RF-SCG: Contactless Seismocardiography via Deep Learning Radars

Presentor: Kechen Liu

Background and origins

1950s: SCG emerged from seismology-inspired techniques. Scientists adapted vibration measurement methods to track cardiac micro-movements via accelerometers.

Space adoption: First used by Russian cosmonauts in 1963 (Vostok 5) to monitor crew health. Since then, it has been used in multiple missions, including International Space Station since 2007.

Why SCG Matters/Clinical Importance of SCG

- SCG ≠ ECG: SCG records the heart's mechanical activity (chest-wall vibrations), while ECG records electrical activity.

- Micro-event timing: SCG can precisely time valve events (AO/AC/MO/MC, IM).

- Current limitation: Standard SCG needs chest-mounted accelerometers (shirt off, supine, strapped sensor), intrusive and usually clinic-only. Bulky setup.

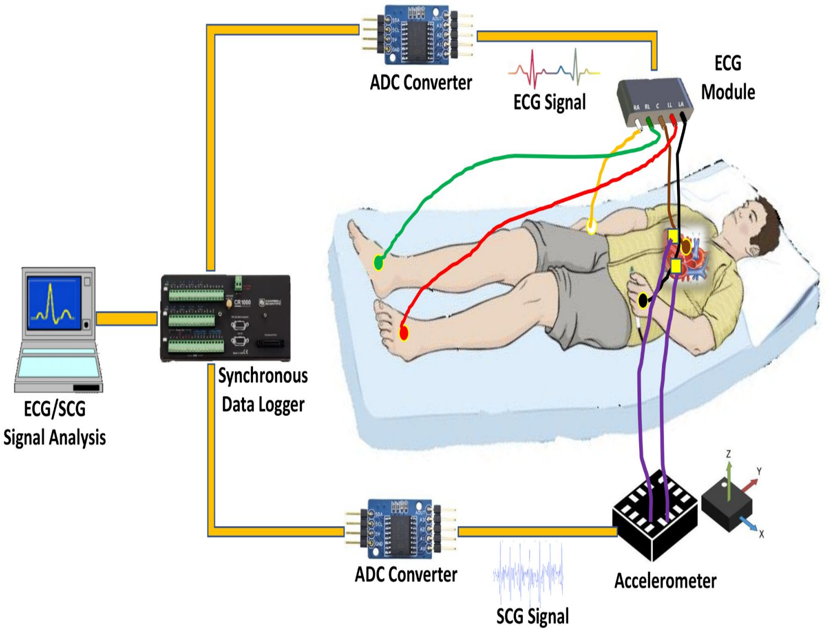

Limitations of Prior Work

Recent advances in RF sensing (including mmWave) demonstrated the ability to capture human vital signs (breathing, heart rate) contactlessly

However, limitations of past RF work:

Mostly restricted to heart rate/period, too coarse for SCG

Attempts at richer recordings often required contact-based sensors (e.g., ECG) or had low accuracy (~58%, near random)

The advantages of RF-SCG

Passive and unobtrusive: Works through clothing, no straps or electrodes needed.

Everyday use: Can enable daily monitoring at home, early warnings for high-risk groups (elderly, neonates, arrhythmia patients).

Emergency use: Quick on-the-spot SCG recording during suspected heart attacks.

Public health: Could support remote monitoring (e.g., infectious patients, burn patients) without physical contact.

2.1 mmWave reflection capture

The mmWave radar transmits signals that bounce off the chest and return with phase shifts related to chest wall vibrations.

Cardiac micro-mechanics (valve motions, contractions) cause very tiny variations in d(t), which are reflected in the phase.

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

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

2.1.2 Hardware Side

Two features in the radar hardware make this usable:

1. **2D Antenna Array**
   - ➢ Radar has multiple antennas arranged in horizontal + vertical arrays.
   - ➢ Each antenna records its own phase stream.
   - ➢ Key: allows focusing on heart's apex reflection vs other parts of chest.
2. **FMCW (Frequency-Modulated Continuous Wave)**
   - ➢ Radar sweeps frequencies in a chirp.
   - ➢ When reflected, different distances → different frequency "buckets".
   - ➢ This separates heart reflections (~30 cm away) from othe objects/walls.

Together: FMCW isolates the right range, antenna array isolates the right direction.

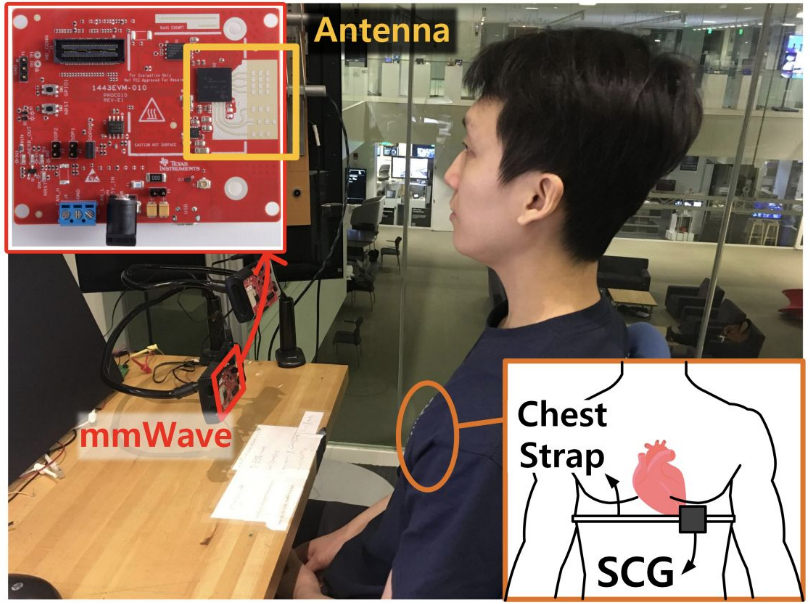

Figure 7—Experimental Environment. The figure shows a subject sitting about 30cm away from the TI IWR1443 mmWave board's antennas. The subject also wears a chest strap and SCG sensor under their T-shirt.

2.2 4D Cardiac Beamforming

Two parallel signal processing chain:

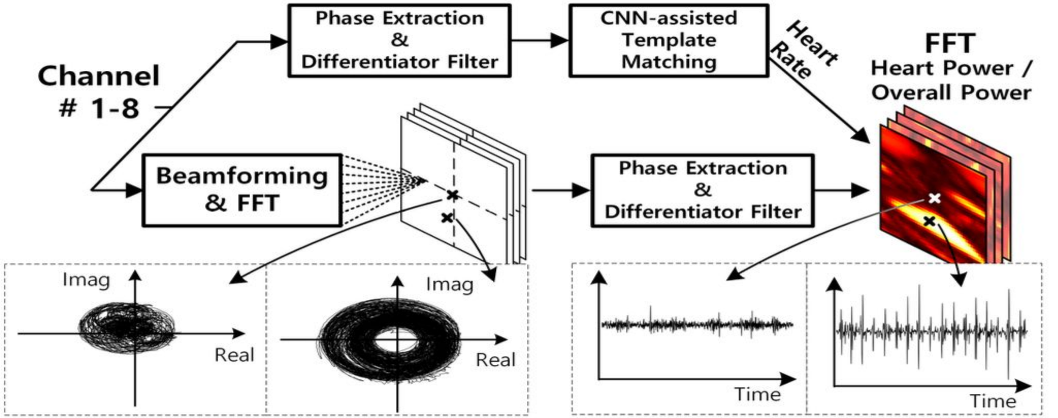

Figure 2—4D Cardiac Beamformer. The 4D cardiac beamformer consists of two parallel signal-processing chains. The first chain extracts a heart rate, and the second chain project eight channel signals into possible directions. Then, it calculates the highest cardiac power ratio with the help of both chains.

- Top chain: heart rate estimation
- 1D Conv = learned heartbeat template.
- MaxPool = peak picker; histogram = robust HR.

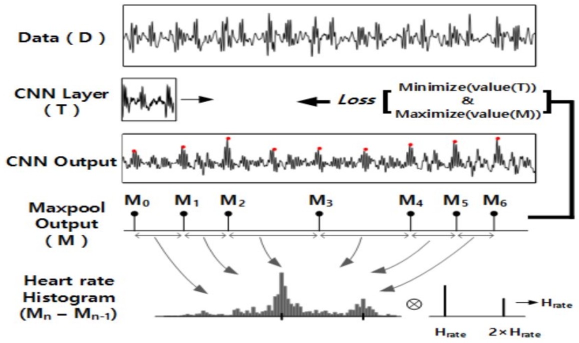

Figure 3—CNN-assisted Template Matching. The figure shows the different stages of RF-SCG's segmentation algorithm. After the *CNN* and *Maxpool* layers, each segment between neighboring *M*'s is a candidate for the hearbeat period (or heart rate). The final stage involves heart rate estimation using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

- Bottom chain: standard beamforming technique to get a a 3D coordinate system (x, y, and intermediate frequency buckets).

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

- Combining the outputs from top chain and button chain to get a 4d map.(ratio)

2.3 RF-to-SCG Translator

They build a **stack of 1D convolutional layers**.

Each layer acts like a **Finite Impulse Response (FIR) filter**.

Stacking them is equivalent to designing a set of **learnable filters** that reshape the radar signal into SCG.

Training:

- Input = radar reflection time-series.
- Ground truth = accelerometer SCG.
- Loss = L2-norm (MSE) between predicted SCG and ground-truth SCG.

Once trained:
No accelerometer needed anymore; Radar alone → SCG waveform.

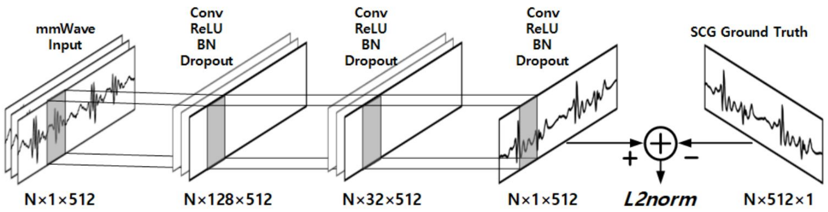

Figure 4—Translation Filter. The figure plots the different layers of RF-SCG's RF-to-SCG Translator. N is the batch size of data set. While the length of one frame is fixed, the number of channels is different across the layers. To translate the mmWave reflection to SCG recordings, every convolution 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.

2.4 SCG Automatic Labeling

They adapt the **U-Net** architecture (normally for images) to **1D time series**.

Input: the predicted SCG waveform.

Output: 5 parallel probability channels (one per fiducial).

Each channel = probability that event occurs at each time step.

Post-processing: threshold the probabilities → mark event locations.

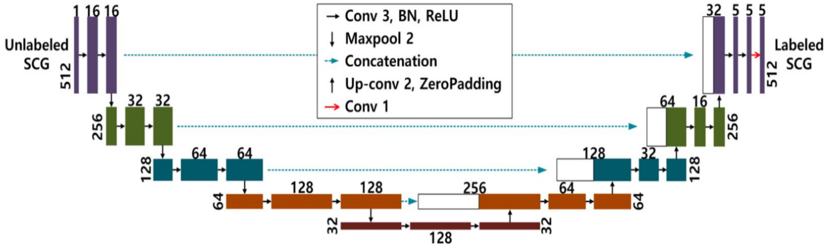

Figure 5—Automatic Labeling. The figure shows the network used for automatic labeling. Each box corresponds to a multi-channel feature map. The arrows denote the different operations. The number of channels is denoted on top of each box. The length of each layer is provided at the left edge of the box.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

3. Implementation

1. Hardware Setup

- **Radar board**: TI **IWR1443BOOST** millimeter-wave radar
  - Operates at **77 GHz** (FMCW chirps).
  - Includes a **steerable 2D antenna array** (horizontal + vertical elements).
- **Subjects**:
  - **21 healthy volunteers**, different days, **regular clothing**.
  - Total: **40,000+ heartbeats** collected
- **Ground truth**:
  - Each heartbeat labeled with **five fiducial points**:
    MC (Mitral Closing), IM (Isovolumetric Contraction), AO (Aortic Opening), AC (Aortic Closing), MO (Mitral Opening).
  - Labels manually inspected for medical protocol adherence

4. Evaluation

RF-SCG is evaluated on:

1. **Reconstructing SCG waveforms** for qualitative review by clinicians.
2. **Timing fine-grained micro-cardiac events** (MC, IM, AO, AC, MO) with clinical-level precision.
3. **Generalizing to unseen subjects**.
4. **Outperforming strong baselines** while being robust to environmental noise and posture changes

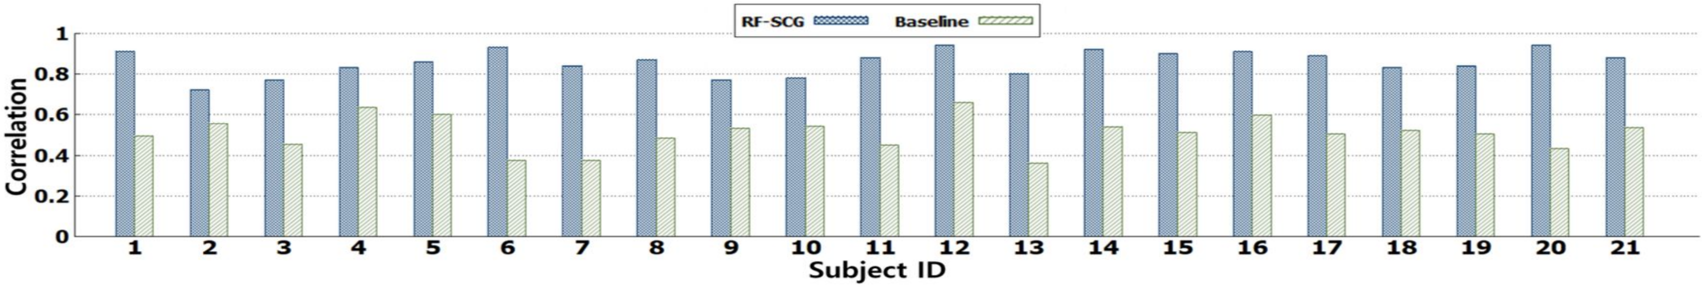

Figure 8—Accuracy in Reconstructing the SCG Waveform. The figure plots the correlation coefficient between the measured radar reflection and the ground truth across different subjects. It compares the median correlation of RF-SCG (in blue) and the baseline (in green), demonstrating that RF-SCG significantly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

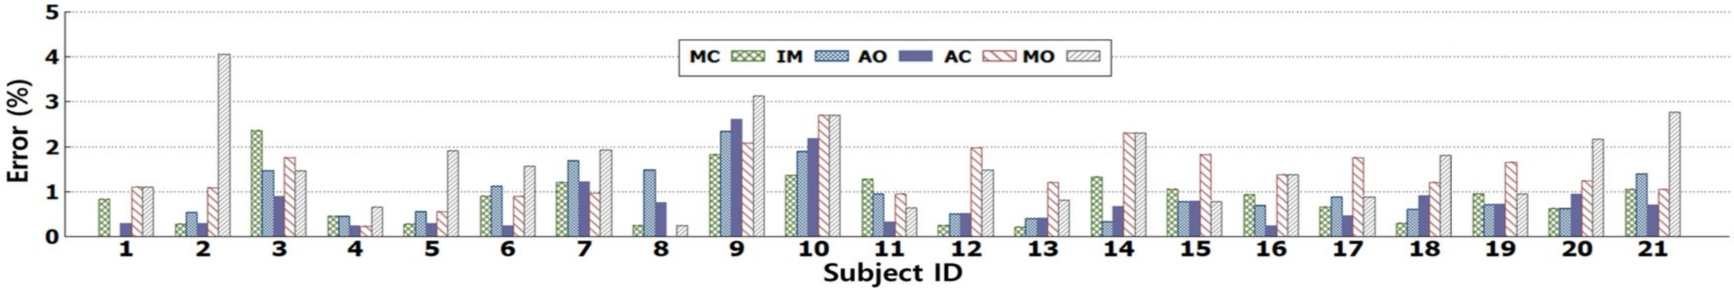

Figure 9—RF-SCG's Overall Accuracy. The figure plots RF-SCG's median error in timing each of the five micro-cardiac events of interest for 21 human subjects. The dashed yellow line denotes the the best-case accuracy achievable using a gold standard cardiac ultrasound (due to its higher quantization error).

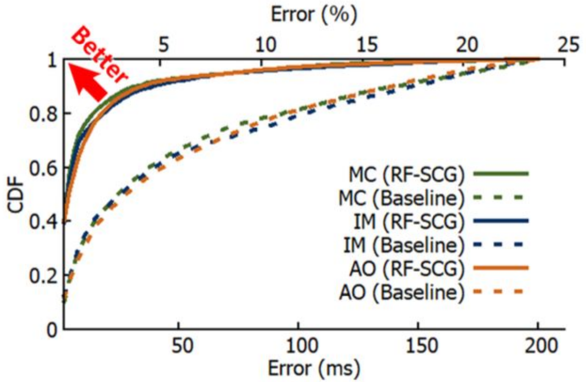

(a) CDF of Accuracies for Systolic Fiducial Points.

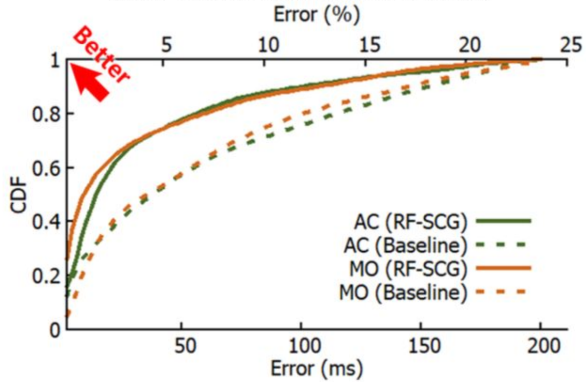

(b) CDF of Accuracies for Diastolic Fiducial Points.

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.

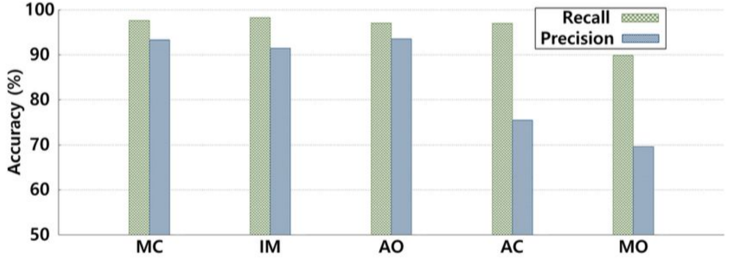

Figure 12—Precision and Recall. The figure plots the precision and recall of RF-SCG's automatic labeling component for each of the five micro-cardiac events of interest.

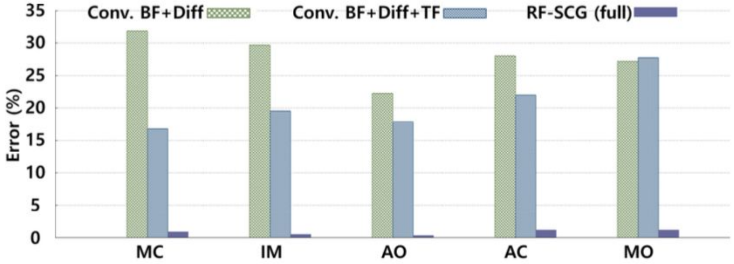

Figure 11—Partial Implementations. The figure plots the accuracy of various partial implementations of RF-SCG for each of the five fiducial points of interest.

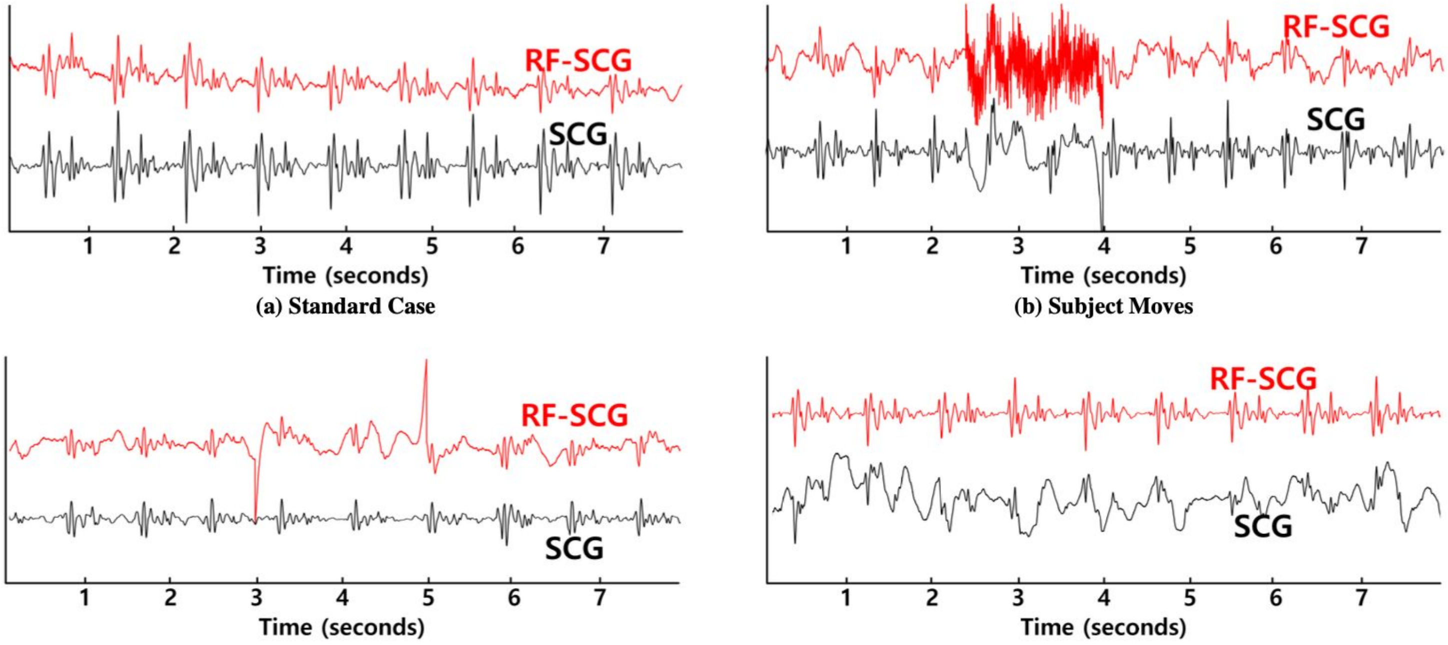

(a) Standard Case

(b) Subject Moves

(c) Nearby Object Moves

(d) Accelerometer Attachment Problem

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.

Let’s check Perusall Comments.

My thoughts on the paper

Hybrid ‘novelty’ = picking the right objective + a pipeline that lets physics do the separation and ML do the translation & semantics.

HR estimate (top chain): Yes with DSP, bandpass + autocorrelation/AMDF, cepstrum, or comb filters. Works, but beat-to-beat variability blurs peaks; the learned template is more tolerant locally.