

Announcements

Please return your
lab kits!

Correction on lab 2:
**use “h100-80”, not
“h100”**

Carnegie Mellon University
Electrical & Computer Engineering

GRADUATE SEMINAR SERIES

OCTOBER 2 | 12 P.M. | Hamerschlag Hall 1107



Friday, October 2, 12 p.m.

Hamerschlag Hall 1107

***Internet of Radar Things: Turning a World of Radios
into One Radar***

[Akshay Gadre](#), Assistant Professor, Electrical and
Computer Engineering Department, University of
Washington



Class Project Overview

18-848 RadarML

Tianshu Huang

Project Requirements

- Must involve a machine learning component
 - **Training and/or fine-tuning is not required**
- Must involve a real radar or wireless system
 - **Must capture and/or deploy on/with real hardware**
 - Demos are encouraged, but not required
- Groups of up to 4
 - **Students with similar ideas / topics should form groups**
 - If there are too many groups, individuals or groups of 2 will be strongly encouraged to merge

Project Resources

Data

- I/Q-1M Dataset (~30h AWR1843Boost, indoor + outdoor + bike)
- Tartan* Pre-Release (~60h AWR1843AOP, indoor + outdoor)

Pre-trained Models

- I/Q-1M (AWR1843Boost):
 - GRT & GRTv2 Model Family
- Tartan* (AWR1843AOP):
 - GRT 3D occupancy model
 - Possible availability of the other GRT family models, GRTv2

Project Resources

Radar Hardware

- AWR1843AOP lab kits
- Limited number of other single-chip radar kits
 - AWR1843Boost, AWR2944, ...

Compute

- Extremely tight due to PSC downtime.
- We can only provide training-class hardware to a limited number of groups.
- Have a inference-only or fine-tune-only backup plan.

Project Ideas


- Real time inference on an edge platform (Jetson Nano, Orin)
- Edge cases & failure analysis
- Interference characterization, detection & mitigation
- Cross-radar fine-tuning
- Generalization to different chirp configurations
- Human pose
- Gesture recognition
- ML-assisted SAR
- Radar-language alignment
- Collect & Train with DDM
- ...

Deadlines & Structure

10/20	Project Proposal
11/12	Progress Report
12/1-12/3	Presentation / Demos
12/4	Final Report

Project Signup Document

<https://docs.google.com/spreadsheets/d/1YA-5bwPVAA-IMyIY3IUczN4UVCAfUvD5d9kV2L7FhKg/edit?usp=sharing>



The screenshot shows a Google Sheets interface. At the top, the title bar reads '18-848 Projects & Teams' with icons for a star, a folder, and a cloud. Below the title bar is a menu bar with 'File', 'Edit', 'View', 'Insert', 'Format', 'Data', 'Tools', 'Extensions', and 'Help'. A toolbar contains various icons for undo, redo, print, copy, paste, zoom (100%), currency, percentage, decimal, thousandths, text color, background color, border, and text wrap. The active cell is G28, containing the formula '=fx'. The spreadsheet has four columns: 'Project', 'Proposal', 'Open for Teammates?', and 'Team Member(s)'. Row 1 is the header. Row 2 contains an example project description, a checked checkbox for 'Project Proposal Template', a 'No' dropdown, and the name 'Tianshu Huang'. Rows 3 through 7 are empty.

	A	B	C	
1	Project	Proposal	Open for Teammates?	Team Member(s)
2	Example Project and a Brief Description. Keep this reasonably short and within a few sentences.	☒ Project Proposal Template	No	Tianshu Huang
3				
4				
5				
6				
7				

Project Proposal

Don't wait until you've finalized your group/project to write the proposal!

- Use a draft proposal to search for group members
- Reach out to projects which you are interested in joining

18-848 Project Proposal

 Project Proposal Template

Limit: 1000 words total.

Status: **Final**

Authors: (tag your group members!)

Reviewers: Tianshu Huang

LGTM:

Revision History:

Date	Description
Sep 30, 2026	Template ready (please change the status from "Final" to "Draft" when you prepare your proposal; then change to "In review" once you submit). Don't forget to clear this revision history as well!

Context & Problem Statement

What are you trying to accomplish? Why are you trying to accomplish this?

Goals & Non-Goals

What specific goals do you have for this project? On the other hand, what aspects are you not planning to pursue?

Proposed Solution or Methodology

If you are seeking to solve a particular problem or build a system, what is your approach and technical solution? On the other hand, if you are seeking to answer a research question, how do you plan to do so?

Alternatives Considered

What other approaches or solutions did you consider, and why did you choose the one you proposed?

Evaluation & Outcomes

How will you evaluate the success of your project? What are the expected outcomes and how will you measure them?

Resources Needed

Briefly describe the required resources. What sensor(s) are needed? What compute resources are needed? What sensor(s) and/or compute resources do you have access to already, and what resources will need to be provided?

Next Class

- Brainstorm ideas
- If you have an idea already, please add it to the project list
- We will leave ~30 minutes to discuss project ideas with each other at the end of class

Learning on Spectrum

18-848 RadarML

Tianshu Huang



Learning for Radar


- Lossy Representations
- Lossless Representations
 - Range-Doppler Geometry
 - 4D Spectrum / GRT
- Complex Data
 - Phase Ambiguity
 - Spatial Signals / GRTv2
 - Time Signals



Learning for Radar

- **Lossy Representations**
- Lossless Representations
 - Range-Doppler Geometry
 - 4D Spectrum / GRT
- Complex Data
 - Phase Ambiguity
 - Spatial Signals / GRTv2
 - Time Signals


Point Clouds




Point Cloud Models

- Generally interchangeable with Lidar methods
- Point cloud models
 - NVRadarNet
 - PointTransformer
- BEV projection methods
 - SWFormer
 - PointPillars

“Flattened” BEV Models



“Flattened” BEV Models







Learning for Radar

- Lossy Representations
- **Lossless Representations**
 - **Range-Doppler Geometry**
 - 4D Spectrum / GRT
- Complex Data
 - Phase Ambiguity
 - Spatial Signals / GRTv2
 - Time Signals


Classical Signal Processing Pipelines



Geometric Doppler




Geometric Doppler




Doppler Carries (Messy) Geometric Information

What about dynamic objects?



Doppler Carries (Messy) Geometric Information



Geometric Doppler

See DART, Appendix A.5 for a (pre-AI!) formalization




Figure 12. Illustration of our integration arc equation derivation.

for velocity vector v . We select p such that its direction is aligned with the positive x axis (which define to be the front of the sensor), i.e.

$$p \propto [1 \ 0 \ 0]^T - \frac{v_x v}{\|v\|_2^2}, \quad (12)$$

and define q accordingly.

Each range-Doppler bin is modeled as the intersection of a sphere with radius r and a cone $\langle v, w \rangle = d$. When intersected with the half-space corresponding with $x \geq 0$, this results in an arc or circle on the p - q plane (or the empty set if degenerate). Since we define the sensor to be facing in the $+x$ axis, and define p to be aligned with $+x$, the valid integration arc will always be $(-\psi, \psi)$ for angle ψ counterclockwise from the $+p$ axis, again relative to the p - q plane.

We now calculate this angle ψ . We begin by calculating the distance γ to the center of the integration arc and its radius $\tilde{r}$ (Fig. 12a). Choosing an arbitrary direction w in our p - q - v basis, the cosine of the angle η between v and w yields

$$\cos(\eta) = \frac{\gamma}{r} = \frac{\langle v, w \rangle}{\|v\|_2 \|w\|_2} \implies \gamma = \frac{rd}{\|v\|_2} \quad (13)$$

using the definition of Doppler d and orthonormality of w . It follows that

$$\tilde{r} = \sqrt{r^2 - \gamma^2} = r \sqrt{1 - \frac{d^2}{\|v\|_2^2}}. \quad (14)$$

Next, we calculate the length of the chord corresponding to the integration arc $\pm\psi$. Let h be the distance from the chord to v (Fig. 12b); h can be specified relative to γ and the angle θ between v and the $+x$ axis (Fig. 12c). It follows that

$$h = \gamma \tan(\theta) = \frac{rd}{\|v\|_2} \left(\frac{v_x}{\sqrt{\|v\|_2^2 - v_x^2}} \right), \quad (15)$$

which finally yields (Fig. 12b)

$$\psi = \arccos\left(\frac{h}{\tilde{r}}\right). \quad (16)$$

We also must account for three degenerate cases:




Figure 13. Annular section A for the revolved solid approximation of Eq. 5. η is the angle to the velocity vector v .

- When $d > \|v\|_2$, the Doppler bin exceeds the velocity (and Eq. 14 becomes imaginary).
- When $h > \tilde{r}$, the entire doppler cone is pointing behind the radar (i.e. the $x > 0$ half-space no longer intersects with the cone; Eq. 16 is undefined).
- When $h < -\tilde{r}$, the entire doppler cone is in front of the radar. In this case, we instead have $\psi = \pi$.

Volume of integration region Expanding on Eq. 4 more explicitly using our integration bound ψ , the radar return $\mathbf{Y}(r_i, d_j, k)$ can be expressed as the circular integral


$$\mathbf{Y}(r_i, d_j, k) \propto \tilde{r} \int_{-\psi \leq \phi \leq \psi} C(r, k; w(\phi)) d\phi \quad (17)$$

for projection w to the integration arc.


While our rendering equation (Eq. 4) makes a “thin-shell” assumption, we must also examine the effect of *bin size* in order to correctly determine which constants to apply to our approximation. Let ΔR and ΔD be the width of each range and Doppler bin. Instead of using a “point” approximation of C , we instead derive the integral based on the revolved solid of an annular section A (Fig. 13). Using again the fact that $\cos(\eta) = d/\|v\|_2$, we arrive at the approximation

$$\begin{aligned} dA &= r_i dr d\theta = \frac{r_i}{\|v\|_2 \sqrt{1 - \frac{d^2}{\|v\|_2^2}}} dr dd \\ &= \frac{r_i^2}{\|v\|_2^2 \tilde{r}} dr dd, \end{aligned} \quad (18)$$


Micro-Doppler



Range–azimuth




Range–Doppler





Micro-Doppler

- People
- Animals
- Wheels
- Propellers
- ...



Micro-Doppler analysis in ISAR - Review and perspectives, V.C. Chen et al, 2009






Learning for Radar

- Lossy Representations
- **Lossless Representations**
 - Range-Doppler Geometry
 - **4D Spectrum / GRT**
- Complex Data
 - Phase Ambiguity
 - Spatial Signals / GRTv2
 - Time Signals


I/Q-1M Dataset

- 29 hours of data
- 1M radar-lidar-camera samples





Generalizable Radar Transformer



GRT Small IQ-1M
bike/ ptbreeze.back
+00:52.93s
radar 001077
lidar 000550
camera 001689

Birds Eye View Occupancy (CFAR)




Birds Eye View Occupancy (GRT)



Birds Eye View Occupancy (Lidar)



Input (4D Radar Spectrum)



Depth (CFAR)



Depth (GRT)



Occupancy (Lidar)



RGB (Ground Truth)



Segmentation (GRT)



Segmentation (RGB)







GRT Depth




Traditional Point Cloud




Lidar



GRT Birds Eye View



GRT Model



Geometric Doppler




Figure 20. Relative test losses binned by the speed of the data collection rig (20 equal 5% bins) on the *bike* subset; due to less available Doppler information, lower speeds are associated with higher losses for each of our tasks.

Table 4. **Test loss for each ablation** (smaller is better) relative to GRT-small trained on our full dataset, along with 95% confidence intervals for the relative differences.

	Ablation	Relative Test Loss
Inputs	Amplitude Only	$+0.04 \pm 0.85\%$
	Angle of Arrival	$+28.0 \pm 2.12\%$
	CFAR Thresholding	$+31.5 \pm 2.38\%$
Doppler	Without Doppler FFT	$+22.5 \pm 1.76\%$
	Slow Time Shuffled	$+23.10 \pm 1.94\%$
Post-Patch Axes	Doppler-Az-El	$+6.22 \pm 1.11\%$
	Range-Az-El	$+6.27 \pm 1.10\%$
	Range-Doppler-Az-El	$+4.18 \pm 1.09\%$
Augmentations	None	$+5.87 \pm 1.32\%$
	Scale, Phase, and Flip Only	$+3.89 \pm 1.19\%$
Separate Models	Indoor Data Only	$+5.76 \pm 1.85\%$
	Outdoor Data Only	$+5.58 \pm 1.26\%$
	Bike Data Only	$+2.77 \pm 1.41\%$

Geometric Doppler




Figure 19. **Sample sequences of three consecutive frames before (top) and after (bottom) stopping at a red light.** Video frames are provided for reference (left), with GRT's semantic segmentation (center left) and depth (center right) outputs across the two sequences along with the range-Doppler spectrum. After stopping, the Doppler spectrum (horizontal axis; right) collapses to a single Doppler bin, resulting in significantly decreased information available to the model. This manifests as noisier (as seen by larger frame-to-frame variations and hallucinations) and less accurate predictions by the model.

Scaling Laws






Figure 6. **Scaling laws for mmWave radar transformers** across `indoor`, `outdoor`, and `bike` test splits. Models and confidence intervals are measured relative to the `small` transformer trained on 10% of the dataset. While our models show weak scaling over model size when trained on our dataset, we see strong log-linear scaling across dataset size of $\approx 20\%$ loss decrease per $10\times$ increase in data size. For variants of this graph with respect to absolute metrics, see App. C.2.

Key Results: Towards Foundational Models?






How does GRT look in retrospect,
10+ years after its initial development?

*1 human year = 7 machine learning years

GRT: What Aged Well?


Overall methodology is robust, and has been independently replicated



GRT: What Aged Well?

Encoder-decoder with positional decoder query

A popular choice in
contemporary
perception
foundation models,
e.g. D4RT (CVPR
Best Paper, 2026)



A. Model Overview & Training Details






Figure 7. **Full D4RT model overview** – We provide a holistic overview of the model together with its inputs and outputs. FC corresponds to a fully connected layer, and PE for positional encoding. See Sec. 2 of the main paper for reference.



GRT: What Aged Well?

Data scaling is highly robust



GRT: What Aged Poorly?

Phase angle is a bad representation!



GRT: What Aged Poorly?


Oops!

Table 4. **Test loss for each ablation** (smaller is better) relative to GRT-small trained on our full dataset, along with 95% confidence intervals for the relative differences.


	Ablation	Relative Test Loss
Inputs	Amplitude Only	$+0.04 \pm 0.85\%$
	Angle of Arrival	$+28.0 \pm 2.12\%$
	CFAR Thresholding	$+31.5 \pm 2.38\%$
Doppler	Without Doppler FFT	$+22.5 \pm 1.76\%$
	Slow Time Shuffled	$+23.10 \pm 1.94\%$
Post-Patch Axes	Doppler-Az-El	$+6.22 \pm 1.11\%$
	Range-Az-El	$+6.27 \pm 1.10\%$
	Range-Doppler-Az-El	$+4.18 \pm 1.09\%$
Augmentations	None	$+5.87 \pm 1.32\%$
	Scale, Phase, and Flip Only	$+3.89 \pm 1.19\%$
Separate Models	Indoor Data Only	$+5.76 \pm 1.85\%$
	Outdoor Data Only	$+5.58 \pm 1.26\%$
	Bike Data Only	$+2.77 \pm 1.41\%$

GRT: What Aged Poorly?

Azimuth-elevation is also a bad representation!



TI AWR1843
Single Chip, 3x4



GRT: What Aged Poorly?

Lab 2 & Lecture 9/15 Applied ML – What's wrong with this plot?




Figure 6. **Scaling laws for mmWave radar transformers** across indoor, outdoor, and bike test splits. Models and confidence intervals are measured relative to the `small` transformer trained on 10% of the dataset. While our models show weak scaling over model size when trained on our dataset, we see strong log-linear scaling across dataset size of $\approx 20\%$ loss decrease per $10\times$ increase in data size. For variants of this graph with respect to absolute metrics, see App. C.2.



Learning for Radar

- Lossy Representations
- Lossless Representations
 - Range-Doppler Geometry
 - 4D Spectrum / GRT
- **Complex Data**
 - **Phase Ambiguity**
 - Spatial Signals / GRTv2
 - Time Signals

Next Class

- Brainstorm ideas
- If you have an idea already, please add it to the project list
- We will leave ~30 minutes to discuss project ideas with each other at the end of class