



TEXAS INSTRUMENTS



Asia 2023

Target Classification on the Edge Using mmWave Radar: A Novel Algorithm and Its Real-Time Implementation on TI's IWRL6432

Muhammet Yanik, Slobodan Jovanovic, Sandeep Rao

Radar Systems and Algorithms R&D, Texas Instruments





What We'll Cover

- Introduction
 - What/why is mmWave radar sensing?
 - Why do we need target classification on the edge using mmWave radars?
- The proposed target classification algorithm details. What is the novelty of this work?
- Real-time implementation on TI's 60GHz low-power mmWave radar IWR6432
- Performance results of the algorithm and the real-time implementation benchmarks
- Summary and conclusions

The mmWave Sensing

- The mmWave sensing typically refers to radar applications that operate in the 60GHz and 77GHz frequency bands.
- Why radar now?
 - Technological advances => smaller size, lower cost, higher performance
 - Automotive remains a key market for these sensors
 - New applications: Building surveillance, factory automation, robotics, automotive, in-cabin sensing, and more!


Safety bubbles



Energy efficient units



Energy efficient Smart A/C



Mobile robots

In-cabin sensing



Smarter devices



Need for Target Classification

- Traditional applications leverage radar strengths in target detection and tracking.
- Emerging applications require target classification as well:


Segment	Use-case	Requirement
Automotive Radar	Road User Classification	High Performance
In-Cabin Sensing	Occupancy Detection and Classification	High Performance, Low Power
Building Automation / Surveillance	Motion Classification (people vs pets, trees etc.)	Low Cost, Low Power, Edge Processing
Elderly Care	Fall Detection, Activity Classification	Low Cost, Low Power, Edge Processing
Gesture Recognition	Automotive Dash Board, Mobile Devices, Household Appliances, Kick-2-Open (Automotive Boot opener)	Low Cost/ Low Power, Edge Processing

Main focus
of this work

- Radar classification can leverage the work done in the context of Vision
 - Machine Learning architectures/frameworks can be re-used
- However, the nature of the signal is different
 - Vision: High spatial resolution. None or limited range/velocity resolution
 - Radar: High range/velocity resolution. Limited spatial resolution
 - Hence, intuition on feasibility, performance, and complexity has to be built from the ground up

Resolution in Radar and μ -Doppler Signatures

Radar images are very sharp in the range-velocity domain but blurred in the spatial domain




	Radar	Vision
Range resolution	Green	Grey
Velocity resolution	Green	Grey
Spatial resolution	Grey	Green
All Weather	Green	Grey
Privacy	Green	Grey

- Radars' native good velocity resolution can be used to capture the unique "velocity signature" of moving targets
- Different targets (bike, pedestrian, pet, drone, bird, car, etc.) can have different "signatures."
- Can be used in target classification/ target filtering.

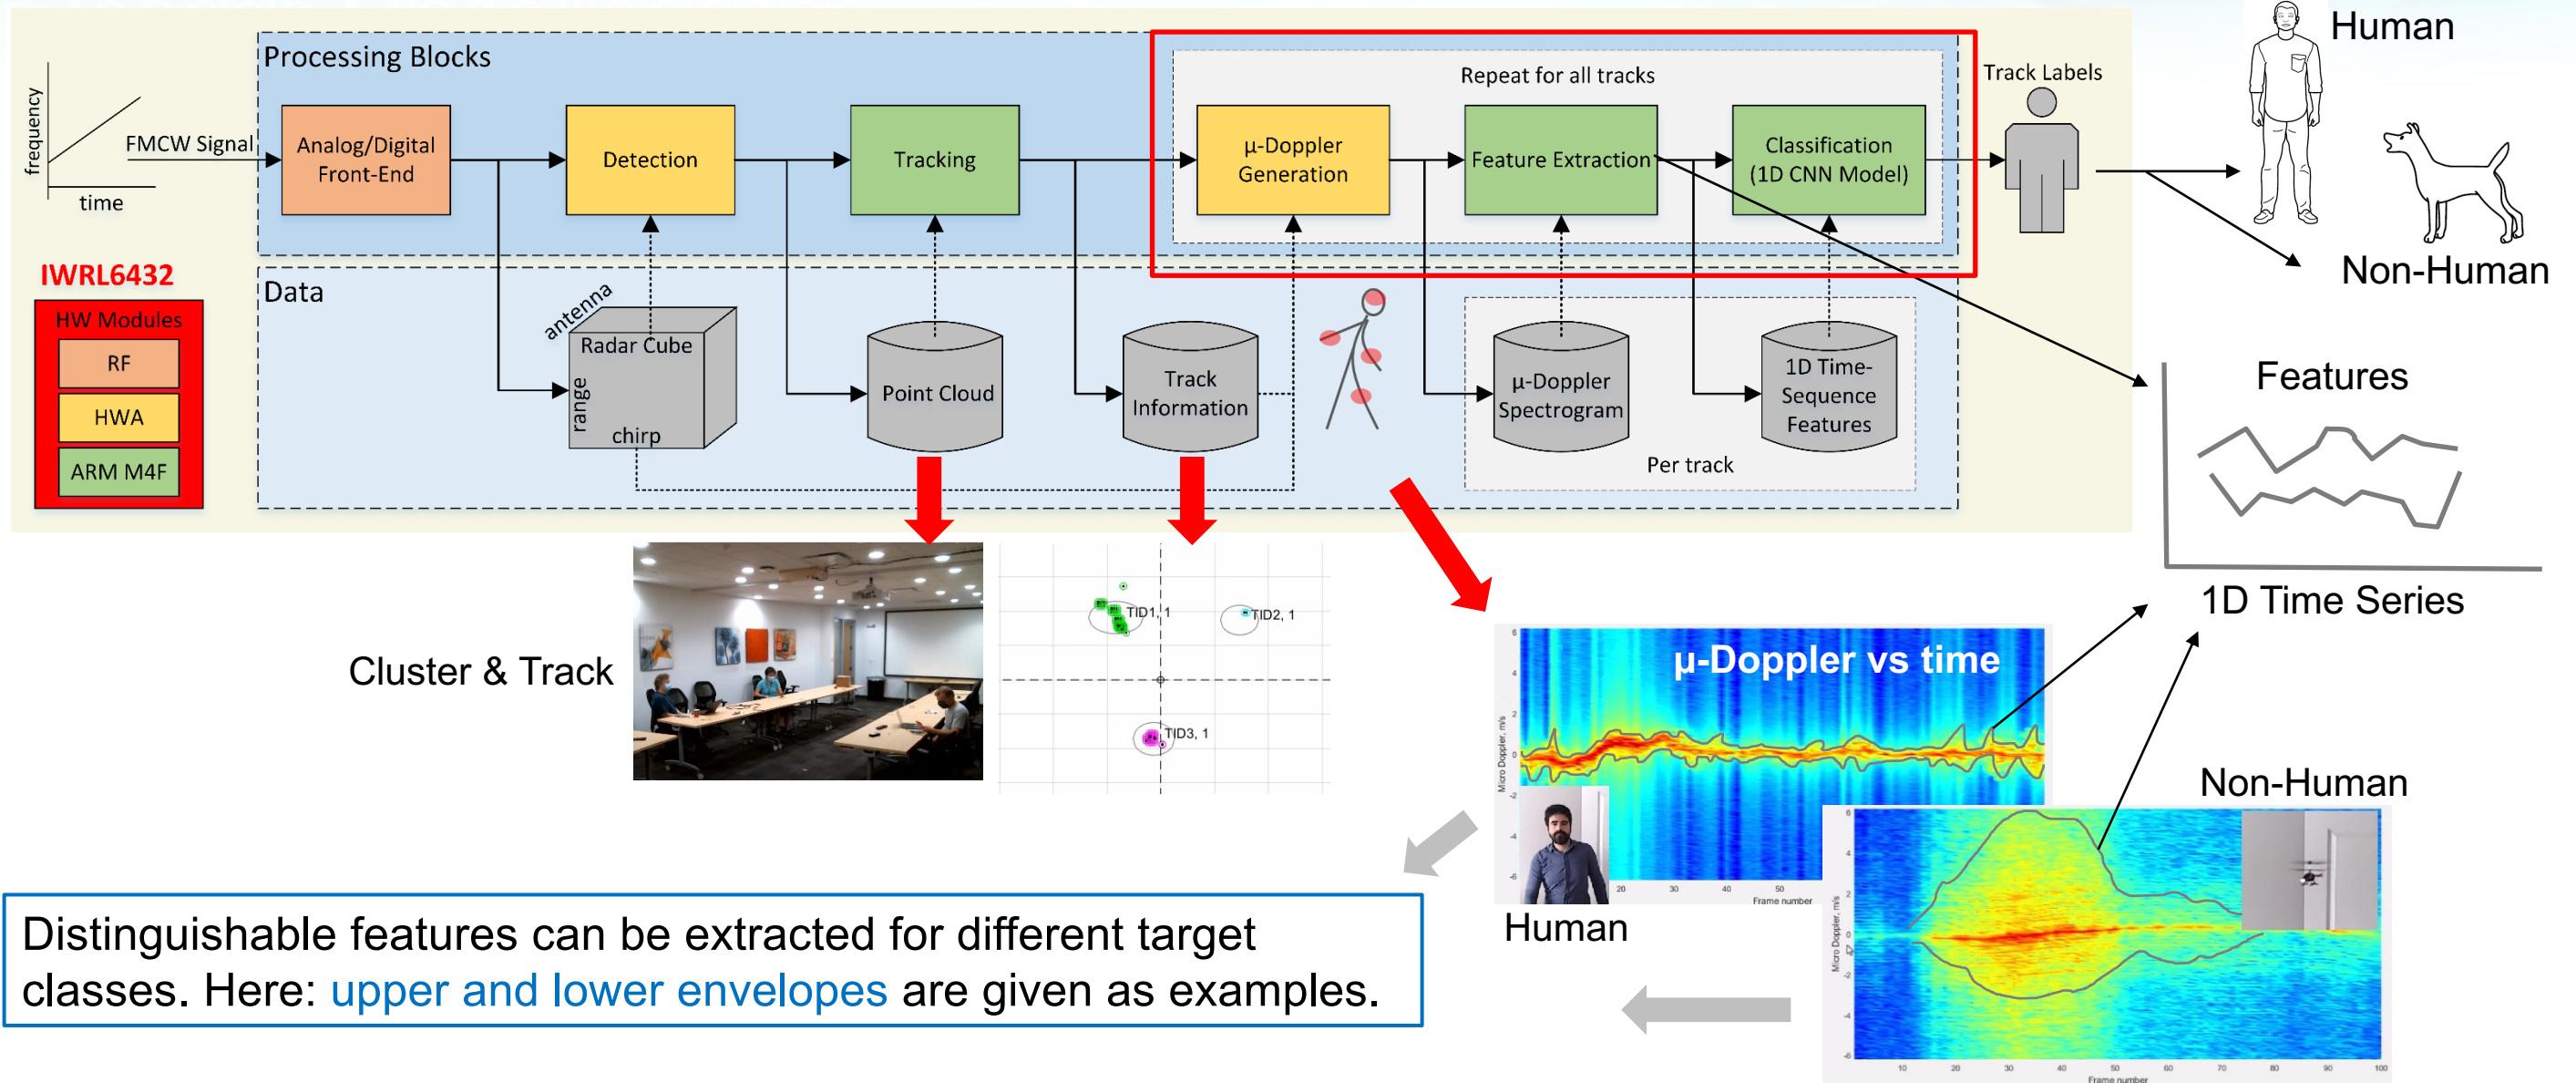
The work in this tutorial relies heavily on these Doppler signatures

Main challenge: How do we create these signatures for multiple objects concurrently, and how do we process them?



Proposed Solution for Target Classification

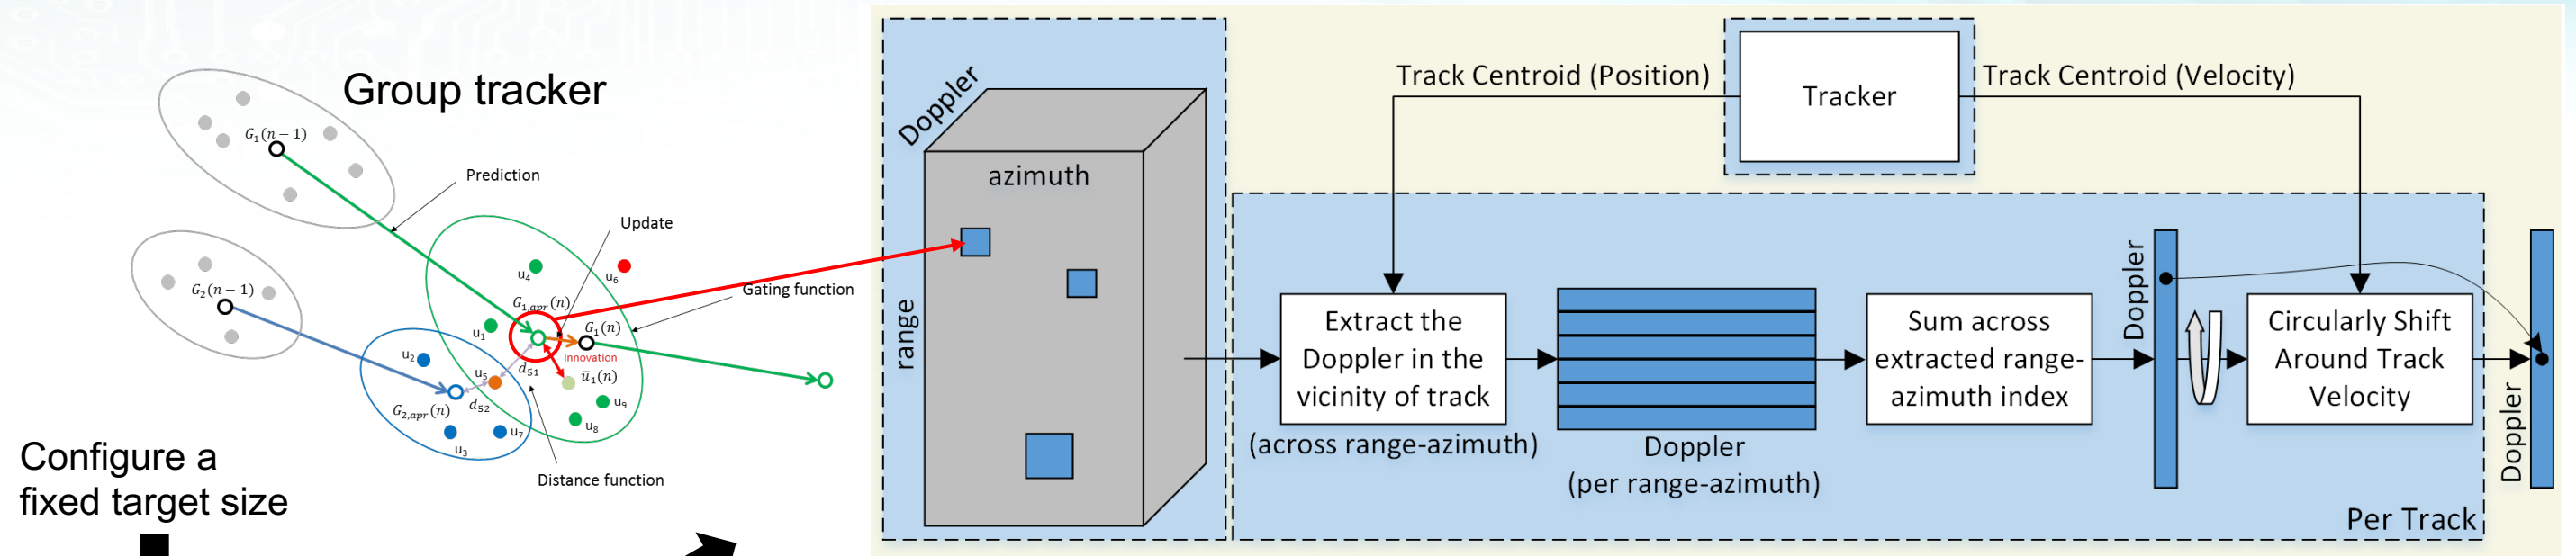
- The end-to-end processing chain to classify multiple target objects: The detection layer generates the point cloud. The tracker layer is run to track multiple objects. Required features are created, and the classifier is run per track.



Distinguishable features can be extracted for different target classes. Here: upper and lower envelopes are given as examples.

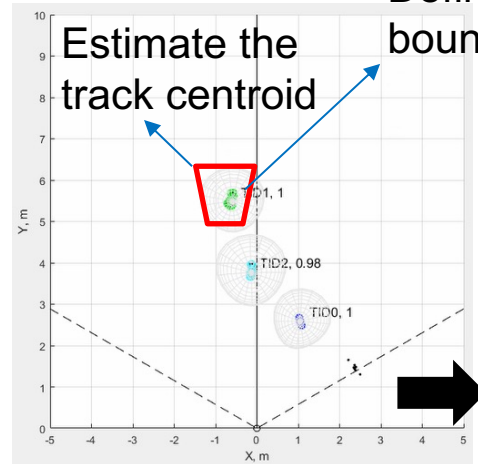
μ-Doppler Generation

- Tracking and spectrogram generation based on referencing the radar cube by the track centroids are central to our focus.



Configure a fixed target size

Main approach:



Define the boundary box

Scenario, 3 targets

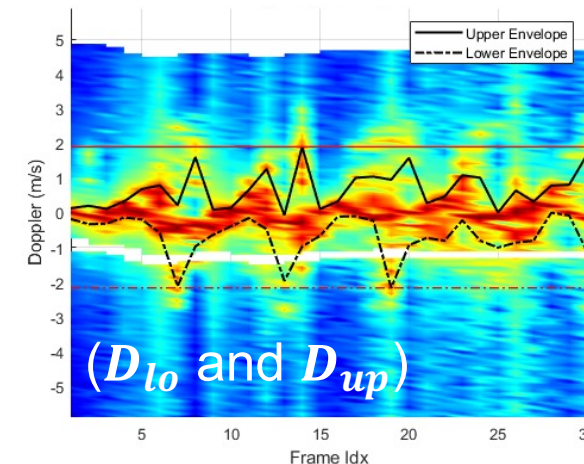


Extract spectrums


4 features are then extracted:

- (1) Normalized bandwidth power (P_{BW})
- (2) Lower Envelope (D_{lo})
- (3) Upper Envelope (D_{up})
- (4) Spectral entropy (H)

The spectrum is normalized (between [0 1]) (to eliminate the effect of SNR)



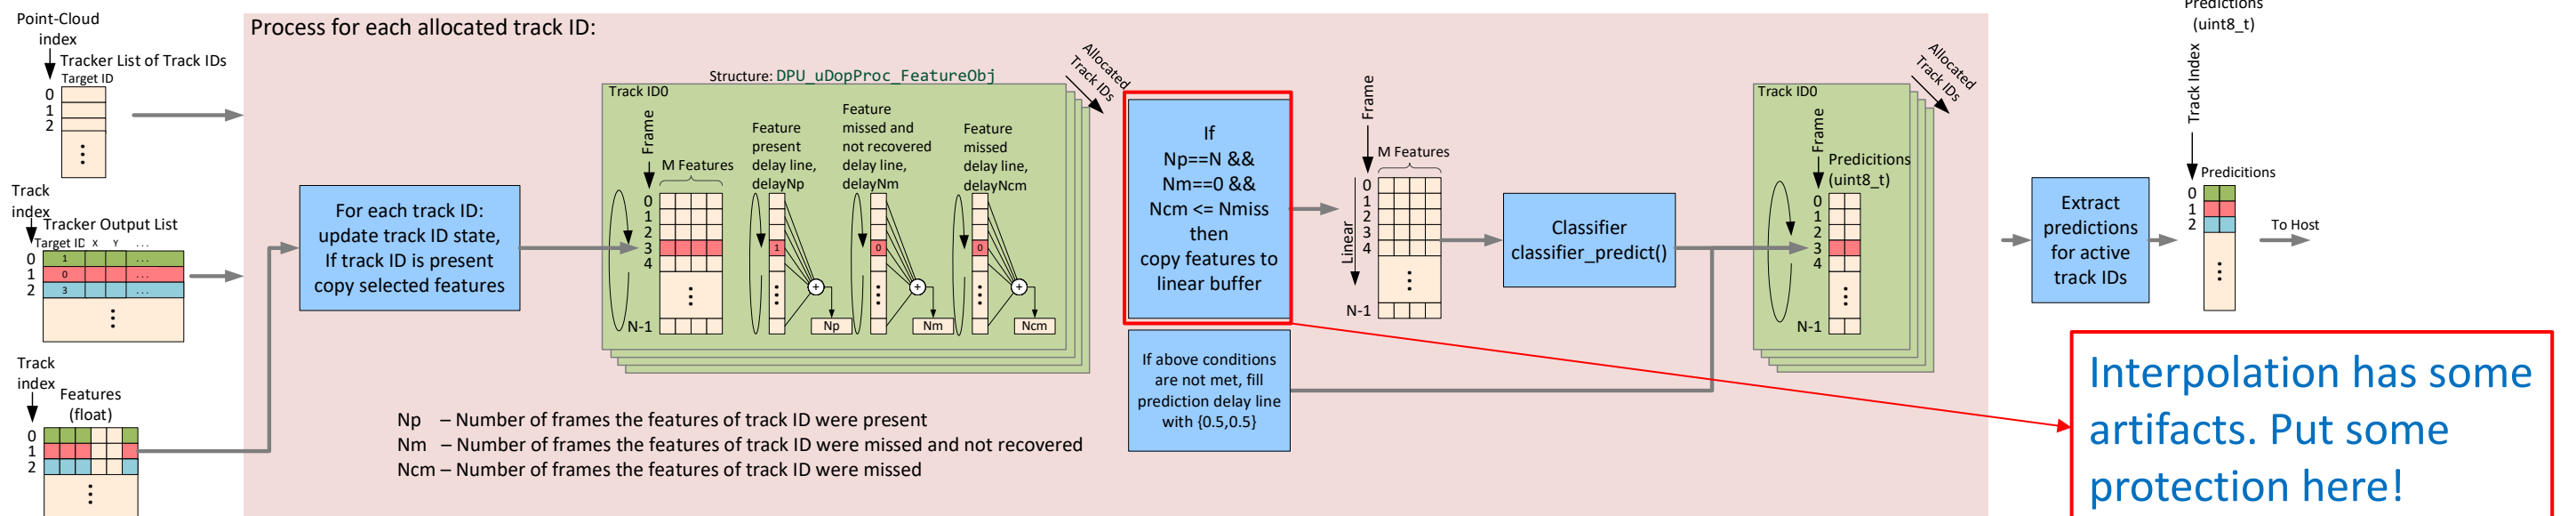
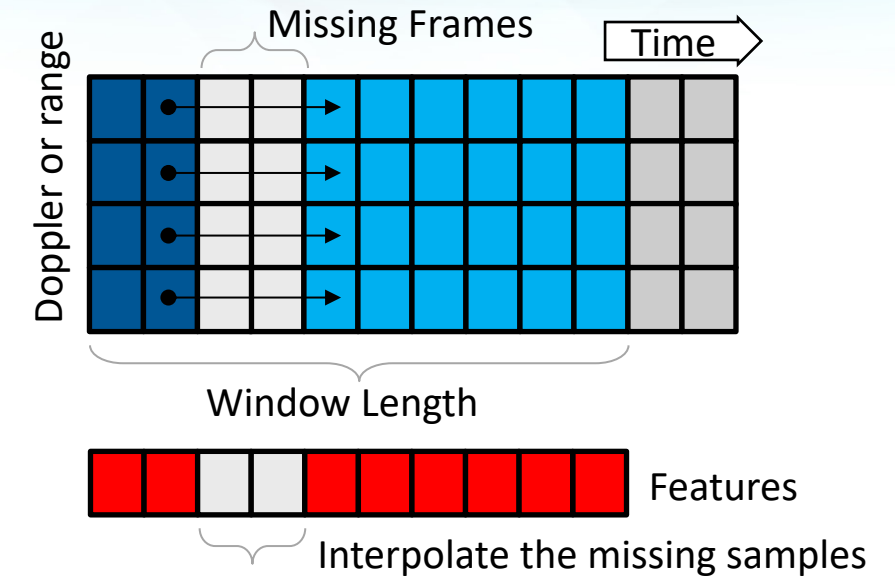
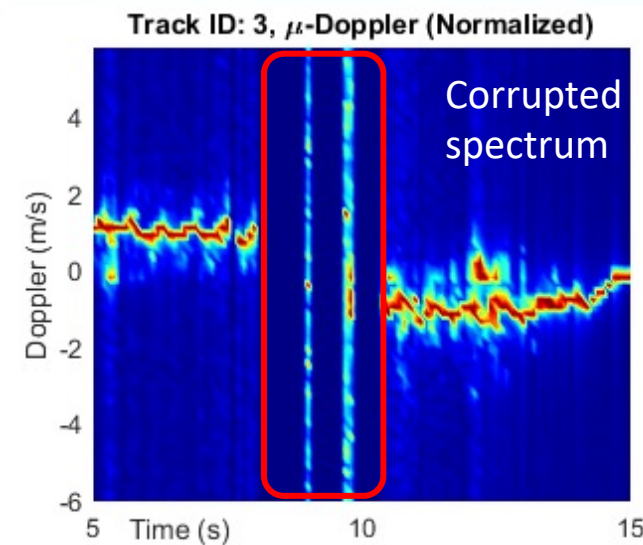

(P_{BW} , and H)



For details: M. E. Yanik and S. Rao, "Radar-Based Multiple Target Classification in Complex Environments Using 1D-CNN Models," *IEEE Radar Conf.*, San Antonio, TX, USA, 2023, pp. 1-6

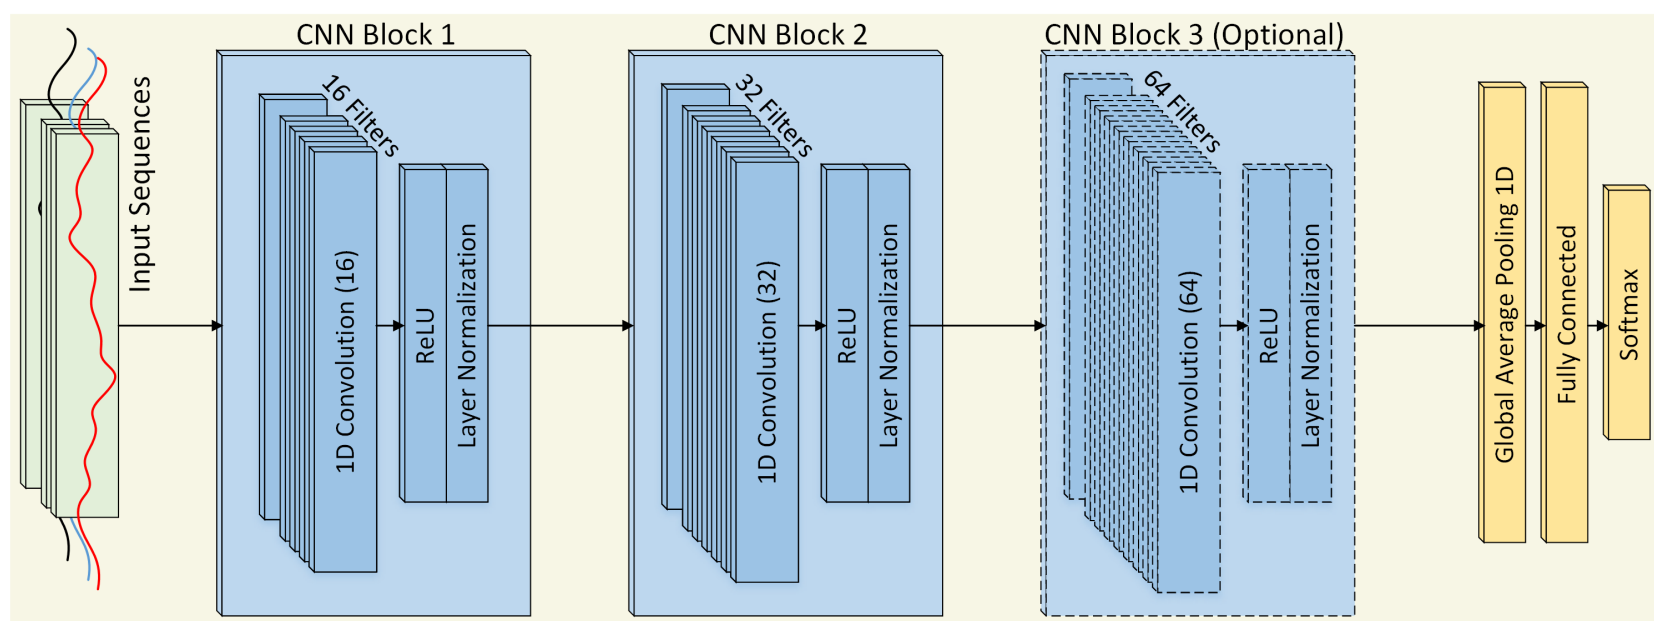
Feature Extraction – Improve Robustness

- If a certain track has no associated points at some frames (it may exit from the scene or obscured by another track), the tracker may still continue to track. In such scenarios, we should avoid updating the feature set because the extracted information **will not be reliable**.



Classification Model

- In the proposed 1D-CNN architecture:
 - 2 blocks of 1D convolution, ReLU, and layer normalization layers are specified.
 - We create 16 and 32 filters for the first and second convolutional layers (with a filter size of 3), respectively.
 - To reduce the output of the convolutional layers to a single vector, a 1D global average pooling layer is added.
 - To map the output to a vector of probabilities, a fully connected layer is specified with an output size of two (matching the number of classes), followed by a softmax layer.



Processing time of the classifier

- On ARM M4F core @160MHz
- Per prediction call (i.e., per track)
- 30-frame window

- 2-layer CNN: 3ms
- 3-layer CNN: 15ms

Memory required by the classifier

- Internal buffer for computations + weights
- 2-layer CNN: 5.7KB + 7.5KB = 13.2KB
- 3-layer CNN: 13.4KB + 32.5KB = 45.9KB

- An optional block of 1D convolution with 64 filters, ReLU, and layer normalization layers (following the first two blocks) is also added to compare the performance of 2 vs. 3 layers CNN models.

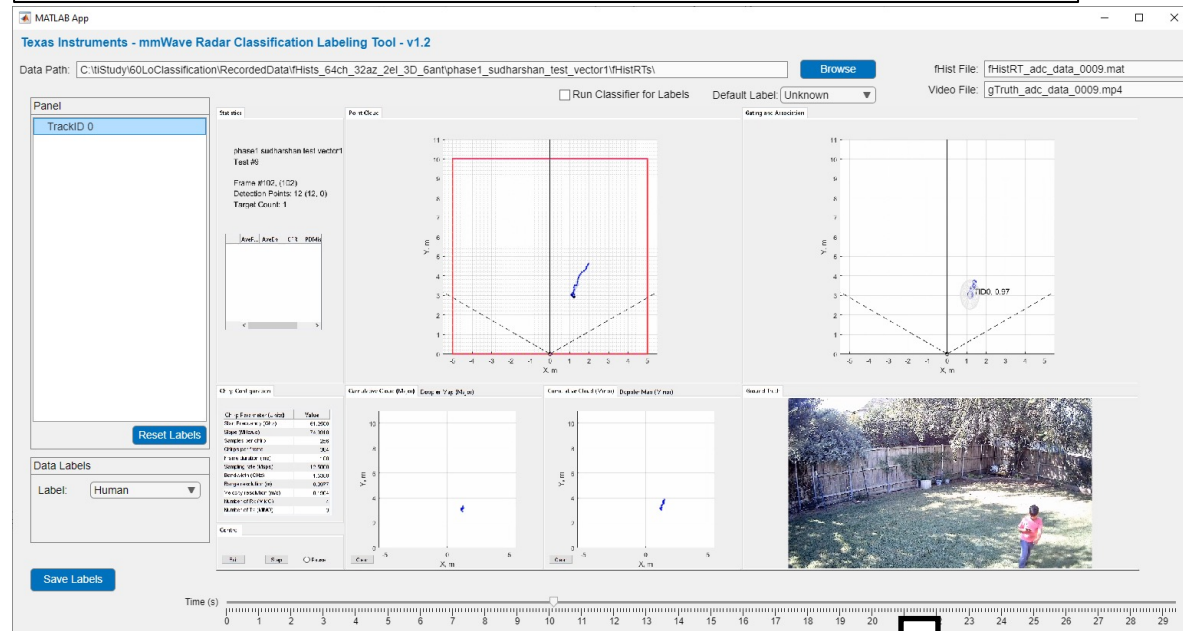
Data Capture and Feature Set Summary

Input the ADC raw data

Point cloud: range, angle, Doppler
Tracker output: position, velocity.

μ -Doppler generation.

Data labeling (with synchronized videos).



A labeling GUI is built in Matlab.

Model/train/evaluate the classifier

Data statistics:

- **310** scenarios in distinct environments.
- Total **125816** frames (**3.5hours**).
- From ~20 different human and numerous non-human targets.


Feature set statistics:

- 2D μ -Doppler maps
- 4 features (time sequence)
- **55079** Observations
 - Block length: 30 frames
 - **25448** human, **29631** non-human

- phase1_empty_no_specific_nonhuman_target
- phase1_human_and_nonhuman_fan_plant
- phase1_human_random_walk_run
- phase1_nonhuman_fan_plant
- phase1_outdoor_dogpark
- phase1_outdoorpark_human_nonhuman_tests
- phase1_sudharshan_test_vector1
- phase1_sudharshan_test_vector2
- phase1_two_human_random_walk_run
- phase2_20220307_dan_test_vector
- phase2_20220308_chris_test_vector
- phase2_20220309_ken_test_vector
- phase2_20220323_muhammad_drone_data
- phase2_20220324_dog_data_with_jackson
- phase2_20220325_dog_data_with_angie
- phase2_20220330_tree_data
- phase2_20220401_dog_data_with_alec
- phase2_20220412_indoor_plant_data
- phase2_20220418_muhammad_frontyard_human_data

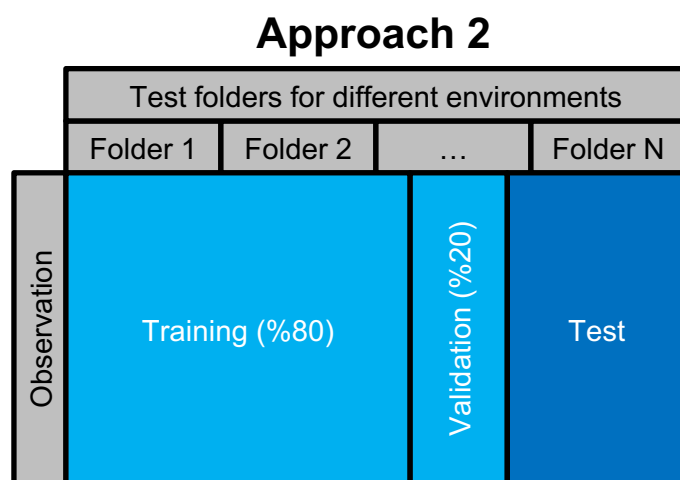
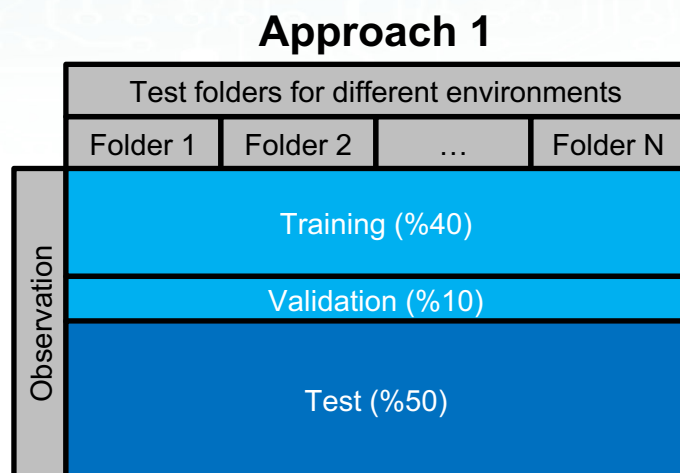
- fHistRT_adc_data_0001.mat
- fHistRT_adc_data_0002.mat
- fHistRT_adc_data_0003.mat
- fHistRT_adc_data_0004.mat
- fHistRT_adc_data_0005.mat
- gTruth_adc_data_0001.mp4
- gTruth_adc_data_0002.mp4
- gTruth_adc_data_0003.mp4
- gTruth_adc_data_0004.mp4
- gTruth_adc_data_0005.mp4

Various people and non-human objects (drone, different dogs, fans, trees, plants, etc.) are available in the data.

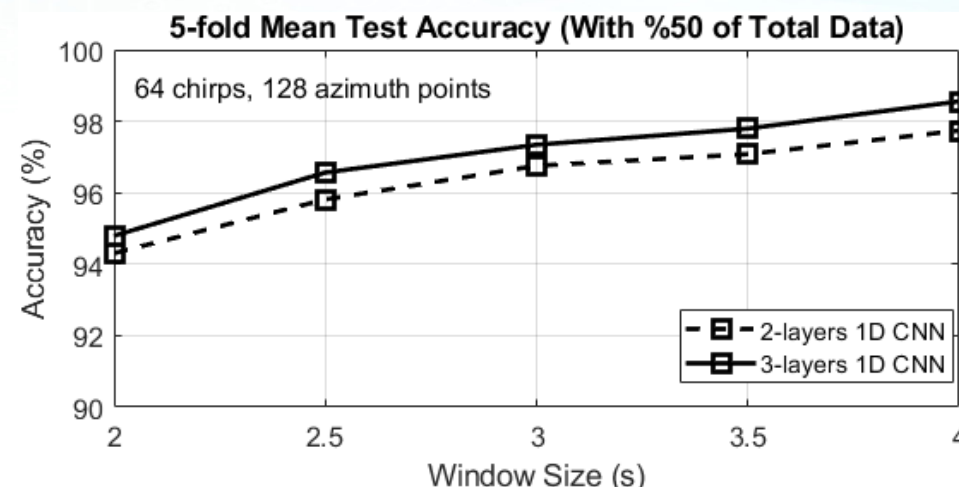


Cross Validation and Performance Results


- The 1st approach uses all the data mixed from different environments. Data split: 40% Train, 10% Validation, **50% Test**.
- The 2nd approach **reserves some folders completely for testing**. The remaining data is used for training + validation



The final classification error is calculated with 5-fold cross-validation.


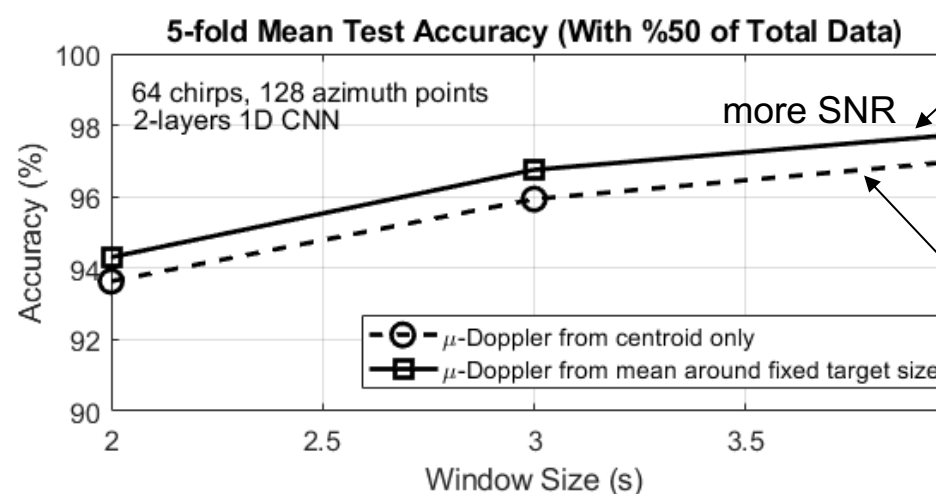


Approach 1 (All the data is mixed)



Approach 2 (Reserved folders for test)

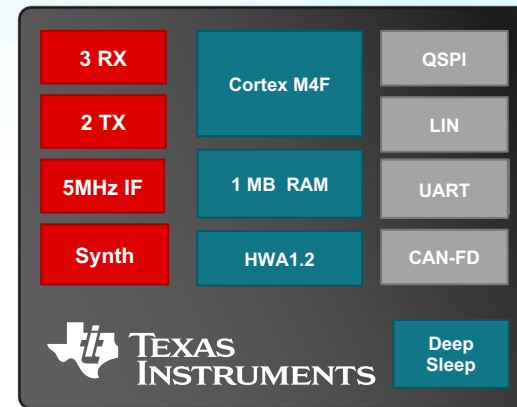
Repeated 5 times with different combinations



Summary: ~97% mean classification accuracy with only 3sec latency (2-layers 1D-CNN)

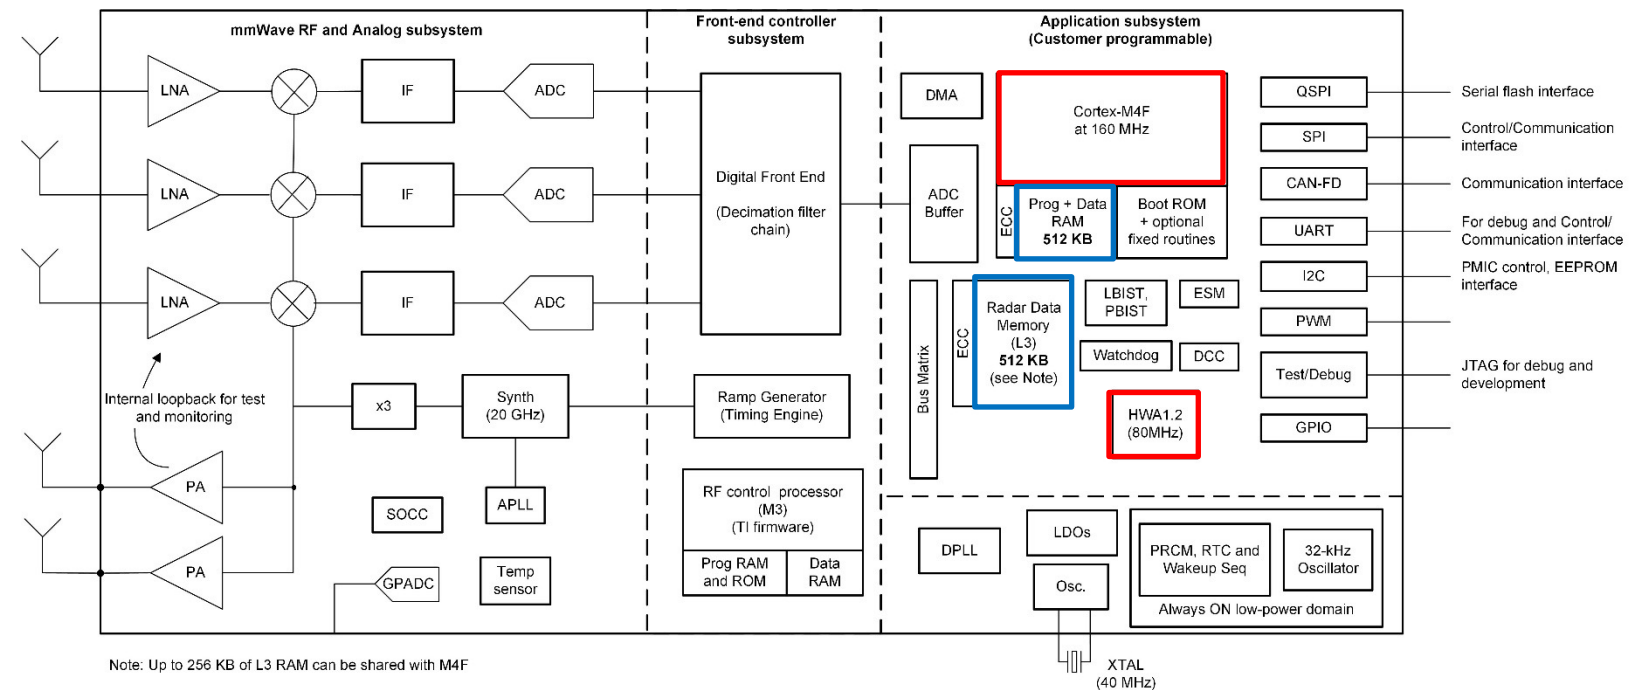
Enable Edge AI – IWRL6432 Overview

IWRL6432: Cost and power-optimized radar on a chip with analog (RF, IF), digital front-end, and processing (M4F+Accelerator) capability in one device.



← The target classification algorithm discussed in this work is implemented on this tiny device to create an edge AI solution.

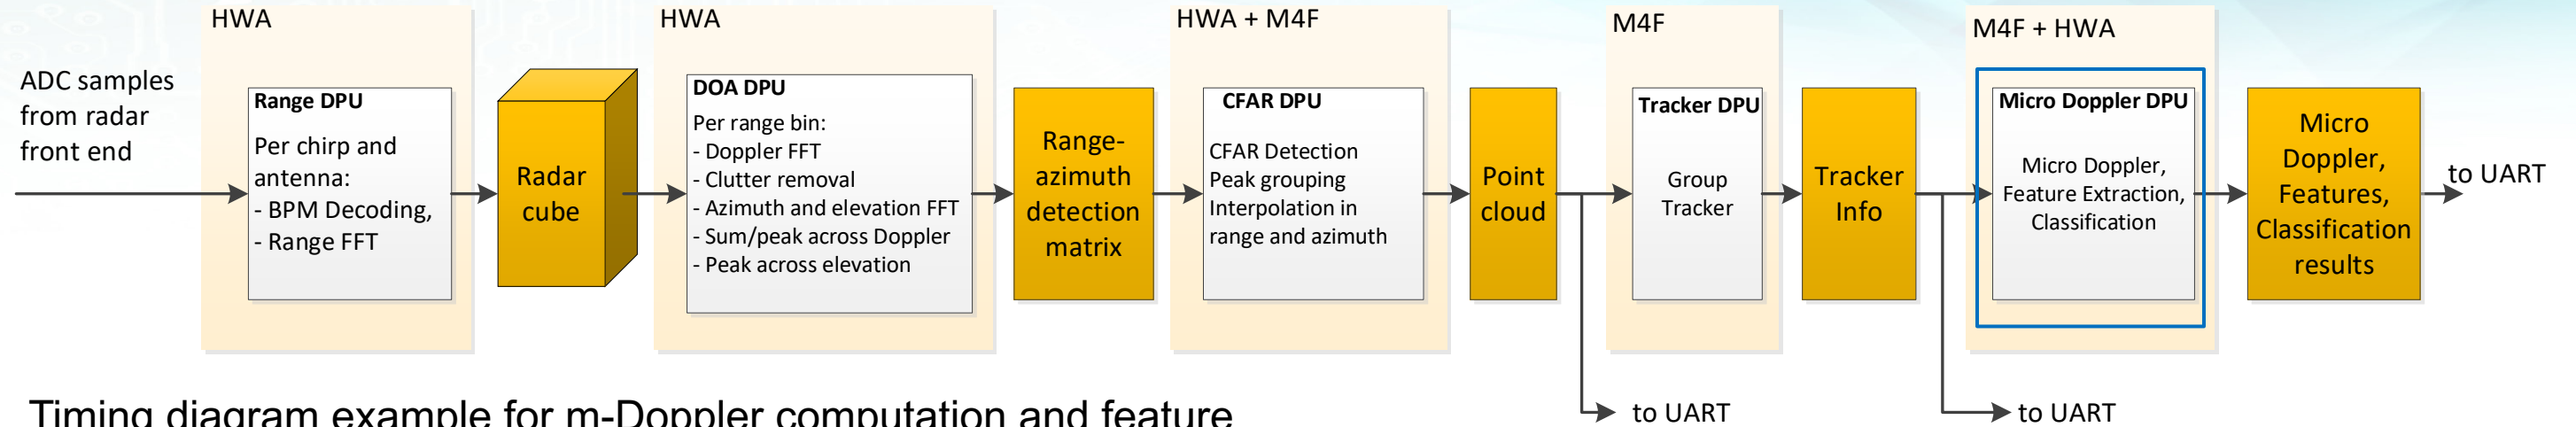
- Integrated transceivers : 3 Rx and 2 Tx.
- Integrated frequency synthesizer: 57GHz to 64GHz
- Processing:
 - ARM Cortex M4F MCU @ 160 MHz
 - Radar hardware accelerator (HWA) @ 80 MHz
- Memory: 1.0 MB RAM
- Interfaces: SPI, UART, I2C, QSPI, GPIOs
- Power: <100mW @ 10Hz (including classifier)
- Package options:
 - ~6.45 x 6.45 mm FCCSP (0.5mm pitch)
 - ~5 x 4.5mm WCSP (0.4mm pitch)




The entire real-time implementation is architected around Cortex M4F and HWA and fits into the 1MB of total memory.

Enable Edge AI – Implementation Details

Detailed view of the signal processing chain on IWR16432. Each block is implemented as a **DPU** and partitioned across **HWA** and **M4F**.



Timing diagram example for m-Doppler computation and feature extraction for three tracking objects

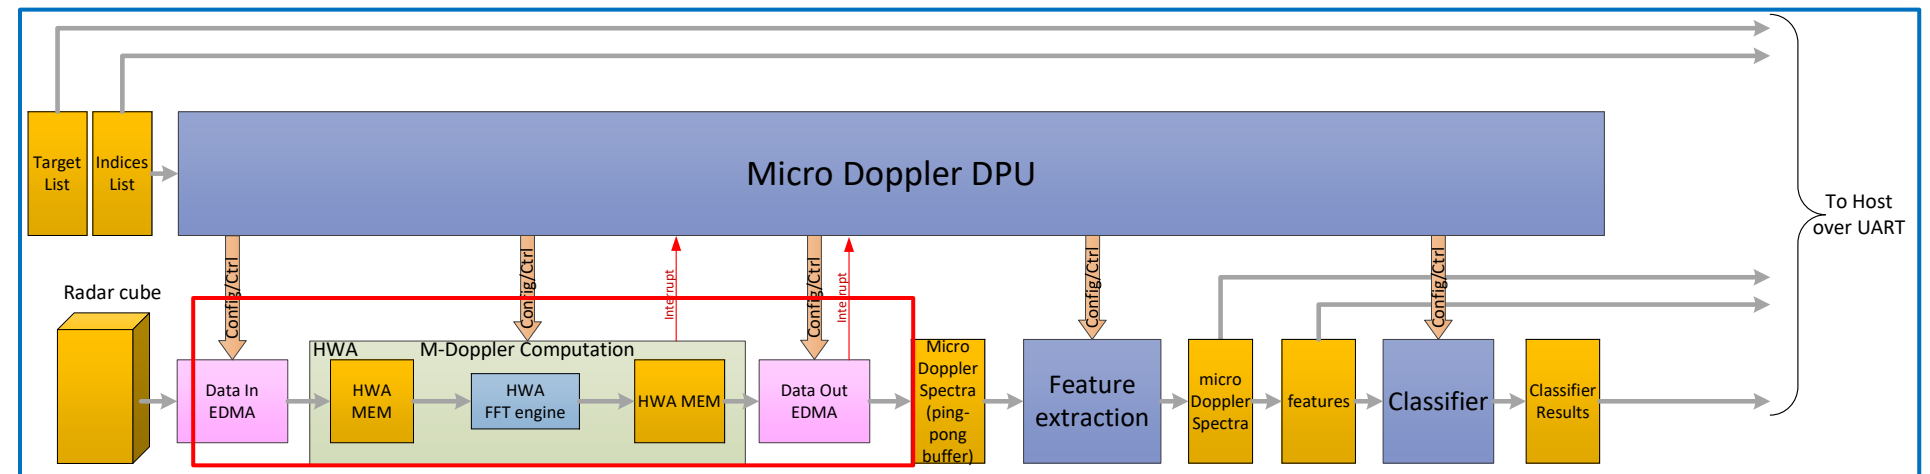


- The entire end-to-end chain takes only **~35ms** for 3 tracks.
- Each additional track requires an additional **~3ms** (tracker + classifier).

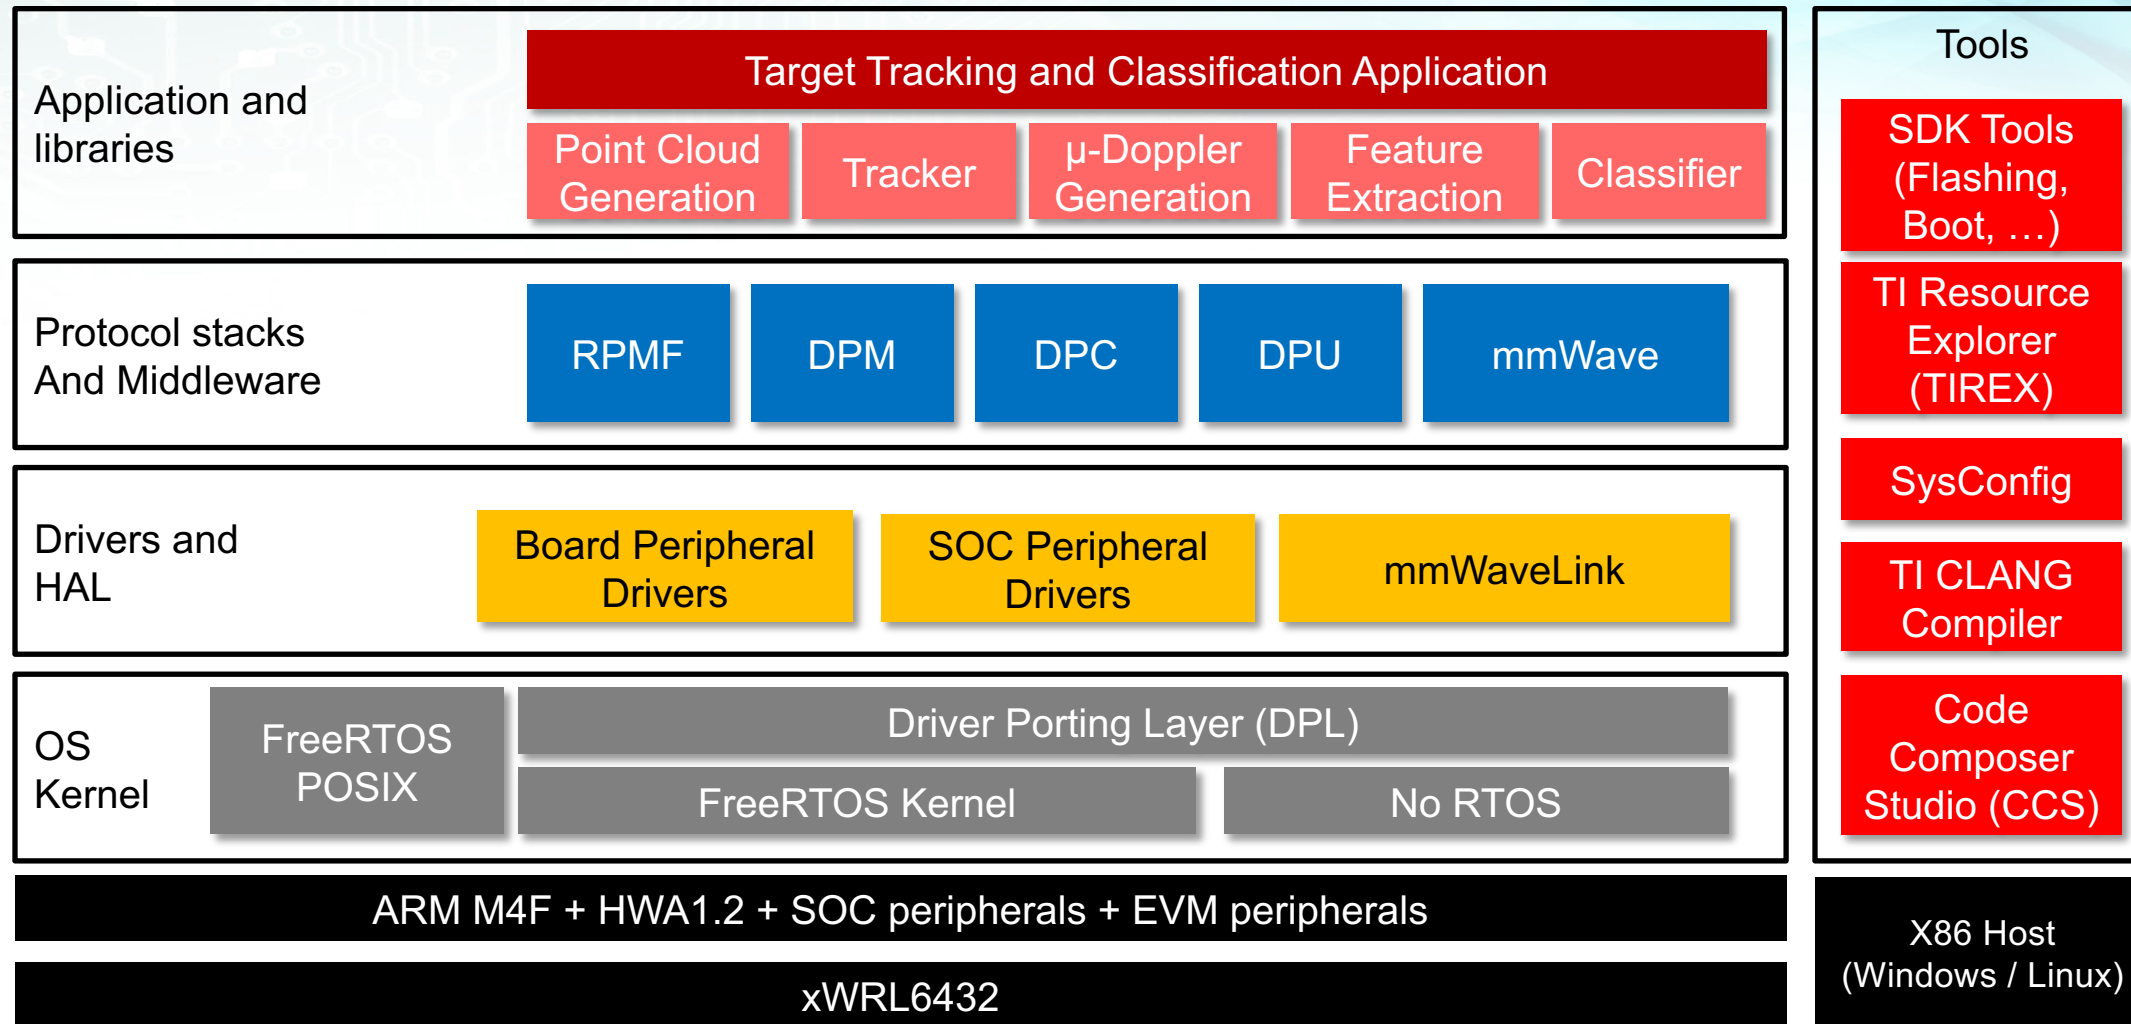
Highly parallelized architecture across processing cores to improve processing efficiency.

Efficient utilization of the HWA engine and EDMA channels is critical to the implementation, making the end-to-end chain work!

Note: See the next page for glossary.



Enable Edge AI – Software Architecture



Glossary


- RPMF: Radar Power Management Framework
- DPM: Data Processing Management
- DPC: Data Processing Chain
- DPU: Data Processing Unit
- SOC: System-on-Chip

Highlights:

- Open Source OS – FreeRTOS
- Simplified, low memory, low latency optimized drivers
- Pre-integrated mmWave components.
- Plug n Play middleware

Available on TI.com!

- Software: [MMWAVE-L-SDK](#)
- Evaluation module: [IWRL6432BOOST](#)



Radar Academy



TIREX

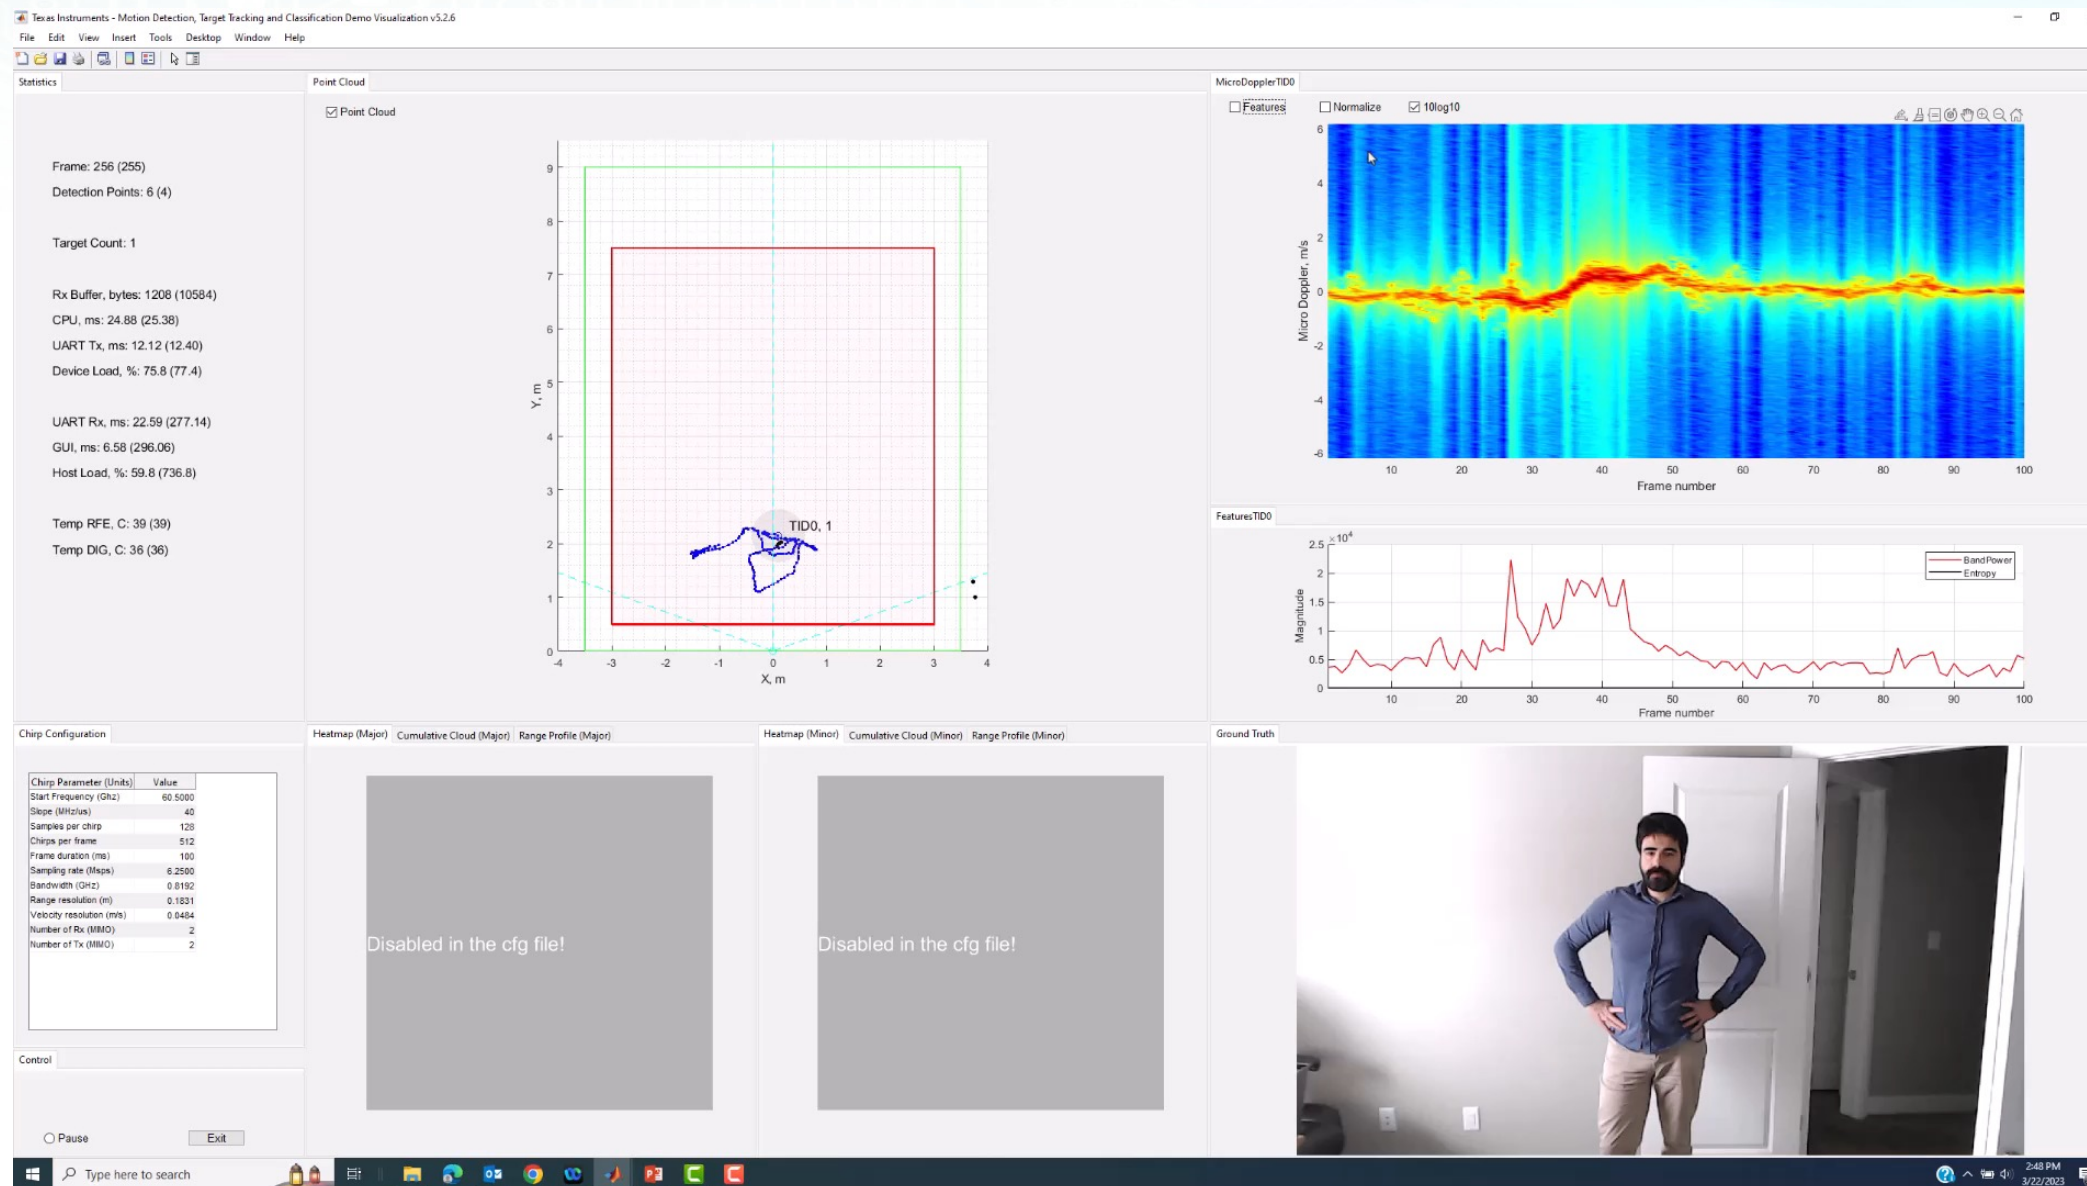


CCS




Real-Time Demo – μ -Doppler Generation

- In the real-time demo on IWR6432, μ -Doppler spectrograms from multiple tracks can be generated and streamed.
- Different ML techniques can be explored using this information without putting effort into lower-level processing.




- The micro-Doppler maps of a human vs. non-human object:



RC helicopter

- This feature can also help in more use cases (pose estimation, fall detection, etc.) when privacy is important.



Summary – Available on TI.com!

Summary: Using the multi-target tracker (MTT) integration, a robust μ -Doppler map generation method is built to support classifying **multiple objects concurrently** (to support more realistic scenarios) with reduced complexity.

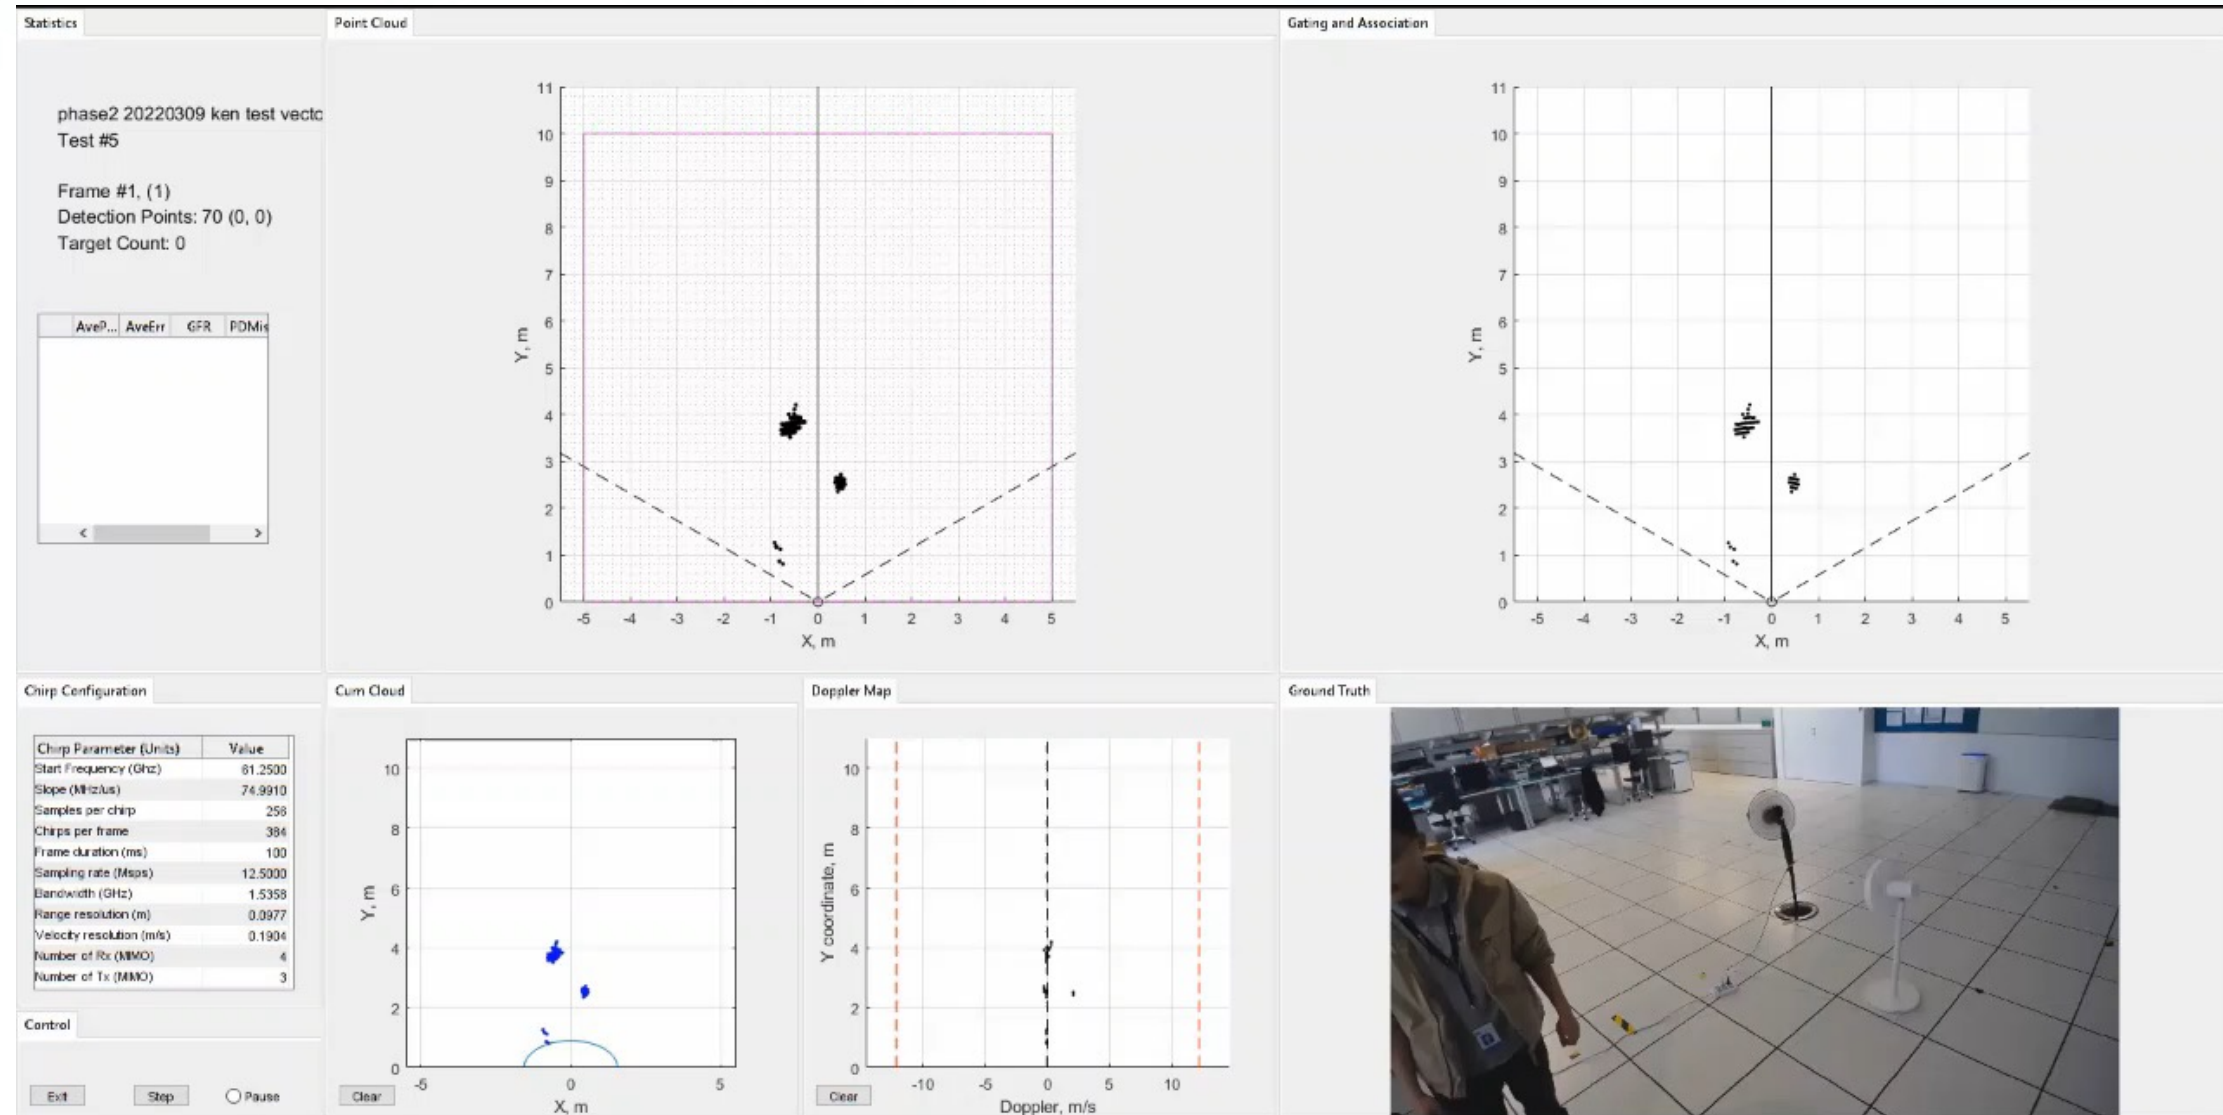
- Reduced the complexity of the classifier with hand-crafted feature extraction.
- Majority of the existing works use μ -Doppler maps directly.

- Enabled **simultaneous** tracking and classification of **multiple targets** in real-world environments.

- Majority of the existing works do not address such multi-object realistic scenarios.

Important facts:

- The machine learning model has never seen this specific scenario in training.





TEXAS INSTRUMENTS



Asia 2023

**Thank You!
Questions?**



Copyright Notice

This presentation in this publication was presented as a tinyML[®] Asia Technical Forum. The content reflects the opinion of the author(s) and their respective companies. The inclusion of presentations in this publication does not constitute an endorsement by tinyML Foundation or the sponsors.

There is no copyright protection claimed by this publication. However, each presentation is the work of the authors and their respective companies and may contain copyrighted material. As such, it is strongly encouraged that any use reflect proper acknowledgement to the appropriate source. Any questions regarding the use of any materials presented should be directed to the author(s) or their companies.

tinyML is a registered trademark of the tinyML Foundation.

www.tinyml.org