. Moreover, the reflected signal and the vibrations are very minute, making it difficult to robustly capture the SCG in the presence of other sources of motion and noise in the environment. This also makes it challenging to associate recorded micro-vibrations with their corresponding fine-grained cardiovascular events. This is why past attempts to wirelessly capture the human SCG either required additional contact-based sensors (e.g., ECG leads [64]) or achieved low accuracy in timing micro-cardiac events [32].

To overcome these challenges, RF-SCG introduces a hybrid architecture that combines signal processing and deep learning to robustly capture SCG recordings from millimeter-wave reflections. The pipeline consists of a series of spatio-temporal filter functions that capture the physics of RF propagation and incorporate constraints from heart physiology. RF-SCG's architecture consists of three main components, which collectively enable it to address the above challenges:

- RF-SCG's first component is a *4D Cardiac Beamformer* (detailed in §3.2) that zeroes in on the reflection coming from the apex of the heart. At a high level, this component combines 3D filters with a time-domain CNN to discover the 3D location of the heart while estimating the heart rate. This component also filters out various sources of noise and interference in space and time.

<sup>1</sup>Note that heart attacks share symptoms with other medical conditions like heartburns or panic attacks, which do not need urgent medical care [30, 49].

- RF-SCG's second component is an *RF-to-SCG Translator* (detailed in §3.3) that aims to learn a transformation the function between wireless reflections off the human chest and standard SCG recordings. Our intuition for why learning such a transformation is possible stems from the fact the radar reflections (used in RF-SCG) and accelerometer measurements (which are used in standard SCG recordings) capture chest vibrations that arise from the same underlying micro-cardiac events. Hence, during the training phase, this component uses recordings from an accelerometer placed at the apex of the heart in order to learn the transformation between radar reflections and standard SCG recordings. Once it has been trained, RF-SCG does not need the accelerometer anymore, and it uses the learned translation filters to transform the reflections to SCG recordings.
- The final component RF-SCG's pipeline performs automatic labeling of the SCG recordings (detailed in §3.4) in order to extract the timing of five fiducial points of interest: mitral valve closing, isovolumetric contraction, aortic valve opening, aortic valve closing, and mitral valve opening. This component modifies and adapts that the U-Net model [42] – which is typically used in Computer Vision to identify salient features in images [18, 36] – to identify the fiducial points of interest in 1D SCGs.

The reasons for adopting the above-principled architecture in our design is multi-fold: first, it allows us to isolate the different sources of errors and develop our model to achieve high accuracy; such flexibility would not be possible by implementing an end-to-end deep neural network as a black box, which may be too risky for critical tasks such as cardiovascular monitoring. And second, by incorporating spatio-temporal filters that reflect domain knowledge from the radar and signal processing communities, we can reduce the training time and the complexity of the deep learning architecture and initialize its tuning parameters with well-informed guesses allowing it to converge faster and to more accurate solutions.

We evaluated our design using the IWR1443BOOST millimeter-wave radar board from Texas Instruments [4]. The board has a steerable antenna array with vertical and horizontal elements and transmits FMCW chirps in the 77 GHz band. We tested RF-SCG on 21 subjects wearing regular clothing on different days and collected over 40,000 heartbeat measurements. Five fiducial points in each heartbeat were hand-labeled and inspected to ensure that the labels follow standard protocols for SCG labeling from the medical literature [7, 12, 44, 48].

Our empirical evaluation demonstrates that RF-SCG can both extract SCG recordings for qualitative review by medical professionals (similar to Fig. 1) as well as precisely time cardiovascular events. The results also show that RF-SCG generalizes well to new (unseen) subjects. Specifically:

- When RF-SCG is tested on subjects that it has not been trained on, it achieves a median error between 0.26%-1.29% in timing each of the five fiducial points.<sup>2</sup> The corresponding timing errors (of 2-12 ms) are within one to two sampling periods (i.e., quantization error) of gold standard cardiac ultrasounds used by medical practitioners in cardiovascular monitoring [41]. Moreover, the 90<sup>th</sup> percentile error of RF-SCG ranges from 3.6%-13.6% across the micro-cardiac events.

<sup>2</sup>The error here is computed with respect to a heartbeat period.

- In order to quantify RF-SCG’s ability to capture inter-subject variability, we also implemented a baseline that learns the *average* heartbeat template across subjects; the baseline then estimates the timing of micro-cardiac events by stretching the template to the heartbeat period. Unlike RF-SCG, the baseline is also provided with the ground-truth heartbeat segmentation to aid it in stretching the template. Our results demonstrate that the baseline achieves median errors of 3.1%-4.6% and 90<sup>th</sup> percentile errors of 17.7%-19.5%.<sup>3</sup> Thus, RF-SCG outperforms the baseline’s median by 3-12× and its 90<sup>th</sup> percentile by 1.3-3.7×.

Our results also demonstrate that RF-SCG is resilient to slight changes in posture (due to its 4D cardiac alignment) and can work correctly in the presence of other objects/humans moving in the environment. However, similar to past designs for RF vitals sensing, RF-SCG requires the user to remain relatively static and close to the millimeter-wave sensor (within a quarter to half a meter) [71]. To deal with scenarios where the user is further away or significantly moving, RF-SCG incorporates algorithms that automatically detect and discard the corresponding signals, enabling it to avoid reporting incorrect measurements (i.e., it favors precision over recall).

**Contributions:** RF-SCG is a system that can capture and reconstruct a human’s seismocardiogram without requiring any contact with the human body. It is based on a hybrid pipeline – of signal processing and deep learning – that consists of a series of learnable spatio-temporal filter functions which incorporate domain knowledge from RF and physiological models. The paper also contributes to a prototype implementation and evaluation, demonstrating median errors of 0.26%-1.29% in timing five key micro-cardiac events.

So far, RF-SCG has been tested on healthy subjects (as per our IRB protocol). As the research evolves, we hope it moves to clinical trials for testing on patients, and that it may be integrated with past proposals that have demonstrated the ability to diagnose and detect various cardiovascular conditions using SCG recordings (obtained from on-body accelerometers) [38, 45, 54]. Such clinical studies are beyond the scope of this paper. We believe this research direction is particularly promising since millimeter-wave radars are already incorporated in millions of manufactured phones and their adoption is likely to grow. Thus, it holds the potential of bringing powerful cardiac monitoring solutions to the hands of lay consumers.

## 2 BACKGROUND & RELATED WORK

Seismocardiography was first studied in the late 1950s by scientists who were inspired by the technology used in seismology to register underground vibration and predict earthquakes [9]. They adapted the technology to measure fluctuations of the cardiac movements using accelerometers.

SCG was first adopted by aerospace programs for monitoring crew health. It was used by Russian cosmonauts in 1963 on board the Vostok 5. Since then, it has been repeatedly used in space missions, and has been onboard the International Space Station since March 2007 as part of the Pneumocard and Sonocard experiments [34].

The medical community has invested significant effort in studying and understanding SCG recordings since their discovery. Clinical

studies have demonstrated that the SCG is more sensitive and specific than the ECG in detecting coronary heart disease during stress exercise testing [63]. Multiple projects have been dedicated to understanding the peaks and valleys of SCG recordings and to mapping these fiducial points to micro-cardiac events [7, 12, 22, 44, 48]. Various studies have demonstrated that SCG recordings can be used to diagnose and monitor various cardiovascular conditions including arrhythmia, myocardial infarction, ischemia, and hemorrhage [21, 37, 38, 43, 45, 47, 50, 51, 53, 54, 68].

Despite its clinical advantages over ECGs, the SCG was largely ignored by the medical community for over two decades due to the overly-cumbersome setups of early versions of the device [23]. Recent advances in miniaturization of accurate accelerometers have renewed interest in SCGs and led researchers to propose and develop a variety of wearable SCG sensors over the past decade [14, 57, 69].

RF-SCG builds on this rich literature and introduces a non-contact approach for recording SCGs. Our approach is motivated by recent advances in RF sensing which have demonstrated the ability to capture human vital signs (specifically breathing and heart rate) [6, 16, 33, 40, 59, 71]. These systems use RF signals (including millimeter wave signals [8, 11, 66]) that bounce off the human body in order to sense minute body movements and map them to human vitals. Prior work, however, remains too coarse and/or inaccurate for the capturing the SCG waveform. In particular, the vast majority of this work is limited to measuring the *average* heart rate or the heartbeat period [6, 32, 65, 71]. Those that can recover richer recordings either require additional contact-based sensors (e.g., ECG leads [64]) or achieve low accuracy (58%, i.e., close to a random guess [32]). Finally, two recent papers have used millimeter-wave radars to time cardiovascular events via wave modeling and signal processing [15, 61]. However, unlike RF-SCG, these systems could only capture one or two fiducial points. Specifically, [61] focuses on recovering the morphology of the two dominant heart sounds (S1 and S2) found in PCG signals, and it achieves correlation coefficients of 80-82% for each of them. More recently, [15] proposed an approach to recover ECG waveforms from radar reflections; however, the system was only evaluated qualitatively on one individual and did not report accuracy numbers.

RF-SCG extends this body of work by introducing new deep-learning-based techniques that enable capturing micro-vibrations *within a heartbeat* in order to reconstruct the entire SCG waveform and to time the associated micro-cardiac events. Quantitatively, in contrast to prior work that aims to time coarser cardiac events, RF-SCG needs to time micro-cardiac events within few milliseconds, i.e., it requires two-to-three orders of magnitude finer granularity than state-of-the-art contactless designs.

## 3 RF-SCG’S DESIGN

RF-SCG consists of three main components: the first is a 4D cardiac beamformer that focuses on reflections coming from the apex of the heart; the second transforms these reflections into SCG recordings; and, the third automatically labels the SCG recordings to extract the timings of micro-cardiac movements. This section describes each of these components in detail.

<sup>3</sup>These numbers are in line with the nominal values from the literature on inter-subject variability [26].

### 3.1 Capturing the mmWave Reflection

Before we dive into RF-SCG's core components, we provide a quick background on how RF-SCG captures millimeter-wave reflections off the human body. To capture these reflections, the sensor transmits a millimeter wave signal, which bounces off the human body and comes back to the device. The device captures these reflections and analyzes them to extract the phase, which is given by the following equation [56]:

$$\phi(t) = 2\pi \frac{d(t)}{\lambda}$$

where  $\lambda$  denotes the wavelength,  $d(t)$  denotes the distance between the sensor and the human body, and  $t$  is time. Since cardiac micro-mechanics (like opening and closing of valves) result in small vibrations on the chest wall, they impact the distance  $d(t)$  between the sensor and the human body. RF-SCG can sense these vibrations in the phase and use them to record the SCG signals.

We note two additional points about how RF-SCG captures RF reflections:

- The sensor incorporates a 2D antenna array with horizontal and vertical components. Each of the receive antennas can compute its own phase  $\phi(t)$ . This 2D array will be used by the beamformer in the next section to focus on the heart's reflection.
- The transmitted radar signal is a Frequency Modulated Carrier Wave (FMCW) signal. FMCW is a well-known radar technique that isolates reflections coming from different ranges into different buckets [5, 31].<sup>4</sup>

The above two methods are well-known techniques, which RF-SCG repurposes for zooming in on the heart's reflection and maximizing the SNR of the received signal as we explain next.

### 3.2 4D Cardiac Beamforming

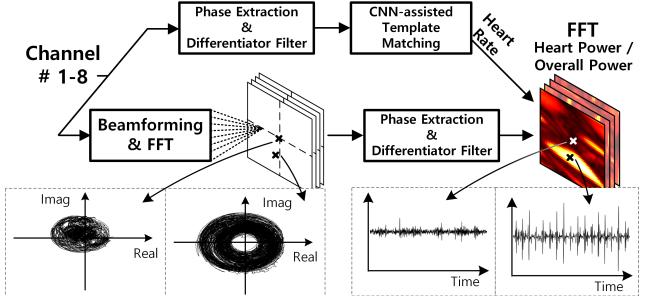
The goal of RF-SCG's first component is to focus on RF signals reflected from the human heart. Obtaining a good signal-to-noise ratio (SNR) for these reflections is important since RF-SCG needs to time micro-cardiac events which are very minute.

At first blush, one might wonder if standard antenna array beamforming can be used to focus on the direction of the reflection coming from the heart. Specifically, since RF-SCG has a 2D antenna array, it can try to use standard beam steering techniques to scan the space for the direction with the strongest reflection. Unfortunately, this standard approach to beamforming might not work (and indeed doesn't work well as we show empirically in §6). This is because the direction of the strongest reflection may not coincide with the heart's apex but rather a different location on the chest.

Instead, RF-SCG must identify the direction that has the strongest heart signal. The challenge in doing so is that it does not know exactly how the heart's reflection looks like in the first place; plus, this reflection is likely to change depending on the user, how they are sitting, and their heart rate. Moreover, the heartbeat signal itself is masked by motion due to breathing, whose magnitude is significantly larger than the micro-vibrations of cardiac movements.

To overcome this challenge, RF-SCG exploits the fact that the heartbeat signal is periodic, and it leverages this periodicity in order to identify the best direction of obtaining the corresponding periodic

<sup>4</sup>We refer the interested reader to the radar literature [31] for more details on FMCW processing.



**Figure 2—4D Cardiac Beamformer.** The 4D cardiac beamformer consists of two parallel signal-processing chains. The first chain extracts a heart rate, and the second chain projects eight channel signals into possible directions. Then, it calculates the highest cardiac power ratio with the help of both chains.

signal. Specifically, it decomposes the task of identifying the best beam with the heart rate into two subtasks. The overall architecture of its *4D Cardiac Beamformer* is shown in Fig. 2 and consists of two main processing chains:

- The top chain is used to robustly extract the heart rate. It does so by formulating the task as a 1D CNN-assisted template matching problem.<sup>5</sup>
- The bottom chain combines beamforming and FFT (via 4D projections)<sup>6</sup> in order to identify the spectral features coming from each point in 3D space. It then uses the estimated heart rate from the top chain in order to identify the correct 3D spatial beam and extract its phase.

These subproblems can be solved sequentially, as we demonstrate in the next two subsections.

#### 3.2.1 CNN-assisted Template Matching

To robustly extract the heart rate from the reflected millimeter-wave signal, RF-SCG formulates the estimation problem as a 1D matching problem. Below, we describe how it solves this problem and why this solution is more robust than simply taking an FFT.

In template matching, the goal is to extract the best template  $T$  that matches a repetitive sequence. In our scenario, the template of interest corresponds to a single heartbeat interval and the sequence is a filtered time series of phase data obtained from the millimeter wave reflection. Formally, our goal is to minimize the error  $e(i)$  between the template  $T$  and repeating segments of the input Data  $D$ :

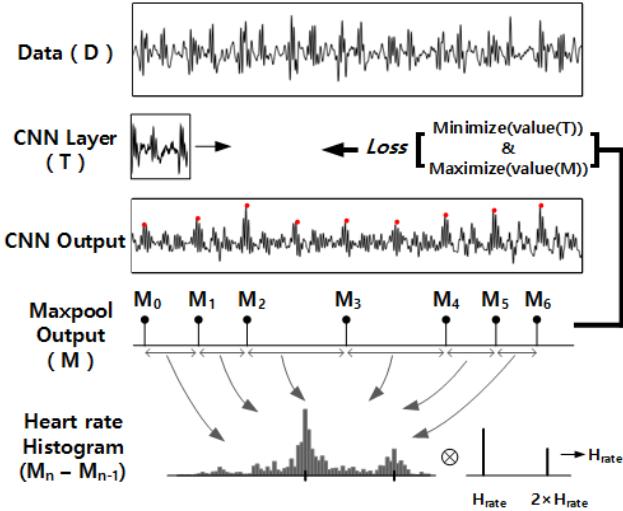
$$e(i) = |D_i - T|^2 = |D_i|^2 + |T|^2 - 2|D_i * T| \quad (1)$$

where  $D_i$  is the  $i^{th}$  segment of the input time-series data which has the same size as  $T$ , and  $*$  denotes the cross-correlation operation. By minimizing the above error function, we can recover the template  $T$ , thus learning the heartbeat template as well as the heartbeat period.

The above optimization problem can be solved using stochastic gradient descent (SGD). To efficiently solve SGD using modern computing tools, we formulate it as a convolutional neural network (CNN). Specifically, according to Eq. 1, we can minimize  $e(i)$  by maximizing the convolution  $|D_i * T|$  while minimizing  $T$ . Intuitively, we can solve this by learning the optimal convolution filter  $T$  while training the neural network to minimize an appropriate loss function that reflects minimizing the error of Eq. 1.

<sup>5</sup>As we show in the next section, this approach is much more robust than simply taking an FFT.

<sup>6</sup>The four dimensions are: 2D beamforming, FMCW, and FFT.



**Figure 3—CNN-assisted Template Matching.** The figure shows the different stages of RF-SCG’s segmentation algorithm. After the *CNN* and *Maxpool* layers, each segment between neighboring  $M$ ’s is a candidate for the heartbeat period (or heart rate). The final stage involves heart rate estimation using a histogram.

Formally, we can represent this problem using the architecture shown in Fig. 3:

- The architecture has a 1D CNN that must learn a template  $T$  as its convolution filter.
- The CNN layer is followed by a *Maxpool* layer, which tries to find the peaks at the output of the CNN layer (i.e., the convolution). Since the peaks arise whenever the template is aligned with a single heartbeat, the number of peaks is indicative of the number of heartbeats.
- The loss function can be formulated as:

$$Loss = \mathcal{L}_2(T) + \lambda \frac{1}{M}$$

where  $\mathcal{L}_2$  is *L2-norm*,  $M$  is the value of each *Maxpool* output element, and  $\lambda$  is a tuning parameter. Mathematically, the first term corresponds to a regularization term for penalizing the loss function to prevent overfitting.

**Exploiting the 1D CNN architecture to extract the heart rate.** Upon convergence, the above architecture would result in minimizing the template matching error (as per Eq. 1). It would also result in a segmentation of the time domain series into individual heartbeats. This can be seen by the peaks (labeled in red) in the CNN output in Fig. 3 as well as in the Maxpool estimates ( $M_0, M_1, \dots, M_6$ ). Moreover, the individual heartbeat period may be derived by taking the difference between any two neighboring *Maxpool* outputs ( $M_n - M_{n-1}$ ).

The last row of Fig. 3 plots the histogram of the maxpool differences. Recall that these histograms correspond to the estimated heartbeat period across all individually measured heartbeats. Interestingly, the histogram has two peaks; these two peaks correspond to both the fundamental and the second harmonic of the heartbeat frequency. Thus, we can use this histogram in order for us to estimate the heartbeat by performing a simple cross-correlation between the histogram and the harmonic impulse function [17].

Few points are worth noting about RF-SCG’s approach for extracting the heartbeat period:

- While the method uses standard machine learning techniques to estimate the heartbeat period, it does not need to be pre-trained since the training process itself can be used to extract the heart rate as described above.
- Naturally, a much simpler approach to extract the heart rate is to use an FFT and identify the peak frequency. In practice, however, this wouldn’t result in an accurate estimate as demonstrated by past work [71]. This is because the heartbeat changes from interval-to-interval, making it infeasible to robustly obtain a sharp peak at the heartbeat frequency. In contrast, since the template matching method relies on the local maximum of the CNN output in Fig. 3, it does not require the response to be perfectly periodic.
- The histogram-based method for extracting heartbeats is highly robust to noise or random sources of error since these would not contribute to or distort the dominant period/frequency.
- In principle, it may be possible to replace RF-SCG’s template matching component by other template matching schemes [62, 71]. In developing RF-SCG, we experimented with various schemes but found the CNN-based approach to be more desirable for three reasons: first, it can be easily implemented using highly optimized machine learning frameworks (e.g., Tensorflow, PyTorch); second, it is linear time; and third, it achieves high robustness and accuracy with minimal parameterization.<sup>7</sup>

### 3.2.2 Beamforming-based SNR Improvement

Now that we have a reliable algorithm to obtain heart rates from periodic signals, we can employ an algorithm that exploits this periodicity in order to extract an optimal steering direction. Fig. 2 shows the overall flow to identify the optimal projection. The input to this architecture are eight channels from our mmWave device. The signals received on each of these channels may be expressed as:

$$m_n = A_n \times e^{j\phi_n(t)} \quad (2)$$

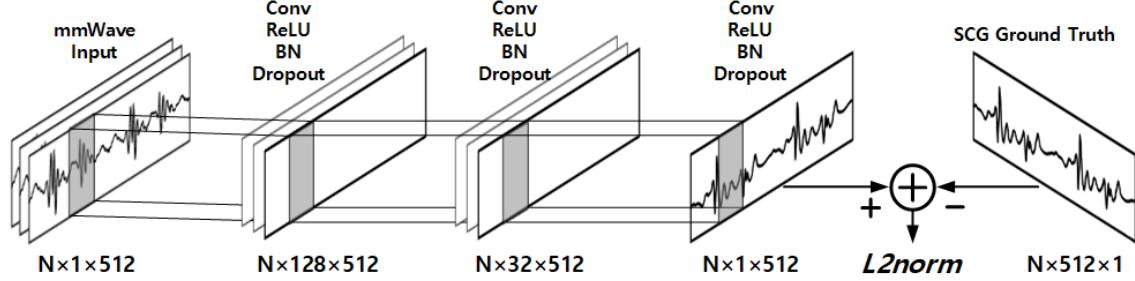
where  $A_n$  denotes the amplitude, and  $\phi_n(t)$  denotes the phase of  $n^{th}$  channel signal. The signals are fed to two different signal-processing chains: the CNN-assisted Template Matching chain and the beamforming chain.

In the first chain, RF-SCG extracts the phase  $\phi_n(t)$  from all the channels, then applies a differentiator filter to each of them. The filter serves two main purposes: first, it suppresses low-frequency respiration signals and boosts the high-frequency heart rate, thus accentuating the feature points related to micro-cardiac movements; the second purpose of this filter is to transform the phased-based distance estimate (as per Eq. 3.1) to an acceleration by designing the filter to operate as a second derivative. Thanks to the differentiator filter, the CNN-assisted Template Matching stage can easily extract histogram from all the channels and obtain a representative heart rate as described earlier.

The second chain of RF-SCG’s 4D Cardiac Beamformer applies a standard beamforming method. Specifically, the signals received on all eight channel are combined using the following equation:

$$BF(x, y) = \sum_{n=1}^8 \exp \left( \frac{-j2\pi}{\lambda} [x_{d,n} \cdot x + y_{d,n} \cdot y] \right) \times m_n \quad (3)$$

<sup>7</sup>The only hyper-parameters we needed to set were the learning rate,  $\lambda$ , and a layer size of *Maxpool* and *CNN*. The main constraint we needed to account for was to ensure that the *Maxpool* and *CNN* layers larger than one heartbeat interval, which can be easily addressed by choosing a sufficiently large number.



**Figure 4—Translation Filter.** The figure plots the different layers of RF-SCG’s RF-to-SCG Translator.  $N$  is the batch size of data set. While the length of one frame is fixed, the number of channels is different across the layers. To translate the mmWave reflection to SCG recordings, every convolution layer functions as an FIR filter. The coefficients of every CNN layer are updated based on the calculated  $L2 - norm$  between the ground truth and translated results.

where  $\lambda$  is the wavelength of transmitted RF signal,  $x_{d,n}, y_{d,n}$  are the coordinates of  $n^{th}$  antenna array,  $x, y$  are projected coordinates, and  $m_n$  is received signal from  $n^{th}$  antenna array. After applying this beamforming, we perform an FMCW projection (by taking an FFT) in order to isolate reflections coming from different distances into different frequency buckets (as described in §3.1).

To sum up, the second chain obtains projections in a 3D coordinate system ( $x, y$ , and intermediate frequency buckets). Each of these projections is a time-domain signal. It is possible to verify the quality of each projected time-domain signal by plotting it on a complex plane. To demonstrate this visually, Fig. 2 plots two potential outcomes of the projection with different projection qualities. The left plot shows messy circles around the origin without a hole at the center; this indicates destructive interference (i.e., the projection did not constructively combine the channels). The right plot shows thick concentric circle, demonstrating that the projection is close to ideal (as per Eq. 2). Subsequently, the projected signals are fed to the next stage which extracts the phase and applies a differentiator filter.<sup>8</sup>

Now that each of the two chains is completed, it is time to combine their outputs. To do this, our beamformer takes the FFT of the projected time-domain signals and computes the power ratio of the heartbeat power to the overall power across coordinates and buckets. Specifically, recall that each time-domain signal (at the output of the second chain) is parameterized by 3 coordinates. Thus, by taking the FFT of each of these signals, we obtain their time-domain spectral properties. Knowing the heart rate estimate from the output of the CNN-based template matching (as per §3.2.1), we can compute the ratio of the heartbeat power to the overall power across coordinates and buckets. This provides us with a 4D heatmap ( $x, y$ , frequency buckets, and power ratio), a sample of which is visualized to the right of Fig. 2. Using this 4D output, we can now identify the optimal 3D location as the one having the highest power ratio (i.e., where the heartbeat’s periodicity is most pronounced).

To demonstrate the value of this projection, Fig. 2 shows the time-domain phase extracted from two different locations in this projection (the two plots at the bottom right of the figure). The phase from a random coordinate in the projection (without loss of generality, the origin) barely shows any periodicity. However, the one obtained from the projection’s peak (rightmost plot) clearly shows periodic signals. This demonstrates the ability of RF-SCG’s 4D cardiac beamformer to focus on and accentuate the heart’s reflections.

<sup>8</sup>The differentiator filter is trained together with the following stages which will be discussed in §3.3. The differentiator filter is initialized with a conventional differentiator filter coefficients and will be trained to maintain its symmetry.

**Dealing with Outliers and Extraneous Movements.** After the beamforming stage, we still need to filter out corrupted data. In particular, various environmental factors (such as a subject’s excessive movement or another object coming in the way between device and subject) may cause sudden peaks or distort the received signal. Not accounting for such factors would degrade the performance of the overall system.

To identify and eliminate such outliers, RF-SCG leverages a statistical analysis called Tukey’s fence. The method identifies thresholds for rejecting outliers based on the quartiles of the time series. To implement this method, the system divides the obtained time series projections into two-second windows. Then, it can calculate the power distribution across all unit windows and extracts the first quartile ( $Q_1$ ) and third quartile ( $Q_3$ ). Based on these values, it defines a threshold  $P_{th}$  as follows:

$$P_{th} = Q_3 + 1.5 \times (Q_3 - Q_1) \quad (4)$$

Then, if the power of any unit exceeds this threshold, it is discarded.<sup>9</sup>

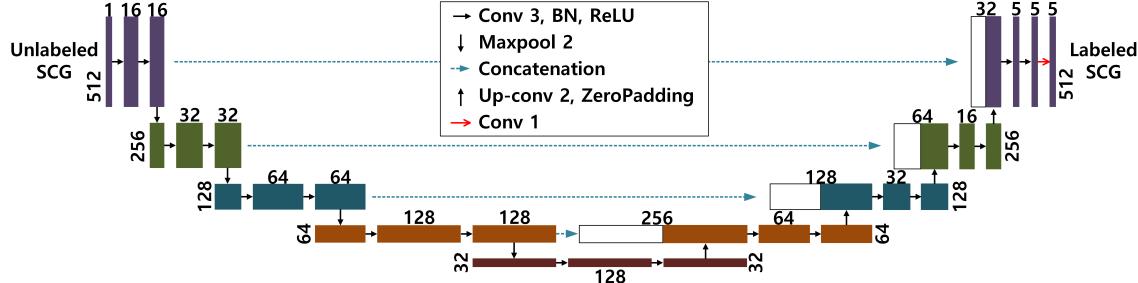
### 3.3 RF-to-SCG Translator

The goal of the translation filter is to translate the wireless reflections into meaningful clinical parameters, more specifically, to translate mmWave reflections to SCG waveforms. The easiest way to do this would be to implement an FIR filter that performs such conversion. In principle, this should be possible since both modalities are based on mechanical movements of the human heart. However, in practice, this is challenging because mmWave reflections and accelerometers are fundamentally different acquisition mechanisms, with different properties, front-ends, and channel characteristics. This makes it difficult to derive a closed-form signal transformation from one modality to the other.

To overcome this challenge, RF-SCG *learns* the translation filter, which we call an RF-to-SCG Translator. In particular, the translator learns the filter’s coefficients from mmWave and SCG data pairs. Since an FIR filter and 1-dimensional convolutional neural network (1D-CNN) have the same mathematical representation, we can think of multiple 1D-CNN layers as multiple FIR filters. This enables us to represent fine-grained cardiovascular movements by the combination of various frequency components from mmWave signals.

Fig. 4 shows the overall architecture of RF-SCG’s translation filter. The input to this filter is the millimeter-wave time-series obtained

<sup>9</sup>In our evaluation, we also applied this same method to reject outliers from the ground-truth time series obtained from an on-body accelerometer. Moreover, if a unit is determined to be an outlier in either RF-SCG or the on-body accelerometer, we eliminated it from both of them for fairness.



**Figure 5—Automatic Labeling.** The figure shows the network used for automatic labeling. Each box corresponds to a multi-channel feature map. The number of channels is denoted on top of each box. The length of each layer is provided at the left edge of the box. White boxes represent copied feature maps.

from the output of the 4D Cardiac Beamformer. The input is organized into  $N$  batch samples, 512 sample points, and 1 channel. Each CNN layer (Conv) is grouped with a rectified linear unit (ReLU) function, batch normalization (BN), and dropout layers (Dropout). ReLU functions allow the nodes to learn more complex features in the data by adding nonlinearity. Batch normalization and dropout layers help the networks generalize to diverse input patterns and prevent overfitting. Every CNN layer has the same length but different channels. To update the coefficients of CNN layers, the network uses the  $L2$ -norm between the final output and actual SCG measurement (obtained from on-body accelerometer during training phase). It is important to note that during testing phase, RF-SCG does not require the on-body accelerometer as it has already learned the translation function during training. Thus, the RF-to-SCG Translator enables it to reconstruct the SCG waveform from millimeter-wave reflections.

### 3.4 SCG Automatic Labeling

So far, we have described how RF-SCG can reconstruct SCG waveforms from RF reflections; such waveforms are useful for assessment and monitoring by expert clinicians. Next, we describe how RF-SCG goes beyond reconstructing the waveform to timing the microcardiac events in these waveforms. To do so, it implements an automatic labeling function for SCG recordings. It is worth noting that past work has made efforts to automatically label SCG recordings (albeit those obtained from on-body accelerometers) [25, 29]; however, these past methods require additional on-body ECGs leads [25] or have demonstrated high error margins [29], making them undesirable for practical/medical use.

In order to achieve high accuracy in labeling wirelessly captured SCG recordings, RF-SCG builds on recent advances in deep learning models. Specifically, it adapts the well-known Unet architecture [42] which is typically used in Computer Vision problems to identify salient features in images, and it modifies this architecture to operate with 1D SCG signals.

RF-SCG's modified Unet architecture is shown in Fig. 5, demonstrating how the architecture's parameters have been optimized for SCG labeling. Here, the height of each rectangle is the number of sample points, and the width is the number of channels for each layer. The output has 5 channels which represent probabilities of the fiducial points of interest are: aortic valve opening (AO), aortic valve closing (AC), mitral valve opening (MO), mitral valve closing (MC), and isovolumetric contraction (IM). If the probability has higher than a certain

threshold value, RF-SCG determines that the point is determined to be the corresponding micro-cardiac event.

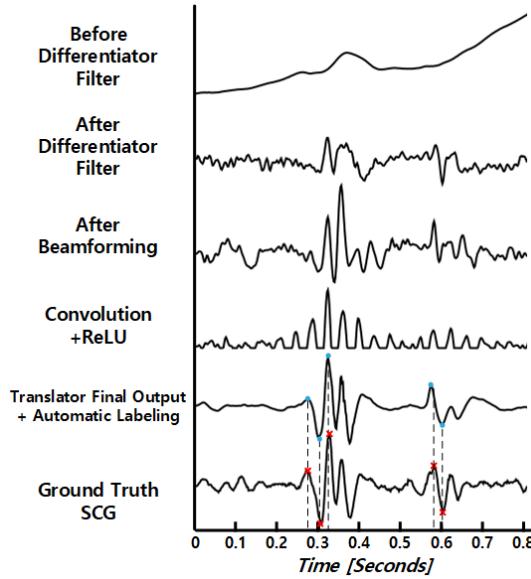
Intuitively, the left layers of the network are encoders and the right layers are decoders. By forward feeding features of each layer directly into subsequent ones (the blue dotted line), the network enables us to annotate fiducial points on the input data. For training, we manually labeled each fiducial point of every SCG data. We refer the reader to [42] for a detailed description of such methods and summarize below the basic information relevant to this paper.

- (1) The first 4 floors on the left side includes two CNNs with Batch-Norm and ReLU functions. These blocks are connected sequentially with MaxPooling layers which extract the maximum value.
- (2) The maximum value is fed to the input of lower floor block containing two CNNs with batch normalization and ReLU function.
- (3) The output from the previous layer through the deconvolution and zero padding layers are concatenated with the output from the layer from the same floor on the left side and is fed to the input of the block that includes CNNs each with the Batch-Norm and the ReLU.
- (4) The output from the previous layer through the deconvolution and zero padding layers is sequentially fed to the input of another 4 floors containing two CNNs each with Batch-Norm and ReLU.
- (5) The output from the previous layer is fed to the input of another CNN with the size of the kind of fiducial points.

As mentioned earlier, the output of this architecture are the probabilities of the fiducial points for each of the samples. Using this probability, RF-SCG can automatically label the SCG waveforms.

### 3.5 Putting It Together

To demonstrate how the entire system fits together, Fig. 6 plots the time-series outputs at different stages of RF-SCG's overall architecture for a sample experimental trial. The topmost plot shows the phase obtained from the that corresponds to the human location (i.e., this is the output of the FMCW processing). The plot contains different sources of noise and interference, especially from breathing. Specifically, the baseline fluctuation (trend) is due to the breathing signal. Thus, after applying a differentiator filter, the breathing signal is removed as shown in the second plot from the top. This plot shows finer variations, but remains quite noisy. The plot below it (i.e., third from the top) is obtained at the output of the 4D cardiac beamformer. Notice that this plot is much cleaner and has sharper



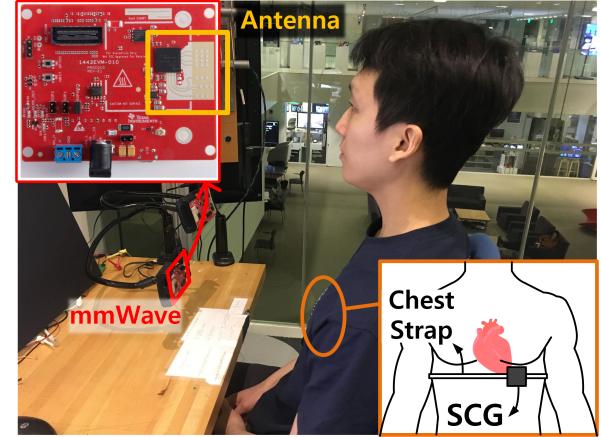
**Figure 6—Overall Signal Processing Flow.** The figure plots the waveforms at different stages of RF-SCG’s pipeline, chronologically from top to bottom. Before the differentiator filter, the signal looks irrelevant to SCG recordings. After the differentiator filter and the beamforming stage, its trend starts showing some micro-cardiac events. The figure then plots one of the channels of the RF-to-SCG translator. Then, it plots the combination of all the translator’s channels, demonstrating the ability to successfully translate the millimeter-wave reflection into SCG recordings. Finally, the five fiducial points are extracted through the automatic labeling function.

peaks and troughs with better SNR. Recall that his signal is fed to the RF-to-SCG translator which performs waveform conversion. The two plots below it show the output at an intermediate stage and the final stage of the translator. Note that the intermediate stage is difficult to interpret, as in most neural networks, but it is interesting to see that it identifies some of the sharp peaks. The plot obtained at the output of the translator (second from the bottom) is very similar to the ground-truth SCG recording (bottom plot). Moreover, the plot shows that RF-SCG is able to correctly label the fiducial points corresponding to the five micro-cardiac events of interest.

## 4 IMPLEMENTATION

**Hardware.** Architecturally, our system combines a mmWave sensor with a real-time data-capture adapter for radar sensing as shown in Fig. 7. The mmWave sensor is an IWR1443BOOST board [4] that operates at 77 GHz. The radar front-end transmits an FMCW signal (with 4 GHz of bandwidth) and consists of two linear array: horizontal (with 3-dB beamwidth of  $\pm 28^\circ$ ) and vertical/elevation (with 3dB-beamwidth of  $\pm 14^\circ$ ). The board incorporates a monolithic implementation of a 3Tx, 4Rx system. In order to maintain high temporal resolution for SCG recordings, we only used 2Tx and 4Rx. The board uses switched transmissions, which enables isolating the received signals. The acquired data is sent to a host PC with the help of a capturing software, mmWave Studios [3].

**Software.** RF-SCG’s software package was implemented in MATLAB R2019b and Python 3.7. Aside from standard FMCW pre-processing and filtering, all of RF-SCG’s software components (§3.2–§3.3) were implemented in Python. mmWave Studio was used to configure parameters of IWR1443BOOST. To enable flexibility in



**Figure 7—Experimental Environment.** The figure shows a subject sitting about 30cm away from the TI IWR1443 mmWave board’s antennas. The subject also wears a chest strap and SCG sensor under their T-shirt.

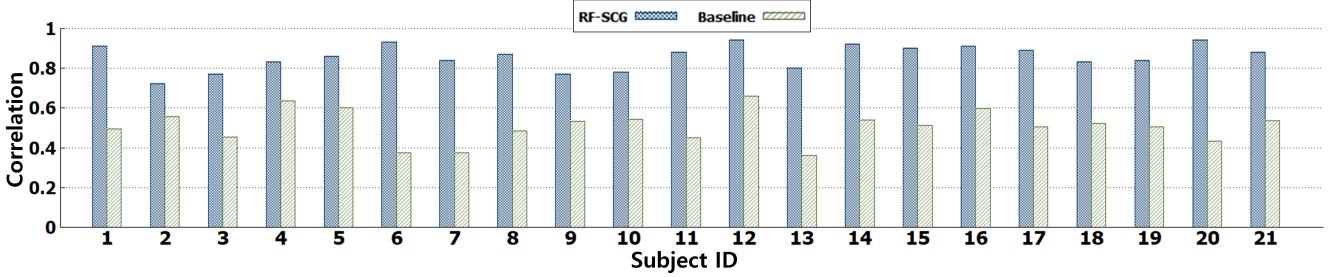
signal processing and acquisition, we wrote a packet sniffing program that captures packets received over Ethernet from the mmWave board in real-time. The extracted phase signals are preprocessed in MATLAB using a bandpass FIR filter (0.2-50 Hz) to mitigate the impact of breathing and high frequency noise.

**CNN-assisted Template Matching.** We set the following hyper-parameters for training the CNN:  $size=600$ ,  $stride=1$ ,  $zero-padding=300$ , and Maxpooling layer  $size=800$ ,  $stride=100$ . We used SGD (Stochastic Gradient Descent) optimizer with  $learning rate=1e-3$ . The processing time of the module is 0.4 sec for 3 minutes data set on a machine with an i7-9700 (3.60GHz) processor, 16GB RAM, RTX 2080.

**Beamforming.** We performed beamforming in post-processing after extracting the individual channels from the radar. The 4D Beamformer identifies the optimal projection vector. Specifically, in the beamforming function, we extracted the power ratio from a  $40 \times 40$  grid representing the  $(x, y)$  2D space, and four range buckets (i.e., FMCW frequency buckets) centered around a region of 30 cm from the sensor. The spectral density of the heart reflection is estimated in the frequency domain with a resolution of 5 beats/min.

**RF-to-SCG Translator.** We adopted three data augmentation methods (overlapping, stretching, noising) from few-shot learning machine learning models [10, 28, 60, 70]. First, we used a sliding window function to divide the input measurements into frames. The window consisted of 512 samples and was shifted by 32 samples to extract overlapping frames. Second, each frame was stretched or squeezed to represent a large number of heart rates (50-140bpm, at 10bpm intervals) using the cubic spline data interpolation [13]. Third, each frame was replicated 10 times, each with different power of random noise. This allowed us to generate more than 1,000 times the original dataset and provided additional resilience to overfitting. To maintain the width of each layer, we used the Adam optimizer and set CNN parameters  $size=[101, 101, 101]$ ,  $stride=[50, 50, 50]$ ,  $batch size = 64$ ,  $learning rate=1e-5$ ,  $stride=1$ ,  $dropout rate=0.2$ .

**SCG Automatic Labeling.** As described in §3.4, we adapted the Unet architecture for automatic labeling of SCG fiducial points, and modified the size of the layers as well as the cropping and



**Figure 8—Accuracy in Reconstructing the SCG Waveform.** The figure plots the correlation coefficient between the measured radar reflection and the ground truth across different subjects. It compares the median correlation of RF-SCG (in blue) and the baseline (in green), demonstrating that RF-SCG significantly outperforms the baseline.

concatenation functions. If an output value has a score larger than 0.5, it is marked as a valid fiducial point. We used the Adam optimizer and set *learning rate* =  $1e-5$ , *beta1* = 0.9, *beta2* = 0.999

**Ground Truth.** Following the standard approach for collecting ground truth measurements [52], our setup employs an on-body accelerometer (KXR94-2283 [2]) connected to a customized processing module (ADS1299EEGFE [1]). The sampling rate is set to 500 samples per second. To synchronize the ADS1299 with the mmWave sensor, we used the DRDY pin of the ADS1299 as a ticker to clock the SYNC\_IN pin of the mmWave sensor. The SCG recording device was only used for training the network. In a real-world use case, the user does not need to wear the device.

## 5 EXPERIMENTAL EVALUATION

**Participants.** To evaluate RF-SCG’s performance, we recruited 21 participants (5 females) between the ages of 21 and 35. None of the participants had reported known pre-existing cardiovascular conditions. Participants wore their daily attire such as T-shirts, blouses, and button downs with different fabric materials. This research protocol has been approved by our IRB.

**Experimental Environment.** We performed our experiments in an open office environment in an office building with standard furniture and in co-existence with other wireless technologies (WiFi, LTE, Bluetooth, etc.). Throughout our experimental trials, other people were walking around in the environment.

During the data collection process, subjects were asked to sit in a chair within 25-50 cm from the sensor as shown in Fig. 7. The mmWave sensors were placed in front of the subjects and aimed at the sternum. The ground truth was collected by attaching an accelerometer to the subject’s chest, slightly to the left of their sternum. The attachment was done using either adhesive tape or an elastic string wrapped around the person’s chest.

We conducted 169 experimental trials, each lasting for around 3 minutes. During these experiments, subjects were asked to breath normally, but remain quasi-static otherwise (i.e., avoid moving their arms and legs). Subjects could sit relatively comfortably and blink or swallow regularly. We collected around 40,000 heartbeats in total. Subject heart rates varied from 55-110bpm. Each heartbeat was manually labeled to time each of the 5 fiducial points of interest. The heartbeat labels were inspected multiple times by different individuals to ensure abidance to standard labeling protocols from the SCG medical literature [7, 12, 44, 48].

## 6 PERFORMANCE RESULTS

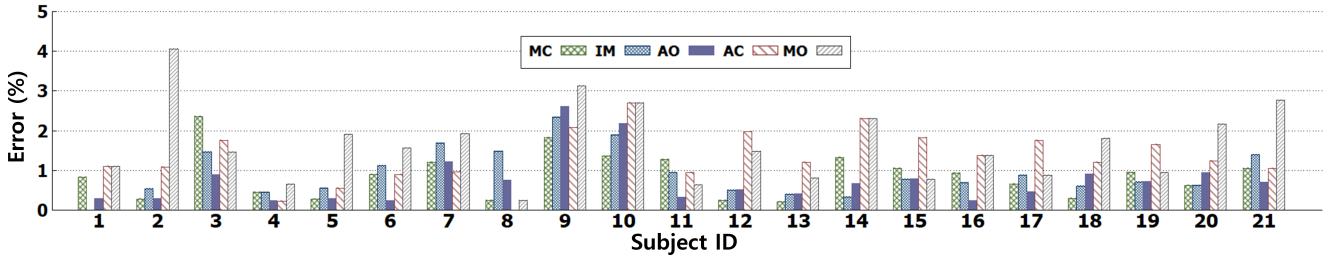
### 6.1 Results across Subjects

#### 6.1.1 Reconstructing the SCG Waveform

First, we would like to understand RF-SCG’s ability to reconstruct the SCG waveform from radar reflections. Specifically, we are interested in evaluating its ability to translate these reflections into recordings that are typically obtained from on-body accelerometers and which have been used previously in medical studies [7, 12, 22, 44, 48]. Moreover, we are interested in evaluating how well it generalizes for unseen subjects. To do so, we divided our dataset into a training set and testing set. For each subject (e.g., subject 1), the model is trained on twenty other subjects (e.g., subjects 2-21) and evaluated on the untrained subject. This 21-fold cross validation ensures that the training and testing are mutually exclusive and that the model is evaluated on subjects it hasn’t been trained on. For each subject, RF-SCG remotely captures the SCG recording and automatically extracts the five fiducial points of interest.

To quantify the similarity between RF-SCG’s continuous waveform output and the ground truth, we computed the Pearson Correlation Coefficient. This coefficient represents how closely two variables move together [39]. We also compared RF-SCG to a baseline implementation that performs standard beamforming and bandpass filtering (detailed in §4), similar to state-of-the-art prior work on contactless cardiac sensing [32, 64]. For each subject, we divided the time series into 1 second intervals and computed the correlation coefficient for each of these intervals.

Fig. 8 plots the median values of the correlation coefficients across subjects for both RF-SCG’s output (in blue) and the baseline (in green). The figure shows that RF-SCG’s correlation coefficient is higher than that of the baseline across all subjects. Moreover, while the conventional method’s highest correlation is 0.66, every subject with RF-SCG has more than 0.72 correlation coefficient. Additionally, for 7 out of 21 subjects, RF-SCG achieves a *very high positive correlation* ( $> 0.9$ ) according to [35]. It is also worth noting that other state-of-the-art work [61] has focused on recovering portions of the PCG waveform (specifically, those corresponding to two of the micro-cardiac events) and demonstrated correlations of 0.80-0.82. This shows that RF-SCG significantly outperforms state-of-the-art work and that RF-SCG’s beamformer and RF-to-SCG Translator enable it to learn representative transformations between the reflected signals and the on-body accelerations. In turn, these results demonstrate that RF-SCG is capable of reconstructing the entire SCG waveform.



**Figure 9—RF-SCG’s Overall Accuracy.** The figure plots RF-SCG’s median error in timing each of the five micro-cardiac events of interest for 21 human subjects. The dashed yellow line denotes the the best-case accuracy achievable using a gold standard cardiac ultrasound (due to its higher quantization error).

### 6.1.2 Timing Micro-Cardiac Events

Next, we would like to evaluate RF-SCG’s ability to precisely time micro-cardiac events from the extracted SCG signals, and how its accuracy varies across subjects. Thus, using the same protocol described above, we divide our measurement dataset into training and testing. This time, instead of using the continuous time-domain waveform obtained at the output of the RF-to-SCG translator, we are interested in evaluating the entire pipeline, including final stage of RF-SCG (i.e., the Automatic Labeling stage). For each detected fiducial point, we computed the error as the time difference between the RF-SCG-estimated timing and the ground truth timing, then normalized by the heartbeat period, as per the following equation:

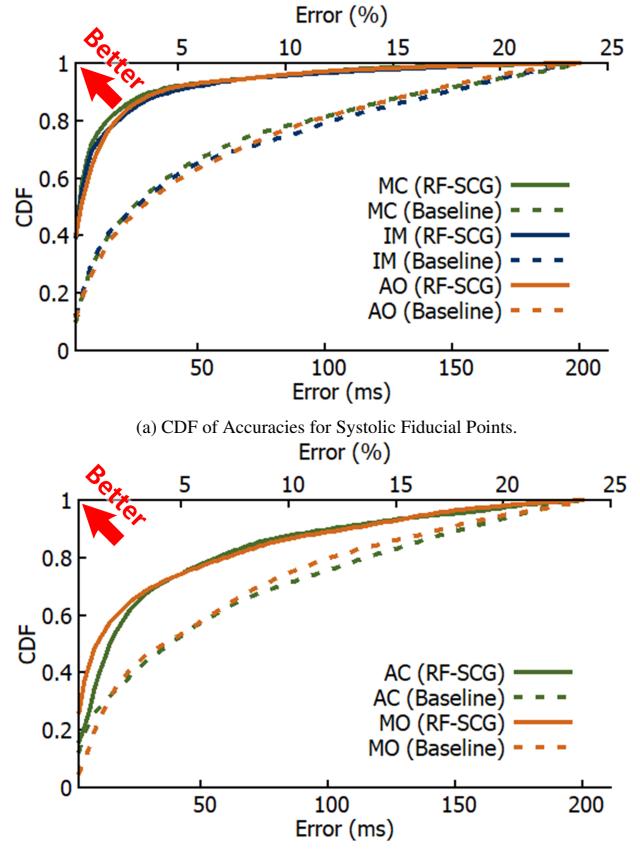
$$\text{Error} = 100 \times \left( \frac{\text{RF-SCG Estimated Time} - \text{Ground Truth Time}}{\text{Heartbeat Period}} \right) \quad (5)$$

Fig. 9 shows RF-SCG’s error across all 21 subjects. Each cluster represents a single subject and the corresponding median accuracy of each fiducial point. We also compared RF-SCG’s accuracy to the accuracy of a cardiac ultrasound device that is typically used as a gold standard for timing micro-cardiac events [41].<sup>10</sup> Specifically, the best-case accuracy is limited by the quantization error of the device which results from its sampling period (around 20 ms) and is computed similarly as per the above equation and denoted by the dashed red line in Fig. 9.<sup>11</sup> We make following remarks:

- Overall, RF-SCG can maintains low error (below 2.5%) across almost all subjects and all fiducial points.
- RF-SCG matches or exceeds the performance of the best-case accuracy for the cardiac ultrasound (dashed yellow line) for all fiducial points in seven out of the ten subjects. Moreover, if one considers only the systolic micro-cardiac movements (MC, IM, and AO), RF-SCG’s accuracy matches or exceeds the gold standard device for all subjects. The higher accuracy for systolic micro-cardiac events over the diastolic movements (AC, MO) is expected since the systolic movements are stronger (they correspond to heart contraction, while the diastolic fiducial points correspond to relaxation which is a weaker movement). This is the same reason why the stronger (lub) sound of the heartbeat comes from the systolic movements and why the corresponding peaks are larger in the SCG shown in Fig. 1 and Fig. 6.
- In some scenarios, RF-SCG’s median accuracy exceeds that best-case accuracy of the gold standard. This is because RF-SCG’s

<sup>10</sup>Note that there are cardiac ultrasounds that have higher sampling rates; however, we take the lowest sampling rate available as it remains to be part of the accepted gold standards.

<sup>11</sup>For the gold standard, the heartbeat period is taken as an average for simplicity.

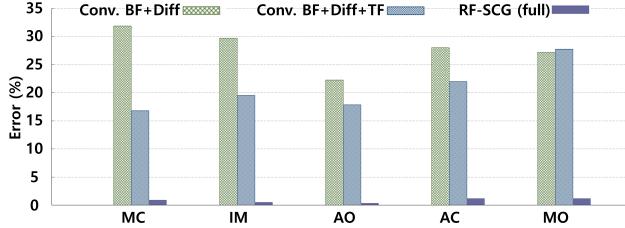


**Figure 10—CDF of Accuracies for Micro-Cardiac Events.** The figure plots the CDF of RF-SCG’s accuracy in timing (a) systolic and (b) diastolic micro-cardiac events. The solid line represents RF-SCG and the dotted line corresponds to the stretched template method.

sampling rate is much higher than that of standard cardiac ultrasound devices.

- The MO of the second subject has the highest error (around 4%). The corresponding error accounts for only two sampling periods in cardiac ultrasound.

Next, we would like to quantify the benefits of RF-SCG’s timing accuracy against a simpler baseline that learns an “average” human heartbeat template. Since state-of-the-art systems can extract at most two fiducial points (as discussed in §2), we developed a new “stretched template” baseline, which operates as follows. After segmenting the recorded signals into individual heartbeats (as per RF-SCG’s beamformer), we stretch each individual heartbeat cycle



**Figure 11—Partial Implementations.** The figure plots the accuracy of various partial implementations of RF-SCG for each of the five fiducial points of interest.

to fit in one second. Then, each subject’s fiducial points are estimated using the remaining subjects’ averaged SCG signals. Even though this baseline is already given the advantage of RF-SCG’s segmentation algorithm, it allows us to investigate the value of the remaining stages in recovering meaningful variations across heartbeats.

Fig. 10 plots the CDF of the error for RF-SCG and the baseline for each of the 5 fiducial points. Each of the CDFs in this figure corresponds to the combined error across all subjects. The plot shows both the absolute error (bottom x-axis) and relative error (top x-axis). We make the following remarks:

- RF-SCG outperforms the baseline across *all* fiducial points. The median improvement varies from around 3× for MO up to 12× for AO. This demonstrates that RF-SCG significantly outperforms the stretched template method and that the system has sufficient accuracy to capture the variability across individuals.
- The median errors in timing the fiducial points (reported as absolute error/percentage per Eq. 5) are: 2 ms/0.26% (MC), 4 ms/0.52% (IM), 4 ms/0.52% (AO), 12 ms/1.55% (AC), 10 ms/1.29% (MO). These errors are 3 – 12× smaller than the variability of individual heartbeats even for healthy individuals (as reported in the appendix).
- Similar to our earlier results, RF-SCG has higher accuracy in timing systolic (MC, IM, AO) than diastolic (AC, MO) fiducial points. This is due to the stronger SNR of systolic vibration (stronger sound of the heartbeat) than diastolic one.

## 6.2 Micro-benchmarks

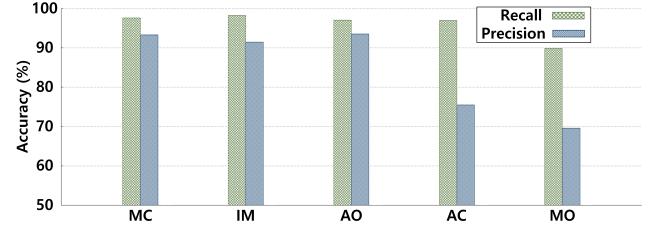
### 6.2.1 Decomposing RF-SCG’s Gains

Next, we would like to quantify the accuracy gains arising from each of RF-SCG’s subcomponents. To do so, we evaluated the accuracy of partial implementations of the overall system: (1) conventional beamforming (Conventional BF) that extracts coordinates at the highest power and a differentiator filter, (2) with translation filter, (3) RF-SCG’s full architecture with its 4D Cardiac Beamforming method and translation filter. RF-SCG’s automatic labeling is applied to all three cases.

Fig. 11 plots the accuracy for each of the above partial implementations. For simplicity, the figure only plots the results for a single subject. Similar to our evaluation in §6.1, the models were trained on all other subjects (i.e., aside from the one used in evaluation).

We make the following remarks:

- Both partial implementation achieve significantly higher errors than RF-SCG’s full implementation. Specifically, the absence of the 4D Cardiac Beamformer increases the sensor’s error to 15%–27% across all fiducial points; and the absence of the RF-to-SCG



**Figure 12—Precision and Recall.** The figure plots the precision and recall of RF-SCG’s automatic labeling component for each of the five micro-cardiac events of interest.

translator further increases its error up to 32%. This demonstrates the importance of each of the sub-components in the system’s overall performance.

- One might wonder why the partial implementations achieve any level of accuracy in timing (i.e., less than 35% error). The reason here is that even the minimally-processed wireless reflection (i.e., upon FMCW and differentiation) has some resemblance to the accelerometer-based SCG.
- Interestingly, the MO fiducial point has almost the same median error for both partial implementations. This shows that when the beamforming is extremely problematic (i.e., produces a very low SNR), the translator cannot help in reducing the error as it cannot learn a meaningful translation.

### 6.2.2 Precision & Recall of Automatic Labeling

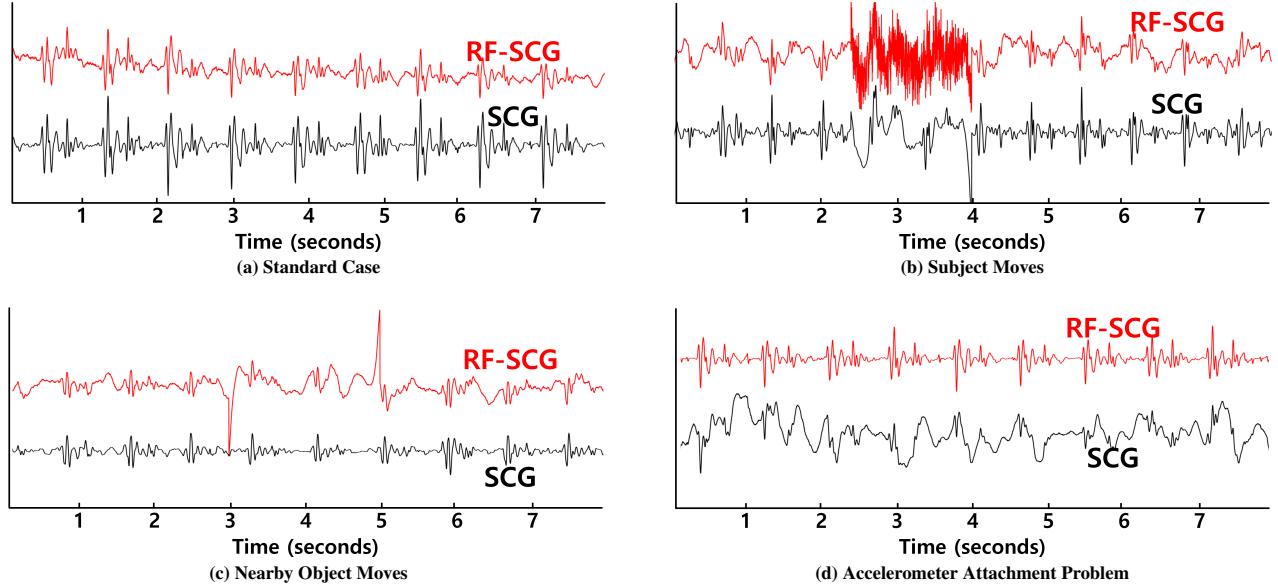
So far, our evaluation has focused on measuring the correlation of the waveform and the temporal errors between an RF-SCG-determined fiducial point and the ground-truth time of the corresponding cardiac movement. Next, we would like to gain more insight into RF-SCG’s ability to detect these fiducial points. We used the same dataset as §6.1 (i.e., testing and training is mutually exclusive on all subjects), and we computed precision and recall metrics as follows:

$$\text{Precision} = \frac{\# \text{ Correct Detections}}{\# \text{ of Reported Detections}}$$

$$\text{Recall} = \frac{\# \text{ Correct Detections}}{\# \text{ of Ground Truth Instances}}$$

Here, we define a detection to be correct if the error between the detected time by RF-SCG and the ground-truth time according to the ground truth (from accelerometer) is within a sampling error of the gold standard cardiac ultrasound. Intuitively, the precision metric scores how well the detected points match the manually labeled points, while the recall metric scores the system’s ability to detect all the manually labeled points in the SCG.

Fig. 12 shows the precision and recall metrics for each fiducial points across all the subjects. (Recall that these metrics differ from the accuracy metric reported earlier). The figure shows that RF-SCG’s precision is around or above 90% across all the fiducial points. However, only three of the fiducial points have recall scores higher than 90%. For diastolic fiducial points, the recall is lower than precision by 20.2 %. Intuitively, this means that RF-SCG cannot detect all the manually-labeled points; however, once it detects a point, it is highly likely to be correct. It is possible to make the automatic labeling component more aggressive by adjusting the loss function of the training regime and threshold of score value to favor recall. However, since the goal of RF-SCG is to sense fiducial points



**Figure 13—Modality Performance and Environmental Problems.** The figure plots the output of RF-SCG and an accelerometer-based SCG when there are environmental problems during the measurement.

within a heartbeat, it may not be necessary to extract all fiducial points from each the heartbeats since the heartbeats repeat over time. This is why we consider precision to be a better metric to evaluate the performance of the labeling function. In that sense, RF-SCG’s labeling component shows very good accuracy. Nonetheless, as the research evolves, it may be desirable to adjust the training parameters differently depending on the cardiovascular condition of interest.

#### 6.2.3 Modality Performance and Environmental Problems

In our final micro-benchmark, we are interested in understanding various sources of error for both RF-SCG and the standard method for obtaining SCG recordings (i.e., using a chest-attached accelerometer). It is important to note that since RF-SCG implements outlier rejection, it is able to eliminate most of the corrupted regions in these scenarios. Thus, we turned off the outlier rejection component of our system here in order to demonstrate the impact of such problems on the recordings.

Fig. 13 shows the output signal for both of these modalities when there are different environmental problems. Fig. 13(a) shows the standard and most common case. Here, the subject remains quasi-static, and there is no interference from surrounding objects. Fig. 13(b) shows the case when the subject rotates their body during the experiment. During the movement, both modalities experience a corrupted region, then go back to the standard case when the subject becomes quasi-static again. Fig. 13(c) shows the signals obtained when there is interference from a nearby object. Here, as the subject was sitting, another user swung their hand between the subject and the mmWave sensor to create interference. Even though RF-SCG’s signal is partially distorted, it still exhibits some of its fiducial points during the period; on the other hand, the on-body accelerometer is not affected at all by this kind of interference. Finally, Fig. 13(d) shows the case when there is an attachment problem in the conventional SCG modality. Since this modality requires the accelerometer to be firmly affixed to subject’s chest (to maintain high SNR), it suffers

when the attachment becomes loose. In contrast, RF-SCG shows very clean signals since it is based on contactless sensing.

## 7 DISCUSSION & CONCLUSION

By enabling contactless cardiovascular recording, RF-SCG marks an important step toward continuous health monitoring. A key benefit of RF-SCG is that it can work with off-the-shelf devices, using millimeter wave radars similar to those already present in millions of phones. Such a system may be incorporated into a desk screen or smartphone (e.g., Pixel 4) to enable continuous opportunistic monitoring. The system also has applications for patients with sensitive skin (e.g., burn patients).

Our evaluation in this paper has focused on healthy individuals. As the research evolves, it would be very interesting to evaluate the system’s accuracy in medical applications, which have already been verified with on-body SCG recordings. Such applications include detecting cardiovascular diseases (such coronary artery disease, ischemia, valvular or structural heart diseases, heart failure) [21, 24, 45, 46, 67] as well as long-term monitoring of patients with cardiovascular conditions (such as arrhythmia) [20, 38]. It would also be interesting to expand the system’s capabilities so that it can capture SCG recordings from further away subjects as well as those who are moving around in their everyday environments.

More generally, we hope that as this line of work evolves, it will help bring passive, long-term cardiovascular monitoring to our everyday environments.

**Acknowledgments.** We thank Manish Singh and Youngjae Min for their help in early developments of the system, and than Dr. Jonathan Ludmir for helpful discussions on cardiovascular monitoring. We also thank the Signal Kinetics group and the anonymous MobiCom reviewers and shepherd for their helpful feedback on the manuscript. This research is supported by an NSF CAREER Award and the MIT Media Lab.

## REFERENCES

[1] Ads1299 performance demonstration kit. <https://www.kionix.com/product/KXR94-2283>. Accessed: 2020-03-25.

[2] Kxr94-2283. <http://www.ti.com/tool/ADS1299EEGFE-PDK>. Accessed: 2020-03-25.

[3] mmwave studios. <http://www.ti.com/tool/MMWAVE-STUDIO>. Accessed: 2020-03-25.

[4] Texas instruments iwr1443boost. <https://www.ti.com/store/ti/en/p/product/?p=IWR1443BOOST>. Accessed: 2020-03-22.

[5] ADIB, F., KABELAC, Z., KATABI, D., AND MILLER, R. C. 3d tracking via body radio reflections. In *Usenix NSDI* (2014).

[6] ADIB, F., MAO, H., KABELAC, Z., KATABI, D., AND MILLER, R. C. Smart homes that monitor breathing and heart rate. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems* (2015), pp. 837–846.

[7] AKHBARDEH, A., TAVAKOLIAN, K., GUREV, V., LEE, T., NEW, W., KAMINSKA, B., AND TRAYANOVA, N. Comparative analysis of three different modalities for characterization of the seismocardiogram. In *2009 Annual International Conference of the IEEE Engineering in Medicine and Biology Society* (2009), IEEE, pp. 2899–2903.

[8] ALIZADEH, M., SHAKER, G., AND SAFAVI-NAEINI, S. Remote heart rate sensing with mm-wave radar. In *2018 18th International Symposium on Antenna Technology and Applied Electromagnetics (ANTEM)* (2018), IEEE, pp. 1–2.

[9] BAEVSKII, R., EGOROV, A., AND KAZARIAN, L. The method of seismocardiography. *Kardiologia* 18 (1964), 87–89.

[10] CHEN, Z., FU, Y., ZHANG, Y., JIANG, Y.-G., XUE, X., AND SIGAL, L. Multi-level semantic feature augmentation for one-shot learning. *IEEE Transactions on Image Processing* 28, 9 (2019), 4594–4605.

[11] CHURKIN, S., AND ANISHCHENKO, L. Millimeter-wave radar for vital signs monitoring. In *2015 IEEE International Conference on Microwaves, Communications, Antennas and Electronic Systems (COMCAS)* (2015), IEEE, pp. 1–4.

[12] CROW, R. S., HANNAN, P., JACOBS, D., HEDQUIST, L., AND SALERNO, D. M. Relationship between seismocardiogram and echocardiogram for events in the cardiac cycle. *American journal of noninvasive cardiology* 8 (1994), 39–46.

[13] DE BOOR, C., DE BOOR, C., MATHÉMATICIEN, E.-U., DE BOOR, C., AND DE BOOR, C. *A practical guide to splines*, vol. 27. Springer-Verlag New York, 1978.

[14] DI RIENZO, M., VAINI, E., CASTIGLIONI, P., MERATI, G., MERIGGI, P., PARATI, G., FAINI, A., AND RIZZO, F. Wearable seismocardiography: Towards a beat-by-beat assessment of cardiac mechanics in ambulant subjects. *Autonomic Neuroscience* 178, 1-2 (2013), 50–59.

[15] DONG, S., ZHANG, Y., MA, C., LV, Q., LI, C., AND RAN, L. Cardiogram detection with a millimeter-wave radar sensor. In *2020 IEEE Radio and Wireless Symposium (RWS)* (2020), IEEE, pp. 127–129.

[16] FANG, B., LANE, N. D., ZHANG, M., BORAN, A., AND KAWSAR, F. Bodyscan: Enabling radio-based sensing on wearable devices for contactless activity and vital sign monitoring. In *Proceedings of the 14th annual international conference on mobile systems, applications, and services* (2016), pp. 97–110.

[17] GALLAGER, R. G. *Stochastic Processes: Theory for Applications*. Cambridge university press, 2013.

[18] HAN, L., LI, X., AND DONG, Y. Convolutional edge constraint-based u-net for salient object detection. *IEEE Access* 7 (2019), 48890–48900.

[19] HILDENBRAND, J. Here's how the Pixel 4's Soli radar works and why Motion Sense has so much potential. Android Central, 2019. <https://www.androidcentral.com/how-does-googles-soli-chip-work>.

[20] HURNANEN, T., LEHTONEN, E., TADI, M. J., KUUSELA, T., KIVINIEMI, T., SARASTE, A., VASANKARI, T., AIRAKSINEN, J., KOIVISTO, T., AND PÄNKÄÄLÄ, M. Automated detection of atrial fibrillation based on time-frequency analysis of seismocardiograms. *IEEE journal of biomedical and health informatics* 21, 5 (2016), 1233–1241.

[21] IFTIKHAR, Z., LAHDENOJA, O., TADI, M. J., HURNANEN, T., VASANKARI, T., KIVINIEMI, T., AIRAKSINEN, J., KOIVISTO, T., AND PÄNKÄÄLÄ, M. Multiclass classifier based cardiovascular condition detection using smartphone mechanocardiography. *Scientific reports* 8, 1 (2018), 1–14.

[22] INAN, O., ETEMADI, M., WIARD, R., GIOVANGRANDI, L., AND KOVACS, G. Robust ballistocardiogram acquisition for home monitoring. *Physiological measurement* 30, 2 (2009), 169.

[23] INAN, O. T., MIGEOTTE, P.-F., PARK, K.-S., ETEMADI, M., TAVAKOLIAN, K., CASANELLA, R., ZANETTI, J., TANK, J., FUNTOVA, I., PRISK, G. K., ET AL. Ballistocardiography and seismocardiography: A review of recent advances. *IEEE journal of biomedical and health informatics* 19, 4 (2014), 1414–1427.

[24] JAVAID, A. Q., DOWLING, S., ETEMADI, M., HELLER, J. A., ROY, S., KLEIN, L., AND INAN, O. T. Quantification of posture induced changes in wearable seismocardiogram signals for heart failure patients. In *2016 Computing in Cardiology Conference (CinC)* (2016), IEEE, pp. 777–780.

[25] KHOSROW-KHavar, F., TAVAKOLIAN, K., BLABER, A. P., ZANETTI, J. M., FAZEL-REZAI, R., AND MENON, C. Automatic annotation of seismocardiogram with high-frequency precordial accelerations. *IEEE journal of biomedical and health informatics* 19, 4 (2014), 1428–1434.

[26] KORTEKAAS, M. C., VAN VELZEN, M. H., GRÜNE, F., NIEHOF, S. P., STOLKER, R. J., AND HUYGEN, F. J. Small intra-individual variability of the pre-ejection period justifies the use of pulse transit time as approximation of the vascular transit. *PloS one* 13, 10 (2018), e0204105.

[27] KORZENIOWSKA-KUBACKA, I., BILIŃSKA, M., AND PIOTROWICZ, R. Usefulness of seismocardiography for the diagnosis of ischemia in patients with coronary artery disease. *Annals of noninvasive electrocardiology* 10, 3 (2005), 281–287.

[28] KRIZHEVSKY, A., SUTSKEVER, I., AND HINTON, G. E. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems* (2012), pp. 1097–1105.

[29] LAURIN, A., KHOSROW-KHavar, F., BLABER, A., AND TAVAKOLIAN, K. Accurate and consistent automatic seismocardiogram annotation without concurrent ecg. *Physiological measurement* 37, 9 (2016), 1588.

[30] LETTER, H. H. *Chest pain: A heart attack or something else?* Harvard Health Publishing, 2019. <https://www.health.harvard.edu/heart-health/chest-pain-a-heart-attack-or-something-else>.

[31] LEVANON, N. Radar principles. *New York, Wiley-Interscience*, 1988, 320 p. (1988).

[32] LI, Y., XIA, Z., AND ZHANG, Y. Standalone systolic profile detection of non-contact scg signal with lstm network. *IEEE Sensors Journal* 20, 6 (2019), 3123–3131.

[33] LIN, F., SONG, C., ZHUANG, Y., XU, W., LI, C., AND REN, K. Cardiac scan: A non-contact and continuous heart-based user authentication system. In *Proceedings of the 23rd Annual International Conference on Mobile Computing and Networking* (2017), pp. 315–328.

[34] MIGEOTTE, P.-F., DE RIDDER, S., TANK, J., PATTYN, N., FUNTOVA, I., BAEVSKY, R., NEYT, X., AND PRISK, G. K. Three dimensional ballisto-and seismo-cardiography: Hij wave amplitudes are poorly correlated to maximal systolic force vector. In *2012 Annual International Conference of the IEEE Engineering in Medicine and Biology Society* (2012), IEEE, pp. 5046–5049.

[35] MUKAKA, M. M. A guide to appropriate use of correlation coefficient in medical research. *Malawi medical journal* 24, 3 (2012), 69–71.

[36] OKTAY, O., SCHLEMPER, J., FOLGOC, L. L., LEE, M., HEINRICH, M., MI-SAWA, K., MORI, K., McDONAGH, S., HAMMERLA, N. Y., KAINZ, B., ET AL. Attention u-net: Learning where to look for the pancreas. *arXiv preprint arXiv:1804.03999* (2018).

[37] PANDIA, K., INAN, O. T., KOVACS, G. T., AND GIOVANGRANDI, L. Extracting respiratory information from seismocardiogram signals acquired on the chest using a miniature accelerometer. *Physiological measurement* 33, 10 (2012), 1643.

[38] PÄNKÄÄLÄ, M., KOIVISTO, T., LAHDENOJA, O., KIVINIEMI, T., SARASTE, A., VASANKARI, T., AND AIRAKSINEN, J. Detection of atrial fibrillation with seismocardiography. In *2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)* (2016), IEEE, pp. 4369–4374.

[39] PEARSON, K. Note on regression and inheritance in the case of two parents. *Proceedings of the Royal Society of London* 58 (1895), 240–242.

[40] PETKIE, D. T., BENTON, C., AND BRYAN, E. Millimeter wave radar for remote measurement of vital signs. In *2009 IEEE Radar Conference* (2009), IEEE, pp. 1–3.

[41] PETROS NIHOYANNOPOULOS, J. K. *Echocardiography*. Springer, 2018.

[42] RONNEBERGER, O., FISCHER, P., AND BROX, T. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention* (2015), Springer, pp. 234–241.

[43] SAHOO, P. K., THAKKAR, H. K., LIN, W.-Y., CHANG, P.-C., AND LEE, M.-Y. On the design of an efficient cardiac health monitoring system through combined analysis of ecg and scg signals. *Sensors* 18, 2 (2018), 379.

[44] SALERNO, D. M. Seismocardiography: A new technique for recording cardiac vibrations. concept, method, and initial observations. *Journal of cardiovascular technology* 9, 2 (1990), 111–118.

[45] SALERNO, D. M., AND ZANETTI, J. Seismocardiography for monitoring changes in left ventricular function during ischemia. *Chest* 100, 4 (1991), 991–993.

[46] SALERNO, D. M., ZANETTI, J. M., GREEN, L. A., MOONEY, M. R., MADISON, J. D., AND VAN TASSEL, R. A. Seismocardiographic changes associated with obstruction of coronary blood flow during balloon angioplasty. *The American journal of cardiology* 68, 2 (1991), 201–207.

[47] SOLAR, B. E., TAEBI, A., AND MANSY, H. A. Classification of seismocardiographic cycles into lung volume phases. In *2017 IEEE Signal Processing in Medicine and Biology Symposium (SPMB)* (2017), IEEE, pp. 1–2.

[48] SØRENSEN, K., SCHMIDT, S. E., JENSEN, A. S., SØGAARD, P., AND STRUIJK, J. J. Definition of fiducial points in the normal seismocardiogram. *Scientific reports* 8, 1 (2018), 1–11.

[49] STAFF, M. C. Heartburn or heart attack: When to worry. Mayo Clinic, 2017. <https://www.mayoclinic.org/diseases-conditions/heartburn/in-depth/heartburn-gerd/art-20046483>.

[50] TADI, M. J., LEHTONEN, E., HURNANEN, T., KOSKINEN, J., ERIKSSON, J., PÄNKÄÄLÄ, M., TERÄS, M., AND KOIVISTO, T. A real-time approach for heart rate monitoring using a hilbert transform in seismocardiograms. *Physiological*

*measurement* 37, 11 (2016), 1885.

[51] TAEBI, A., AND MANSY, H. A. Grouping similar seismocardiographic signals using respiratory information. In *2017 IEEE Signal Processing in Medicine and Biology Symposium (SPMB)* (2017), IEEE, pp. 1–6.

[52] TAEBI, A., SOLAR, B. E., BOMAR, A. J., SANDLER, R. H., AND MANSY, H. A. Recent advances in seismocardiography. *Vibration* 2, 1 (2019), 64–86.

[53] TAVAKOLIAN, K. *Characterization and analysis of seismocardiogram for estimation of hemodynamic parameters*. PhD thesis, Applied Science: School of Engineering Science, 2010.

[54] TAVAKOLIAN, K., BLABER, A. P., AKHBARDEH, A., NGAI, B., AND KAMINSKA, B. Estimating cardiac stroke volume from the seismocardiogram signal. *CMBS Proceedings* 33 (2010).

[55] TAVAKOLIAN, K., BLABER, A. P., NGAI, B., AND KAMINSKA, B. Estimation of hemodynamic parameters from seismocardiogram. In *2010 Computing in Cardiology* (2010), IEEE, pp. 1055–1058.

[56] TSE, D., AND VISWANATH, P. *Fundamentals of wireless communication*. Cambridge university press, 2005.

[57] WAHLSTRÖM, J., SKOG, I., HÄNDEL, P., KHOSROW-KHavar, F., TAVAKOLIAN, K., STEIN, P. K., AND NEHORAI, A. A hidden markov model for seismocardiography. *IEEE Transactions on Biomedical Engineering* 64, 10 (2017), 2361–2372.

[58] WANG, H., ZHANG, D., MA, J., WANG, Y., WANG, Y., WU, D., GU, T., AND XIE, B. Human respiration detection with commodity wifi devices: do user location and body orientation matter? In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (2016), pp. 25–36.

[59] WANG, X., YANG, C., AND MAO, S. Phasebeat: Exploiting csi phase data for vital sign monitoring with commodity wifi devices. In *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)* (2017), IEEE, pp. 1230–1239.

[60] WANG, Y., YAO, Q., KWOK, J. T., AND NI, L. M. Generalizing from a few examples: A survey on few-shot learning. *ACM Computing Surveys (CSUR)* 53, 3 (2020), 1–34.

[61] WILL, C., SHI, K., SCHELLENBERGER, S., STEIGLEDER, T., MICHLER, F., FUCHS, J., WEIGEL, R., OSTGATHE, C., AND KOELPIN, A. Radar-based heart sound detection. *Scientific reports* 8, 1 (2018), 1–14.

[62] WILL, C., SHI, K., WEIGEL, R., AND KOELPIN, A. Advanced template matching algorithm for instantaneous heartbeat detection using continuous wave radar systems. In *2017 First IEEE MTT-S International Microwave Bio Conference (IMBIOC)* (2017), IEEE, pp. 1–4.

[63] WILSON, R. A., BAMRAH, V. S., LINDSAY JR, J., SCHWAIGER, M., AND MORGANROTH, J. Diagnostic accuracy of seismocardiography compared with electrocardiography for the anatomic and physiologic diagnosis of coronary artery disease during exercise testing. *The American journal of cardiology* 71, 7 (1993), 536–545.

[64] XIA, Z., SHANDHI, M. M. H., INAN, O. T., AND ZHANG, Y. Non-contact sensing of seismocardiogram signals using microwave doppler radar. *IEEE Sensors Journal* 18, 14 (2018), 5956–5964.

[65] XIA, Z., SHANDHI, M. M. H., LI, Y., INAN, O., AND ZHANG, Y. The delineation of fiducial points for non-contact radar seismocardiogram signals without concurrent ecg. *IEEE Journal of Biomedical and Health Informatics* (2020).

[66] YANG, Z., PATHAK, P. H., ZENG, Y., LIRAN, X., AND MOHAPATRA, P. Monitoring vital signs using millimeter wave. In *Proceedings of the 17th ACM International Symposium on Mobile Ad Hoc Networking and Computing* (2016), ACM, pp. 211–220.

[67] YAO, J., TRIDANDAPANI, S., WICK, C. A., AND BHATTI, P. T. Seismocardiography-based cardiac computed tomography gating using patient-specific template identification and detection. *IEEE Journal of Translational Engineering in Health and Medicine* 5 (2017), 1–14.

[68] ZAKERI, V., AKHBARDEH, A., ALAMDARI, N., FAZEL-REZAI, R., PAUKKUNEN, M., AND TAVAKOLIAN, K. Analyzing seismocardiogram cycles to identify the respiratory phases. *IEEE Transactions on Biomedical Engineering* 64, 8 (2016), 1786–1792.

[69] ZANETTI, J. M., AND TAVAKOLIAN, K. Seismocardiography: Past, present and future. In *2013 35th annual international conference of the IEEE engineering in medicine and biology society (EMBC)* (2013), IEEE, pp. 7004–7007.

[70] ZEILER, M. D., AND FERGUS, R. Visualizing and understanding convolutional networks. In *European conference on computer vision* (2014), Springer, pp. 818–833.

[71] ZHAO, M., ADIB, F., AND KATABI, D. Emotion recognition using wireless signals. In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking* (2016), pp. 95–108.

## A SUBJECT VARIABILITY

We extracted the standard deviation of each fiducial point from each subject, and added it to Table 1. Note that this is after normalization by stretching, demonstrating that even for individual subjects, there is sufficient variation in the fiducial points.

	MC(ms)	IM(ms)	AO(ms)	AC(ms)	MO(ms)
<b>S1</b>	20.05	27.95	25.35	41.55	35.2
<b>S2</b>	20.55	27.9	30.6	33.75	36.25
<b>S3</b>	19.15	25.75	27.45	40.25	32.95
<b>S4</b>	25.05	31.25	26.95	40.8	35.15
<b>S5</b>	22.3	28.15	32.35	36.2	37.1
<b>S6</b>	19.05	28.95	29	34.15	38.4
<b>S7</b>	22	23.55	24.75	40.2	32.45
<b>S8</b>	26.55	24.75	27.3	34.4	38.5
<b>S9</b>	22.65	29.35	25.75	38.2	39.35
<b>S10</b>	23.75	28.45	28.7	37.05	32.3
<b>S11</b>	20.05	28.6	23.9	33	37.2
<b>S12</b>	22.6	25.15	30.8	41.95	33.75
<b>S13</b>	26.9	23.95	23.85	36.15	36.1
<b>S14</b>	28.1	26.1	29.15	33.25	40.65
<b>S15</b>	21.1	31.7	23.95	37.1	39.45
<b>S16</b>	26.75	26.15	25.9	42.3	37.55
<b>S17</b>	20.15	31.35	26.45	34.75	35.6
<b>S18</b>	28.7	30.5	23.65	34.15	32
<b>S19</b>	20.75	27	24.65	34.5	32.8
<b>S20</b>	26.4	30.75	29	40.3	34.2
<b>S21</b>	19.45	24.9	23	38.65	38.35

Table 1—Standard deviation of each fiducial point from each subject