MEDUSA

Scalable Multi-View Biometric Sensing in the Wild with Distributed MIMO Radars

Paper Review by Raheem Idowu

09/30/25

Outline

What is MEDUSA? (System and results)

Why was it made? (Motivation and background)

How does it work? (Challenges and contributions)

What’s next? (Limitations and future work)

Perusall Comments

Outline

What is MEDUSA? (System and results)

Why was it made? (Motivation and background)

How does it work? (Challenges and contributions)

What’s next? (Limitations and future work)

Perusall Comments

What is MEDUSA

Scalable **Multi-View** Biometric **Sensing in the Wild** with **Distributed MIMO Radars**

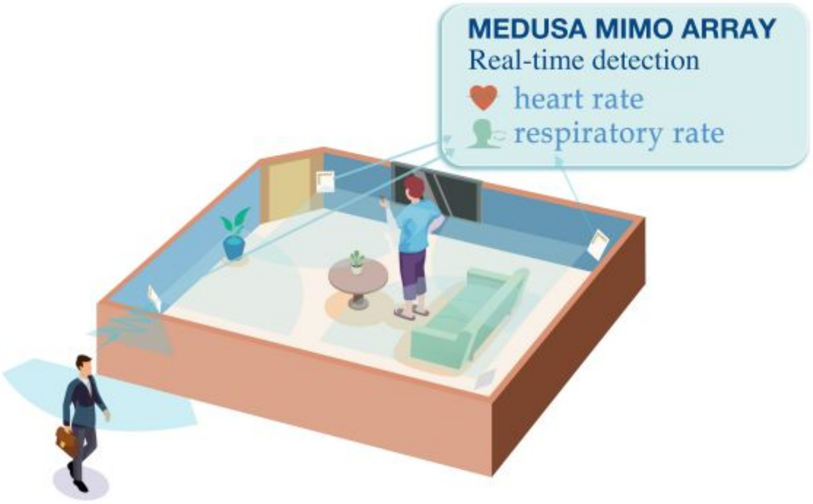

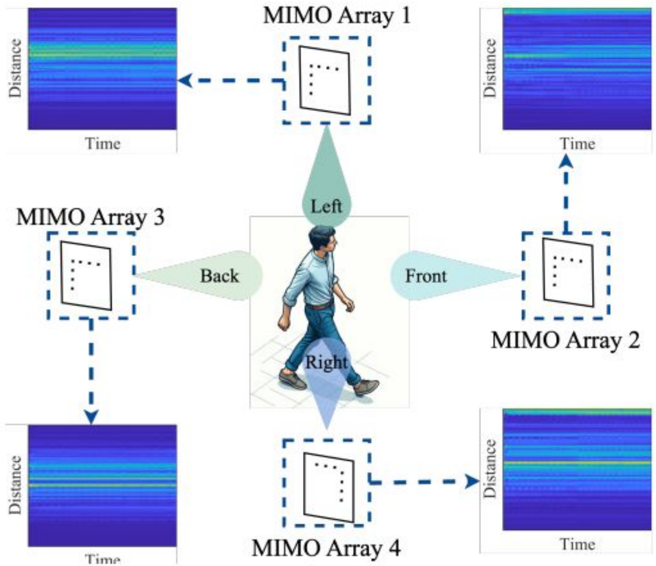

System Overview

4 **radars** in 4 corners of the room (spatially distributed)

- Multi-view sensing

Each radar is 4x4 MIMO with 8 **antenna elements**

4x4 MIMO system= 4 Rx and 4 Tx antennas

- Equals 16 virtual antennas

Radar data → Computer, Radars sync wirelessly

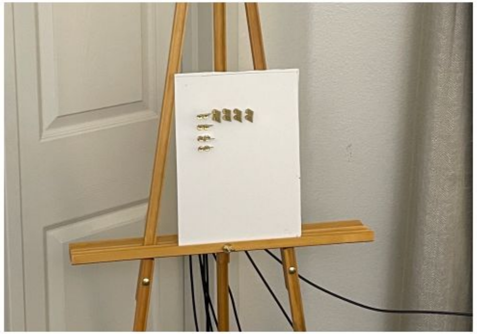

(d) Antenna board 4×4

MIMO UWB Radar

MIMO = multiple input multiple output

Improve sensing by using multiple antennas

UWB = ultra wideband

>500 Mhz of bandwidth

low power, short range, high accuracy

used in iPhone Precision Finding

Results

Heart rate and respiration rate

Accuracy +20% vs. prior baselines

Most benefit when:

- Subject is facing away
- Subject is moving
- Obstacles in the environment

Enables “sensing in the wild”

Real time! (22.94ms median latency)

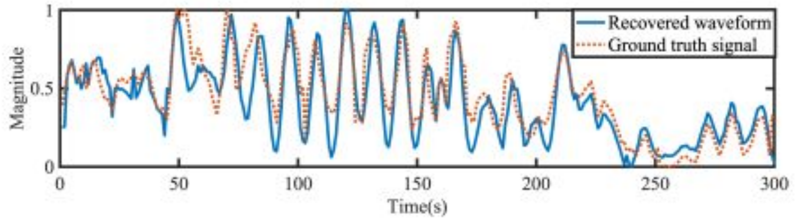

(a) Respiration waveform reconstructed for a moving target in LoS

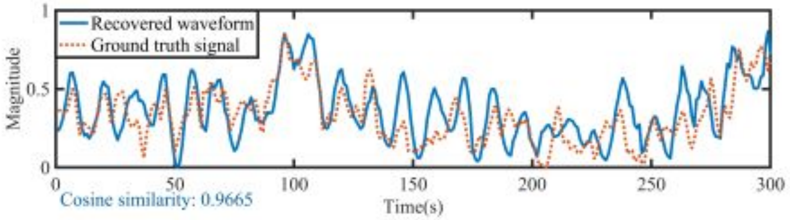

(b) Respiration waveform reconstructed for a static target in NLoS

Figure 19: Respiration waveforms reconstructed by MEDUSA.

Baseline and Ground Truth

Baseline: Novelda UWB (resp) and TI mmWave IWR1443 (heart)

S = Single View MIMO Radar (16 x 16)

Ground Truth = Vernier's breathing belt and heart rate sensor

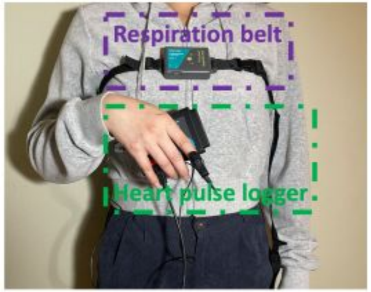

	Medusa Platform	IWR1443BOOST	AWR2243 Cascade
Frequency Band	6.5-9.5GHz	76-81 GHz	76-81 GHz
TX/RX	16TX/16RX	3TX/4RX	12TX/16RX
Azimuth Array	256 element virtual array	12 element virtual array	86 element virtual array
Max Angular Resolution	0.448 degree	9.53 degree	1.4 degree
Min Spacing Separation	0.039m at 5m	0.841m at 5m	0.122m at 5m
Frame Rate (FPS)	50 – 200	10	5

Table 1: Comparison of MEDUSA platform, TI IWR1443BOOST, and TI AWR2243 Cascade.

Results

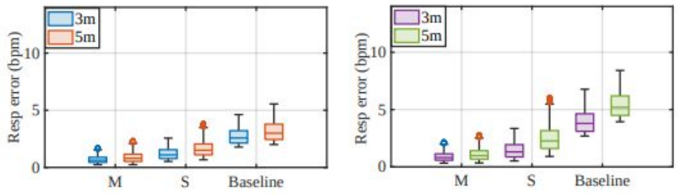

(a) Front (b) Left and Right

Figure 20: LoS Static targets: BPM with different orientations. Comparison of MEDUSA ('M'), Co-located single-view MIMO radar ('S'), and Baseline.

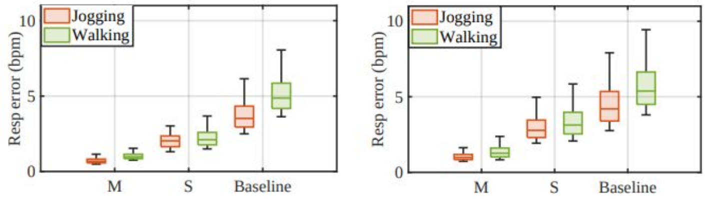

(a) Jogging and Walking in LoS (b) Jogging and Walking in NLoS

Figure 22: Respiration errors (bpm) for targets with movements in LoS and NLoS.

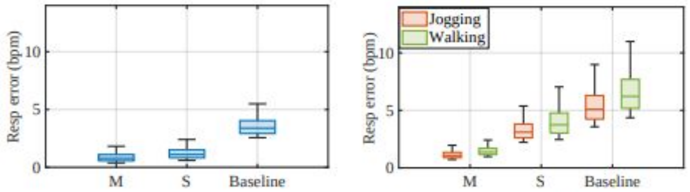

(a) Static (SS) (b) Jogging and walking

Figure 24: Respiration errors (bpm) in LoS of Static and Moving targets in the untrained environment.

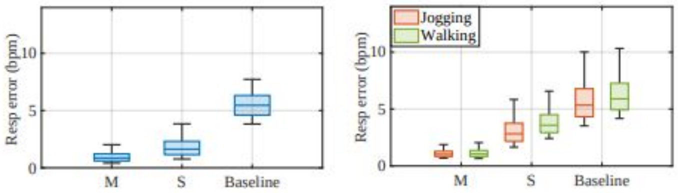

(a) Static (SS) (b) Jogging and walking

Figure 25: Respiration errors (bpm) in NLoS of Static and Moving targets in the untrained environment.

Outline

What is MEDUSA? (System and results)

Why was it made? (Motivation and background)

How does it work? (Challenges and contributions)

What’s next? (Limitations and future work)

Perusall Comments

Motivation

More views = “robustness to blockage and movement”

Movement = blockage due to body parts

Same number of antennas can be

co-located (16x16) = more SNR gain

distributed (2x 8x8) = more spatial diversity gain

Idea: distribute antennas for in-the-wild sensing!

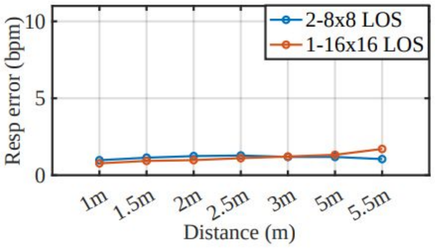

(a) LOS

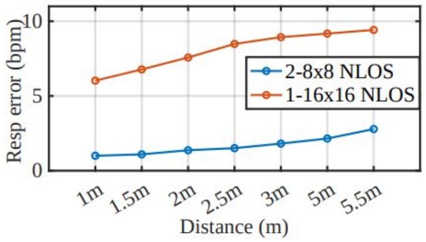

(b) NLoS

Outline

What is MEDUSA? (System and results)

Why was it made? (Motivation and background)

How does it work? (Challenges and contributions)

What’s next? (Limitations and future work)

Perusall Comments

Challenges

1 - Balancing SNR gain with spatial diversity gain

2 - Synchronization between the radars

3 - Fusing data and recovering vital sign signals

Challenges

1 - Balancing SNR gain with spatial diversity gain

2 - Synchronization between the radars

3 - Fusing data and recovering vital sign signals

Challenge 1: Balancing SNR & spatial diversity

Balancing SNR and spatial diversity gain

Two extremes with 256 virtual antennas

- 1 16x16 radar
- 16 1x1 radars

Contribution: 4 4x4 radars sweet spot

Experimentally determined

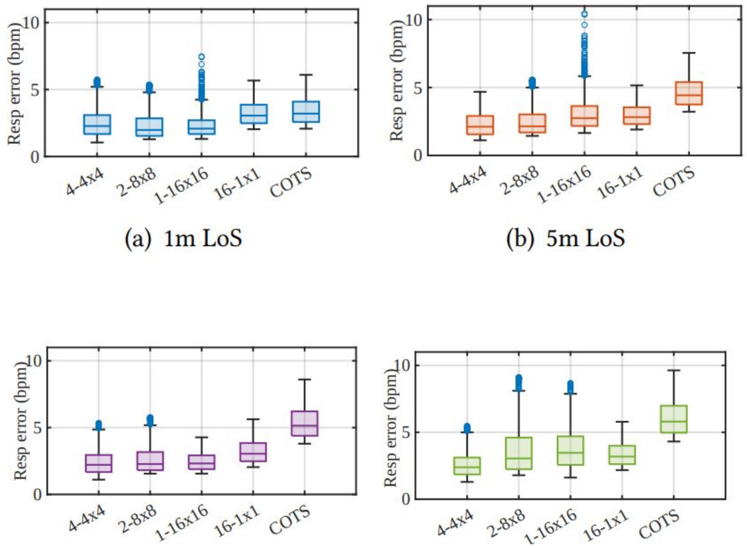

(a) 1m LoS (b) 5m LoS

(c) 1m NLoS (d) 5m NLoS

Challenge 1: Modularity of MEDUSA

Determined **after** making MEDUSA thanks to **modularity**

MEDUSA is modular:

- As many radars as you want
- Each radar max 16 x 16

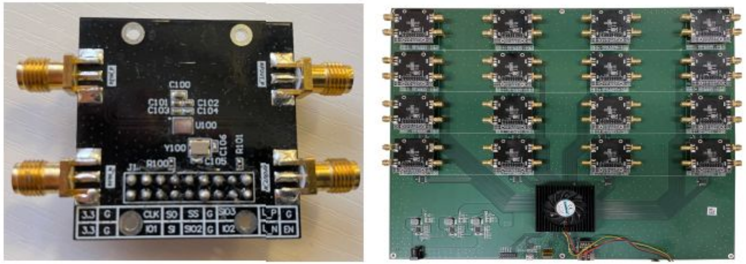

(a) Daughterboard (b) Base board

Radar Hardware:

- Baseboard (FPGA) connects to up to 16 daughterboards
- Each daughterboard (UWB SoC) 1 RX and 1 TX

Challenges

1 - Balancing SNR gain with spatial diversity gain

2 - Synchronization between the radars

3 - Fusing data to recover vital sign signals

Challenge 2: Synchronization between radars

Need a coherent system, otherwise interference and corruption

Contribution: Wireless Sync Channel with Software Defined Radio

One clock server baseboard transmits clock signals

All other baseboards listen to it and adjust based on it (Phase Locked Loop)

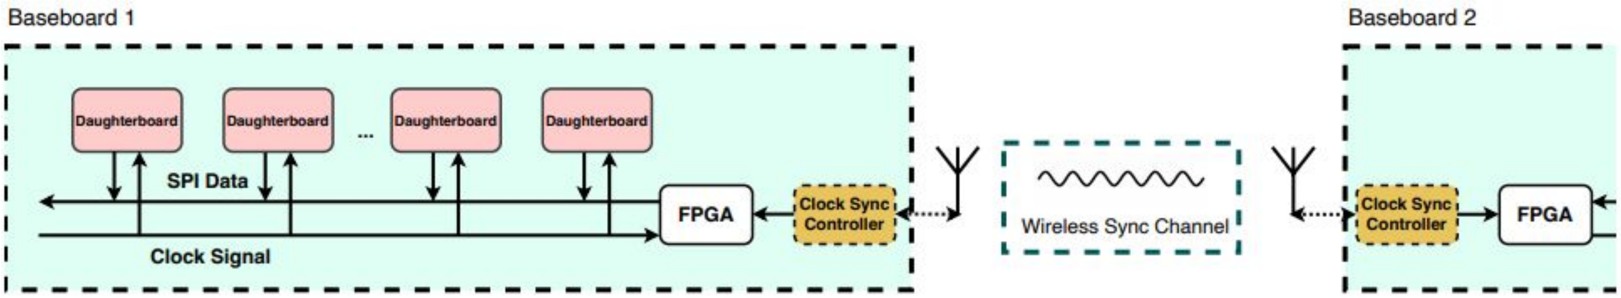

Challenge 2: Evaluation

Achieves median 0.25Hz CFO in LoS and 0.3Hz CFO in NLoS

CFO = carrier frequency offset, (N)LoS = (non) line of sight

802.11 (WiFi) specifies at most ±40ppm of nominal frequency

0.25Hz is <0.00004 ppm of Medusa nominal frequency (6.5GHz)

NLoS in this context is if radars are blocked from each other (?)

Challenge 2: Medusa Hardware Summary

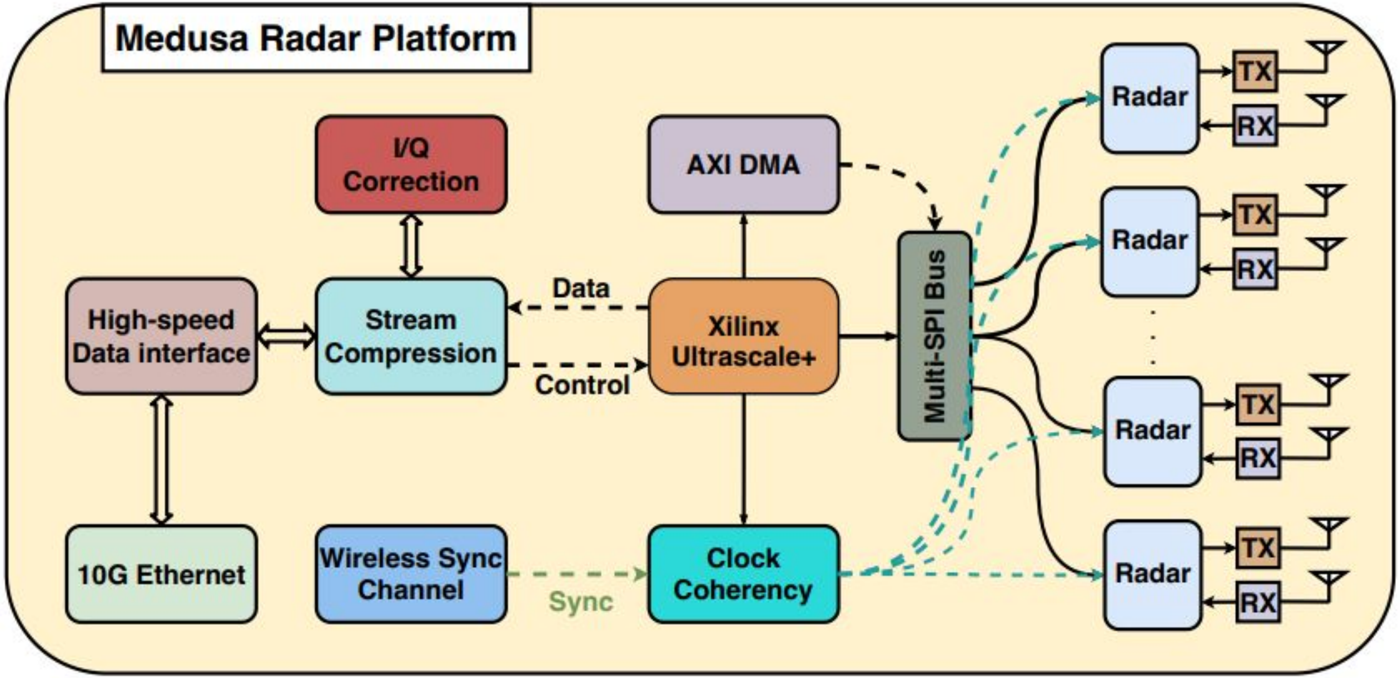

Challenges

1 - Balancing SNR gain with spatial diversity gain

2 - Synchronization between the radars

3 - Fusing data and recovering vital sign signals

Challenge 3: Fusing and recovering vital signs

UWB antenna continuously sends pulses

Each pulses gives amplitude and phase of distance (186 bins)

Extraction is modelled as nonlinear **ICA**

Independent component analysis problem

$$S_k(t) = f([x_1(t), \ldots, x_N(t)]^T) = [s_1(t), \ldots, s_N(t)]^T$$

Recover source signals from a nonlinear mix

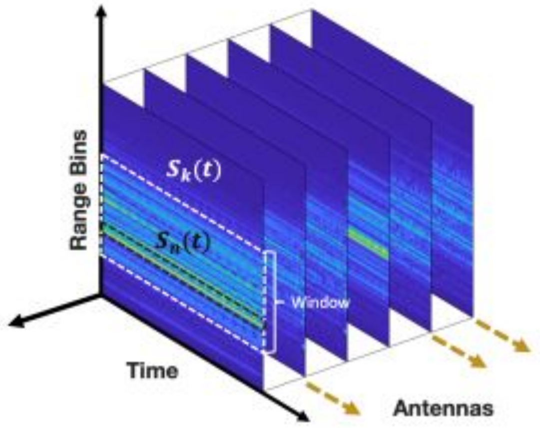

One radar

Challenge 3: ICA

Key assumption of ICA: statistical independence between components

Very useful: (video)

- image steganography[32]
- optical Imaging of neurons[33]
- neuronal spike sorting[34]
- face recognition[35]
- modelling receptive fields of primary visual neurons[36]
- predicting stock market prices[37]
- mobile phone communications[38]
- colour based detection of the ripeness of tomatoes[39]
- removing artifacts, such as eye blinks, from EEG data.[40]
- predicting decision-making using EEG[41]

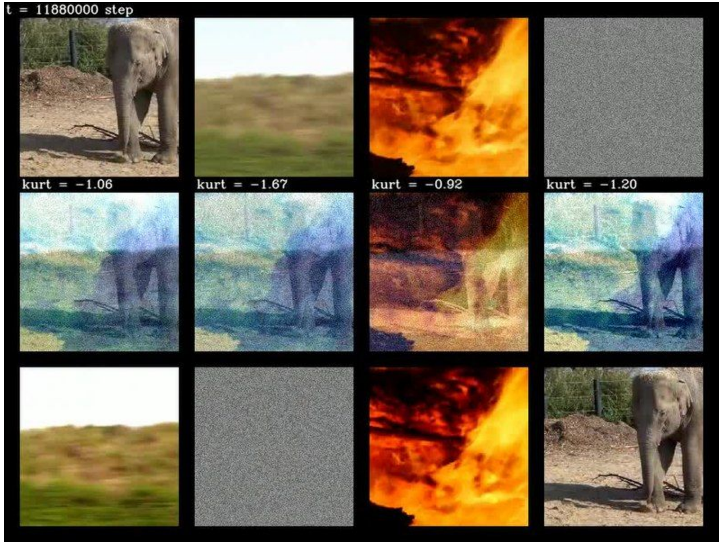

Challenge 3: How to solve nonlinear ICA

Unsupervised (!) contrastive learning = train encoder + classifier

- Goal: Determine if a two points in a signal are in the same “segment”

This network can then be used to approximately solve nonlinear ICA

Paper is unclear on details but seems very powerful

Unsupervised Feature Extraction by Time-Contrastive Learning and Nonlinear ICA

Aapo Hyvärinen and **Hiroshi Morioka**

Nonlinear Independent Component Analysis for Principled Disentanglement in Unsupervised Deep Learning

Aapo Hyvärinen[1*], Ilyes Khemakhem[2], and Hiroshi Morioka[3]

Challenge 3: Contrastive learning (one radar)

Unsupervised learning = avoid inaccurate ground truth sensors

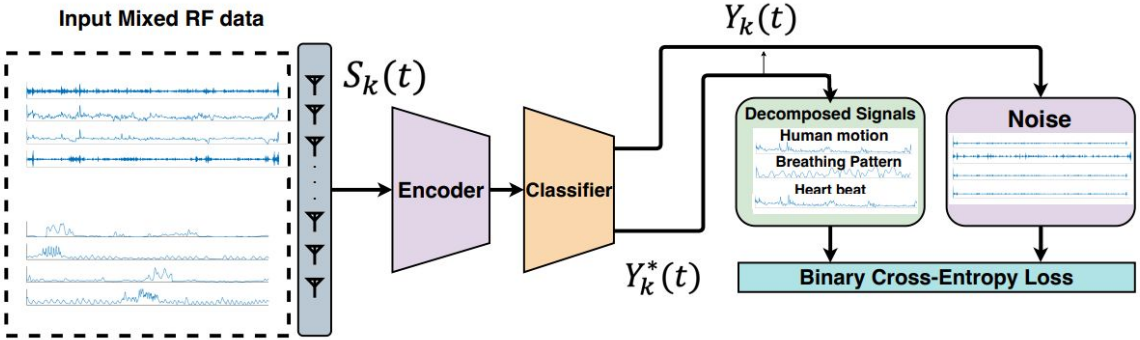

$$Y_k(t) = \begin{pmatrix} S_k(t) \\ S_k(t-T) \end{pmatrix}, \quad Y_k^*(t) = \begin{pmatrix} S_k(t) \\ S_k(t-\delta) \end{pmatrix}$$

Challenge 3: Multiple radars

Multi-head attention layer

Learns to weight & fuse signals from radars based on

- Estimated location of target
- Estimated occlusions
- Estimated orientation

Weighted signals sent to contrastive network

$$[Z_1(t), \ldots, Z_W(t)] = \mathbf{A}([Y_1(t), \ldots, Y_M(t)]),$$
$$[Z_1^*(t), \ldots, Z_W^*(t)] = \mathbf{A}([Y_1^*(t), \ldots, Y_M^*(t)]).$$

Challenge 3: Attention Layer Demonstration

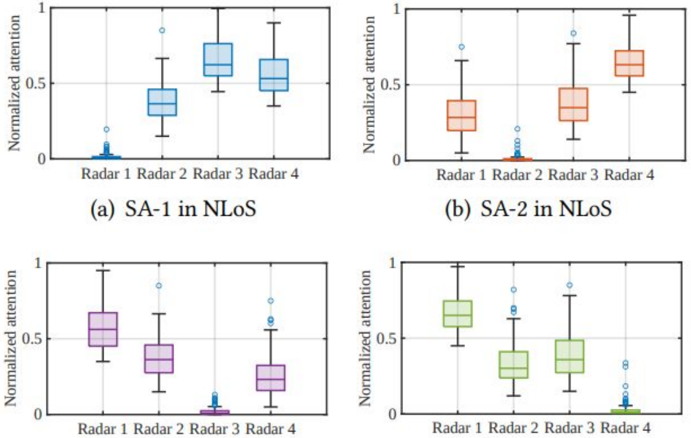

(a) SA-1 in NLoS

(b) SA-2 in NLoS

(c) SA-3 in NLoS

(d) SA-4 in NLoS

Challenge 3: Vital Sign Extraction

ICA outputs source signals

Determine relevant source signals with:

- Respiratory Rate Variability (RRV)
- Heart Rate Variability (HRV)

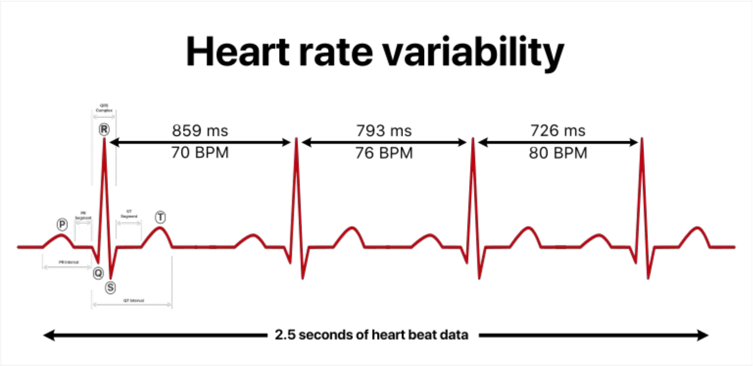

Essentially heuristics

Might defeat healthcare purpose? Relies on “normal variations in heart rate”

Challenge 3: Medusa Model Summary

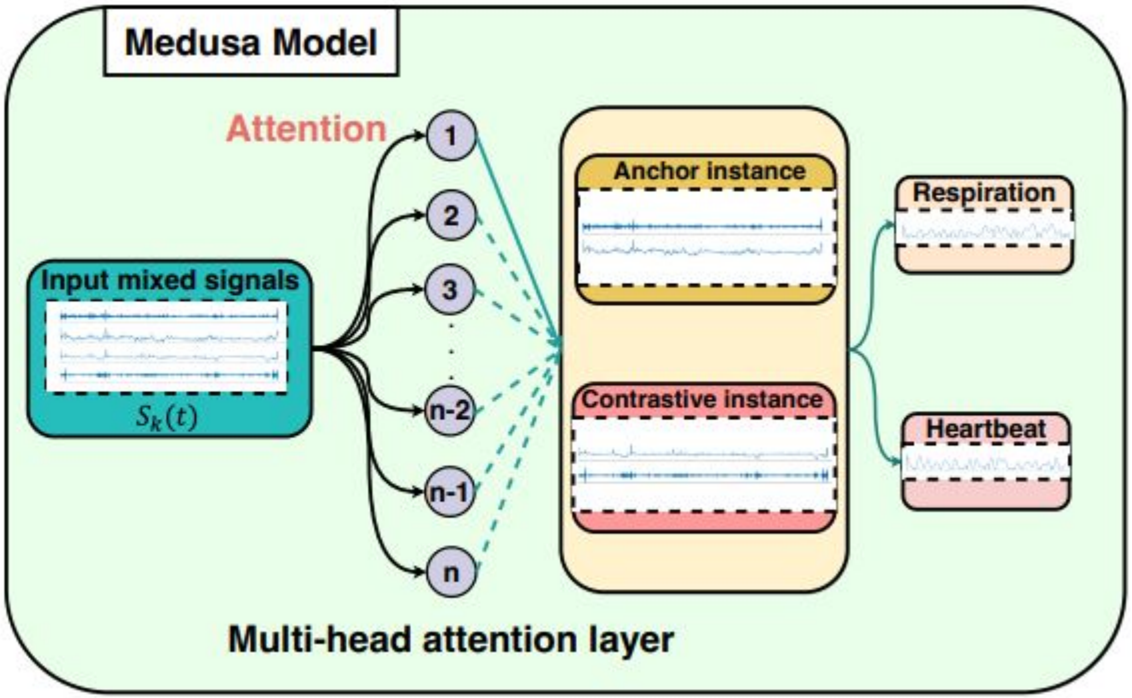

Outline

What is MEDUSA? (System and results)

Why was it made? (Motivation and background)

How does it work? (Challenges and contributions)

What’s next? (Limitations and future work)

Perusall Comments

Future Work: Multi-person tracking

ICA output = all source signals in the scene

Multiple vital sign signals for each person

Almost as accurate (0.92 vs. 0.96)

Future work:

Localization and tracking to match

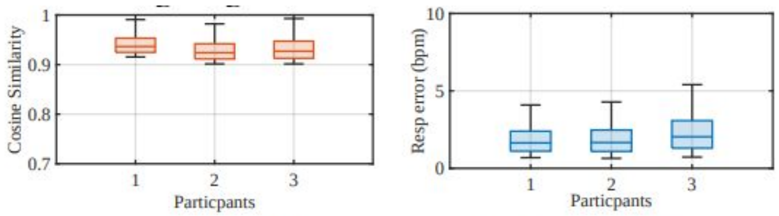

(a) Cosine similarity of the recovered respiration waveform

(b) Respiration Error

Figure 26: Cosine similarity and BPM of respiration rate detection for multiple targets.

Outline

What is MEDUSA? (System and results)

Why was it made? (Motivation and background)

How does it work? (Challenges and contributions)

What’s next? (Limitations and future work)

Perusall Comments