



L11: Learning for Spectrum (Part 2)

18-848 RadarML


Tianshu Huang



Learning for Radar

- Lossy Representations
- Lossless Representations
 - Range-Doppler Geometry
 - 4D Spectrum / GRT
- **Complex Data**
 - **Dealing with Complex Data**
 - Spatial Signals / GRTv2

Recap: range-Doppler spectrum




Spectrum is not a 1:1 match for Images

	Radar spectrum	Camera image
Point features	Most objects	Usually noise
Empty bins	Unoccupied bins	Only in the dark
Noise model	Additive <u>Gaussian</u> (thermal)	Mostly Poisson (shot); Gaussian in low light
Materials	Mostly specular	Mostly diffuse


ADC / fine signal: gaussian $\xrightarrow{\text{FFT}}$ gaussian

Recap: Range vs Phase




round trip



<u>77-81 GHz</u>	{	2π of phase	$= \lambda / 2 = c / 2f$	$\approx 1.875 \text{ mm}$
		1 range bin	$= c / 2B$	$\approx 37.5 \text{ mm}$




Phase can be Ambiguous







Neighboring Bins are not necessarily "Coherent"

Coherent: data/signals with consistent phase



$\angle \text{bin 1} \perp \angle \text{bin 2}$


Recap: Doppler vs Phase




"far field"

Recap: Angle vs Phase


$1.5\lambda = 6\text{mm}$



Broadside




Angled



End-fire

Complex Radar Signals Carry High Frequency Information

- Neighboring bins can have very different phase
- Smooth operations which add/average phasors which are out of phase can cancel
- Small shifts can scramble phase



“Leftover” phase

$$\varphi = \boxed{4\pi R / \lambda} + \boxed{\varphi_{hw}} + \boxed{\varphi_{sp}} + \boxed{\varphi_{target}}$$

Round-trip phase **Hardware delay** **Signal processing** **Target reflectivity**

Constant for a given sensor system Constant for each target

“complex reflectivity”

Real (Magnitude) vs Complex Signals

If complex signals are hard to work with, can we just flatten them into real signals?

Pros:

- Real spectrum behaves more like images
- Can be downscaled/upscaled etc

Cons:


- Complex reflectivity, sub-wavelength position are permanently lost
- Angular information is lost unless there are many (infinite) angular bins



Learning for Radar



- Lossy Representations
- Lossless Representations
 - Range-Doppler Geometry
 - 4D Spectrum / GRT
- Complex Data
 - Dealing with Complex Data
 - **Spatial Signals / GRTv2**

GRTv2: Leverage the Spatial Signal



Spatial vs Angular Spectrum

- Zero-pad the angular spectrum => logarithmic improvements?
- **Conclusion: the model is not learning to use phase across antenna**



Why Spatial Spectrum?


Pros:

- Theoretically most dense representation
 - $\Rightarrow$ less memory, “narrower” network
- Remove a potentially biased layer of signal processing
- Build a stronger understanding of learning on spectrum
 - $\Rightarrow$ one layer closer to learning on time signal

Cons:


- Harder for models to learn / more data-hungry
- Doesn't directly transfer across different devices
- Unintuitive and tricky to work with

SSL: Contrastive Learning




SSL: Self-Distillation

- Very popular due to large open-source investments by Meta (DINO)
- Expensive and hard to tune



V-JEPA 2






SSL: (Masked) Reconstruction

- Less popular than self-distillation
- Still used in some SOTA models (e.g., D4RT)




Figure 1: **VideoMAE** performs the task of masking random cubes and reconstructing the missing ones with an asymmetric encoder-decoder architecture. Due to high redundancy and temporal correlation


SSL Works!




SSL Eliminates the Impact of Signal Processing



GRT @ 24h (100%)




GRTv2 @ 24h (100%)




Camera Reference




Lidar Ground Truth




GRT @ 24h (100%)




GRTv2 @ 24h (100%)




Camera Reference




Lidar Ground Truth



GRT @ 24h (100%)




GRTv2 @ 24h (100%)




Camera Reference



Lidar Ground Truth




GRT @ 24h (100%)




Camera Reference



GRTv2 @ 24h (100%)



Lidar Ground Truth





Paper Critique & Presentation

18-848 RadarML

Tianshu Huang

Paper Topics



RadarML @ CMU

Syllabus

Assignments

Assignments

Overview

Paper List

Lab 1

Lab 2

Project

Paper Critique Reading List

Why Learning on Spectrum?

RAD Tensor Detection

ICCVW 2019

Vehicle Detection With Automotive Radar Using Deep Learning on Tensors

1/ How to present a paper

1. What does the paper do?

- Describe how the technique / methodology works.
- How was it evaluated? **How well does it work?**

2. What is new?

- How does this work **differ** from prior works?
- Is anything about the proposed method and/or results **surprising**?

3. Why is it hard?

- Why did no one else do this previously?

2/ How to critique a paper

1. How does the work relate to applications and other domains?

- Other than researchers working on similar problems, **who cares** about this particular work?

2. How does the work relate to prior (and future) work?

- Not just what is different, but **why** is it different?
- What is the broader trajectory of the research direction?

3. What are the authors **not** saying?

- Is anything conveniently missing or unexplained?
- Are there any obvious things that the paper failed to do?

Your Job

Read your assigned paper, as well as other papers in your session

- Be prepared to answer technical and discussion questions
- Understand how your paper relates to the others

Prepare a 20-minute presentation for your assigned paper

- Address the key questions about your paper
- Plan for an additional 10 minutes of discussion



Sign Ups

- Link will be posted via Canvas & sent via email after class
- 18 students and 18 papers
- First come first served; feel free to trade with other students
- Calendar dates are subject to change

Next Class

Justin Chan:

Intelligent acoustic sensing systems at scale

Abstract: In this talk, I will present our work on intelligent acoustic sensing systems across two areas. First, in equitable healthcare, I will describe a suite of systems that bring health sensing to ubiquitous smartphones and hearables for detecting ear infections using only a smartphone and a paper cone, enable low-cost newborn hearing screening using inexpensive earphones, and brings micro-cardiac monitoring to all hearables. Second, in augmented audio, I will describe our work designing a bio-inspired ear that uses acoustic microstructures to give smartphones directional hearing, as well as our work designing active noise cancellation for smart glasses. Through these examples, I will demonstrate new computational and sensing techniques that generalize across hardware and work in real-world environments.