DEEP LEARNING ON RF DATA

Adam Thompson | Senior Solutions Architect

March 29, 2018

NVIDIA

AGENDA

Background Information – Signal Processing and Deep Learning

Radio Frequency Data Nuances

Complex Domain Representations and Applications

The Case for GPUs

Deployment

SIGNAL PROCESSING AND DEEP LEARNING REVIEW

SIGNAL PROCESSING PRIMER

Definitions and Applications

Signals can be broadly defined as a medium for transmitting information from one place to another

Signal processing is concerned with the manipulation of signals to exploit imbedded information to achieve a certain goal

Applications include feature detection, geolocation, demodulation, emitter tracking, amplification, and filtering among others

Wireless communication is a major component of the signal processing domain and features applications ranging from AM/FM radio to WiFi to RADAR

SPECTRAL CONSIDERATIONS

Definitions and Applications

Limited resource: with increasing popularity of wireless communication devices, the wireless spectrum has become congested

Certain frequencies are physically more desirable than others and the rise of spread spectrum communication

Spectral **limitations** include multipath, noise, and interfering signals

Motivation for both **signal identification** and **spectrum awareness**

Classical signal processing approaches are susceptible to false alarms and are often difficult to scale with emerging technologies

DEEP LEARNING REVIEW

Definitions and Applications

2012 – AlexNet fostered the ‘big bang’ in Deep Learning based on positive results with the ImageNet competition

Powered by NVIDIA Graphics Processing Units (GPUs) and massive amounts of labeled data

Training - generate a mapping between known input data and known labels

Inference - expose data unseen by the network for identification or classification

Traditional applications in imagery, video, and text

MARRIAGE OF DEEP LEARNING AND RF DATA

SIGNAL IDENTIFICATION

Learn features specific to a desired emitter

Fits into many existing RF dataflows

Success in high noise, high interference environments

ANOMALY DETECTION

Facilitates in discovery

Early warning system for defense and commercial applications

Enforce FCC regulations

SCHEDULING

Automatic recognition of free communication channels

Provide a basis for effective signal transmission or reception

RADIO FREQUENCY DATA

RADIO FREQUENCY DATA

Domains, Considerations, and Limitations

Raw RF signal data is complex valued and traditionally split into the inphase (I) and quadrature (Q) channels

Phase is important for signal processing and RF applications

Standard deep learning networks are not constructed for complex-valued data and, historically, work best on images

No large, commercial, labeled dataset like ImageNet exists for RF data

Complex data can be represented in multiple domains and typically represent time and frequency varying features

RF DATA DOMAINS

FM Collection – 90.1MHz, 1.8MHz Bandwidth

SPECTROGRAM

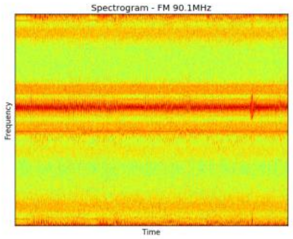

RAW I/Q

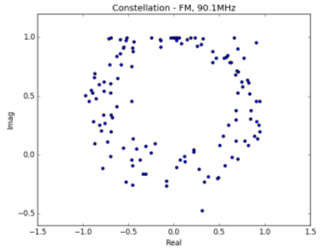

MAGNITUDE/PHASE

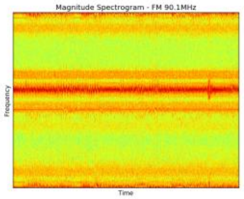

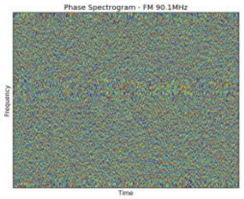

OTHERS

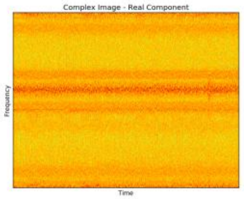

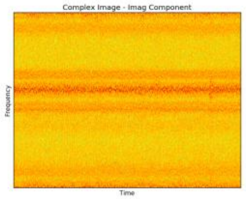

WITH THESE CHOICES, WHAT SHOULD I USE?

SPECTROGRAM DOMAIN

SPECTROGRAM APPROACHES

Overview

Historically most popular domain for RF deep learning research

Discards phase information and is most effective at signal identification

Makes use of standard image domain networks and is a candidate for transfer learning

Demands that the signal footprint is unique and easily separated from the RF environment by an experienced operator

Candidate for image segmentation techniques

DEMONSTRATION

KickView Corporation

Classification of simulated Linear Frequency Modulation (LFM) signals co-existing with noise and interference

Standard GoogLeNet model trained on a Tesla V100 with 30 epochs and 7,500 labeled images yielded the following confusion matrix on a test set of 2,000 images

Training time was 7 minutes and 43 seconds

	Neg	Pos	Accuracy
Neg	990	10	99.0%
Pos	5	995	99.5%

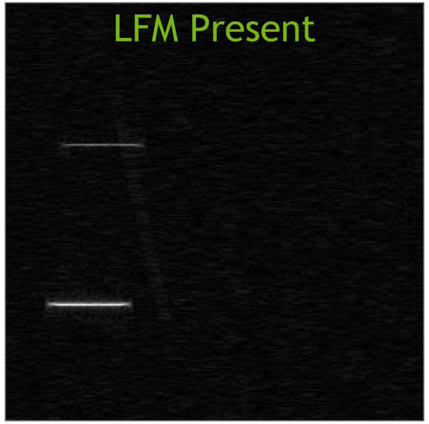

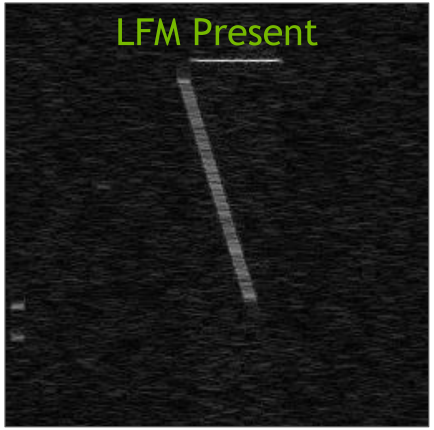

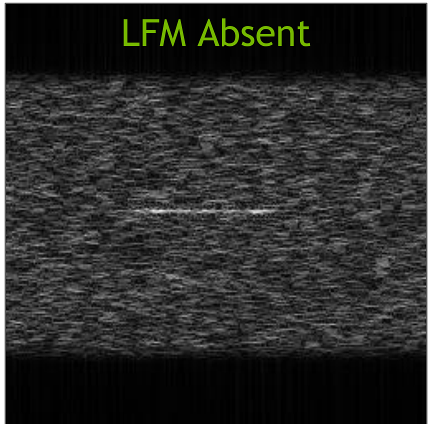

SPECTROGRAM SEMANTIC SEGMENTATION

Overview

Semantic Segmentation is the process of assigning labeled classes on a pixel-by-pixel basis

Commonly used in self driving automobiles and the remote sensing communities

Attempts to provide the true meaning of a given scene

For RF applications, can learn the duration of the transmission, operating frequency, and other emitter specific characteristics such a drift

Research overlaps with medical imaging

DEMONSTRATION

Semantic Segmentation

Manually labeled data by creating a boxed mask highlighting relevant signal energy

1000 training images and 100 validation images using a fully convolutional U-Net architecture shows initial promising results

Trained on a V100 with 30 epochs in 20 minutes and 24 seconds

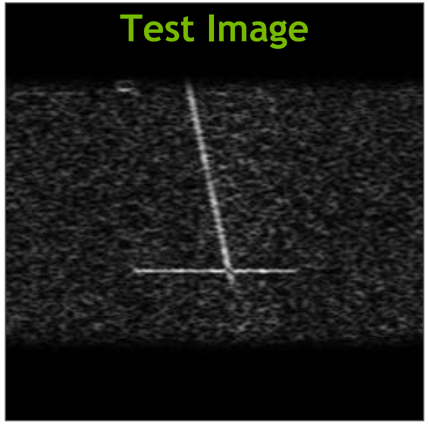

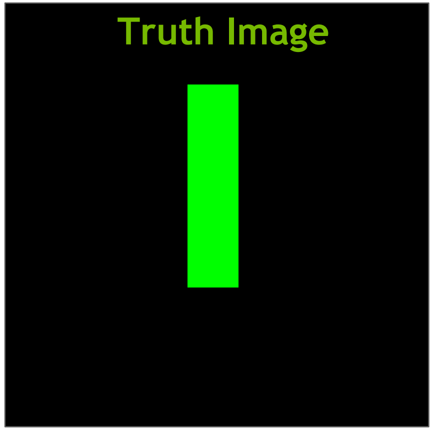

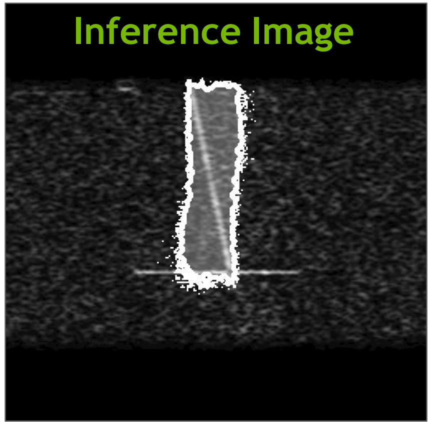

I/Q DOMAIN

I/Q APPROACHES

Overview

Allows deep learning to be applied to the sensor level and can facilitate real time decisions

Preserves phase information which is important in both demodulation and RADAR applications for determining characteristics about the target

Active research on modulation recognition by Tim O'Shea and DeepSig using simulated and OTA data

Training occurred with 120,000 synthetic examples using the ResNet architecture and a TitanX GPU (60 seconds/epoch) - 94% accuracy on simulated data and 87% on OTA

O'Shea et al.: Over the Air Deep Learning Based Radio Signal Classification - https://arxiv.org/pdf/1712.04578.pdf

COMPLEX IMAGE DOMAINS AND OTHERS

COMPLEX IMAGE DOMAIN AND BEYOND

Customer Success Story

KickView OFDM signal detection with simulated data – 20MHz, IEEE 802.11g

Using complex image domain: 2 channel (real and imaginary) stacked outputs of a polyphase channelizer

90% classification accuracy on full band transmissions down to -5.5dB (below noise floor)

90% classification accuracy on partial band transmissions (5MHz) down to 0.5dB

https://blog.kickview.com/deep-learning-meets-dsp-ofdm-signal-detection/

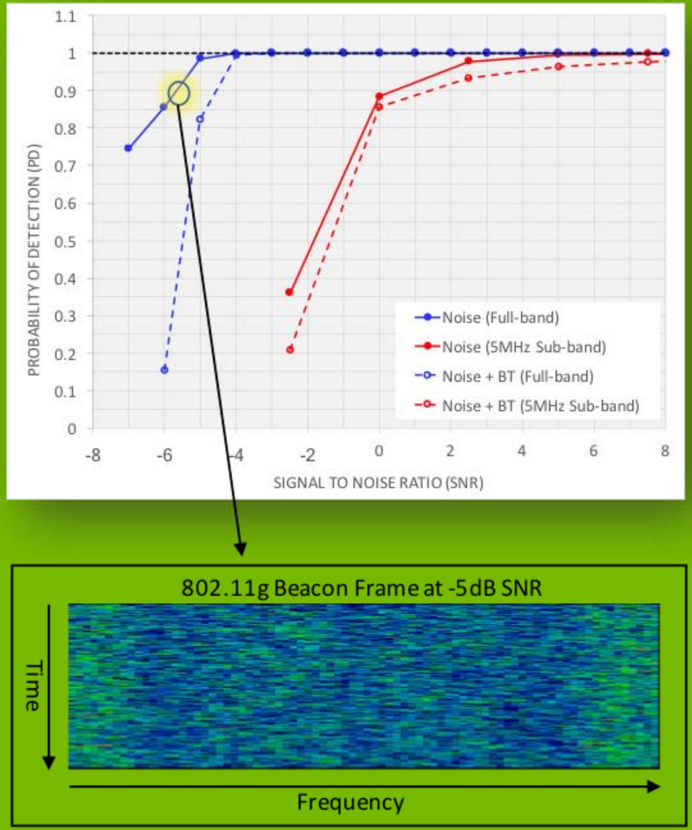

COMPLEX IMAGE DOMAIN AND BEYOND

Multiple data representations and pre-processing techniques specific to complex functions have not been explored in literature

Suggestions for research include I/Q spectral plots, N-dimensional tensors, and others

Desire to find apples-to-apples comparisons when defining a dataset and network architecture

Need for an open, collected dataset!

THE CASE FOR GPUS

THE CASE FOR GPUS

Optional subtitle

Signal processing applications consume a ton of data and real time processing is desired

Traditional signal processing techniques (filtering, windowing, Fourier analysis, eigenvalue decomposition) rely on dense linear algebra

Beyond High Performance Computing, GPUs necessitate fast training and inferencing and have the capability for field deployment

Support all major deep learning frameworks

DEPLOYMENT

EMBEDDED GPU SPECIFICATIONS

	JETSON TX1	JETSON TX2
GPU	Maxwell	Pascal
CPU	64-bit A57 CPUs	64-bit Denver 2 and A57 CPUs
Memory	4 GB 64 bit LPDDR4 25.6 GB/s	8 GB 128 bit LPDDR4 58.4 GB/s
Storage	16 GB eMMC	32 GB eMMC
Wi-Fi/BT	802.11 2x2 ac/BT Ready	802.11 2x2 ac/BT Ready
Video Encode	2160p @ 30	2160p @ 60
Video Decode	2160p @ 60	2160p @ 60 12 bit support for H.265, VP9
Camera	1.4Gpix/s Up to 1.5Gbps per lane	1.4Gpix/s Up to 2.5Gbps per lane
Mechanical	50mm x 87mm 400-pin Compatible Board to Board Connector	

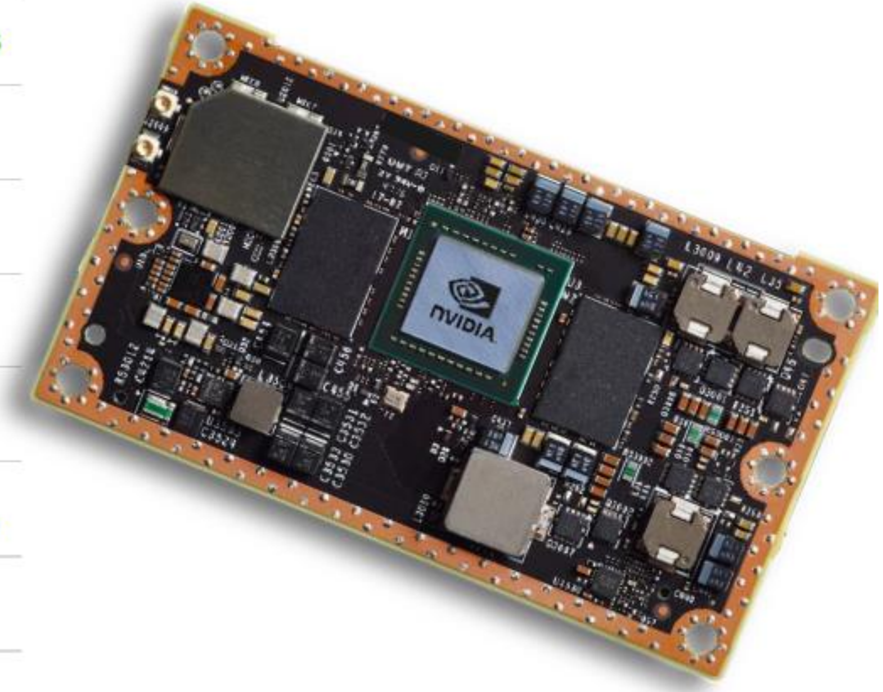

DEEPWAVE

AIR-T Hardware Solution

Software defined radio (SDR) designed for deep learning applications

Placing AI at the edge to process high bandwidth data in real time (> 1GB/s)

Includes FPGA for latency cognizant signal capture

Tegra series embedded GPU

NVIDIA