MATLAB EXPO

임베디드 시스템 적용을 위한 AI 개발

신행재 부장, 매스웍스코리아

MathWorks®

Edge AI innovates many industries!

Hardware Constraints

Memory: 32GB
@ 10TFLOPs

Memory Footprint / Compute

Memory: 8GB
@ GFLOPs

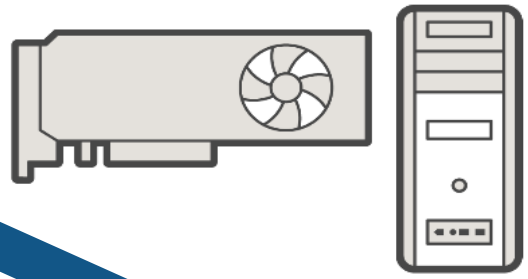

Memory: 1kB~1MB
@ MFLOPs

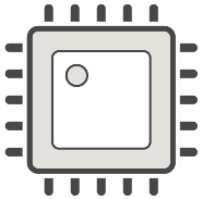

Power Consumption

* FLOPS: floating-point operations per second

What is “Edge” (Embedded) AI?

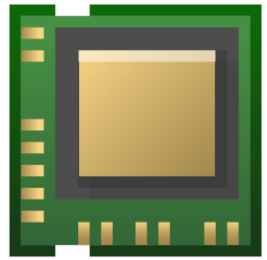

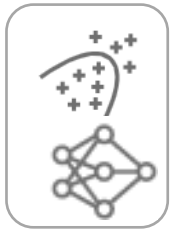

How can I make my AI smaller?

Data Scientist

The chip has only 500 KB memory – make that smaller

Embedded Software Engineer

Why is Edge AI (Model Compression) difficult?

AI is often big

Knowledge Gap

Model Compression Workflow

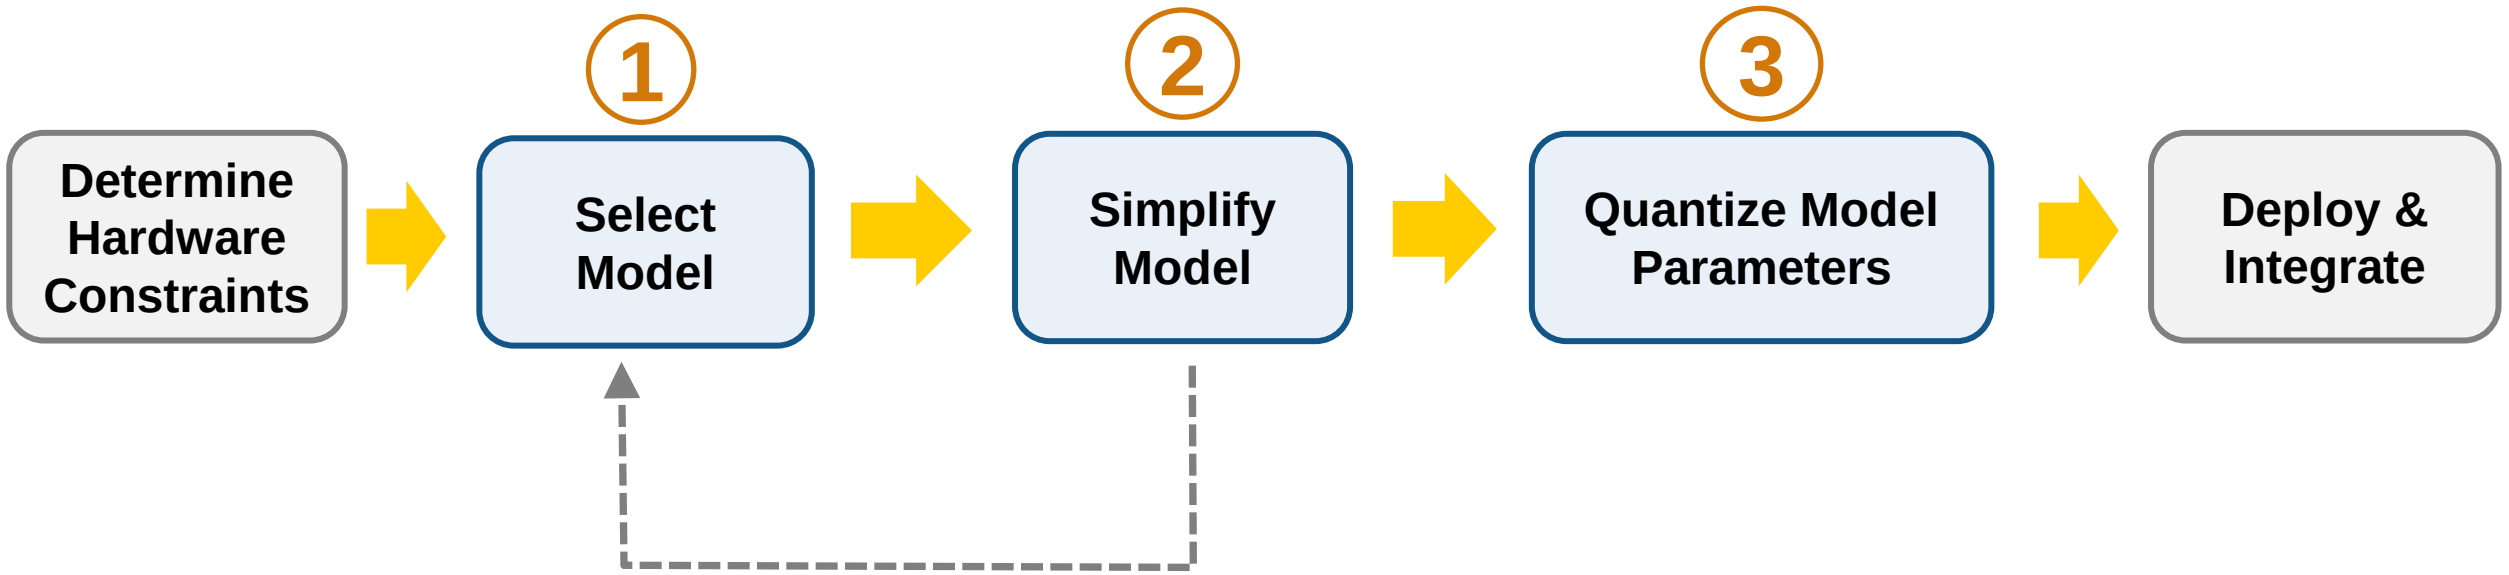

Compressing Machine Learning

Step 1 Size aware model selection

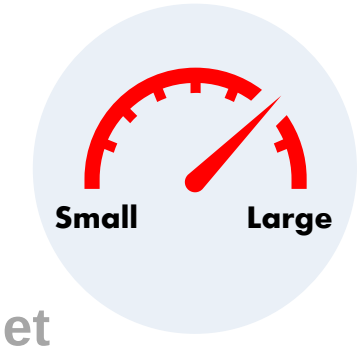

Size /
Execution Time

Deep Neural Net

Gaussian Process

Kernel SVM

Ensembles

Shallow Nets

Decision
Tree

Linear
Model

Accuracy
on Complex tasks

Model Compression Workflow for Machine Learning

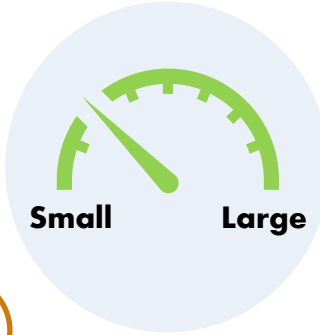

1

Classification / Regression Learner

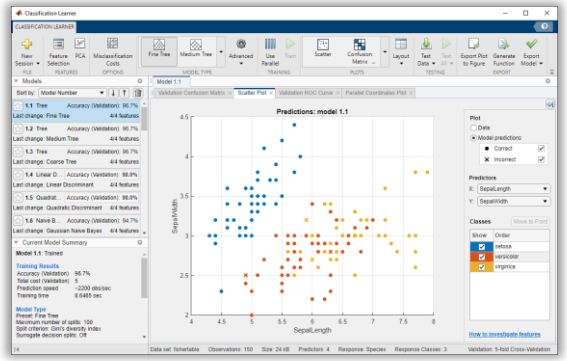

2

In-App Feature Selection

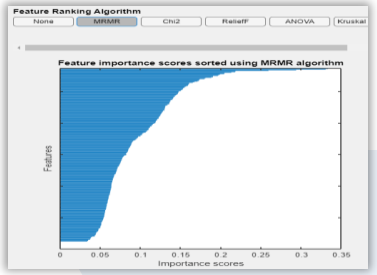

Bayesopt

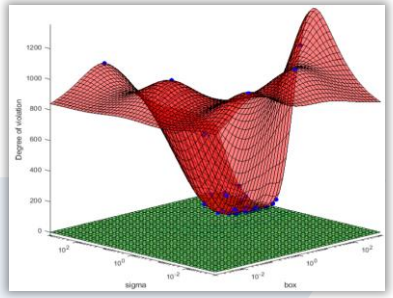

3

Fixed Point Designer / Native Simulink Block

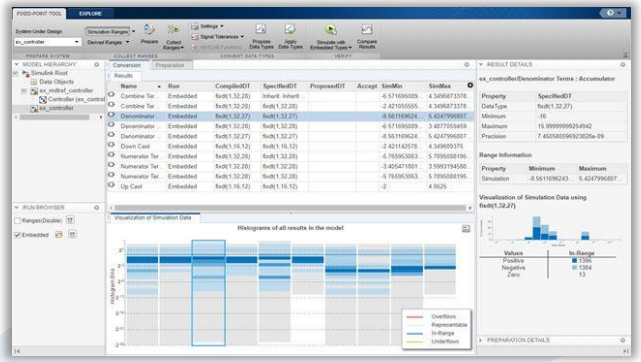

Simplify Model

- Select Features
- Tune Hyper-parameters

Determine Hardware Constraints → Select (Initial) Model → Simplify Model → Quantize Model Parameters → Deploy & Integrate

Demo: Embedding AI in an intelligent Hearing Aid

0.5 to 256 kB
on-chip memory

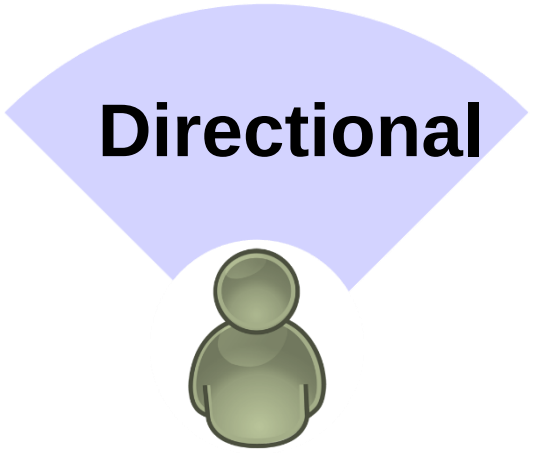

Demo: Fit Machine Learning for Intelligent Hearing Aid

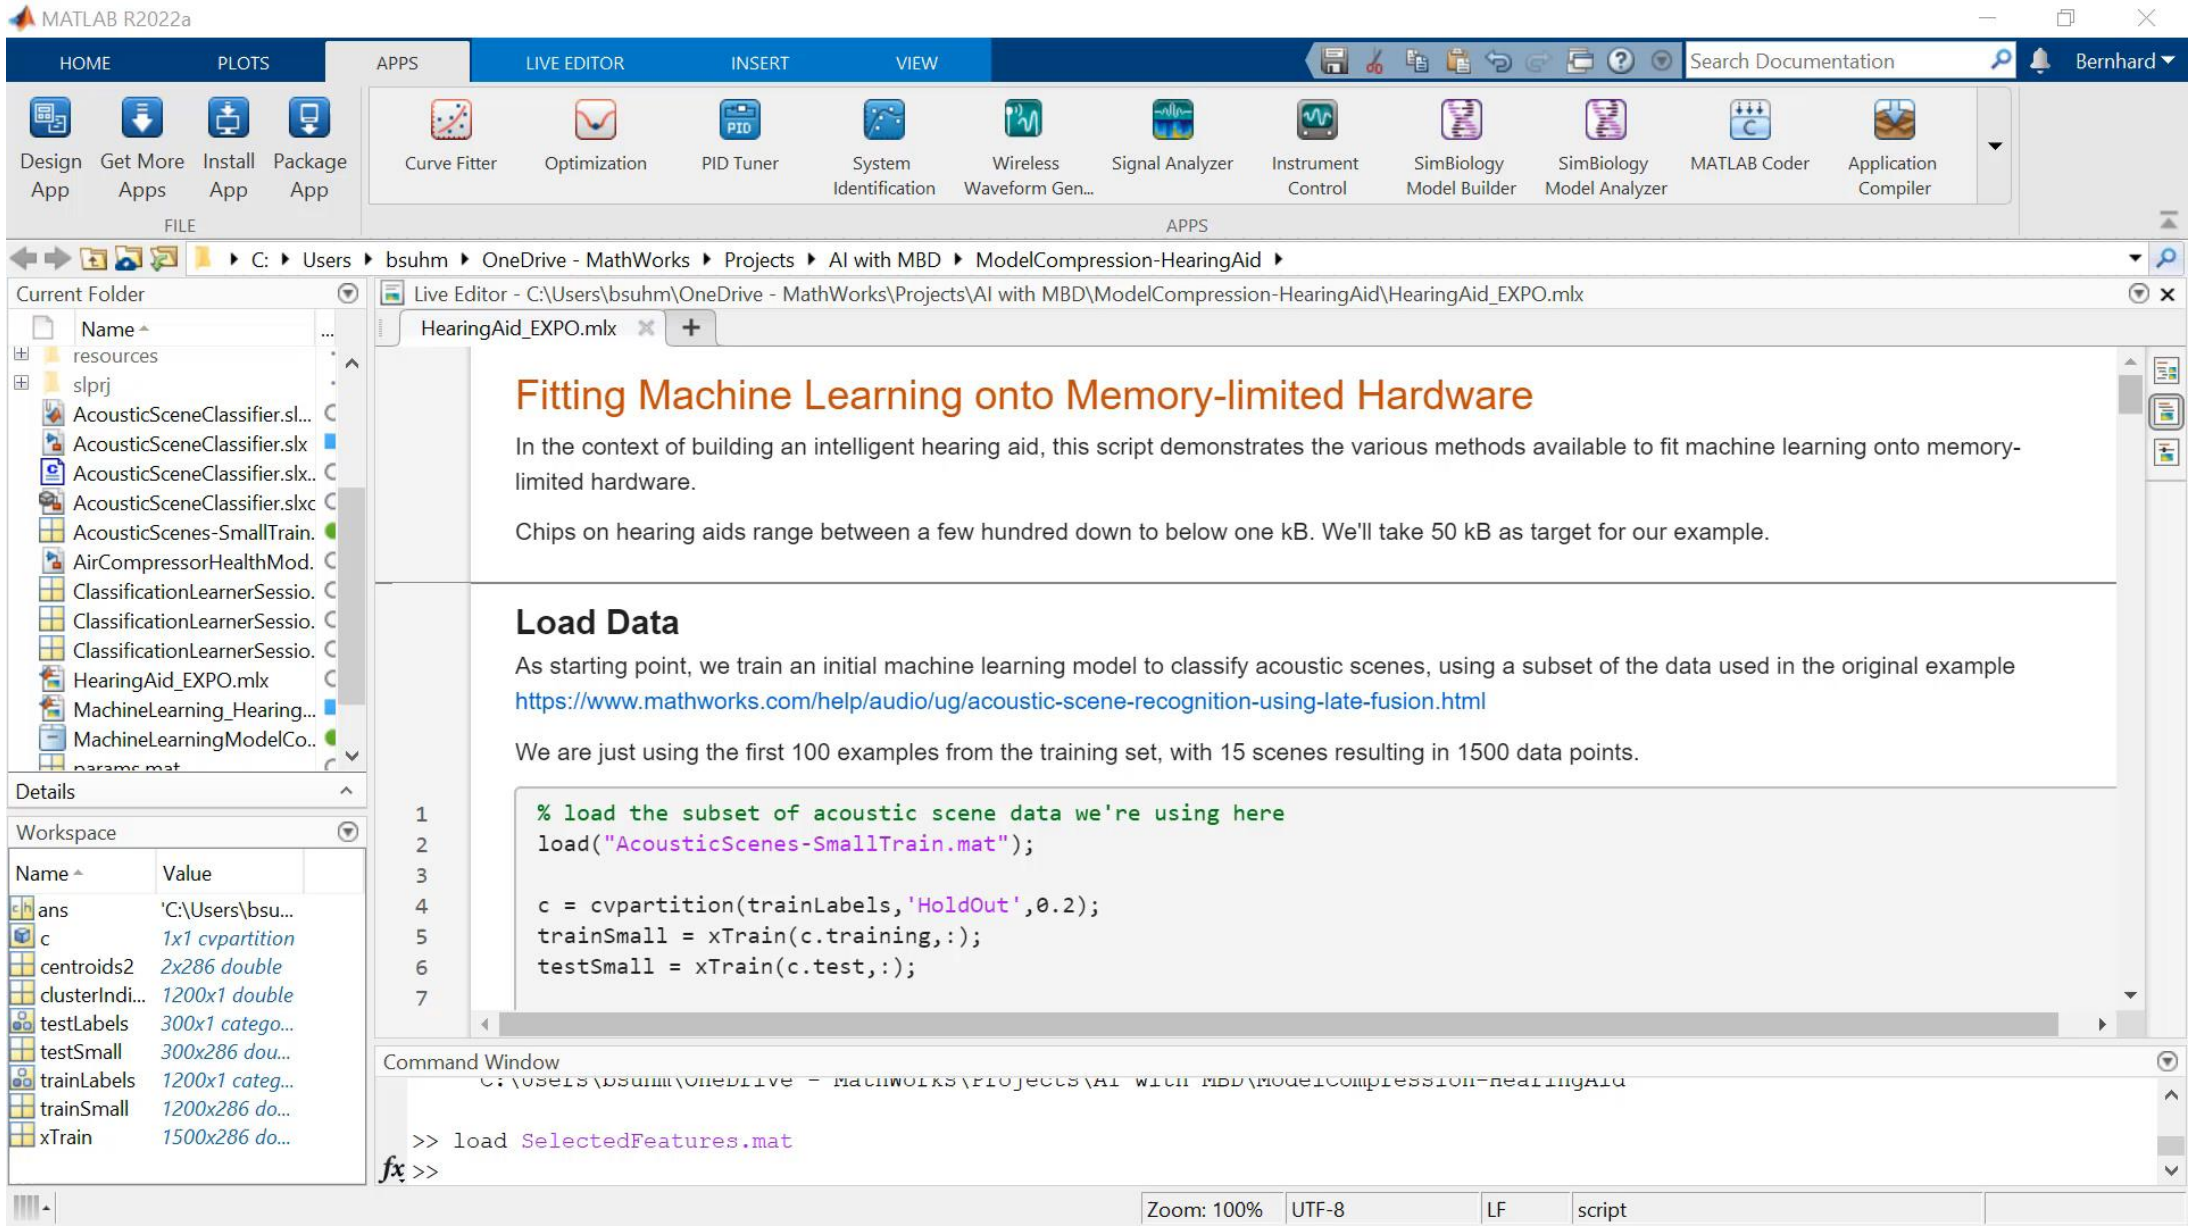

Machine Learning Demo Size Reduction by factor 20

Target Hardware
Size: 50 kB

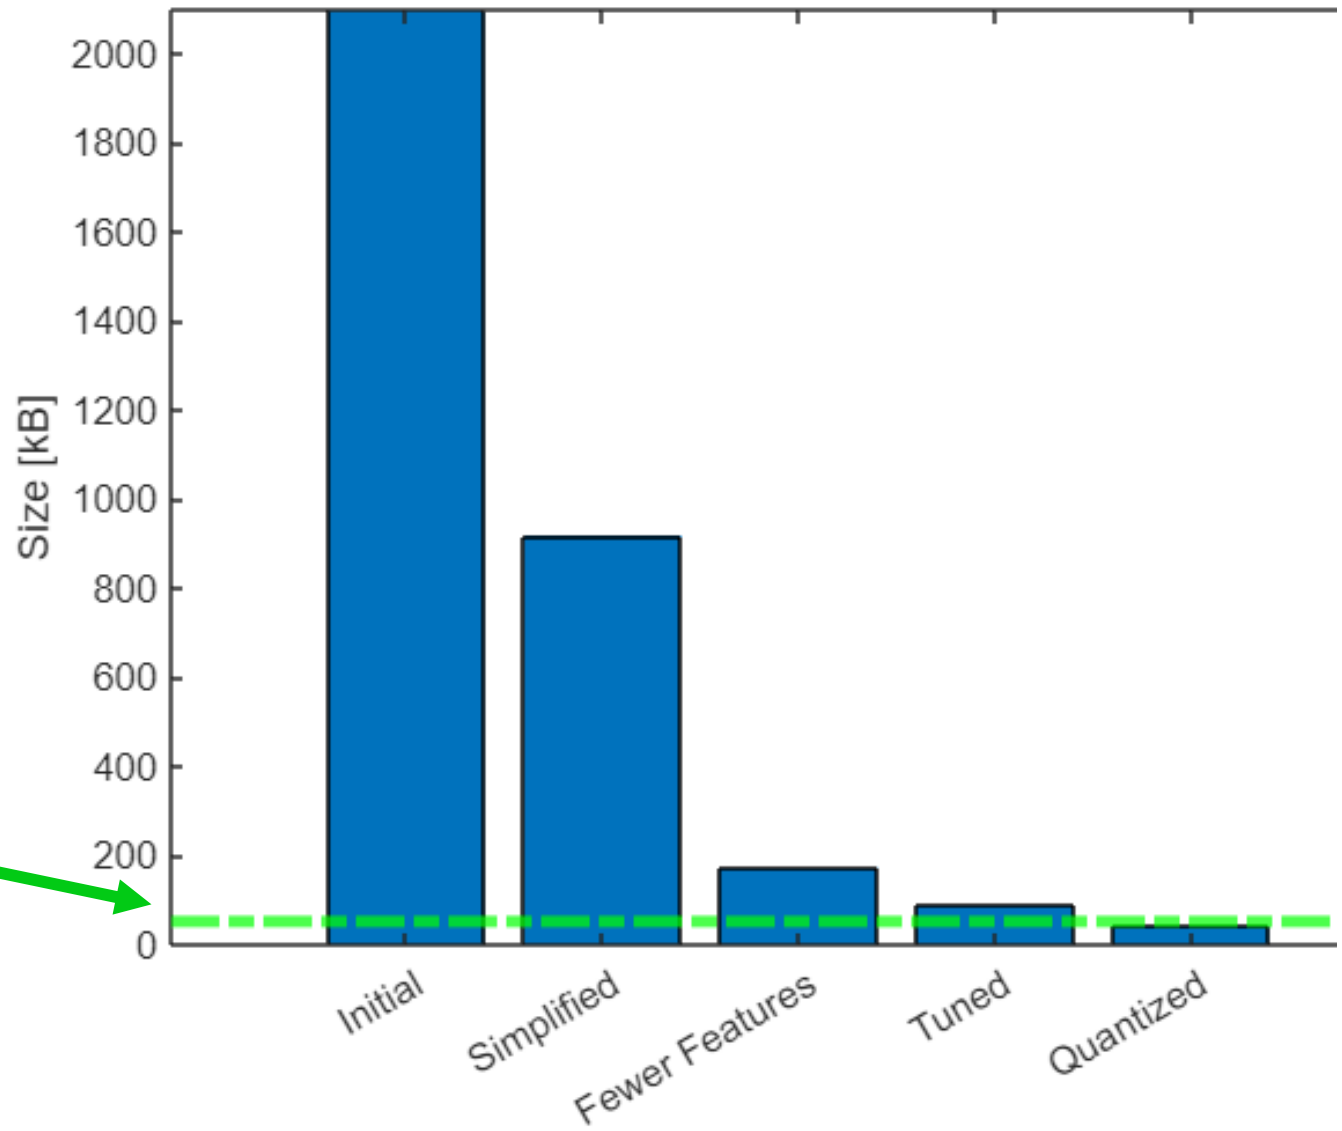

Compressing Deep Learning

Step 1 Size aware model selection

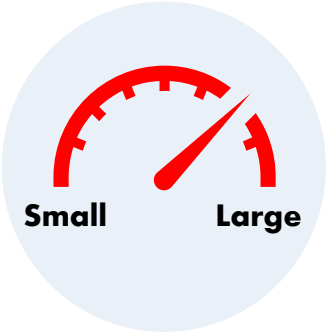

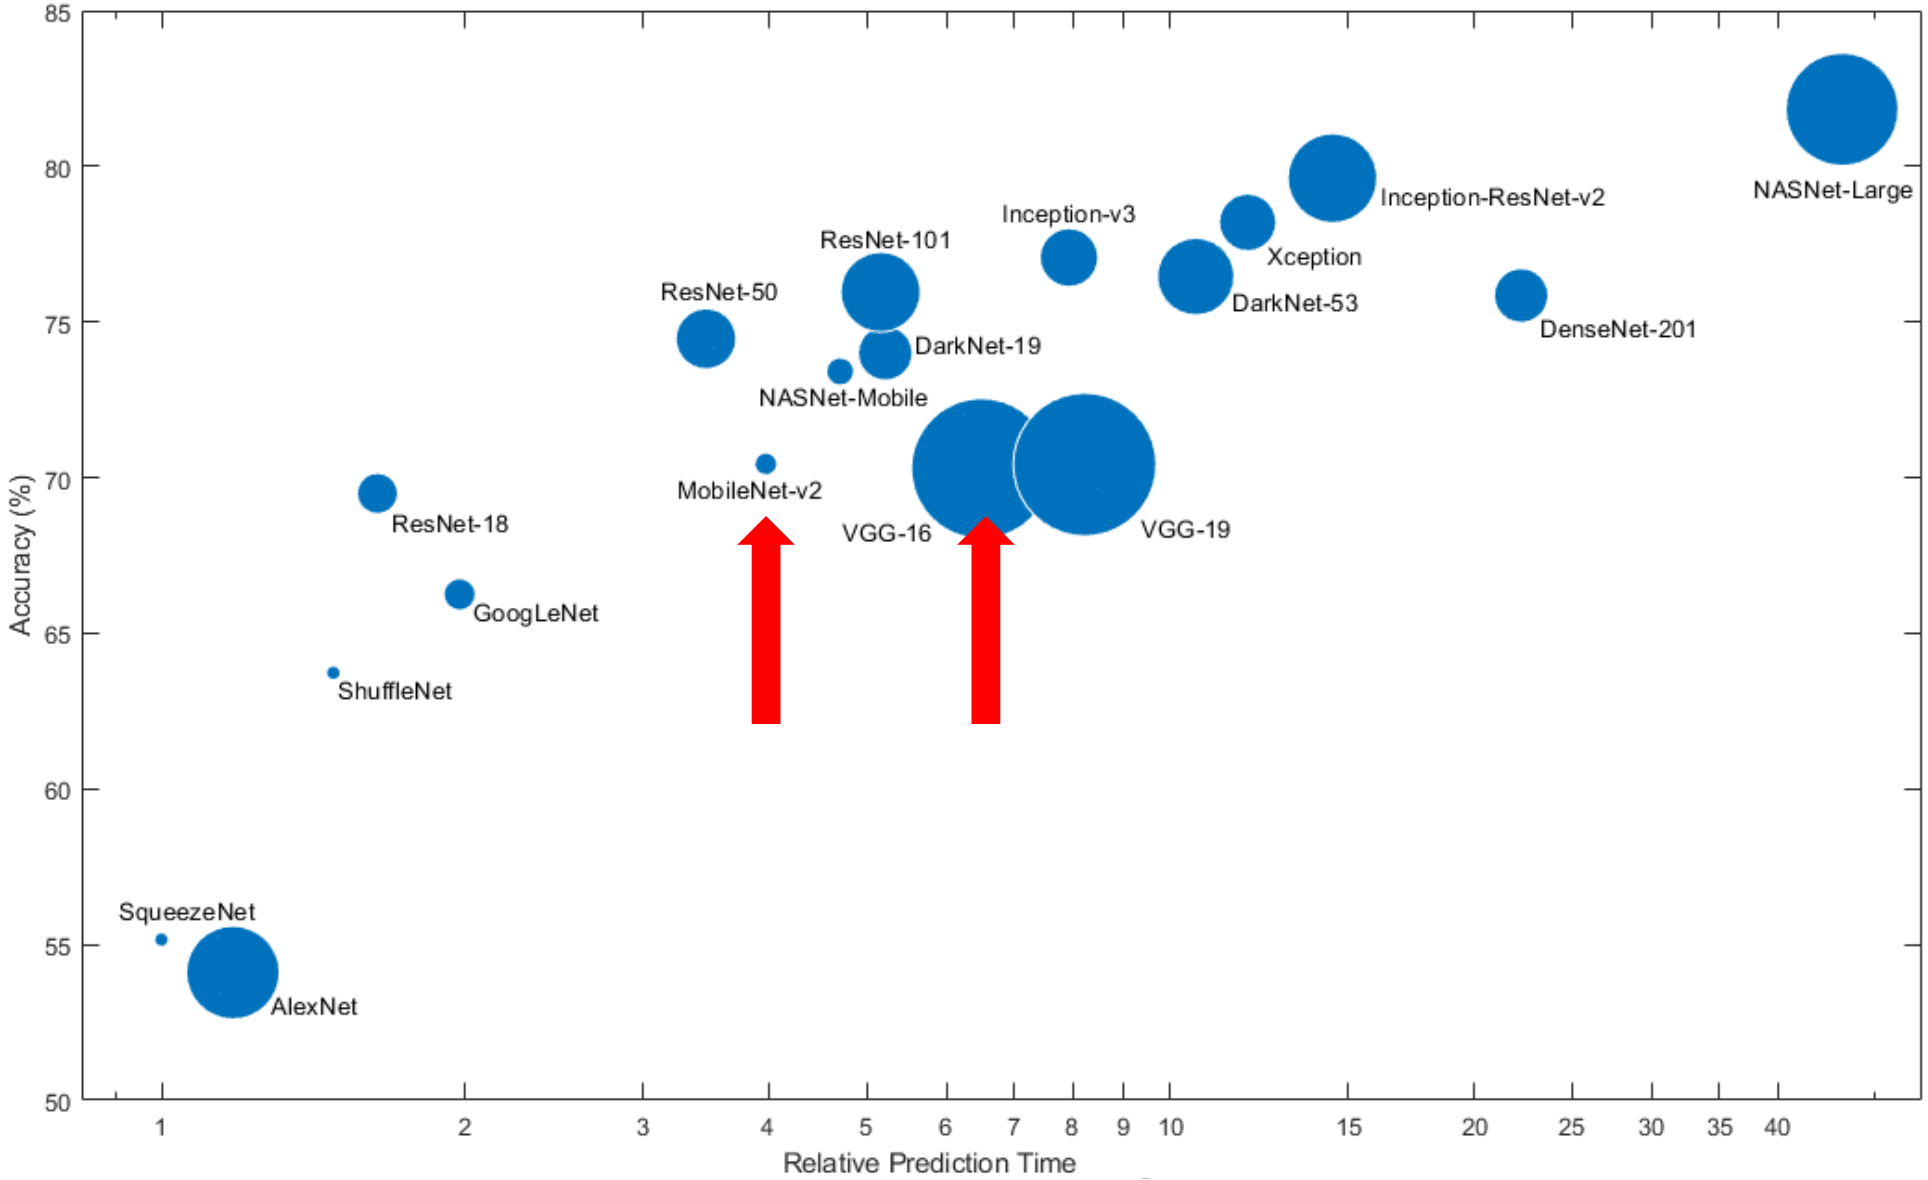

Step 2 Smart pruning

Remove **unimportant** parts of the network

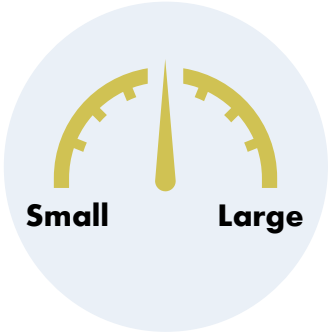

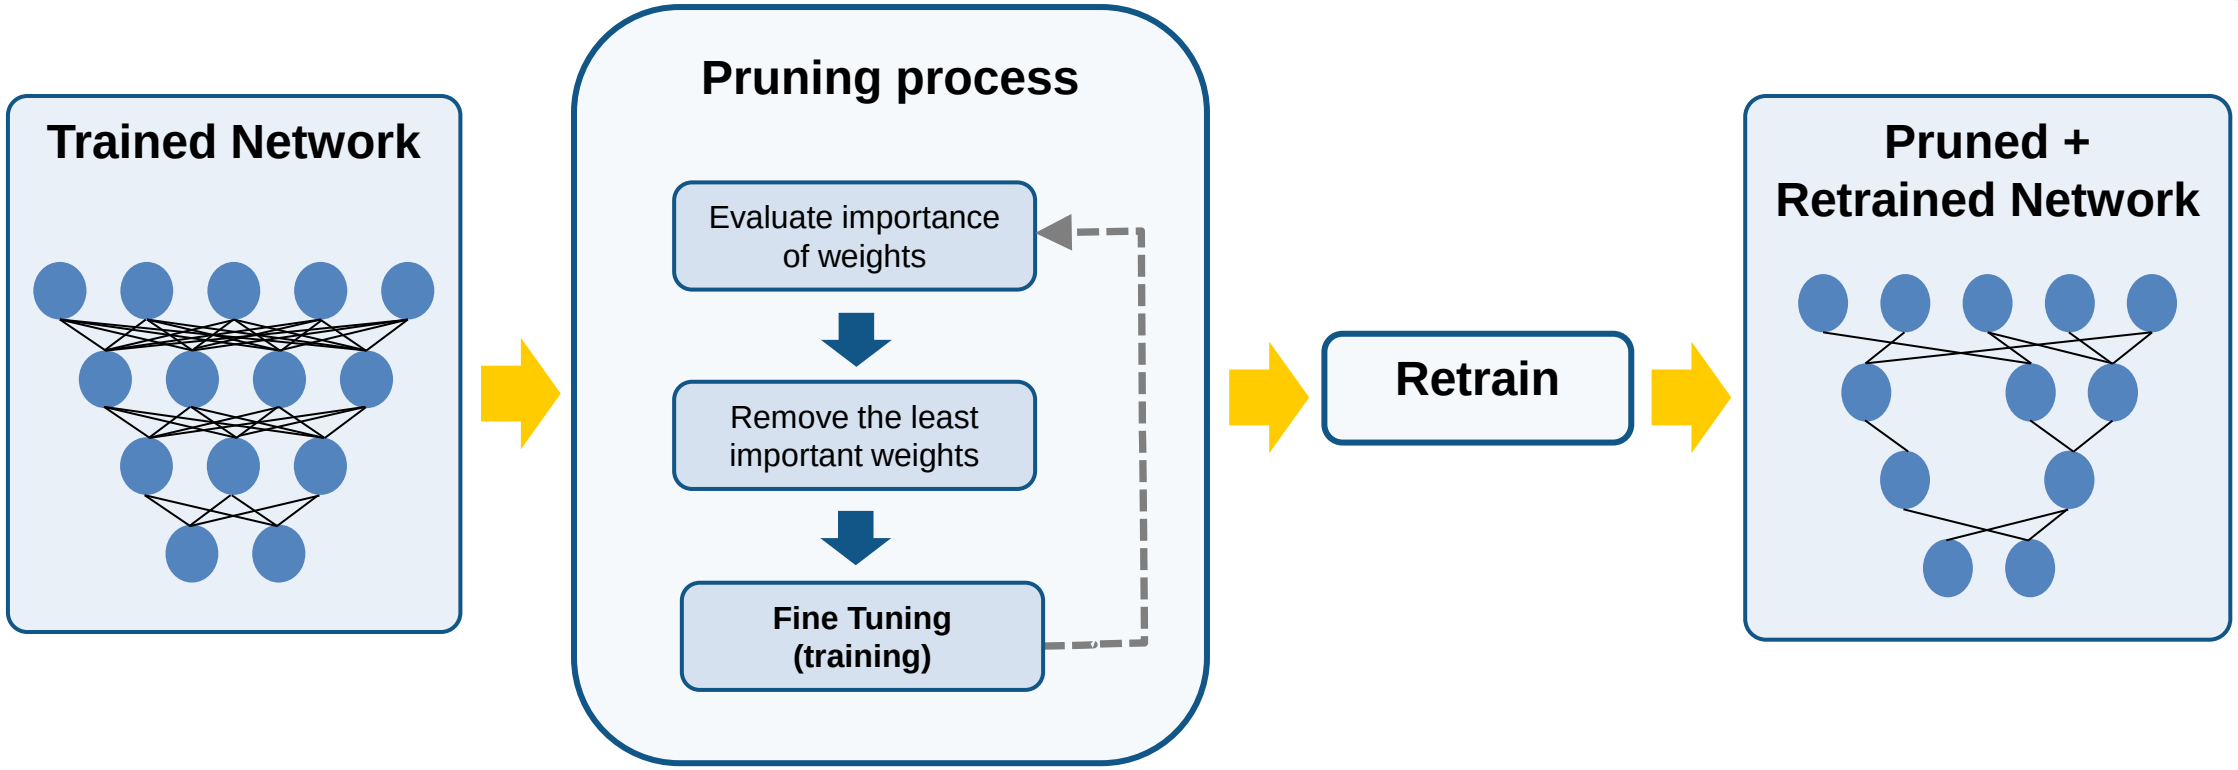

Step 3 Quantize your model

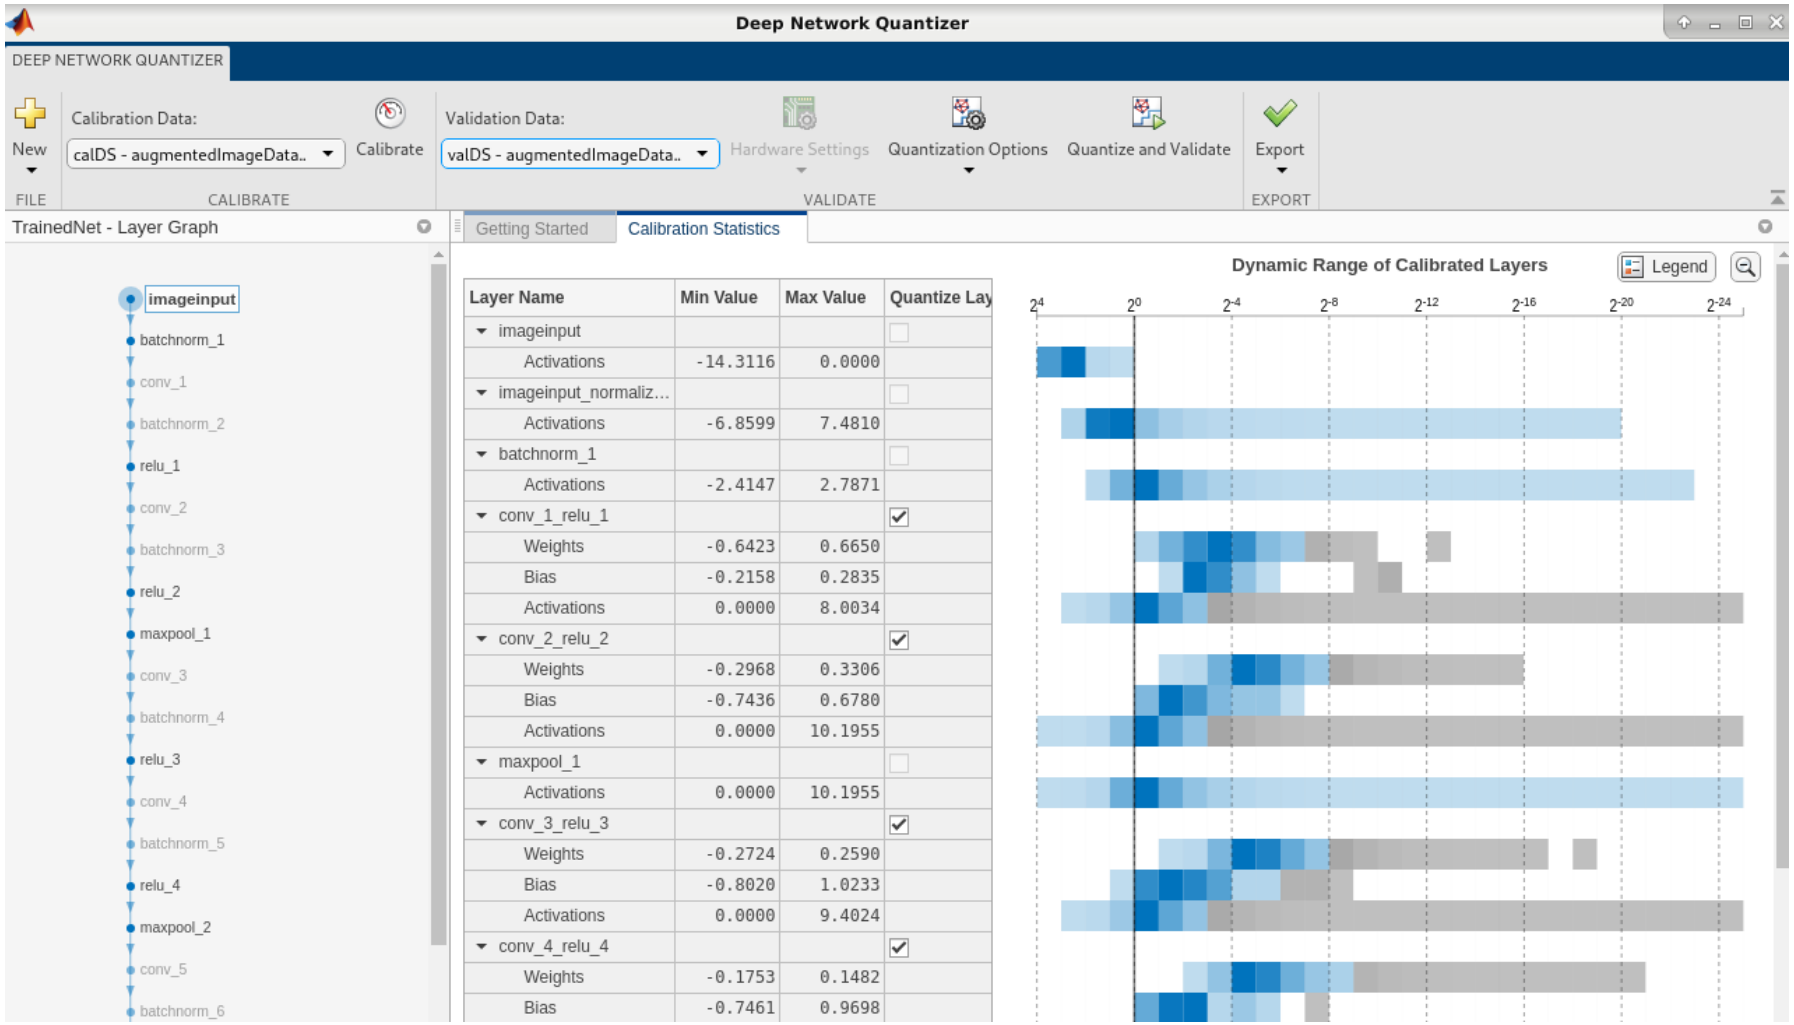

Deep Learning Demo: Scene classification

Classify 10 classes

More difficult problem → more complex model

Functionality for Compressing Deep Neural Nets

1 Deep Network Designer

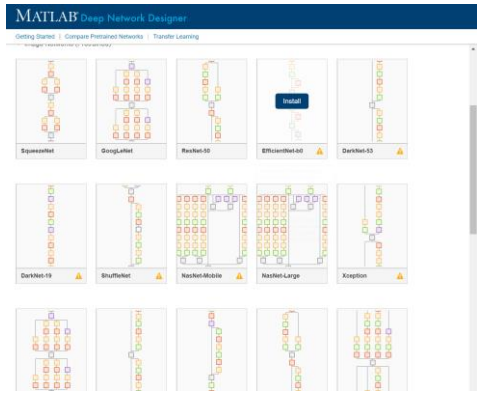

2 Taylor Pruning

```
taylorPrunableNetwork(net)
```

3 Deep Network Quantizer

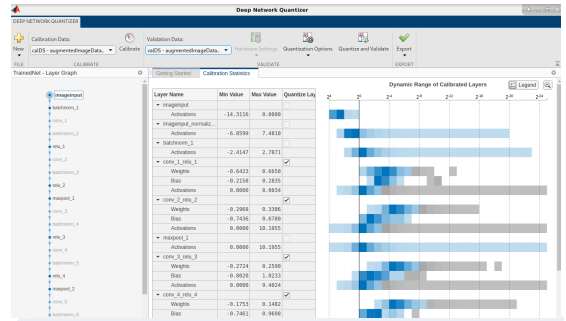

Determine Hardware Constraints → **Select Model** → **Simplify Model** (**Prune Deep Neural Network**) → **Quantize Model Parameters** → **Deploy & Integrate**

MATLAB R2022a

HOME PLOTS APPS LIVE EDITOR INSERT VIEW

Search Documentation

Emelie

New Open Save Compare Print Export FILE

Go To Find Bookmark NAVIGATE

Text Normal B I U M TEXT

Code Control Task Refactor CODE

Run Section Section Break Run and Advance Run to End SECTION

Run Step Stop RUN

Scene identification demo

Current Folder

Name

- PruneQuantizeDemo
  - Helper Functions
    - accuracyVisualization.mlx
    - analyzeNetworkMetrics.mlx
    - assemblePruneLayerGraph.mlx
    - buildDatasetforCNN.m
    - initializeTrainingPlots.mlx
    - modelAccuracy.mlx
    - modelLossPruning.mlx
    - numConvLayerFilters.mlx
    - preprocessMiniBatchTraining...
    - pruneAmountVisualization.mlx
    - PruningLoop.mlx
    - quantizationAccuracyVisualiz...
  - Trained Networks
    - dlquantizePruned.mat
    - prunableNet.mat
    - retrainedPrunedDAGNet.mat
    - trained10classNetwork.mat
  - data.mat
  - Emelie_ASC_Compression10Cla...
  - trainingOptionsRetrain.mlx
- Scene identification demo
  - Helper Functions
  - videos

accuracyVisualization.mlx (Live Script)

No details available

Live Editor - Emelie_ASC_Compression10ClassesDNN.mlx *

Variables - prunedNet

Emelie_ASC_Compression10ClassesDNN.mlx * | LIVEAcousticSceneRecognitionUsingLateFusionExample.mlx | accuracyVisualization.mlx

Step 1: Select Model

1

Select Model

Load original trained CNN model and dataset

```
1  load('trained10classNetwork');
2  load('data')
```

Note: Sounds have been converted to spectograms

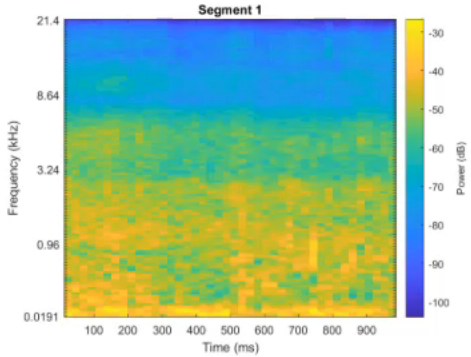

Command Window

Workspace

Name Value

Zoom: 100% UTF-8 LF script Ln 1 Col 32

Deep Learning Demo Size Reduction by factor 5

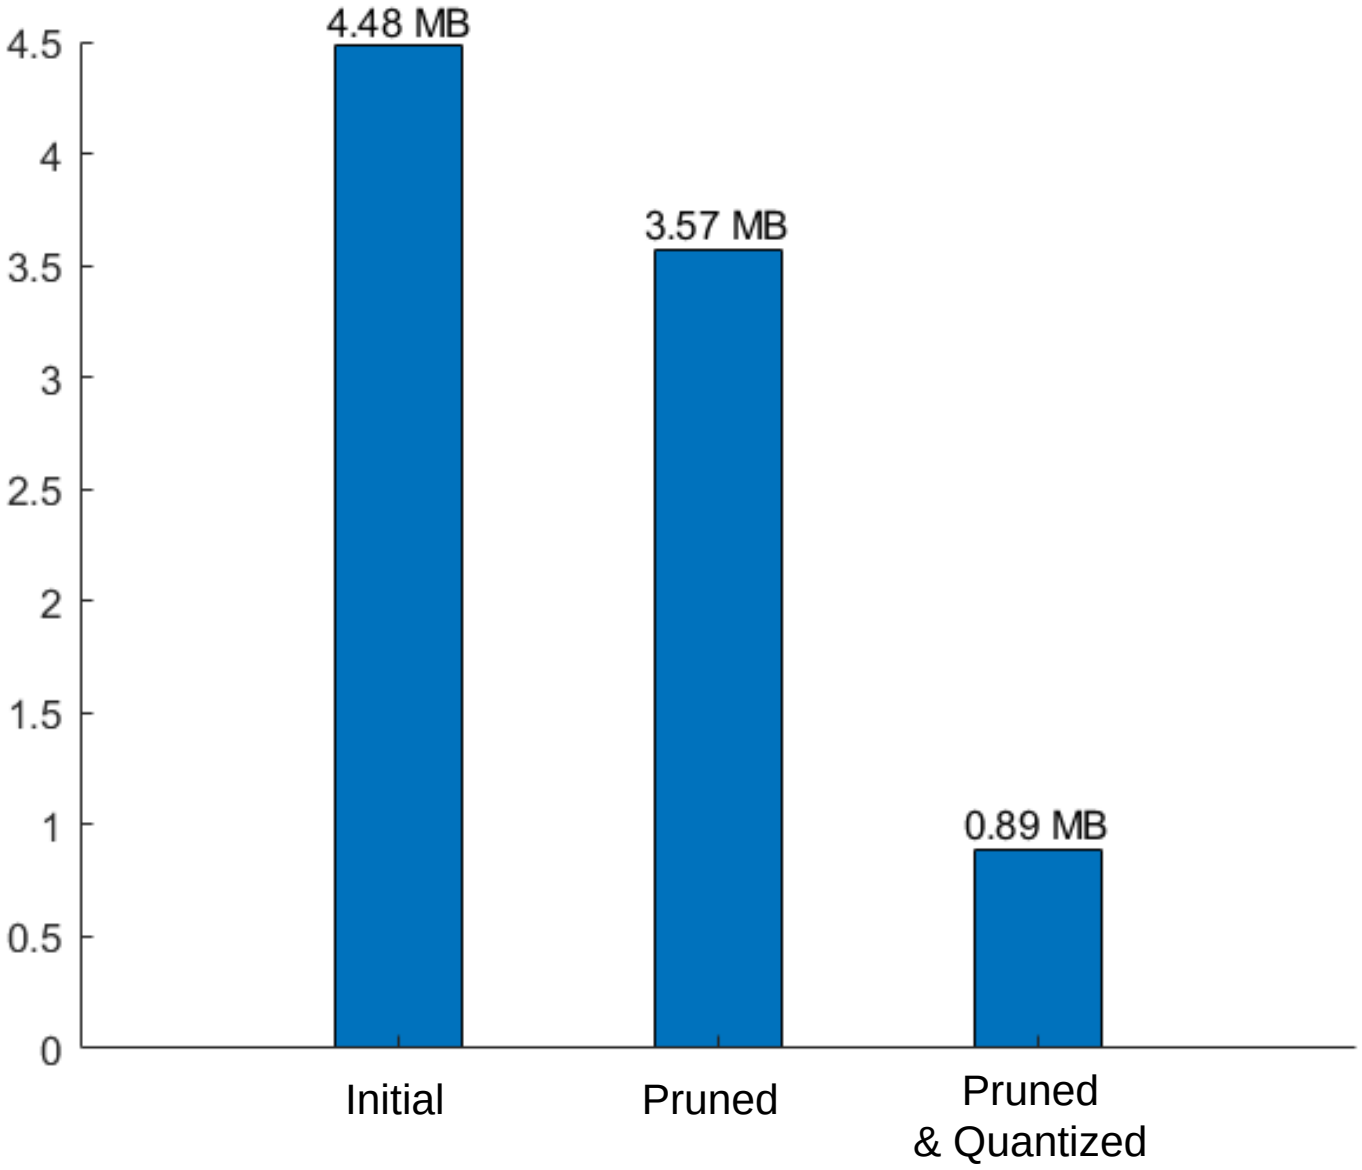

One Codebase – Many Embedded Deployment targets

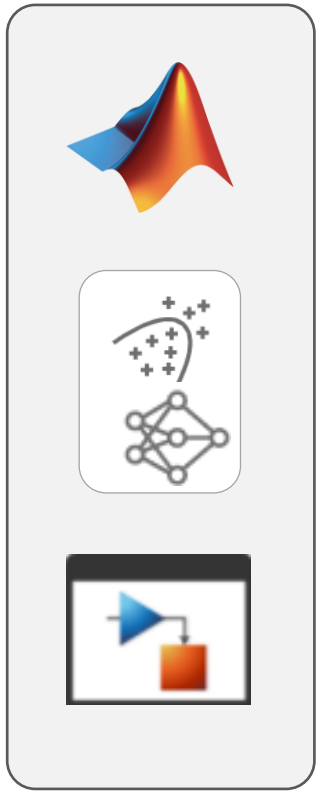

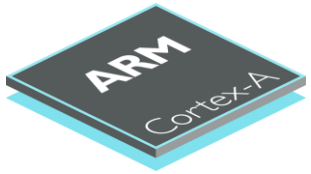

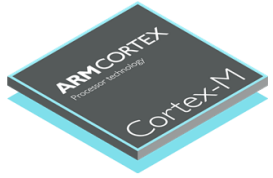

Zynq is a registered trademark of Xilinx Corporation
Intel logo is a registered trademark of Intel Corporation

One Codebase – Many Embedded Deployment targets

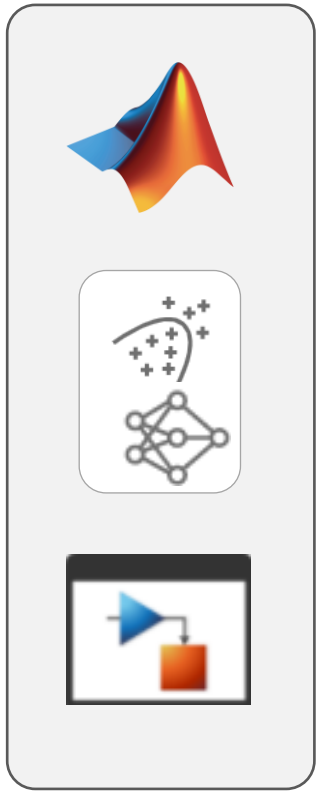

Code Generation

GPU

NVIDIA

CPU

ARM Cortex-A

intel inside XEON

AMD

FPGA

XILINX

ALTERA now part of Intel

Micro-controller

ARM CORTEX Cortex-M

NXP

MICROCHIP

ARDUINO

PLC

SIEMENS

MITSUBISHI ELECTRIC

Rockwell Automation

Schneider Electric

Zynq is a registered trademark of Xilinx Corporation
Intel logo is a registered trademark of Intel Corporation

Conclusions

You can fit AI for many applications onto limited hardware

MathWorks tools make fitting AI models on constrained hardware a lot easier

Same high-level Workflow for any type of AI

Learn More

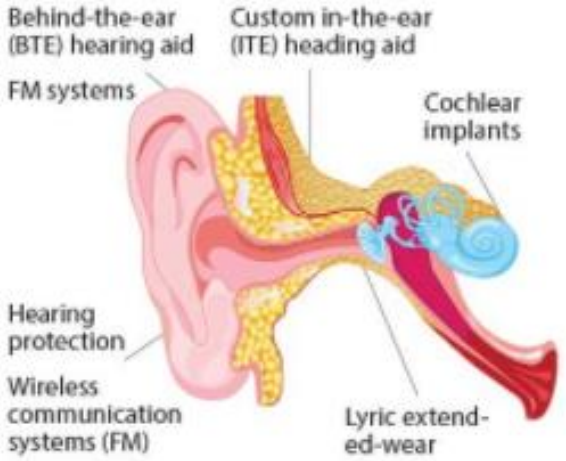

Hearing Implant using MBD

Autonomous Tractor

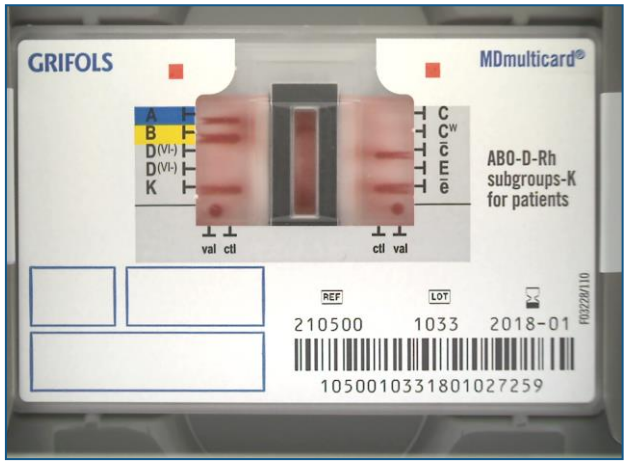

Card to Classify Blood Type

IDNEO

To get your started:

Learn about Embedded Deployment

Quantization of classification SVM (Doc)

Deploy Hand-Gesture Classifier onto Arduino (Doc)

Generate C/C++ Code from Simulink (Video)

Quantizing a Deep Neural Network (Video)

MATLAB EXPO

Thank you

MATLAB EXPO