MATLAB EXPO 2019

AI Techniques in MATLAB for Signal, Time-Series, and Text Data

Paola Jaramillo
Application Engineer

The Use of Deep Learning is Growing Across Industries

Aerospace, Defense and Communications

Communication devices, security, smart cities

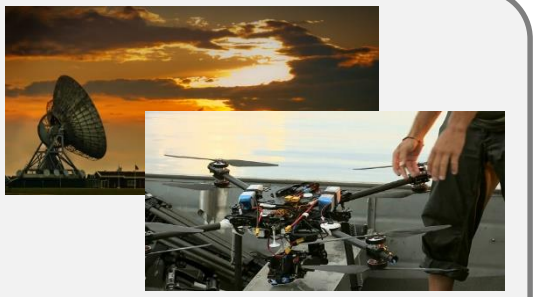

Multi-standard communications receivers, drone recognition

Consumer Electronics and Digital Health

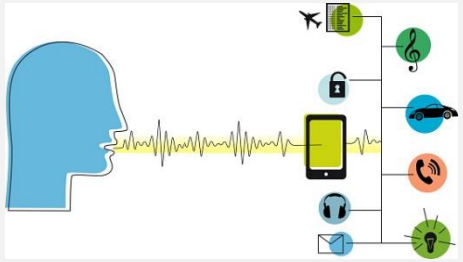

Voice assistants

Digital health

Automotive

Voice control enabled Infotainment

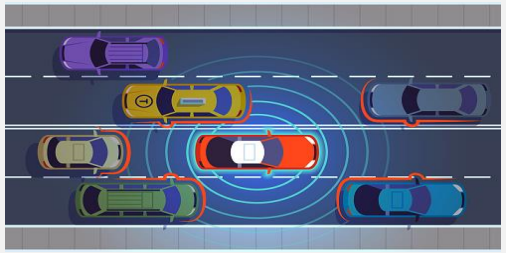

Sensor processing, automated driving

Industrial Automation

Condition monitoring

Predictive maintenance

Deep Learning Workflow

Create and Access Datasets

Data sources

Simulation and augmentation

Data Labeling

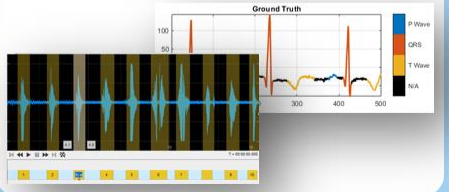

Preprocess and Transform Data

Pre-Processing

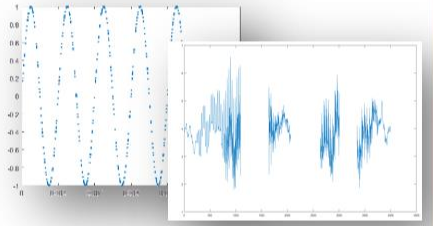

Transformation

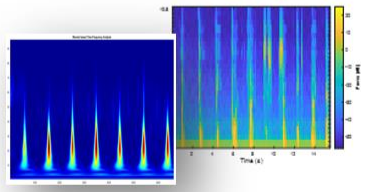

Feature extraction

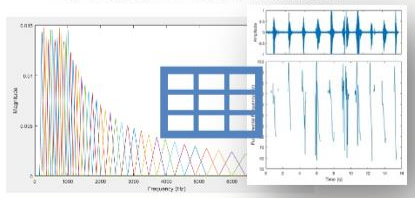

Develop Predictive Models

Import Reference Models/ Design from scratch

Hardware-Accelerated Training

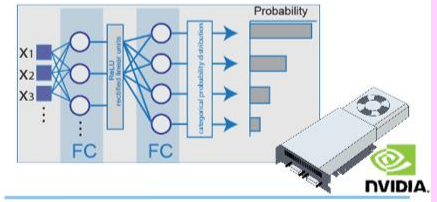

Analyze and tune hyperparameters

Accelerate and Deploy

Desktop Apps

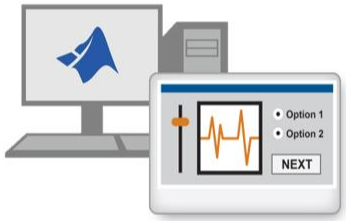

Enterprise Scale Systems

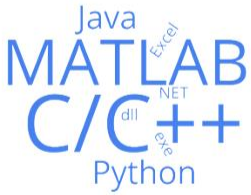

Embedded Devices and Hardware

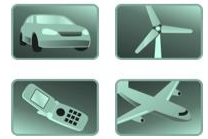

Deep Learning Workflow Challenges – Signals, Time-Series, Text

Deep learning models only as good as training data

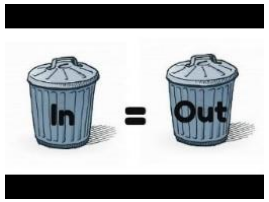

Time consuming data labeling

Limited data availability

Complex deep learning architectures

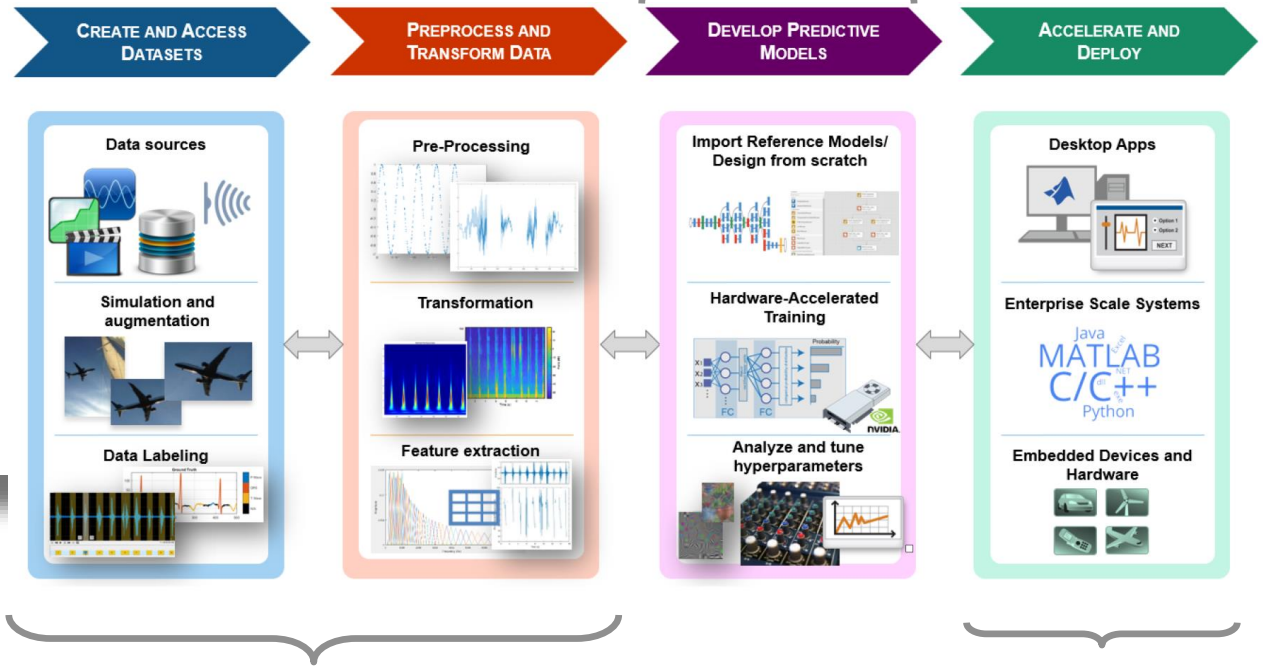

Domain-specific data processing desirable

Limited reference research

Deployment and scaling to various platforms

Agenda

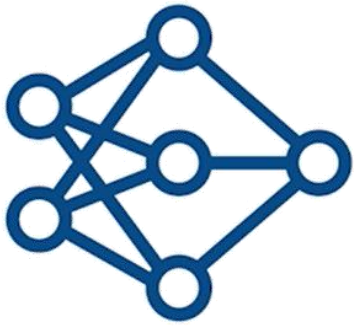

- Deep Learning – Basic ideas

- Deep Learning Model Development for Signals, Time-Series, and Text
  - Data
  - Processing and transformation
  - Model design and optimization
  - Acceleration, prototyping, and deployment

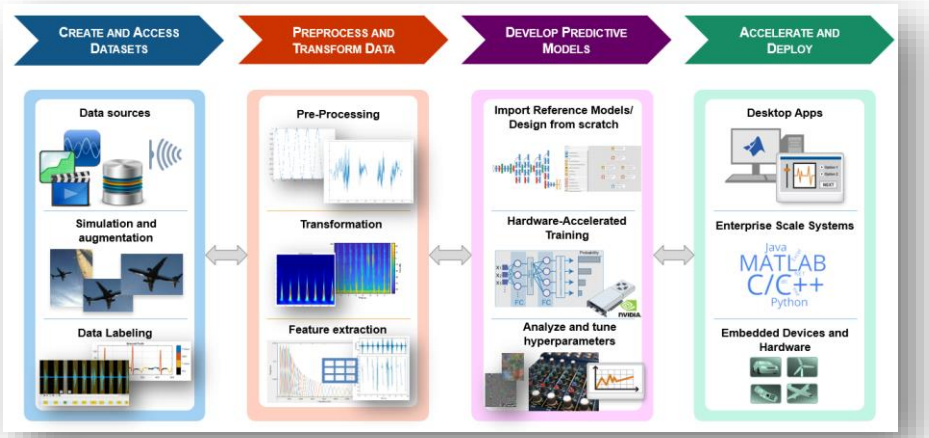

- Conclusions

Application Examples Using MATLAB – Speech and Audio

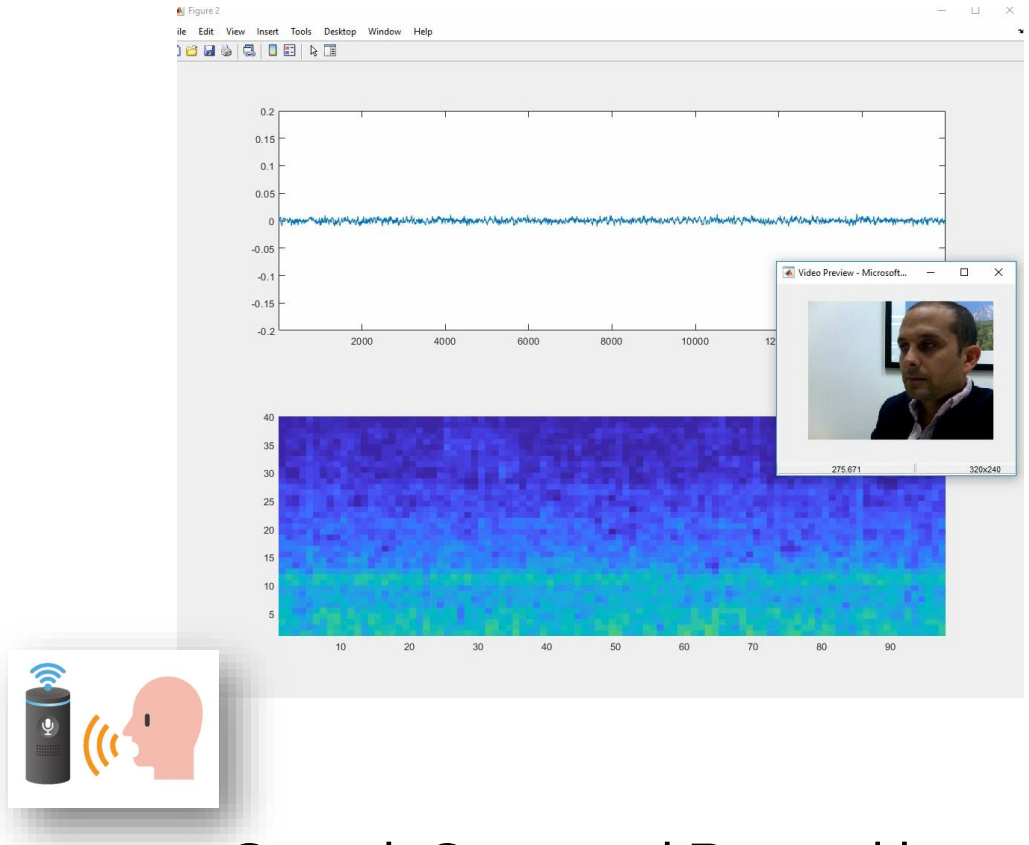

Speech Command Recognition

https://www.mathworks.com/help/deeplearning/examples/deep-learning-speech-recognition.html

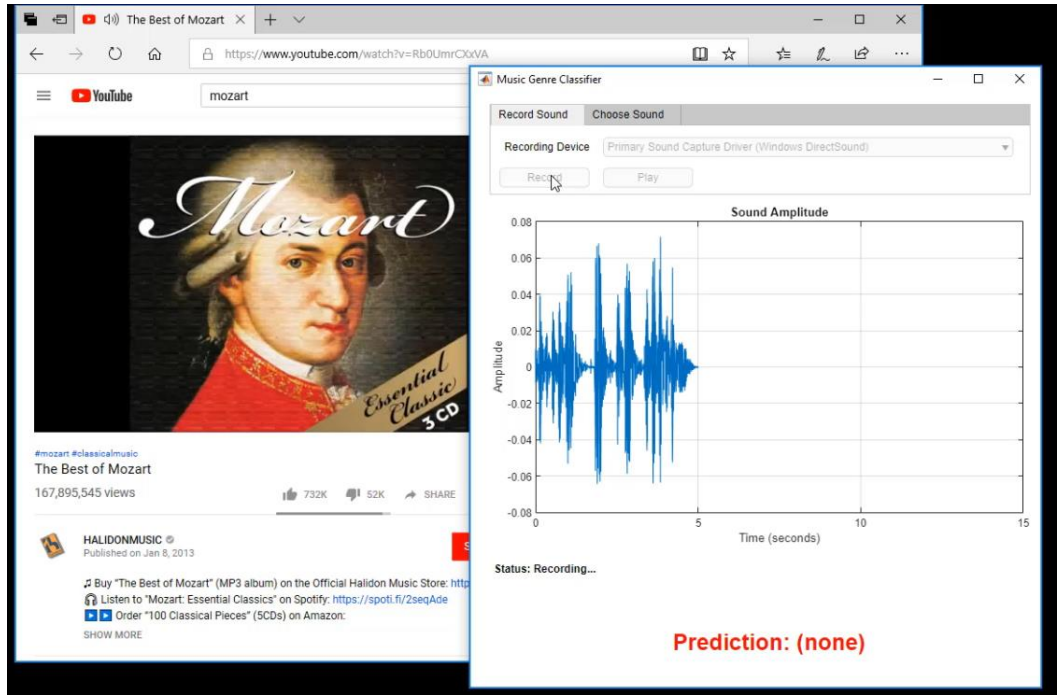

Music Genre Classification

https://www.mathworks.com/help/audio/examples/music-genre-classification-using-wavelet-time-scattering.html

What is Deep Learning?

Deep learning is a type of machine learning in which a model learns from examples.

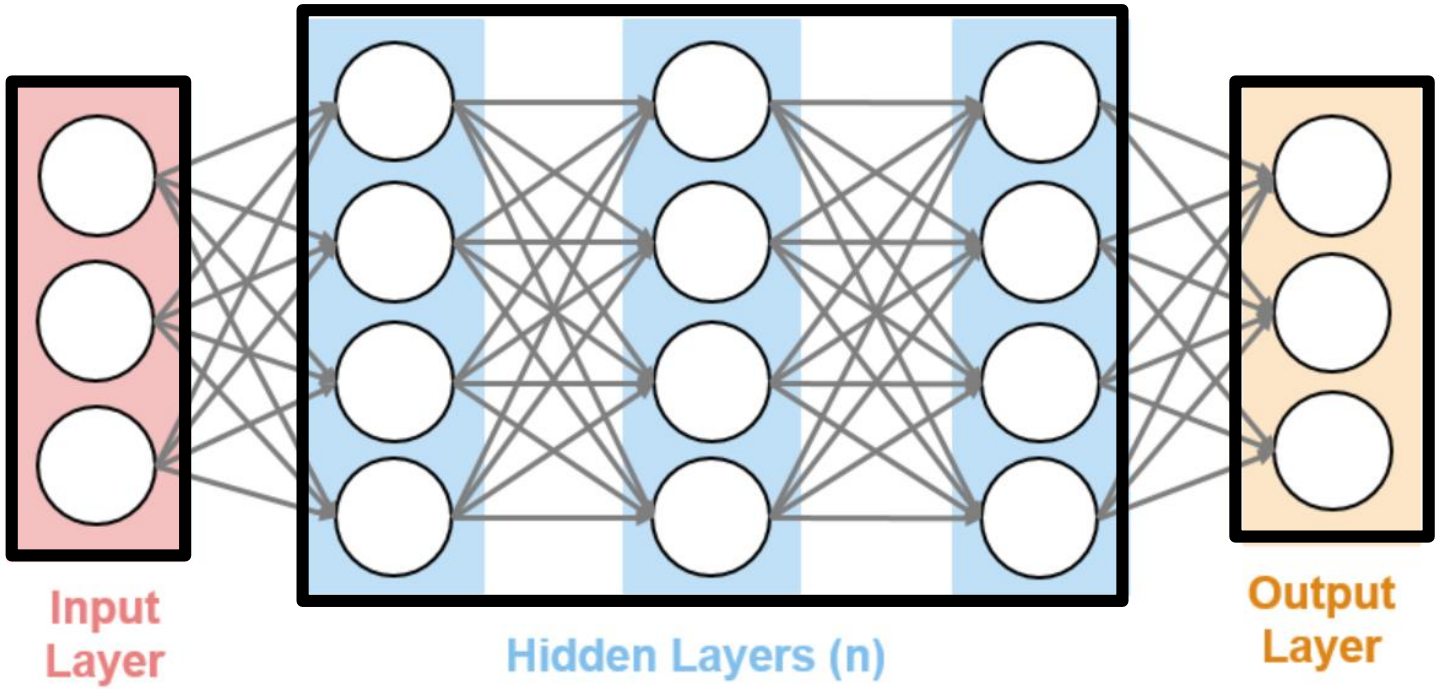

Common Network Architectures – Signals, Time-Series, Text

Convolutional Neural Networks (**CNN**)

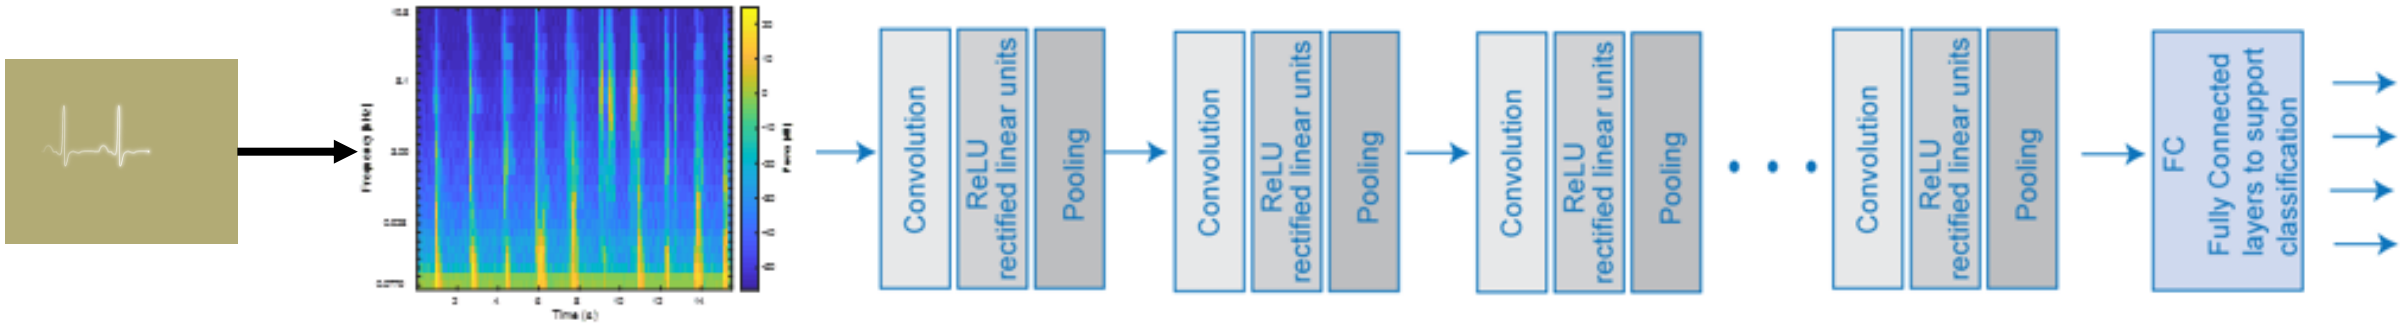

Time-Frequency Transformation

Long Short-Term Memory (**LSTM**) Networks

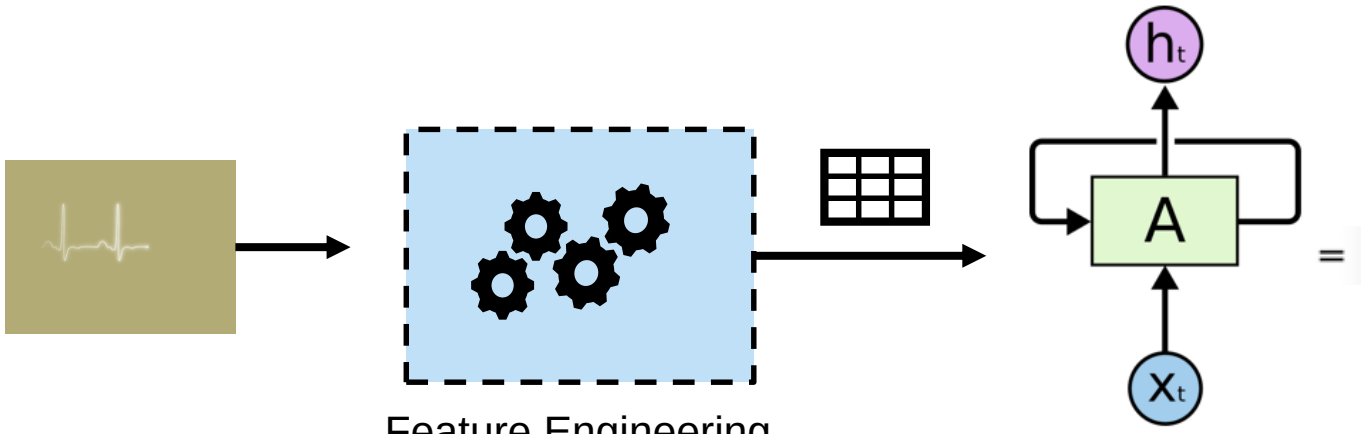

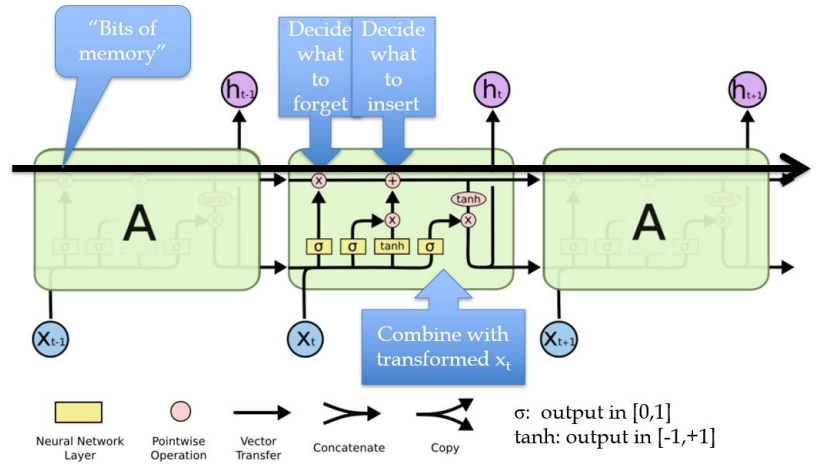

Feature Engineering

Agenda

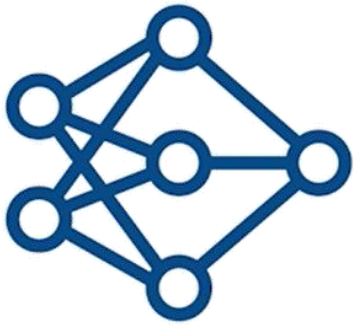

- Deep Learning – Basic ideas
- **Deep Learning Model Development for Signals, Time-Series, and Text**
  - Data
  - Processing and transformation
  - Model design and optimization
  - Acceleration, prototyping, and deployment

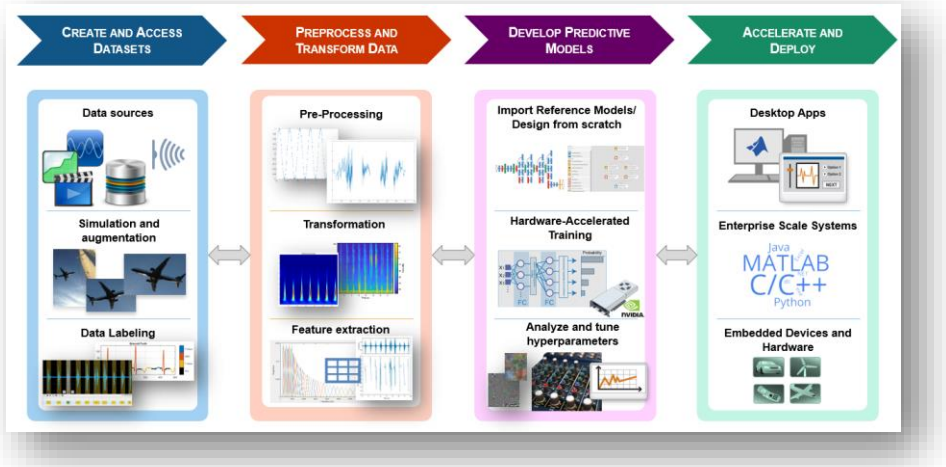

- Conclusions

Agenda

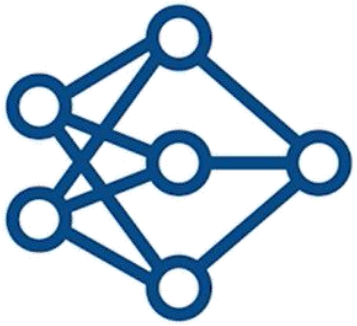

- Deep Learning – Basic ideas

- Deep Learning Model Development for Signals, Time-Series, and Text
  - **Data**
  - Processing and transformation
  - Model design and optimization
  - Acceleration, prototyping, and deployment

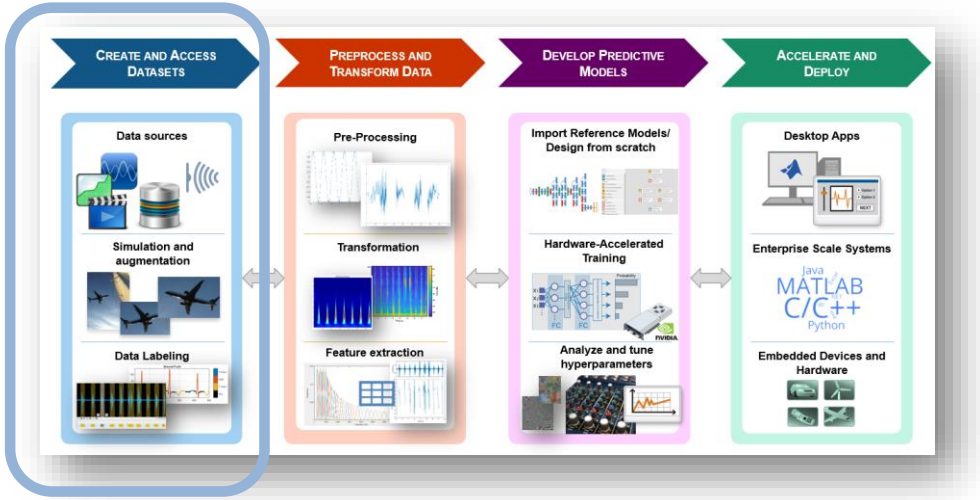

- Conclusions

Current Investments – Models vs. Data

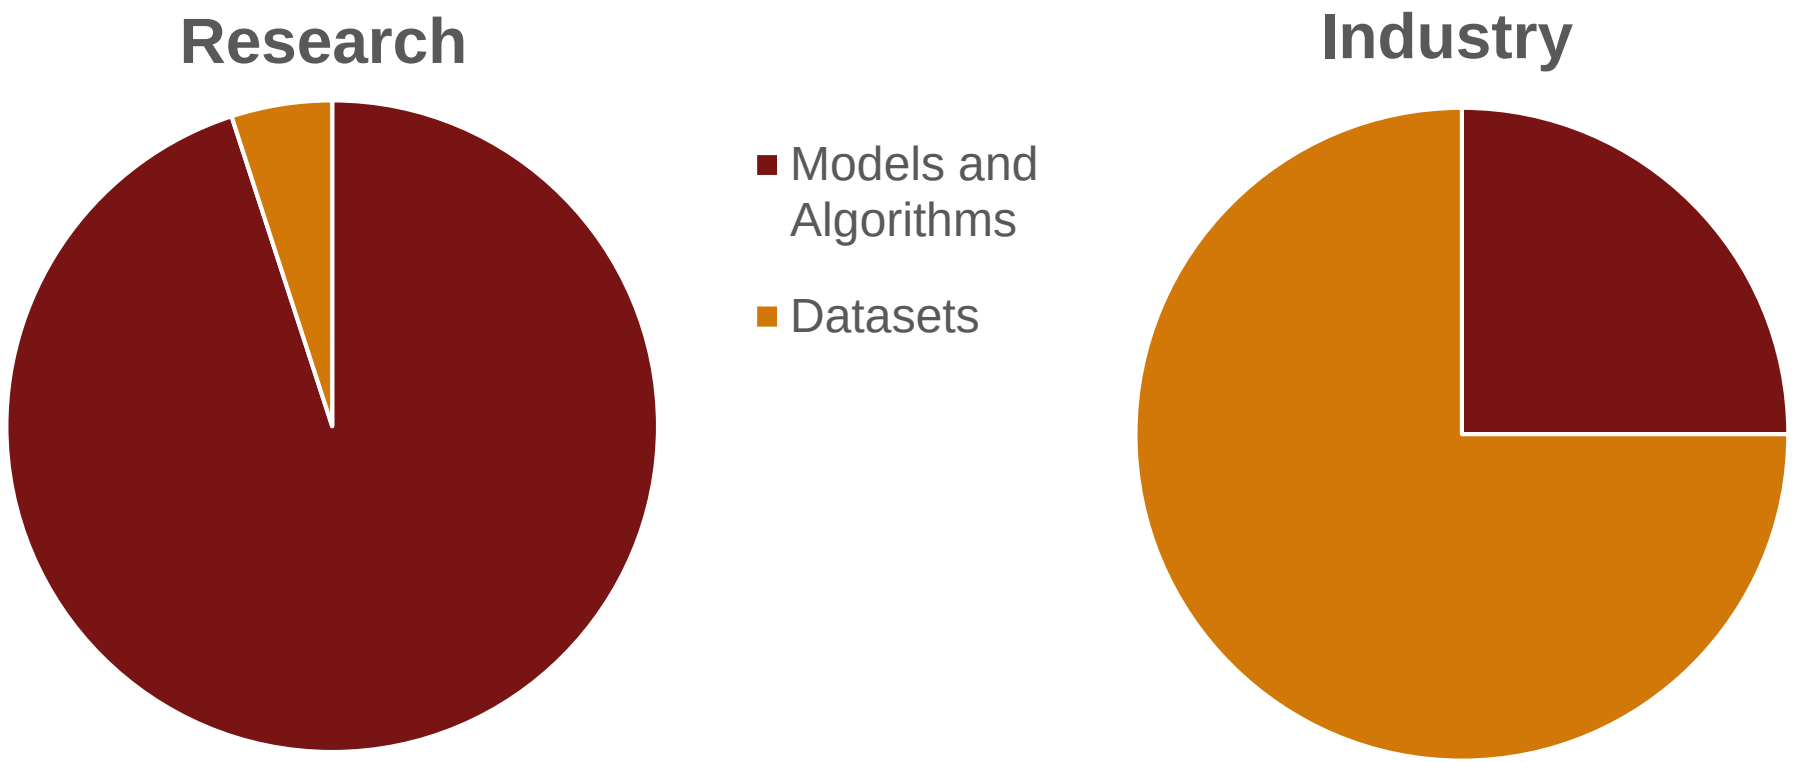

From "Troubleshooting deep neural networks" (Josh Tobin et al., Jan 2019)

Efficient Handling Large Datase

How to navigate, index, read (al

`audioDatastore` `imageDatastore`

`fileDatastore` Custom Datastores

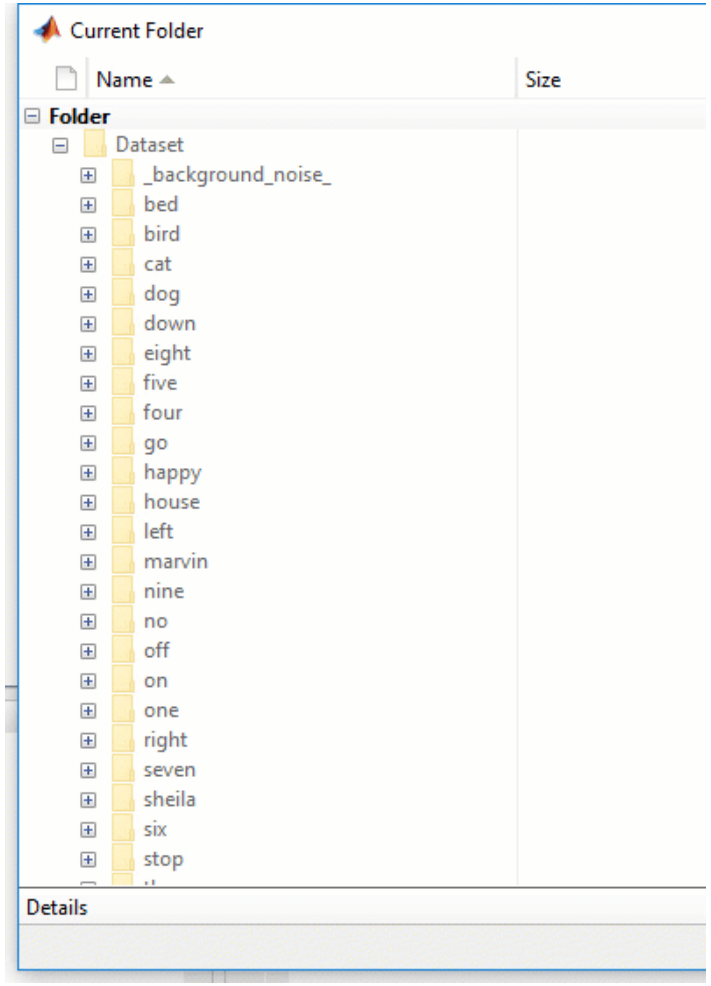

How to...

- Build a list of all data and labels?
- Review basic statistics about available data?
- Select data subsets without nested `for` loops, `dir`, `ls`, `what`, ... aplenty?
- Jointly read data and labels?
- Automatically distribute computations?

Minimize Time Consuming Labelling of Signals

```
patientID = 1;
signalVals = getSignal(QTData,patientID);
labelVals = getLabelValues(QTData,patientID,'WaveformLabels_Chan1');

displayWaveformLabels(signalVals(1,1:1000),labelVals.Value(1:1000))
```

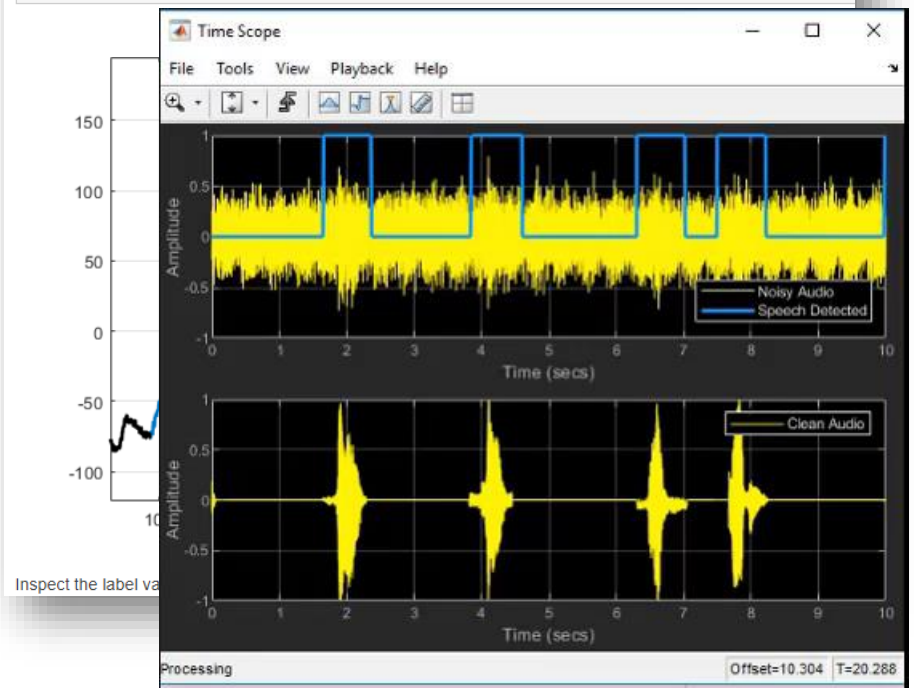

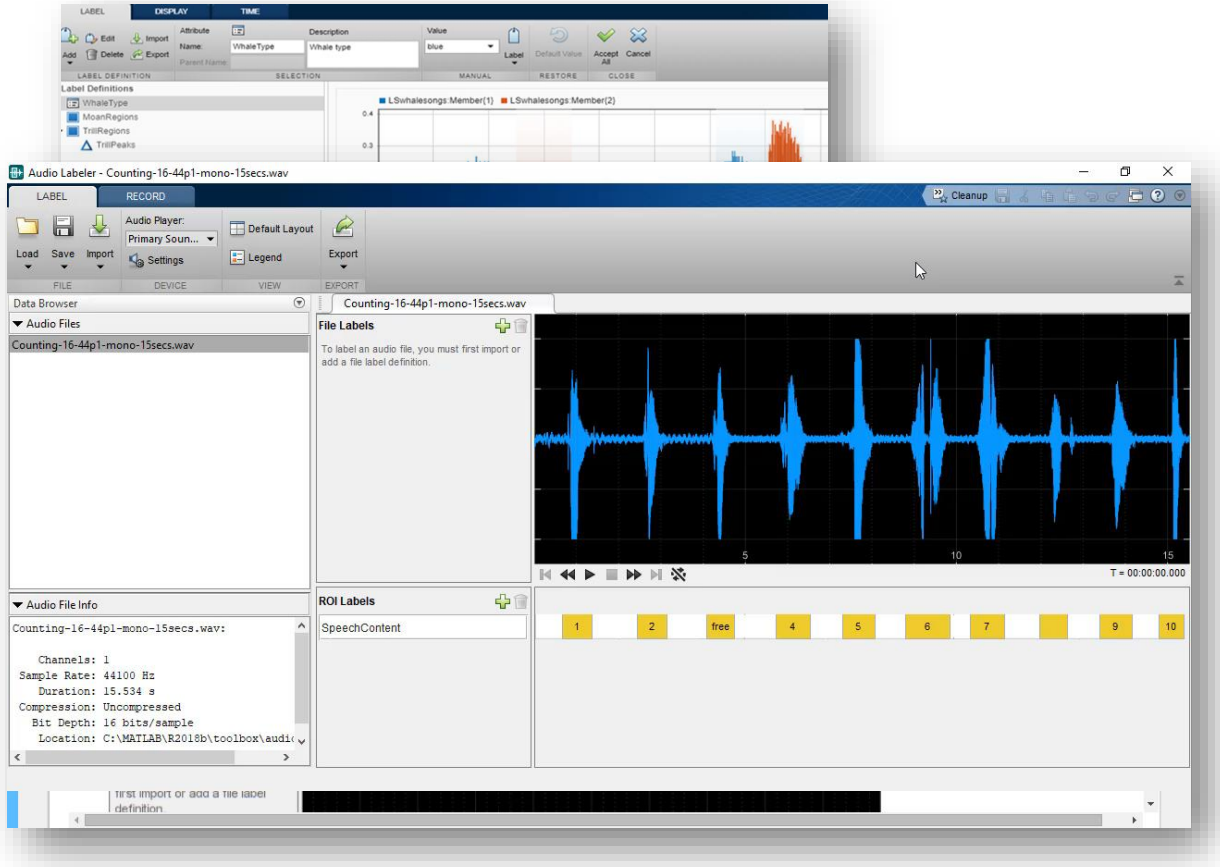

- Programmatically…
- … or via Apps

What if Available (Good) Data Isn't Enough?

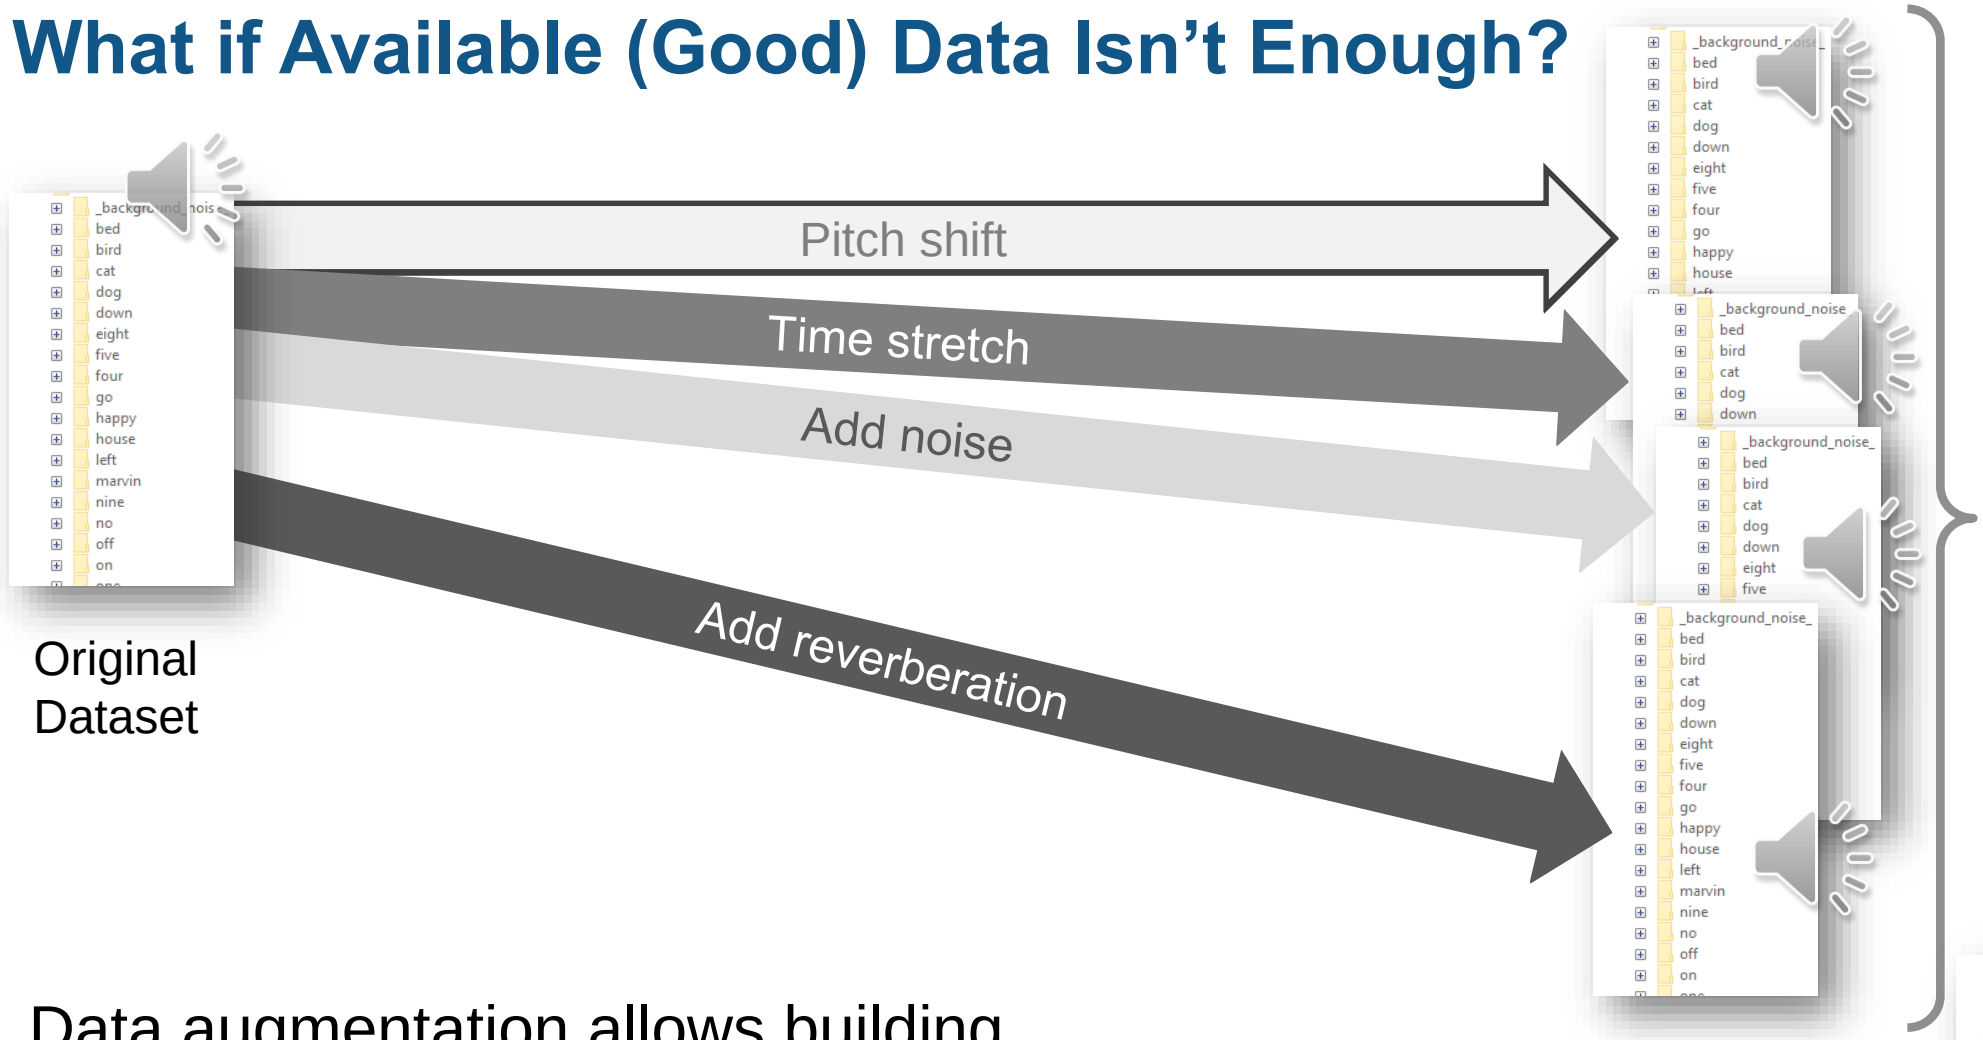

Data augmentation allows building more complex and more robust models

What if Recording and Labelling Real-World Data is Impractical or Unreasonable? - Simulation is Key!

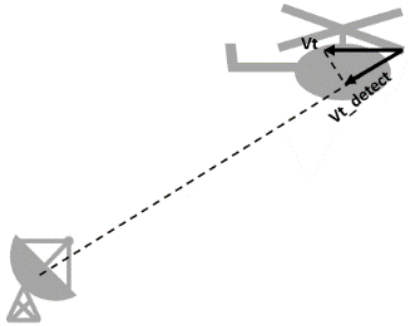

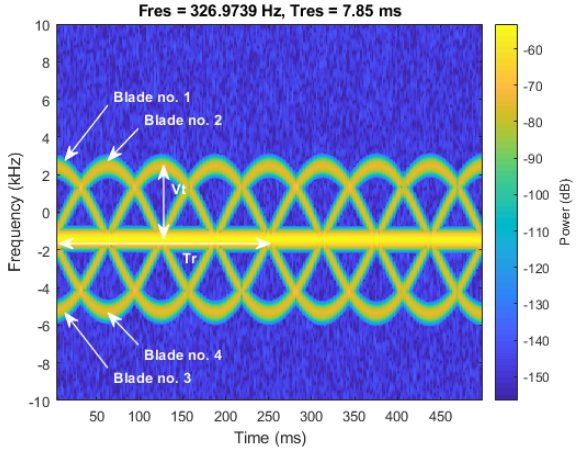

Radar Target Simulation

Micro-Doppler Analysis

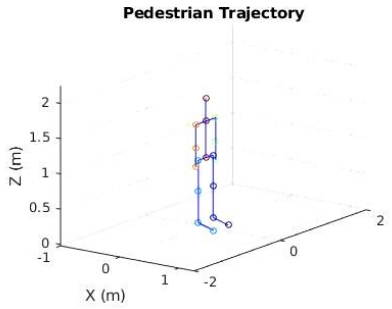

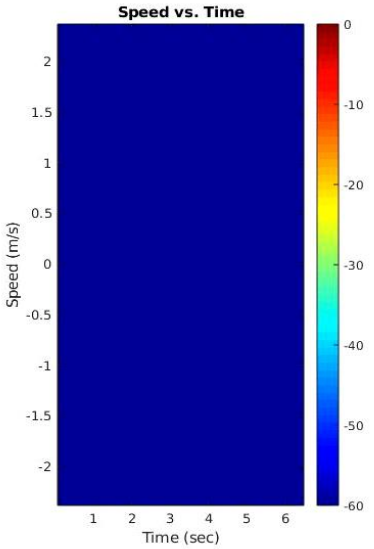

Agenda

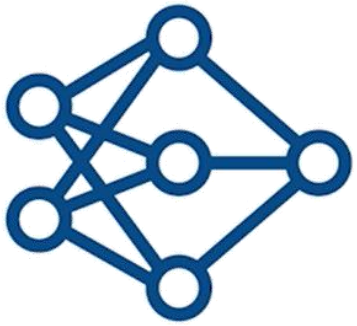

- Deep Learning – Basic ideas

- Deep Learning Model Development for Signals, Time-Series, and Text
  - Data
  - **Processing and transformation**
  - Model design and optimization
  - Acceleration, prototyping, and deployment

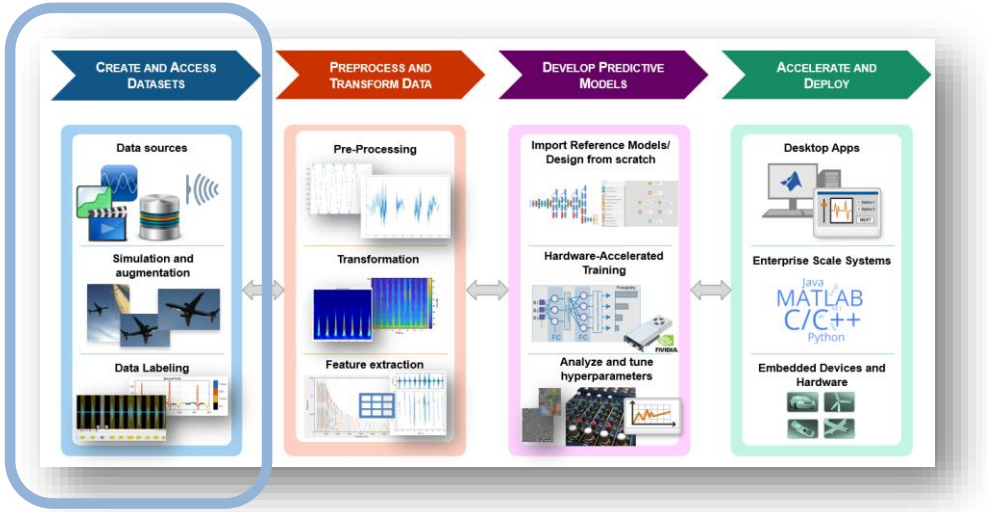

- Conclusions

Common Network Architectures – Signals, Time-Series, Text

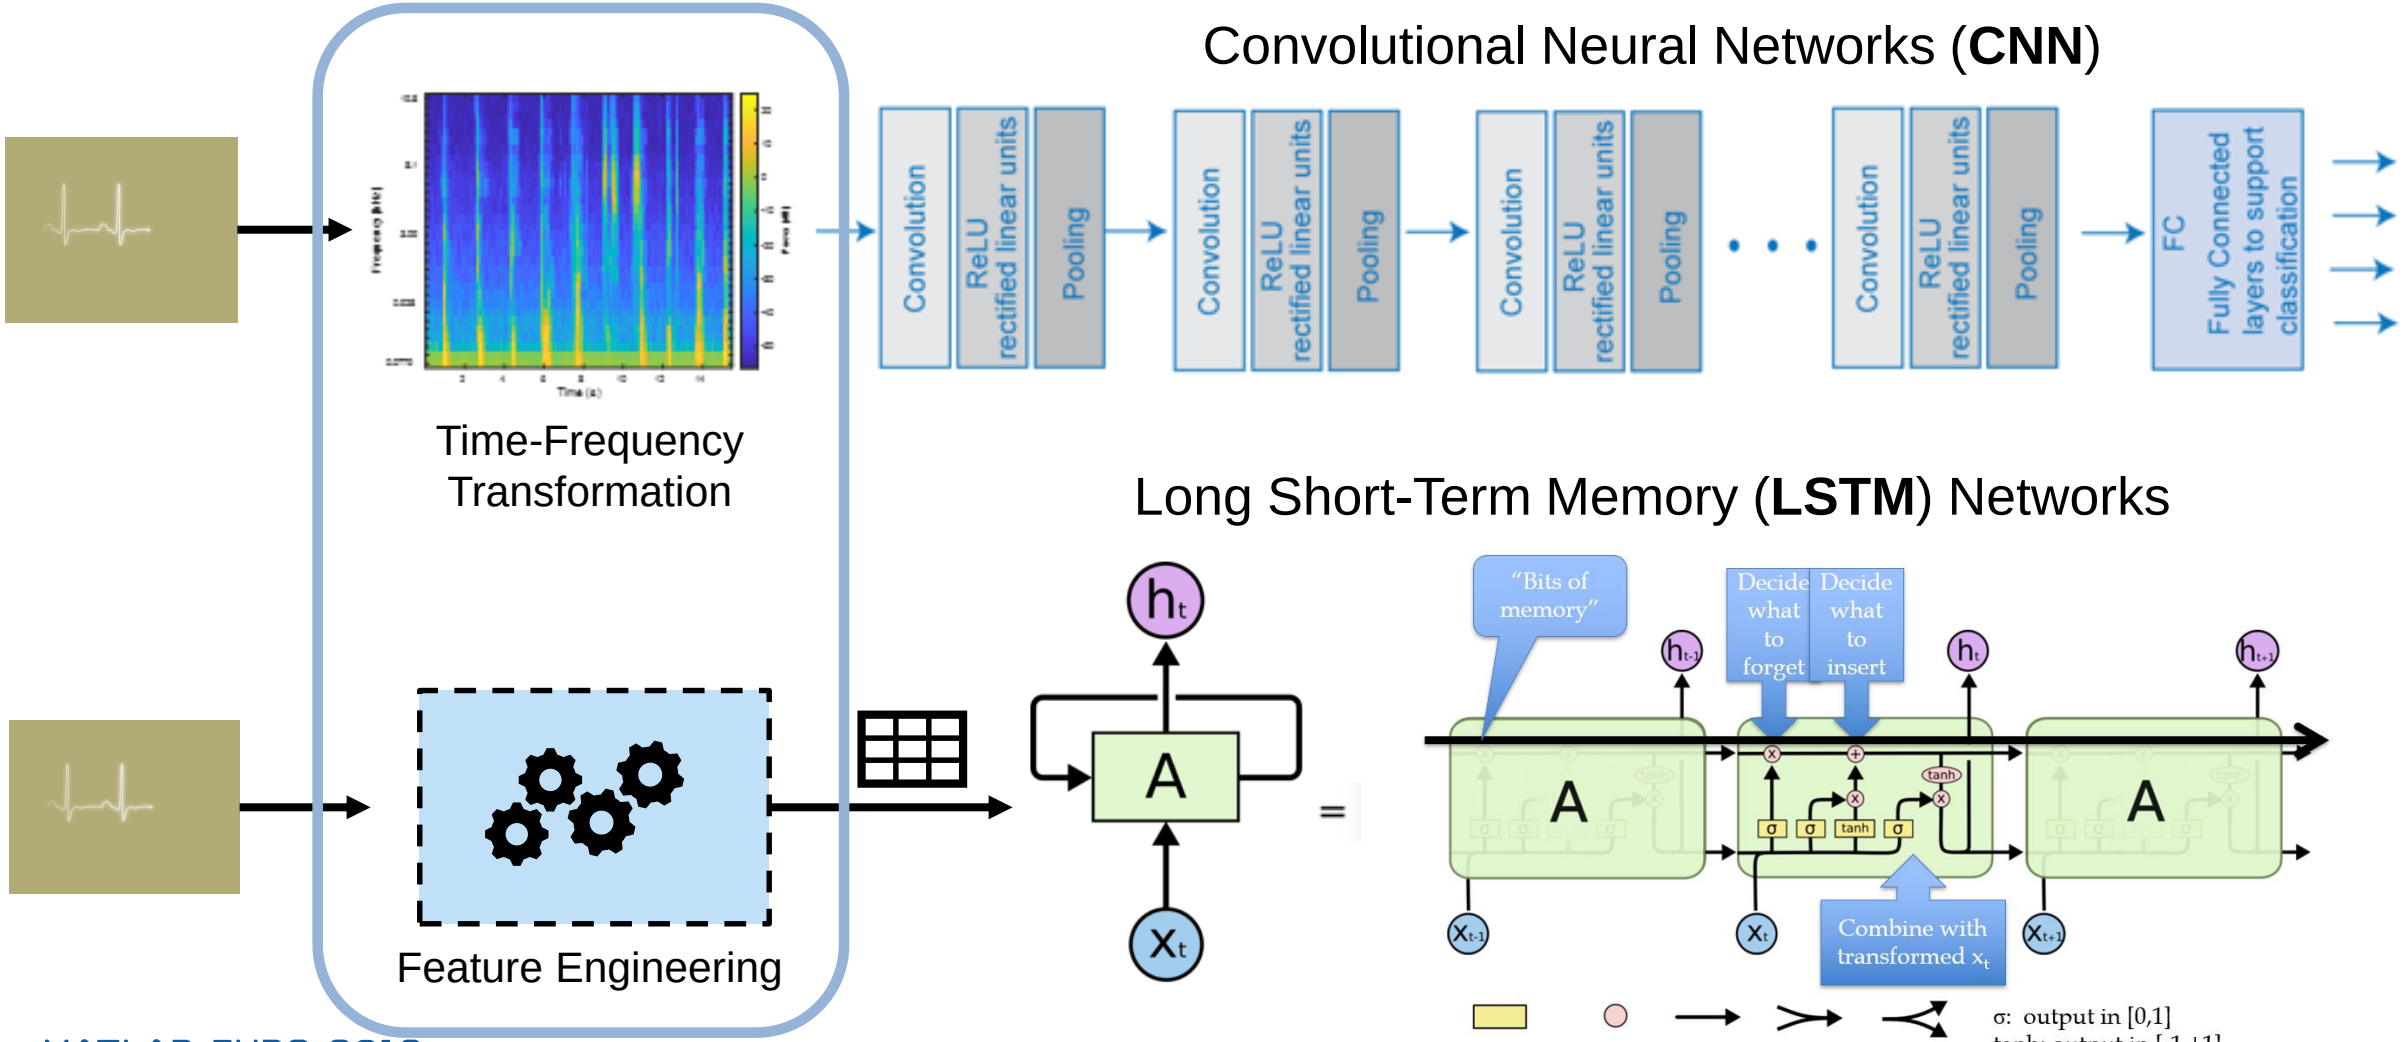

Time-Frequency Transformations

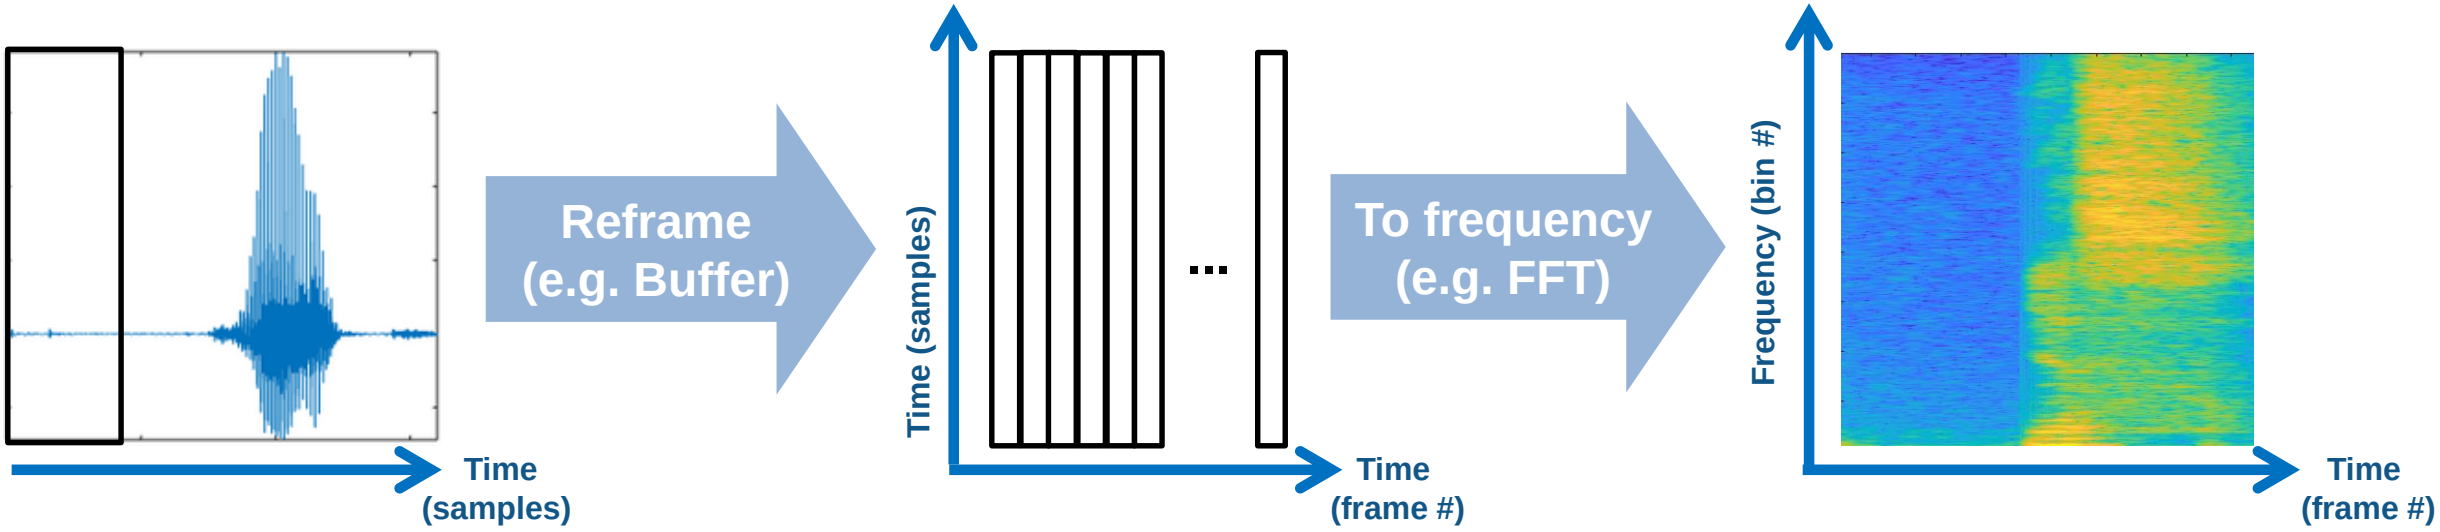

Easiest to understand and implement

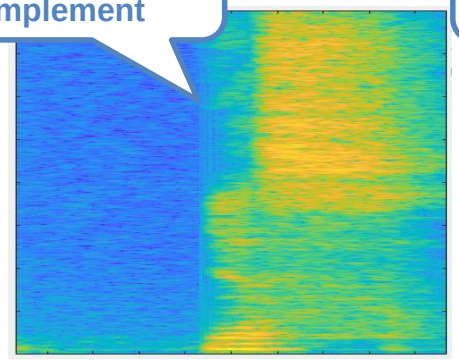

Basic spectrogram

More compact for speech & audio applications

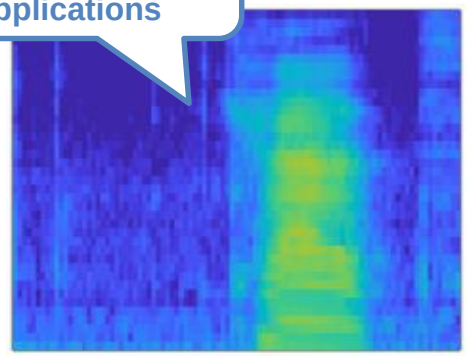

Perceptually-spaced (e.g. Mel, Bark) Spectrogram

Best resolution for non-periodic signals

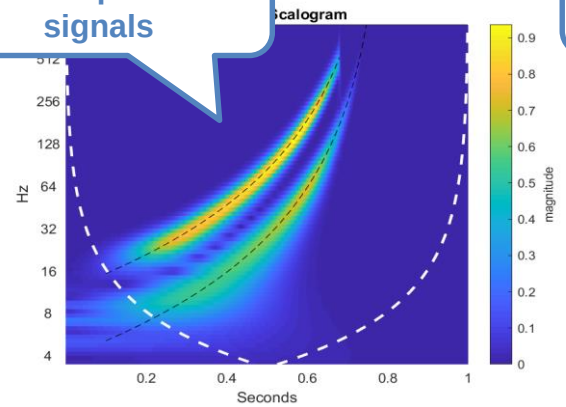

Wavelet scalogram

Better resolution at low frequencies

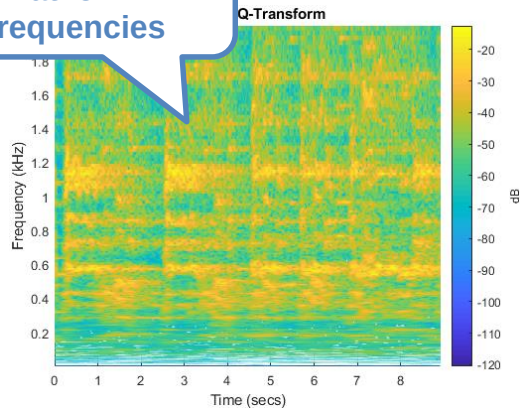

Constant Q transform

Extracting Features from Signals: Application-Agnostic Examples

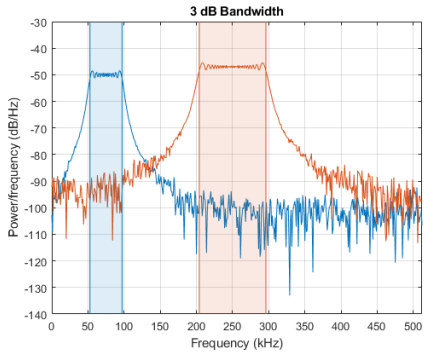

BW measurements

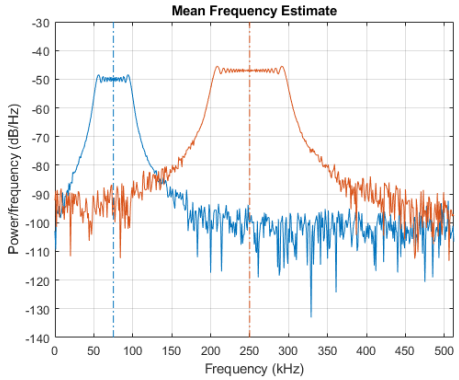

Spectral statistics

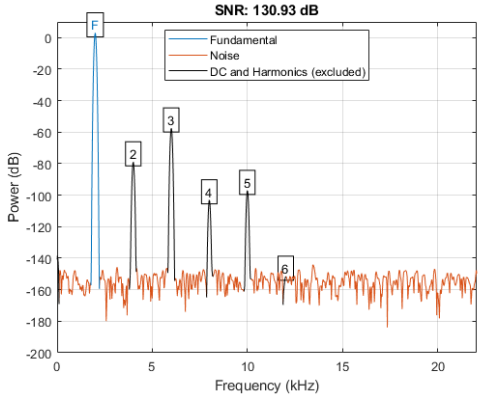

Harmonic analysis

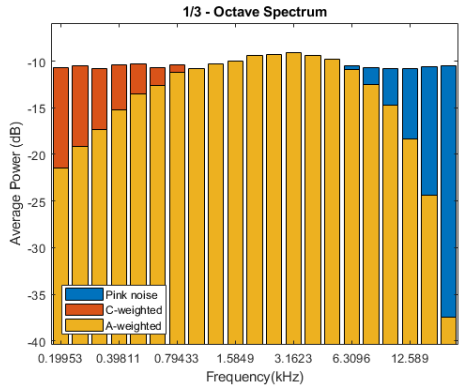

Octave spectrum

Frequency domain

Time domain

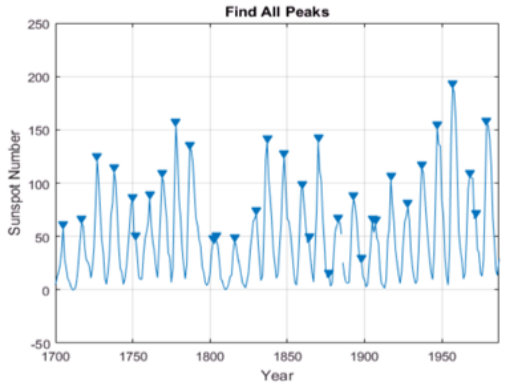

Find peaks

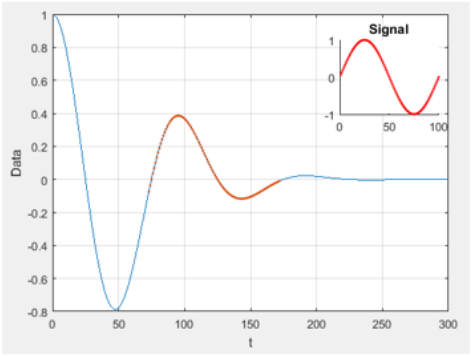

Find signal patterns

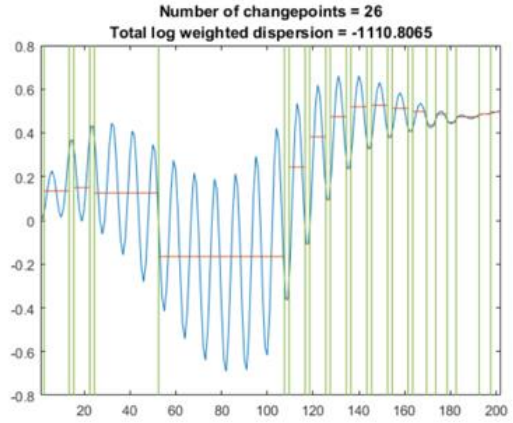

Detect change points

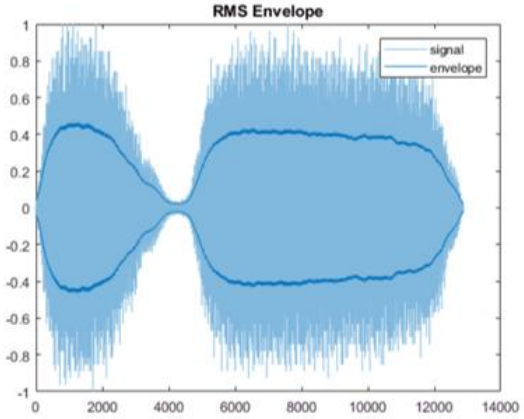

Find signal envelope

Domain-Specific Features and Transformations – Examples

Speech and Audio

- MFCC
- GTCC
- MDCT
- Pitch, harmonicity
- Spectral shape descriptors
- ...

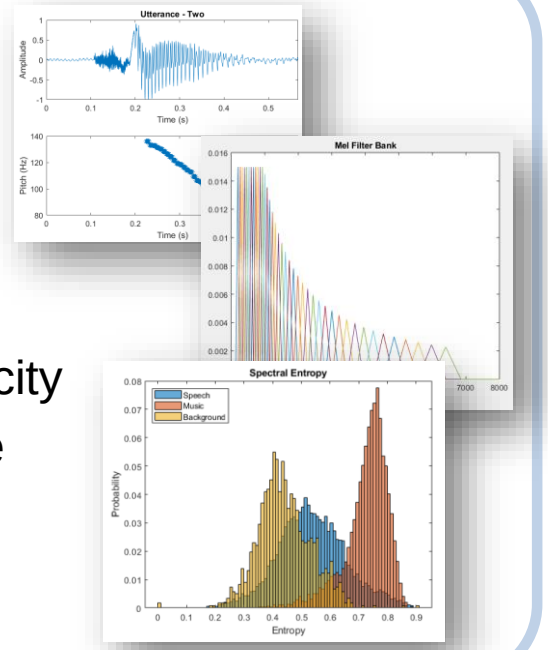

Navigation and Sensor Fusion

- Orientation
- Height
- Position

from

- Acceleration, angular velocity
- Magnetic field
- GPS reading

- Multi-object tracking
- …

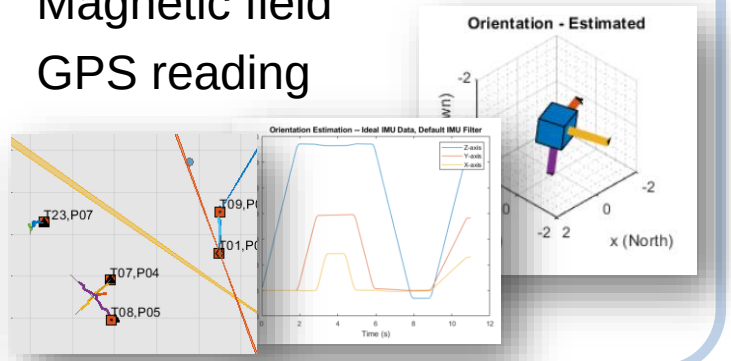

Text Analytics

- Train Word Embeddings
- Word2Vec
- Topic Modeling
- ...

Radar

- Micro-Doppler analysis
- Range-Doppler processing
- Synthetic aperture imaging
- Spectral analysis
- Waveform ambiguity
- ...

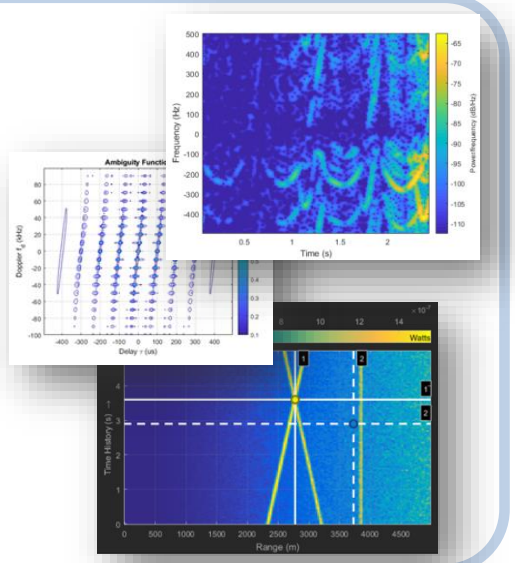

Agenda

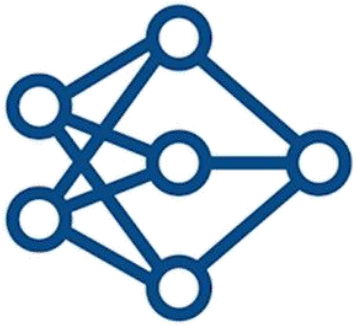

- Deep Learning – Basic ideas

- Deep Learning Model Development for Signals, Time-Series, and Text
  - Data
  - Processing and transformation
  - **Model design and optimization**
  - Acceleration, prototyping, and deployment

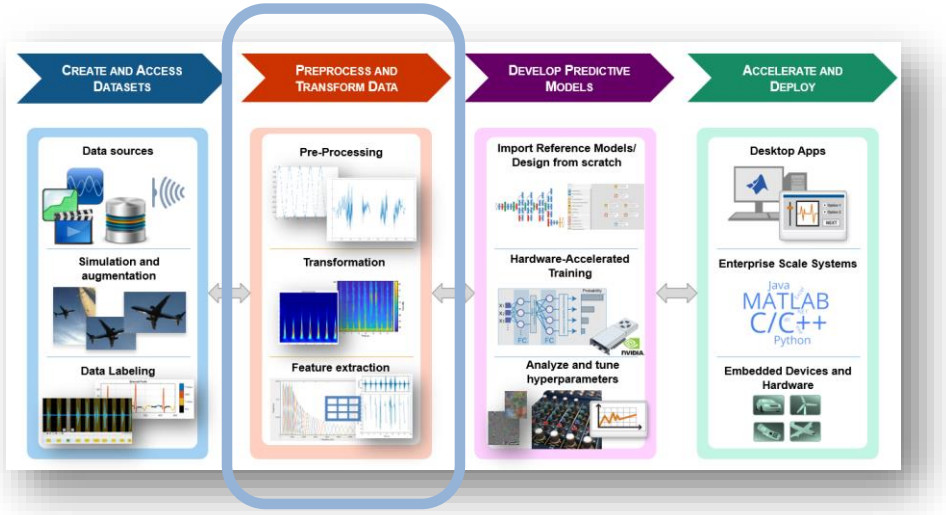

- Conclusions

Deep Learning Model Design and Optimization

Deep Network Designer

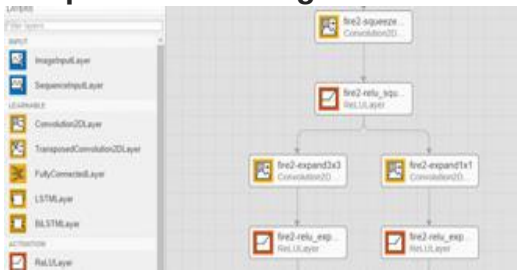

Model Exchange

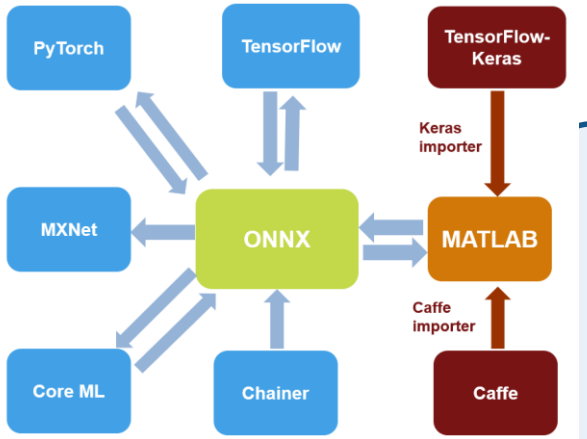

Pre-trained Networks

Places365GoogLeNet
MobileNet-v2
Xception
DenseNet-201
SqueezeNet
ResNet-18
ResNet-50
ResNet-101
Inception-v3
Inception-ResNet-v2
VGG-16
VGG-19
GoogLeNet
AlexNet

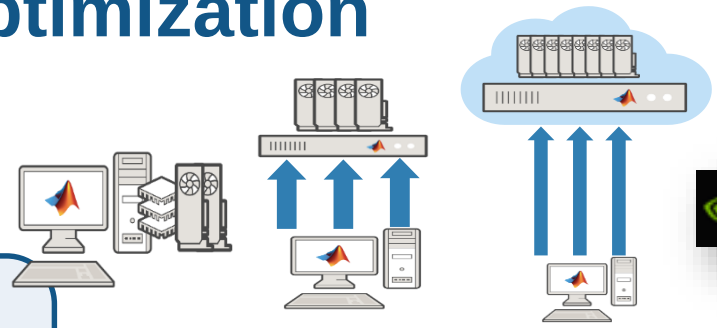

Accelerate Training

NVIDIA GPU CLOUD

MATLAB as a container on NGC

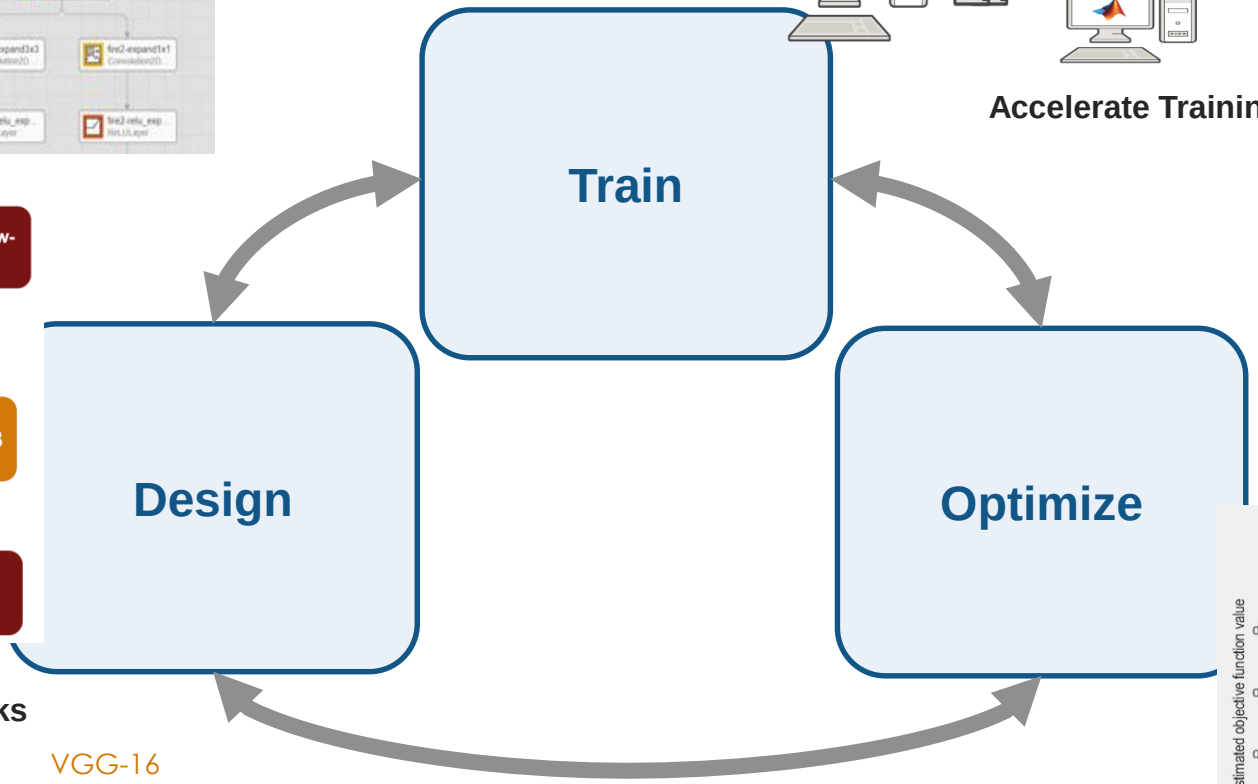

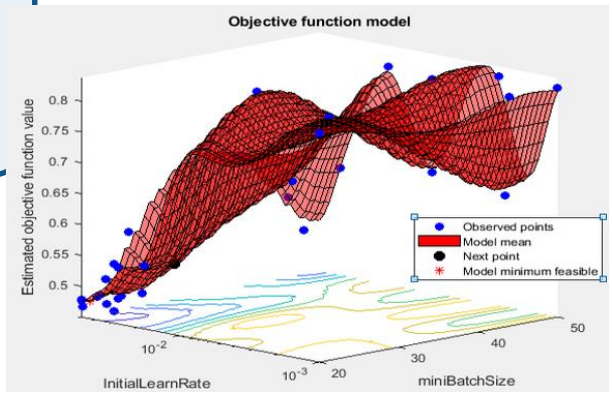

Bayesian Hyperparameter Optimization

CNN Network for Audio Classification

Speech Command Recognition Using Deep Learning

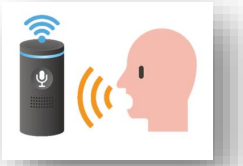

audioDatastore

sheila
six
stop
three
tree
two
up
wow
yes

Google speech command dataset

auditorySpectrogram

Time-Frequency Transform

1. Train

Stop
Up
Yes
⋮
<none>

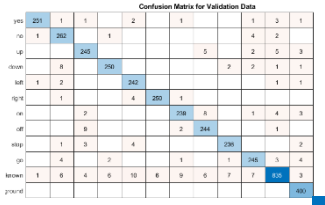

95.41%

audioDeviceReader

auditorySpectrogram

2. Predict

"Up"

LSTM Network for Audio Classification

Speaker Gender Recognition Using Deep Learning

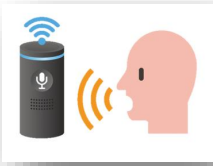

audioDatastore

Training Data

train/sample-000000.mp3
train/sample-000001.mp3
train/sample-000002.mp3
train/sample-000003.mp3
train/sample-000004.mp3
train/sample-000005.mp3
train/sample-000006.mp3
train/sample-000007.mp3

[Mozilla Common Voice Dataset]

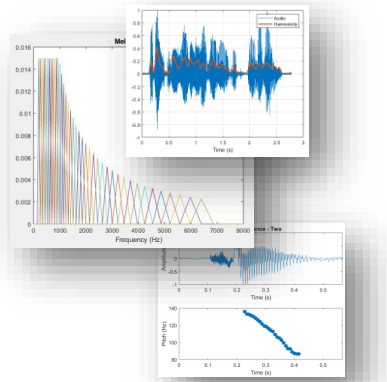

Time-varying features:

- MFCC
- Pitch
- ...

1. Train

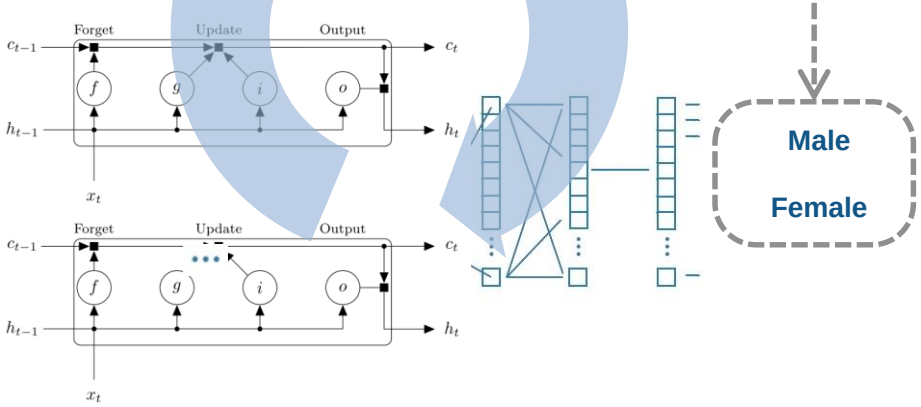

Male

Female

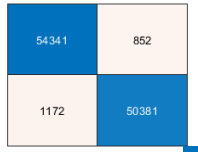

97.2 %

Test Data

train/sample-000000.mp3
train/sample-000001.mp3
train/sample-000002.mp3
train/sample-000003.mp3
train/sample-000004.mp3
train/sample-000005.mp3
train/sample-000006.mp3
train/sample-000007.mp3

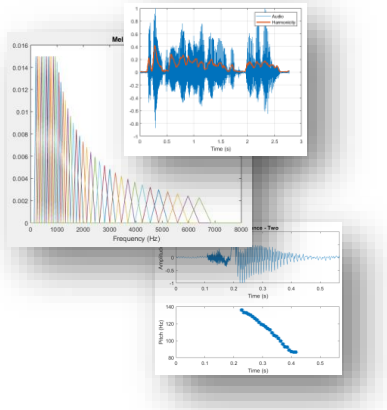

2. Predict

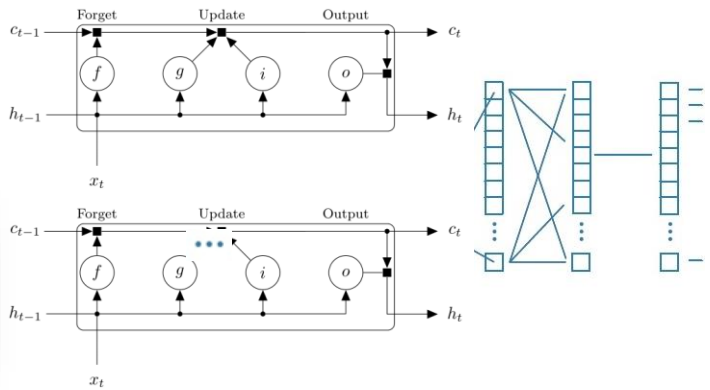

"Male"

Agenda

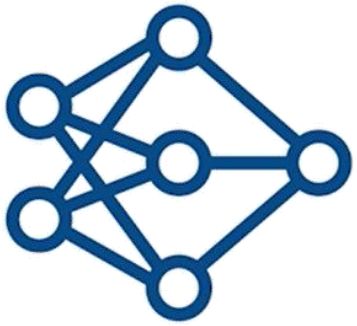

- Deep Learning – Basic ideas

- Deep Learning Model Development for Signals, Time-Series, and Text
  - Data
  - Processing and transformation
  - Model design and optimization
  - **Acceleration, prototyping, and deployment**

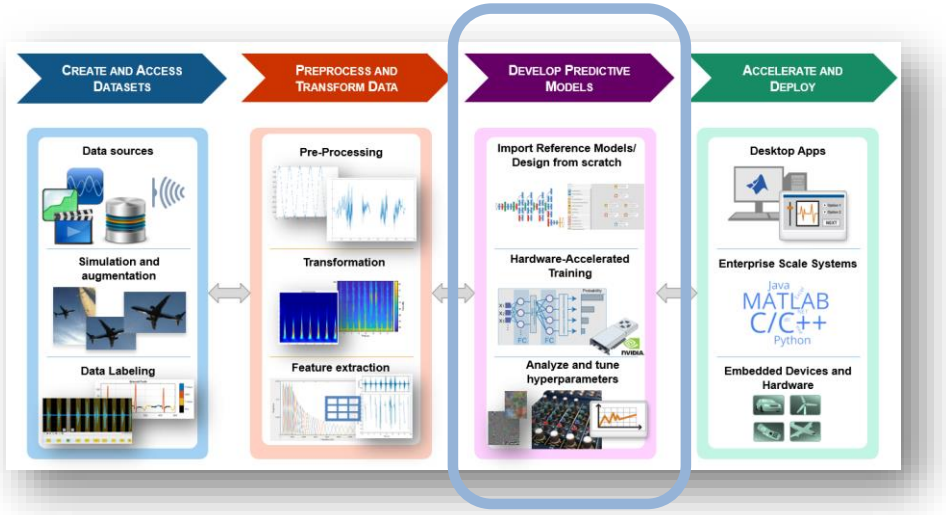

- Conclusions

With GPU Coder, MATLAB is fast

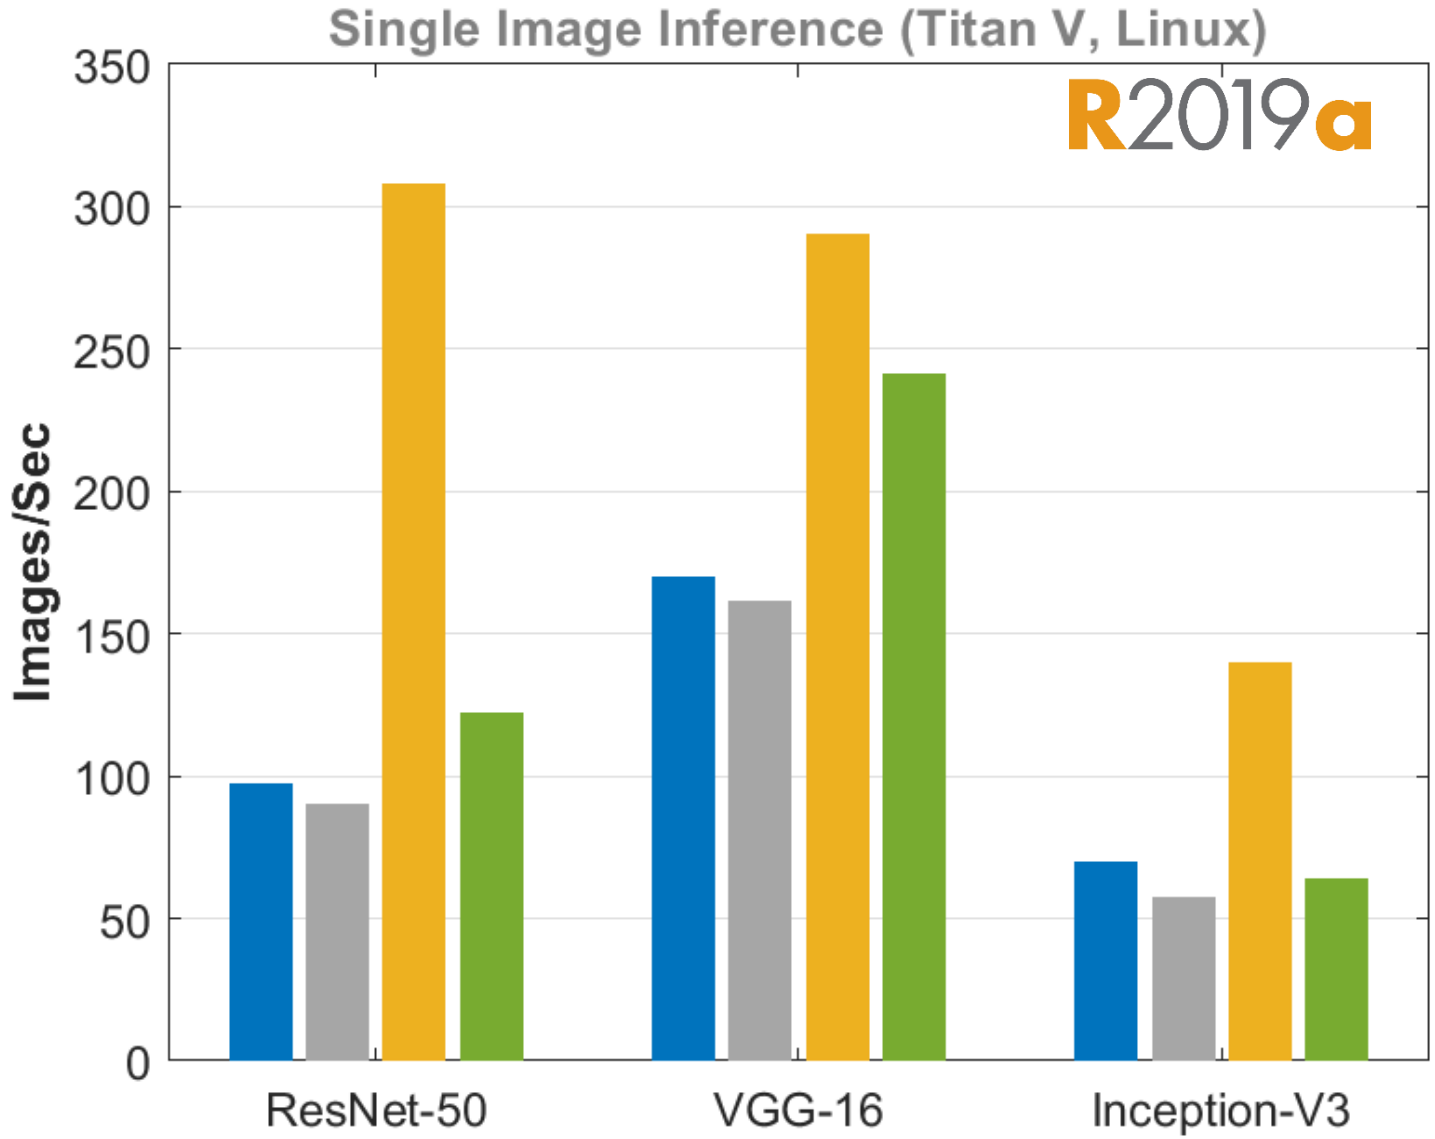

GPU Coder is faster than TensorFlow, MXNet and Pytorch

- TensorFlow
- MXNet
- GPU Coder
- PyTorch

Intel® Xeon® CPU 3.6 GHz - NVIDIA libraries: CUDA10 - cuDNN 7 - Frameworks: TensorFlow 1.13.0, MXNet 1.4.0 PyTorch 1.0.0

Deep Learning from Idea to Product

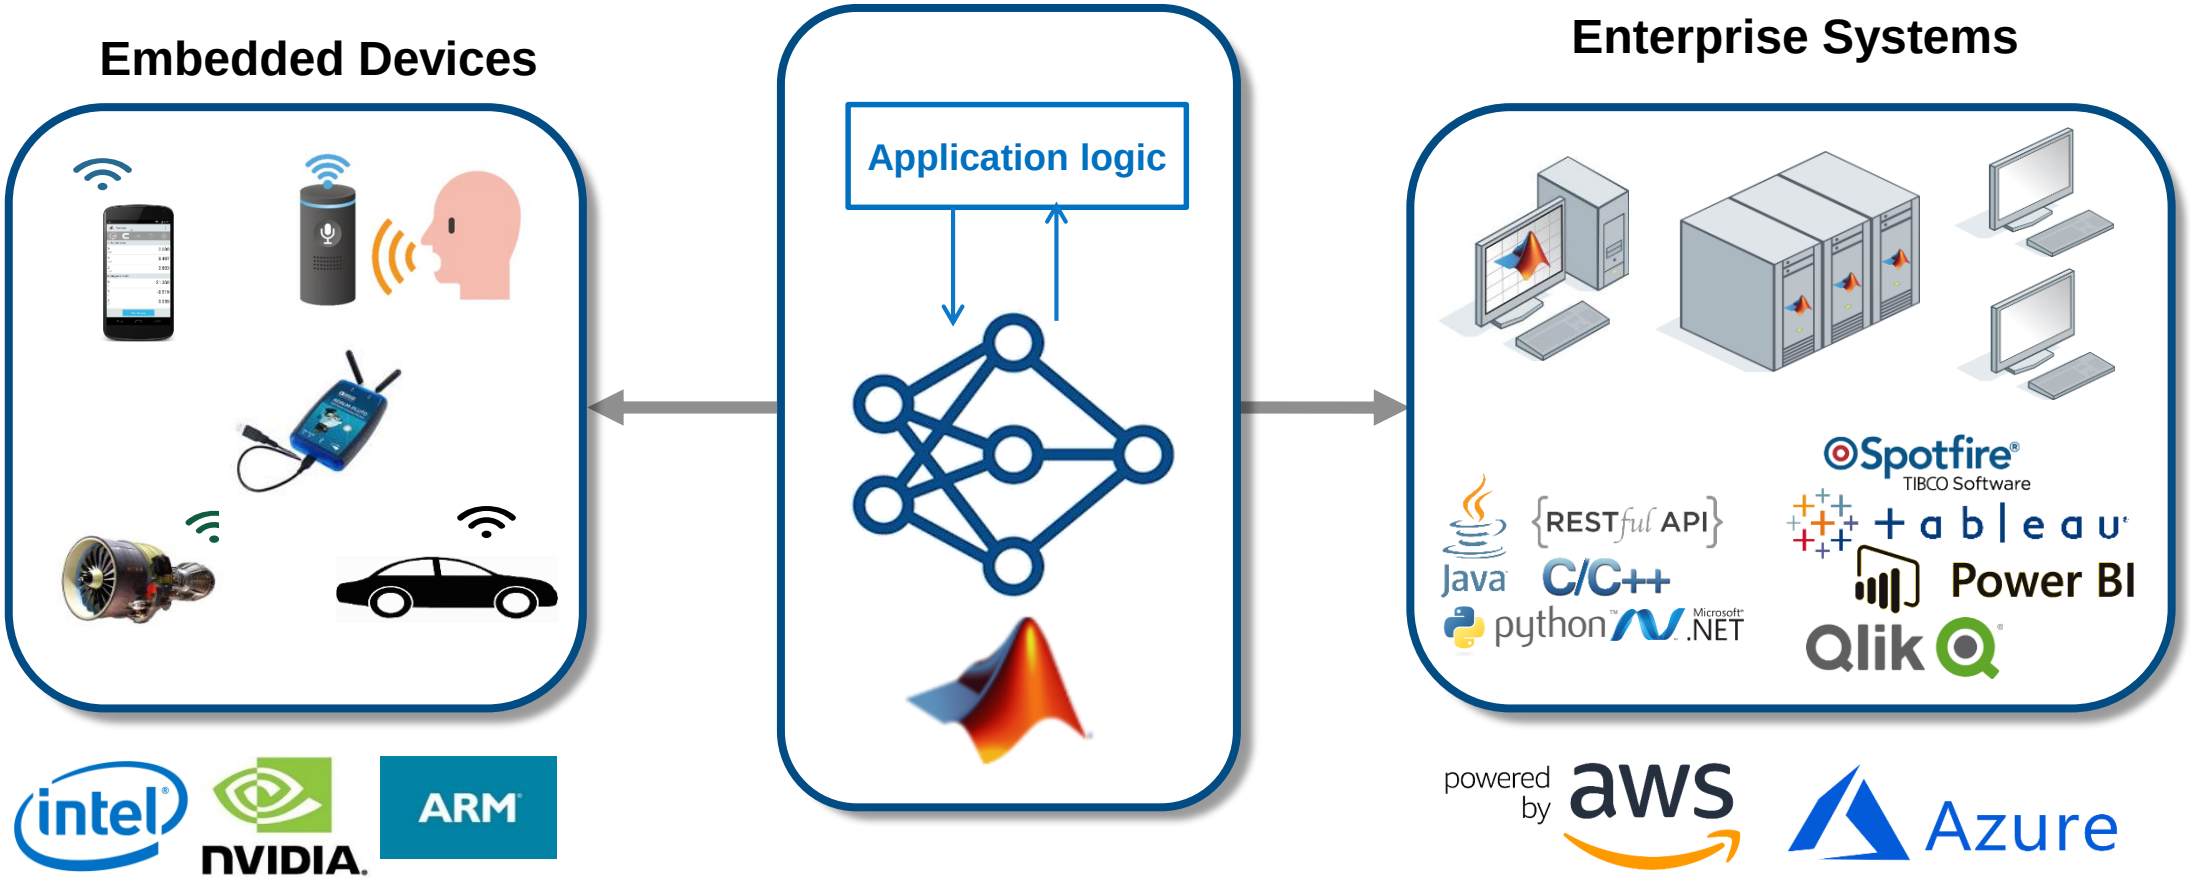

Agenda

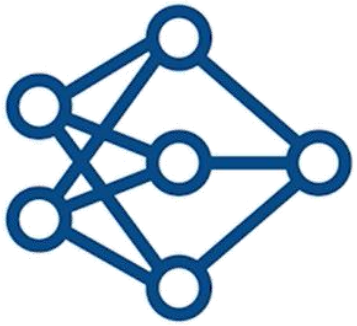

- Deep Learning – Basic ideas

- Deep Learning Model Development for Signals, Time-Series, and Text
  - Data
  - Processing and transformation
  - Model design and optimization
  - Acceleration, prototyping, and deployment

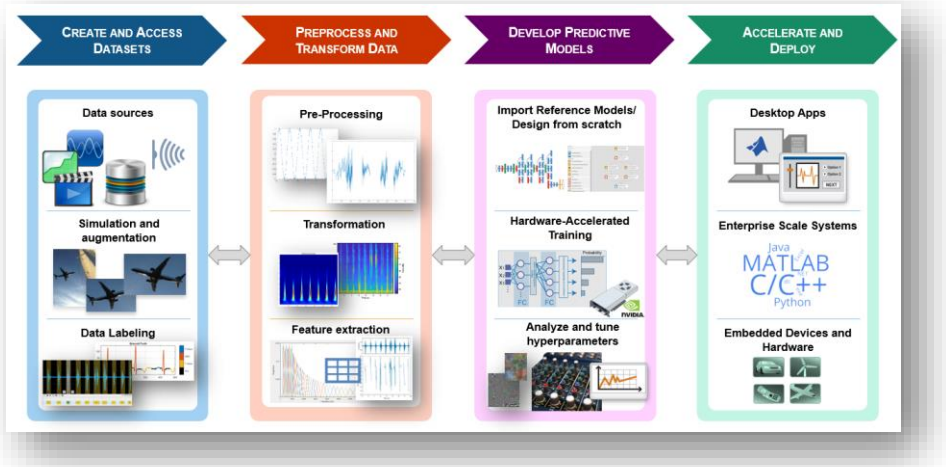

- **Conclusions**

Deep Learning Workflow Challenges – Signals, Time-Series, Text

Deep learning models only as good as training data

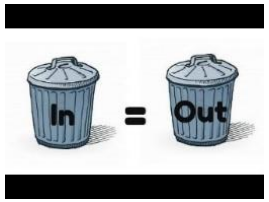

Time consuming data labeling

Limited data availability

Complex deep learning architectures

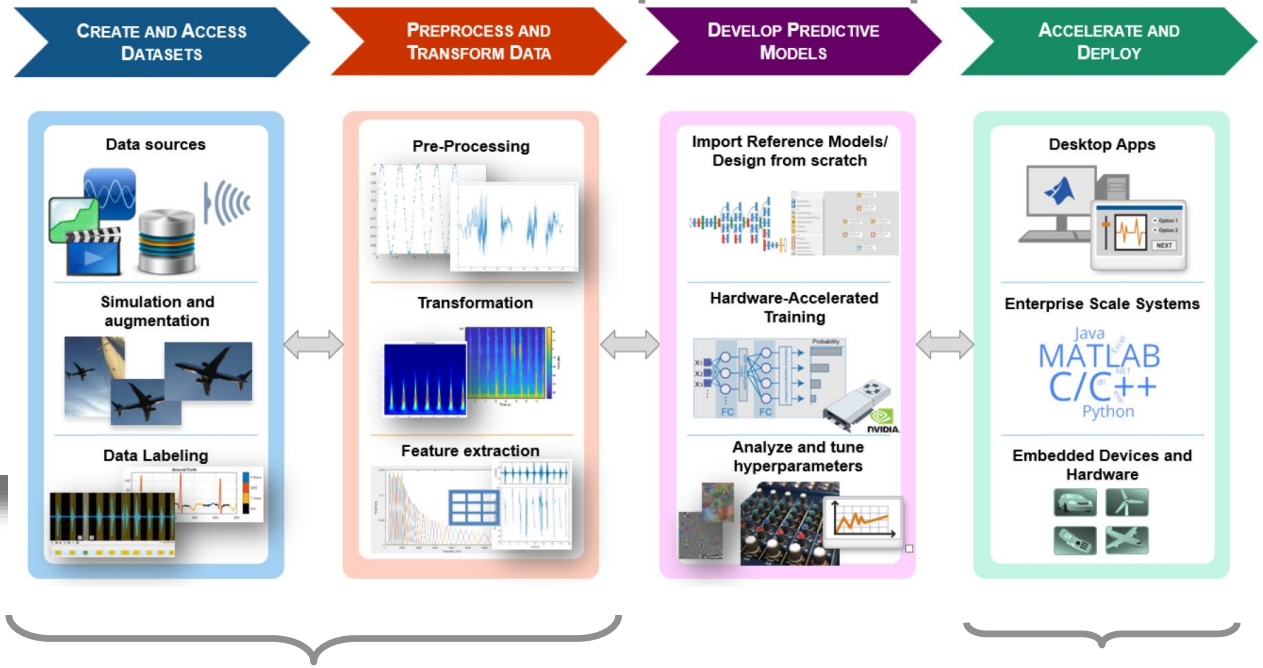

Domain-specific data processing desirable

Limited reference research

Deployment and scaling to various platforms

Deep Learning Workflow MATLAB Strengths - Signals, Time-Series, Text

Data-labeling Apps and examples

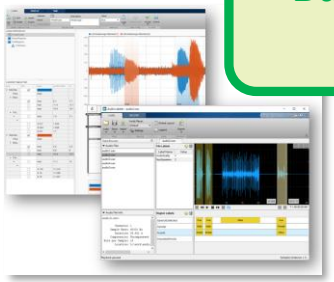

Augmentation and simulation algorithms

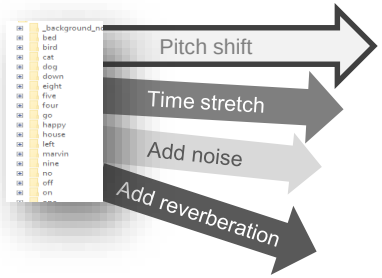

Deep Network Designer and analyzer

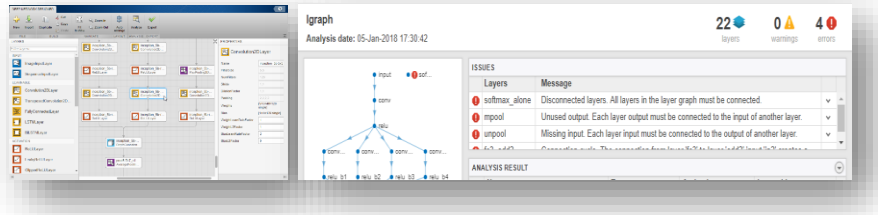

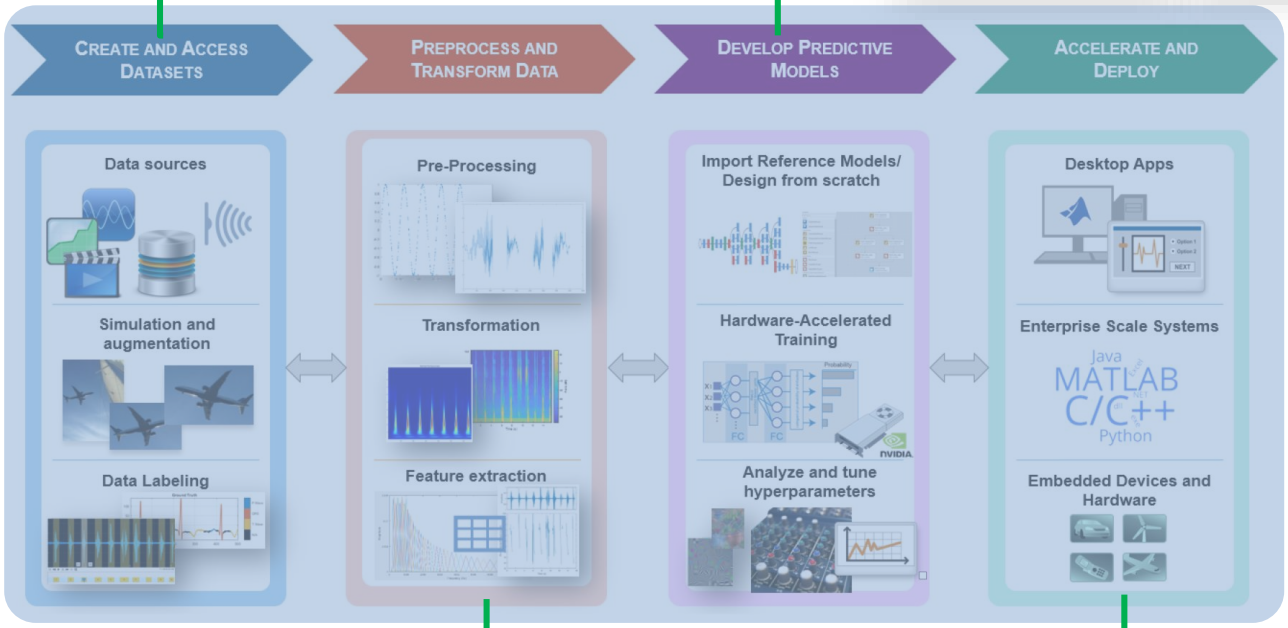

Collaboration in the AI ecosystem

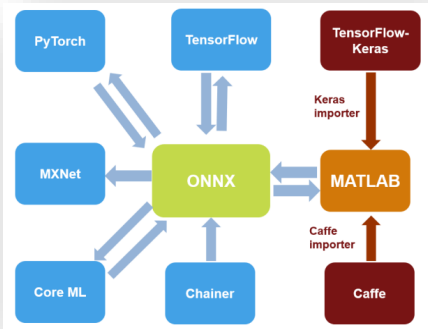

Application-specific algorithms and tools

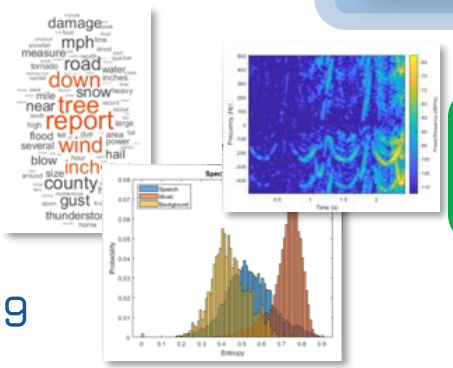

Deployment and Scaling to various platforms

Related Training

- Time series analysis
  - Signal Preprocessing and Feature Extraction for Data Analytics with MATLAB
- Deep learning (including non-vision applications)
  - Deep Learning with MATLAB
- Code generation
  - MATLAB to C with MATLAB Coder

https://nl.mathworks.com/services/training.html