



Image Recognition for Space Applications using Deep Learning on FPGAs & SoCs

Lunar Crater Detection

Stephan van Beek

svanbeek@mathworks.com

*European Technical Specialist
SoC/FPGA/ASIC Design Flows*

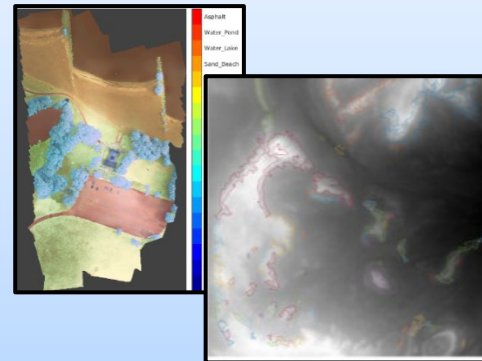


Artificial Intelligence on Embedded Devices

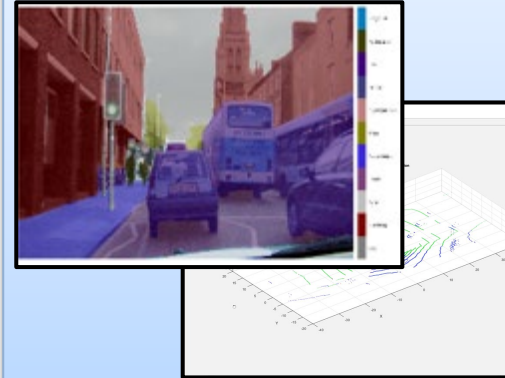
Satellite Navigation



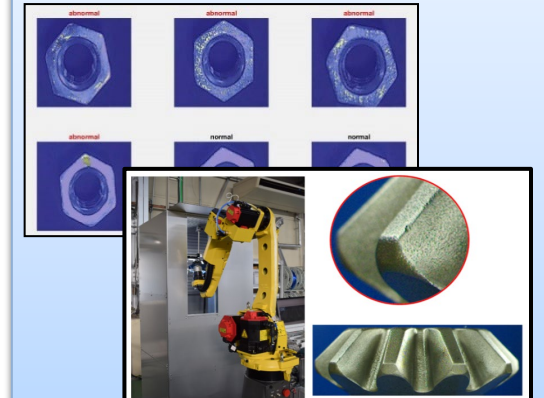
Airborne Image Analysis



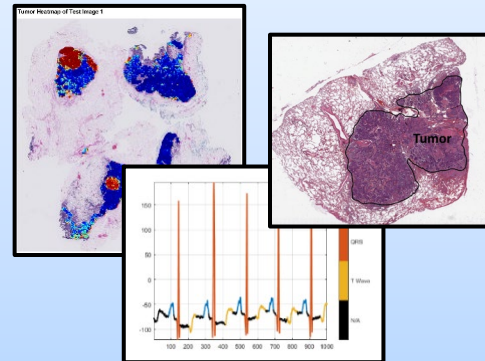
Autonomous Driving



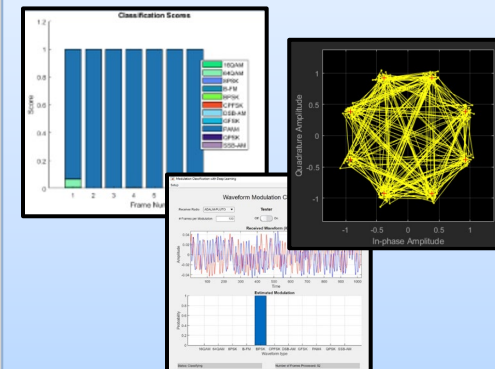
Industrial Inspection



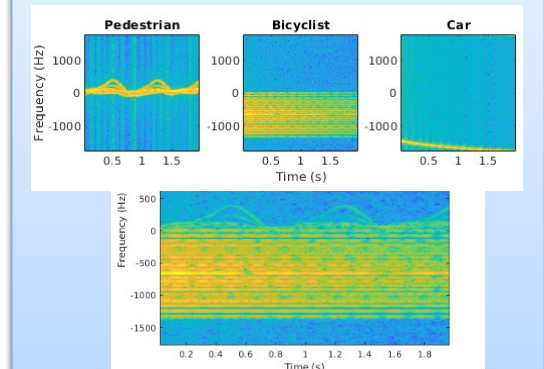
Medical Image Analysis



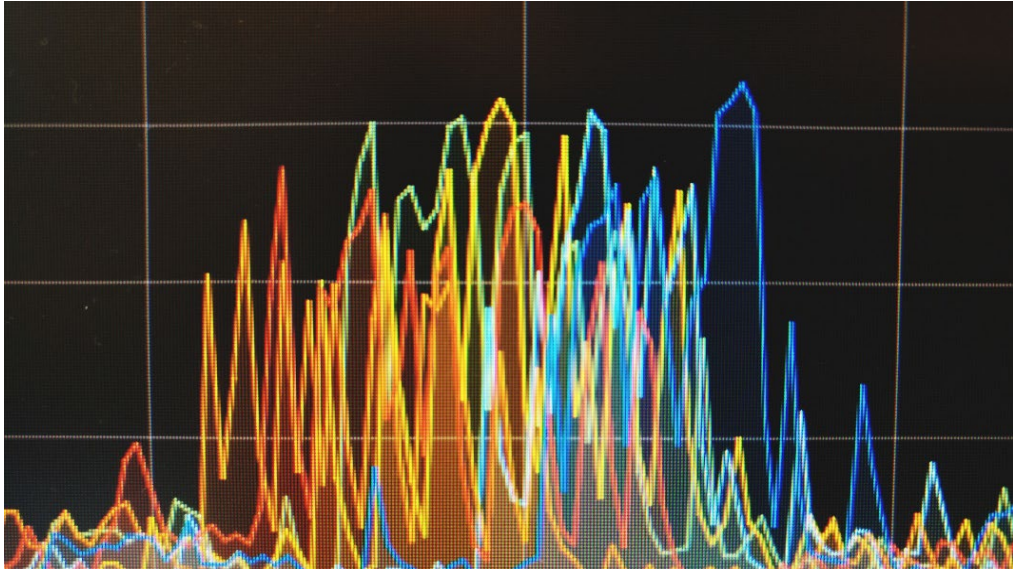
Wireless Modulation Classification



Radar Signature Classification



Machine learning has been deployed on ground segment applications for several years → now moving into space



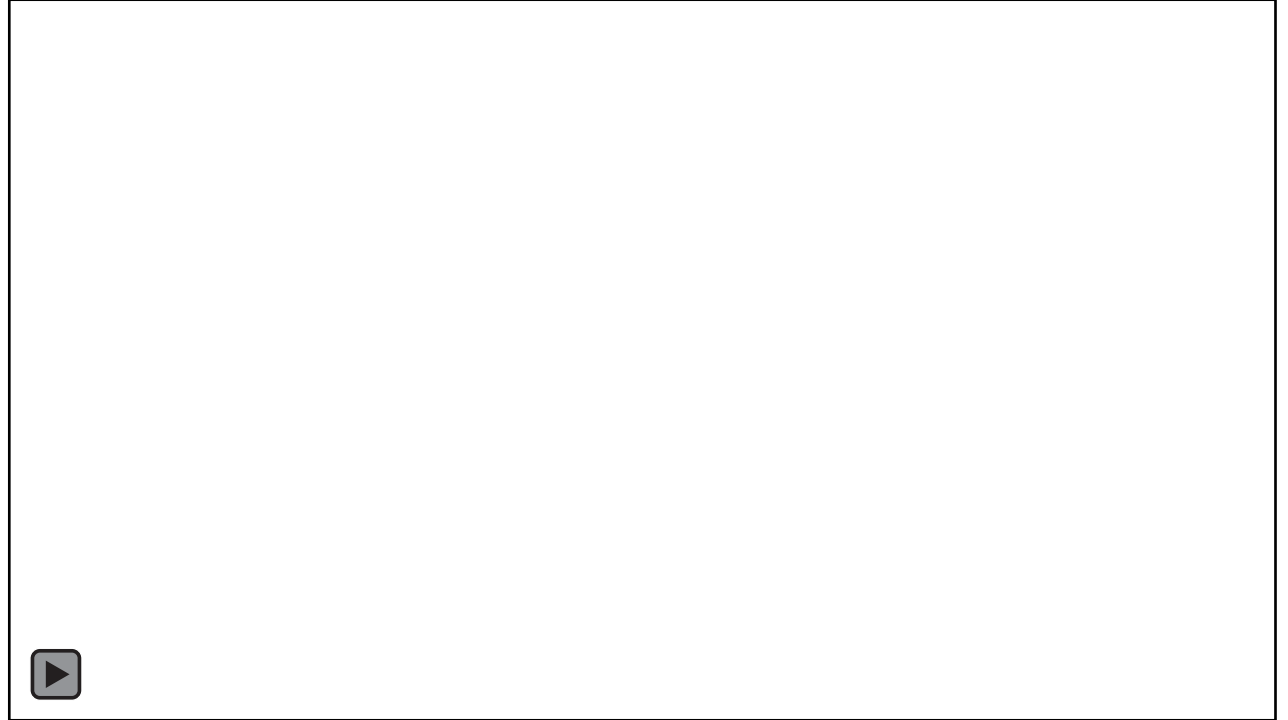
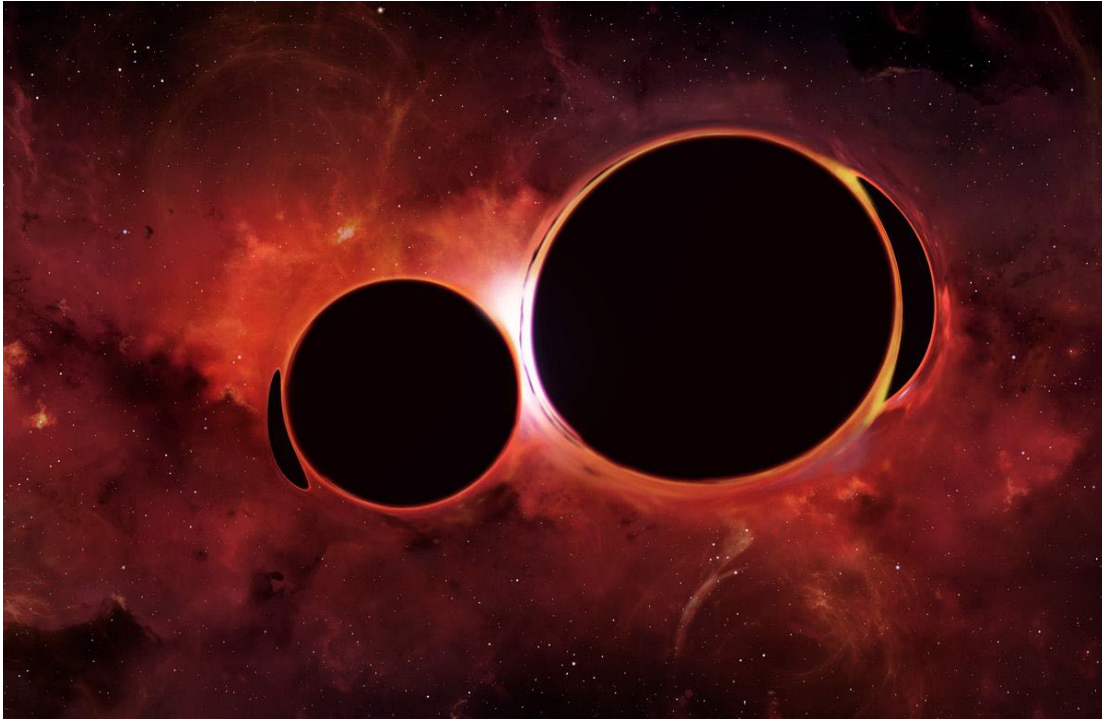
Telemetry Outlier Detection



Geospatial Analytics

Deep Learning Helps Detect Gravitational Waves

Hunting for Black Holes with Artificial Intelligence



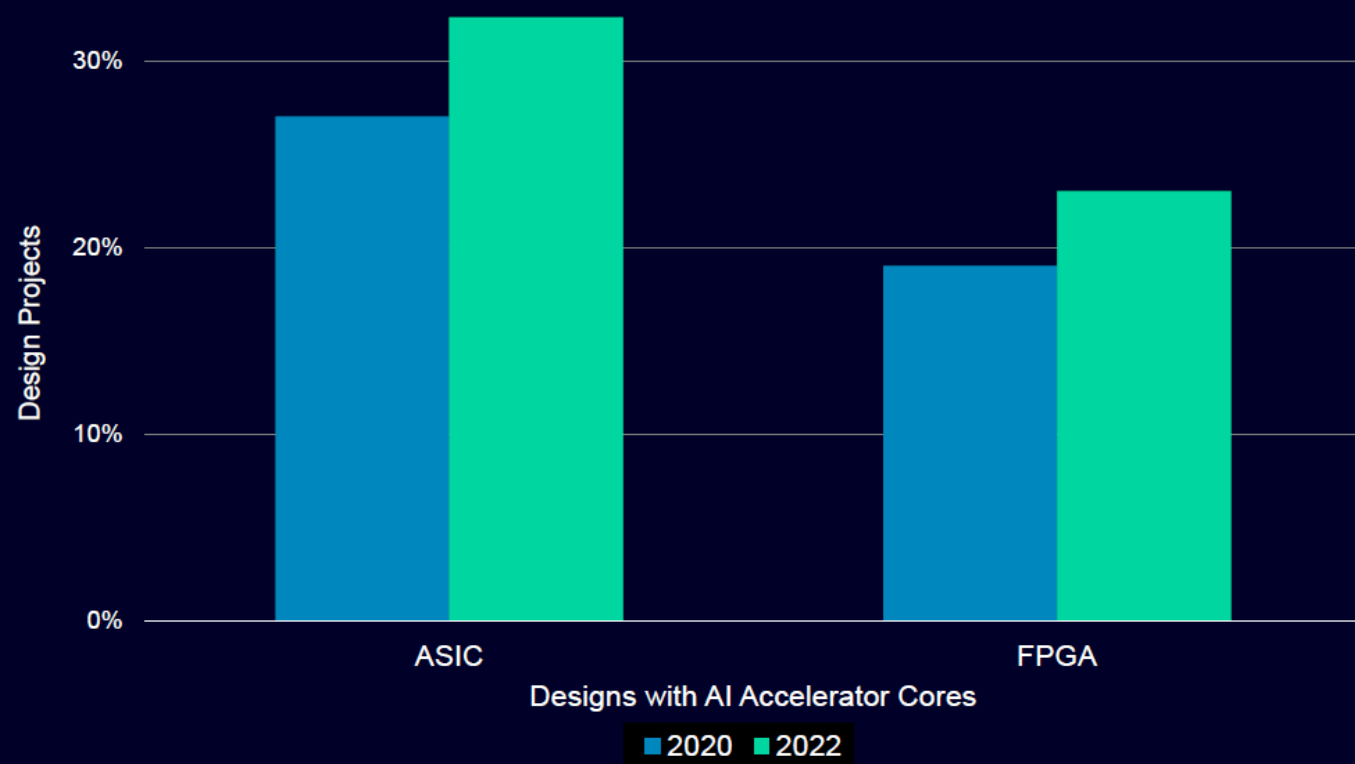
Max Planck Institute used AI and laser interferometry to detect gravitational waves caused by space-time distortions in our solar system.

Industry Trends

Designs with AI accelerator cores increasing

32%
ASICs with AI Cores

23%
FPGAs with AI Cores

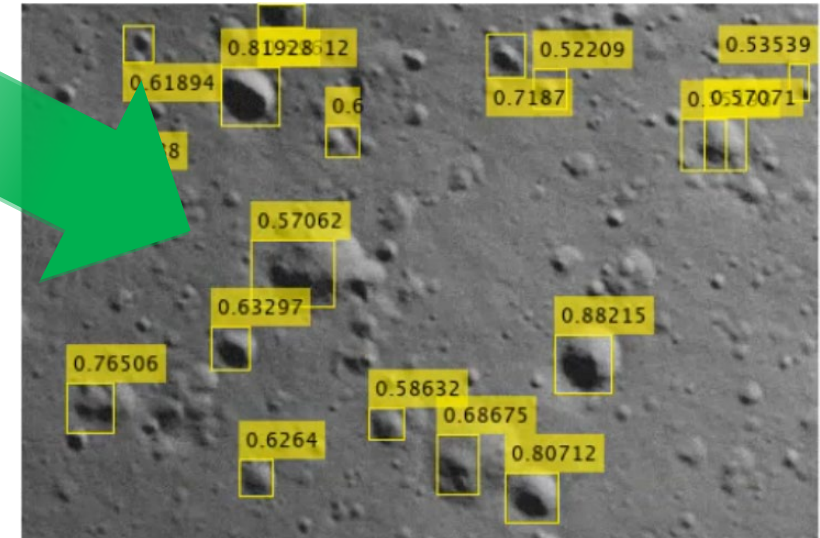
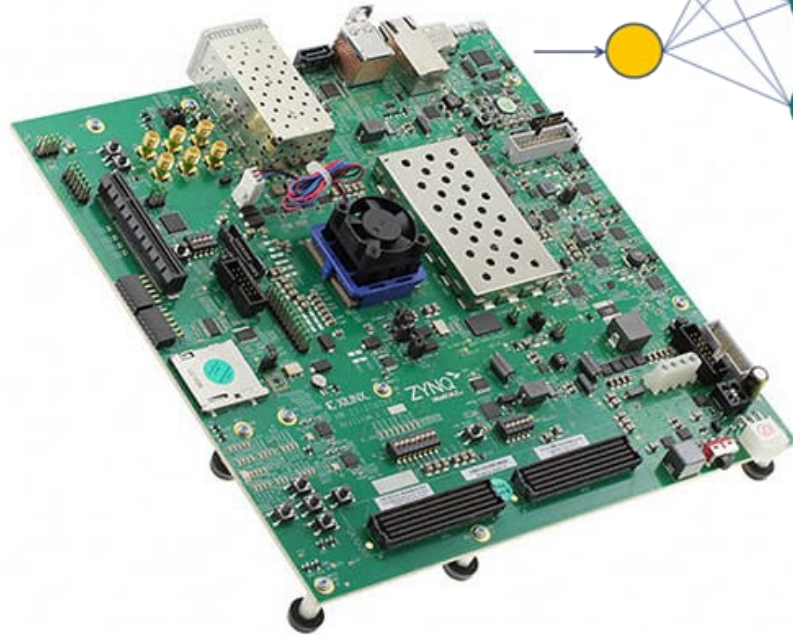
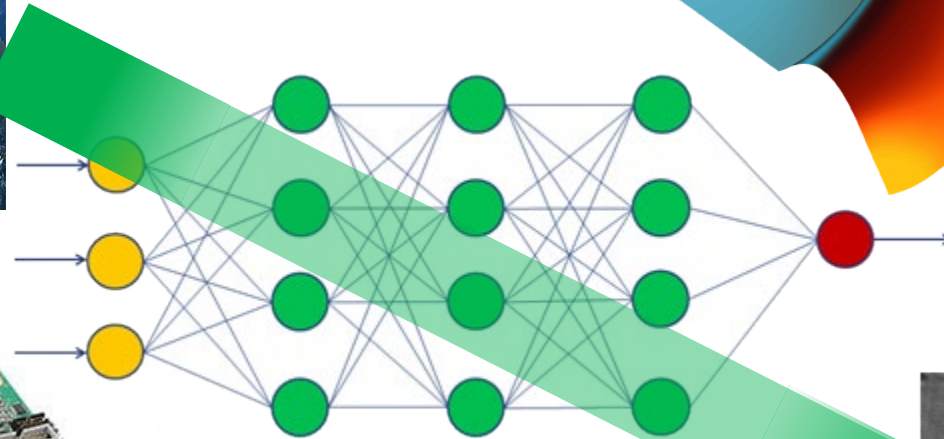
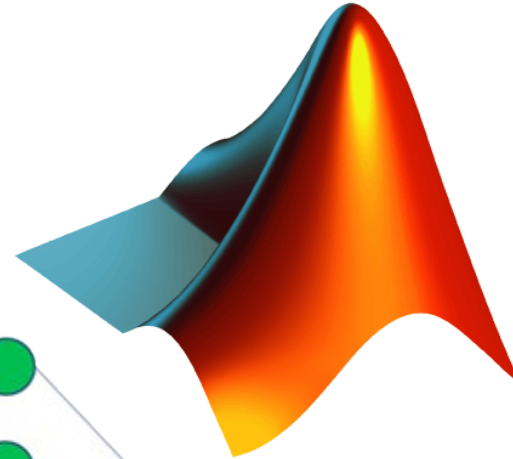


Source: Wilson Research Group and Siemens EDA, 2022 Functional Verification Study

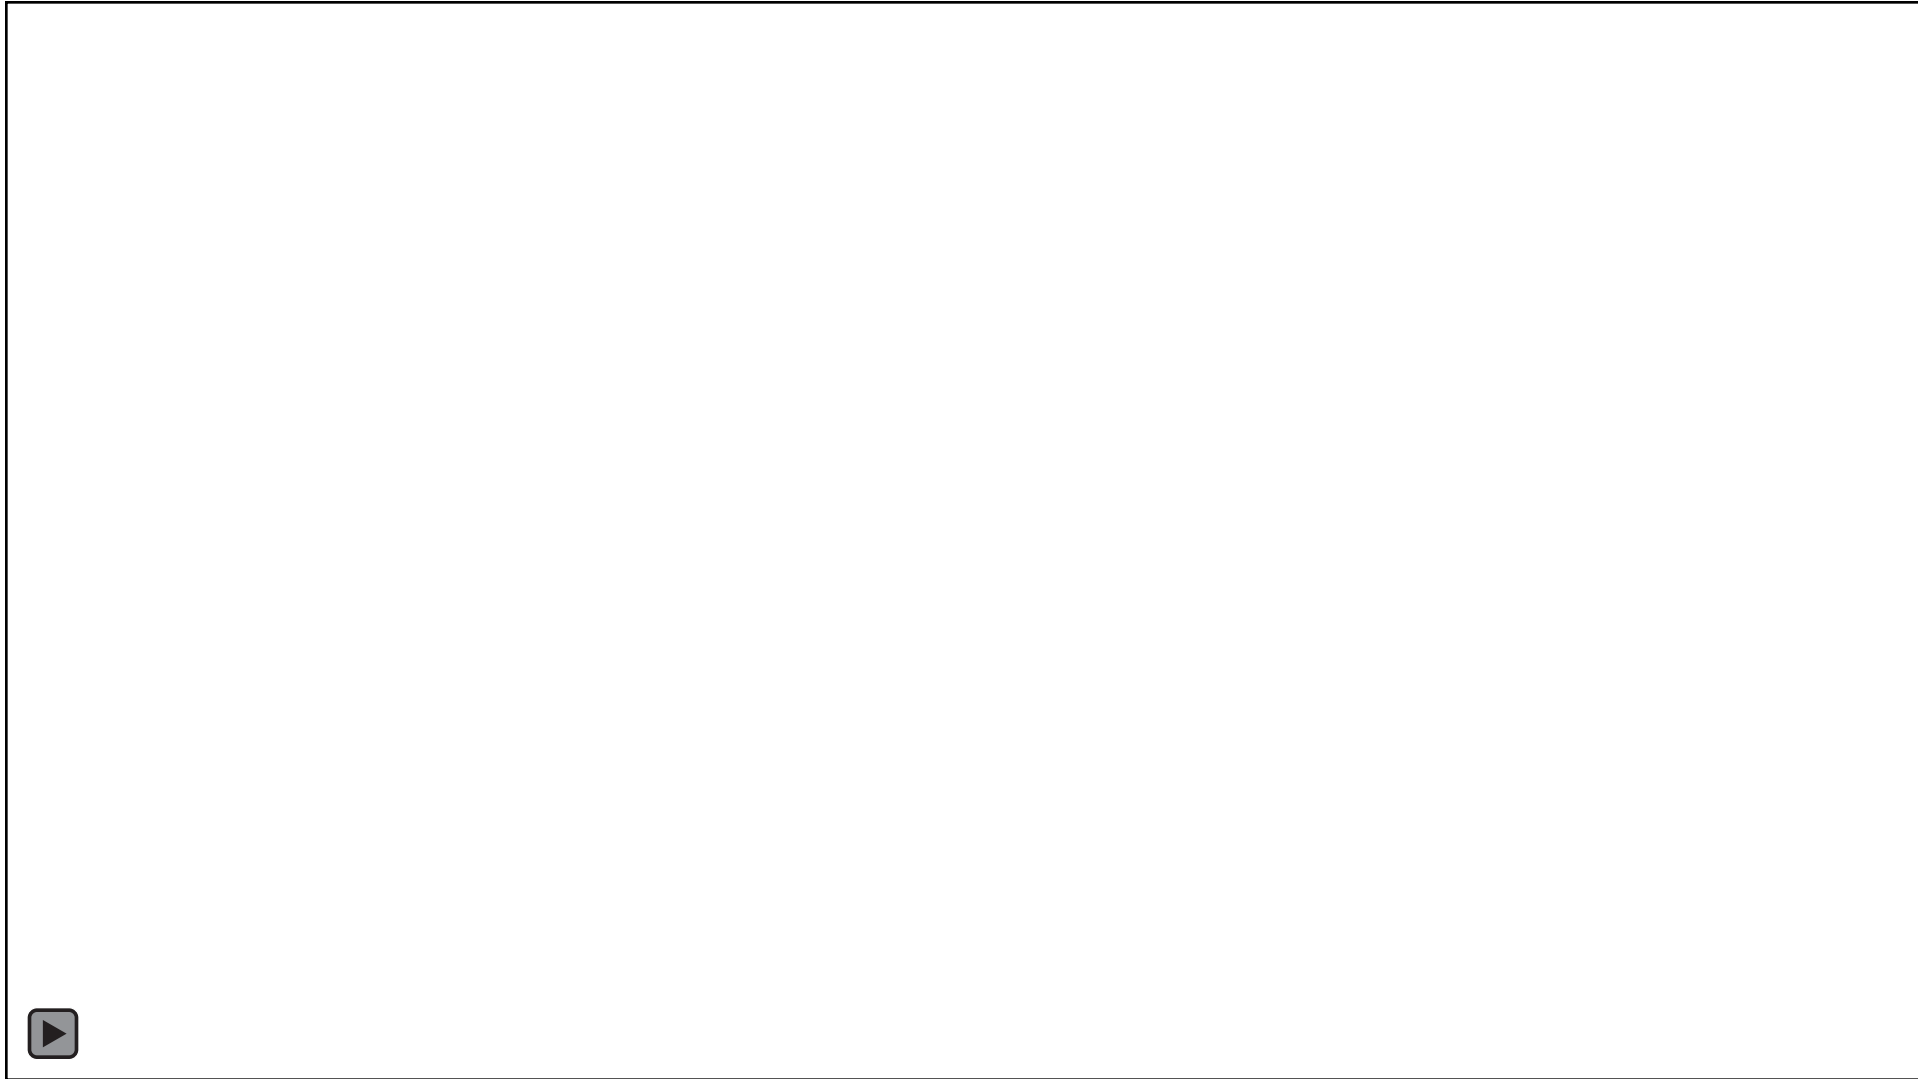
Unrestricted | © Siemens 2022 | Siemens Digital Industries Software | 2022 Functional Verification Study

SIEMENS

Let's have a look at an example: Lunar Crater Detection

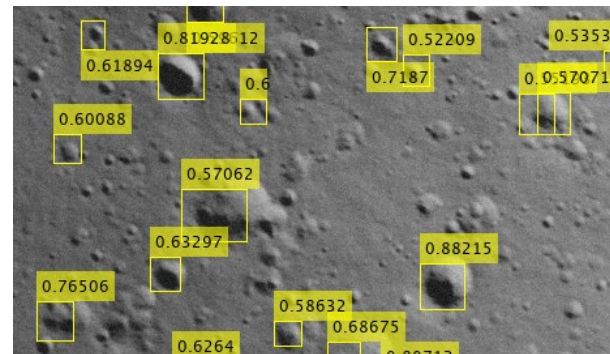
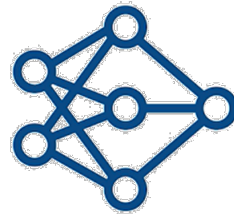
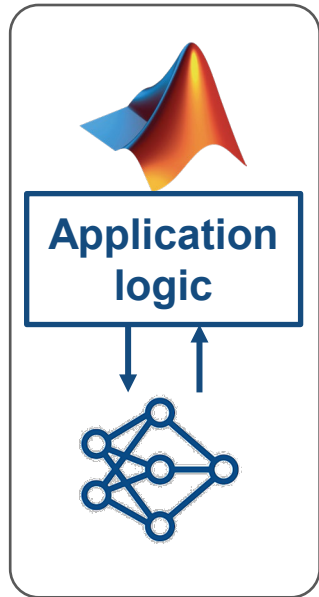


Object Detection with MATLAB

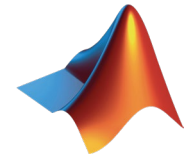


Images and video generated by PANGU and courtesy of the University of Dundee and STAR-Dundee.

Crater Detection Example



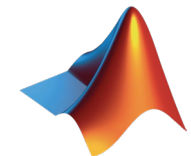
Pre-processing:
Extract regions and
resize



Inference: Predict
using trained network



Post-processing:
Annotate and label



Unique Platform Differentiators

MATLAB is an Interoperable and Integrated Platform for AI-Driven Systems

Data Preparation



Data cleansing and preparation



Human insight



Simulation-generated data

AI Modeling



Model design and tuning



Hardware accelerated training



Interoperability

Simulation & Test



Integration with complex systems



System simulation



System verification and validation

Deployment



Embedded devices



Enterprise systems



Edge, cloud, desktop

Spend less time preprocessing and labeling data

Synchronize disparate time series, filter noisy signals, automate labeling of video, and more.

Data Preparation



Data cleansing and preparation



Human insight




Simulation-generated data




Use labeling apps for deep learning workflows like semantic segmentation

Start with a complete set of algorithms and pre-built models


AI Modeling



Model design and tuning



Hardware accelerated training



Interoperability

Algorithms

Machine learning
Trees, Naïve Bayes, SVM...

Deep learning
CNNs, GANs, LSTM, MIMO...

Reinforcement learning
DQN, A2C, DDPG...

Regression
Linear, nonlinear, trees...

Unsupervised learning
K-means, PCA, GMM...

Predictive maintenance
RUL models, condition indicators...

Bayesian optimization

Pre-built models

Image classification models
AlexNet, GoogLeNet, VGG,
SqueezeNet, ShuffleNet, ResNet,
DenseNet, Inception...

Reference examples

Object detection
Vehicles, pedestrians, faces...

Semantic segmentation
Roadway detection, land cover
classification, tumor detection...

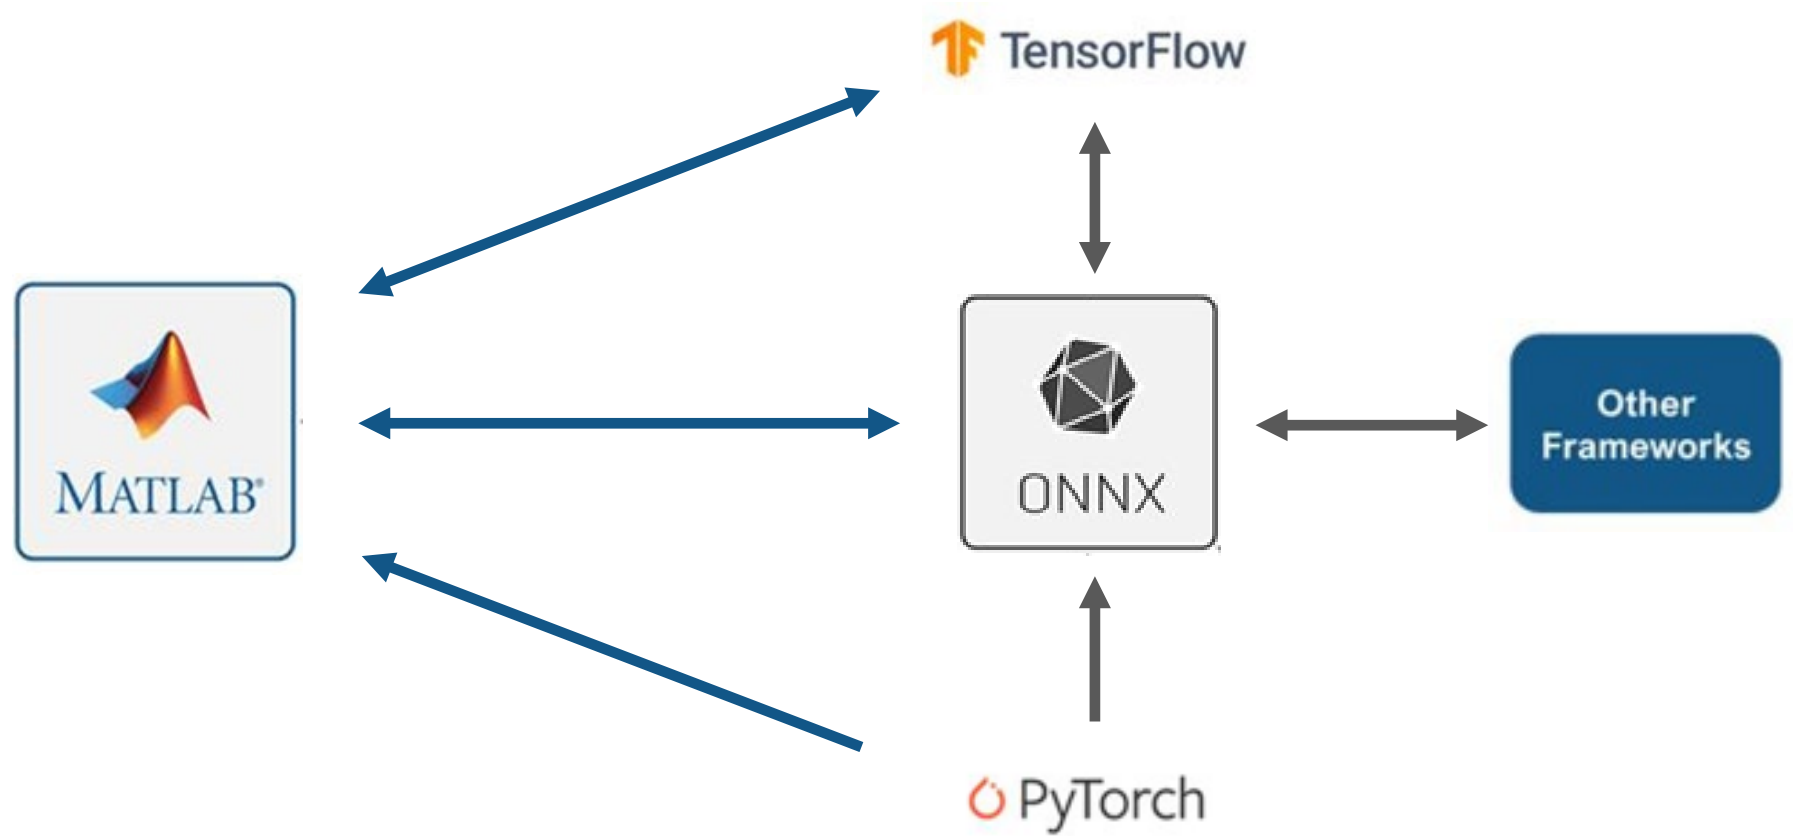
Signal and speech processing
Denoising, music genre recognition,
keyword spotting, radar waveform
classification...

...and more...

Interoperability with Python Based Frameworks

AI Modeling

- Model design and tuning
- Hardware accelerated training
- Interoperability



Increase productivity using Apps for design and analysis

Use MATLAB Apps to design deep learning networks, explore a wide range of classifiers, train regression models, train an optical character recognition model, and more.

AI Modeling



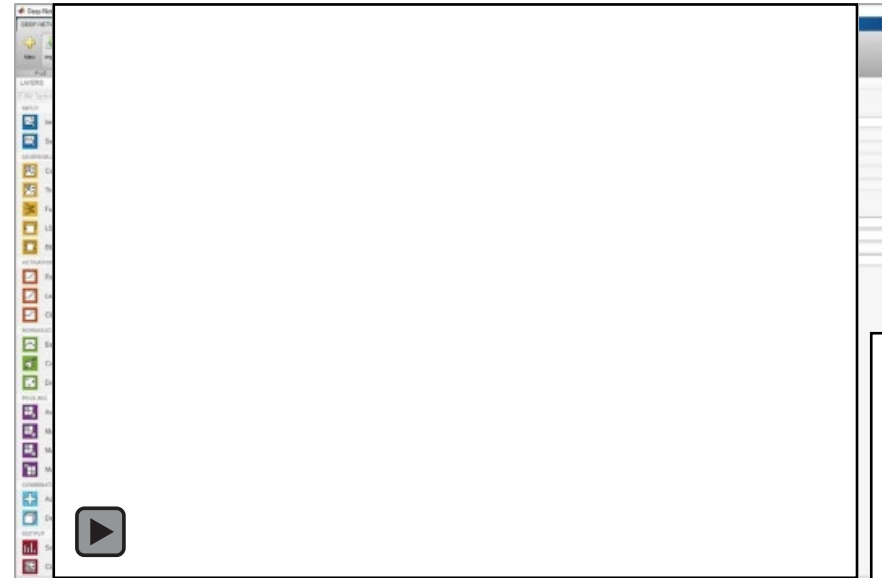
Model design and tuning



Hardware accelerated training



Interoperability



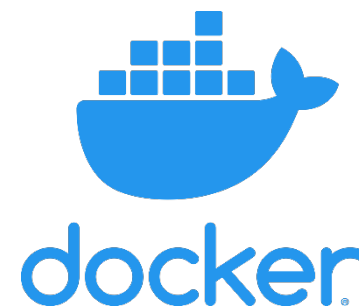
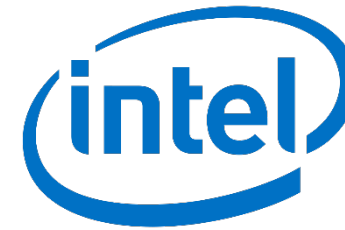
Deep Network Designer app to build, visualize, and edit deep learning networks






Experiment Manager app to manage multiple deep learning experiments, analyze and compare results and code

Hardware acceleration and scaling are critical for training

MATLAB accelerates AI training on GPUs, cloud, and datacenter resources without specialized programming.



AI Modeling

- 
Model design and tuning
- 
Hardware accelerated training
- 
Interoperability

FPGA is a good choice for lower power deep learning applications

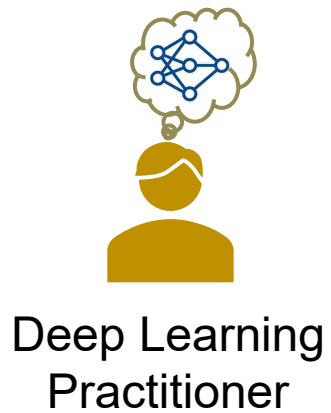
	GPU	ARM	FPGA	ASIC
Speed	High	Low	High	High
Power Consumption	High	Low	Low	Lowest
Engineering Cost	Medium	Low	Medium	High

Deployment

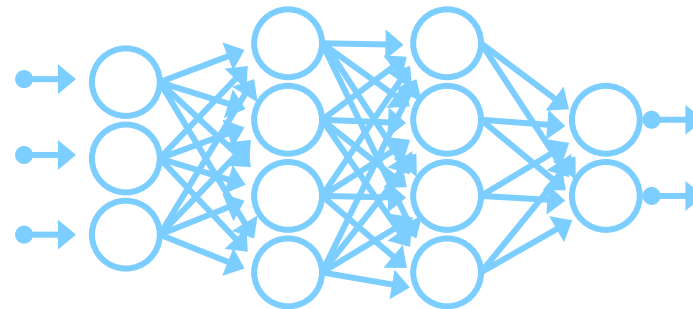
- Embedded devices
- Enterprise systems
- Edge, cloud, desktop

- Low Latency
- High speed I/O connectivity

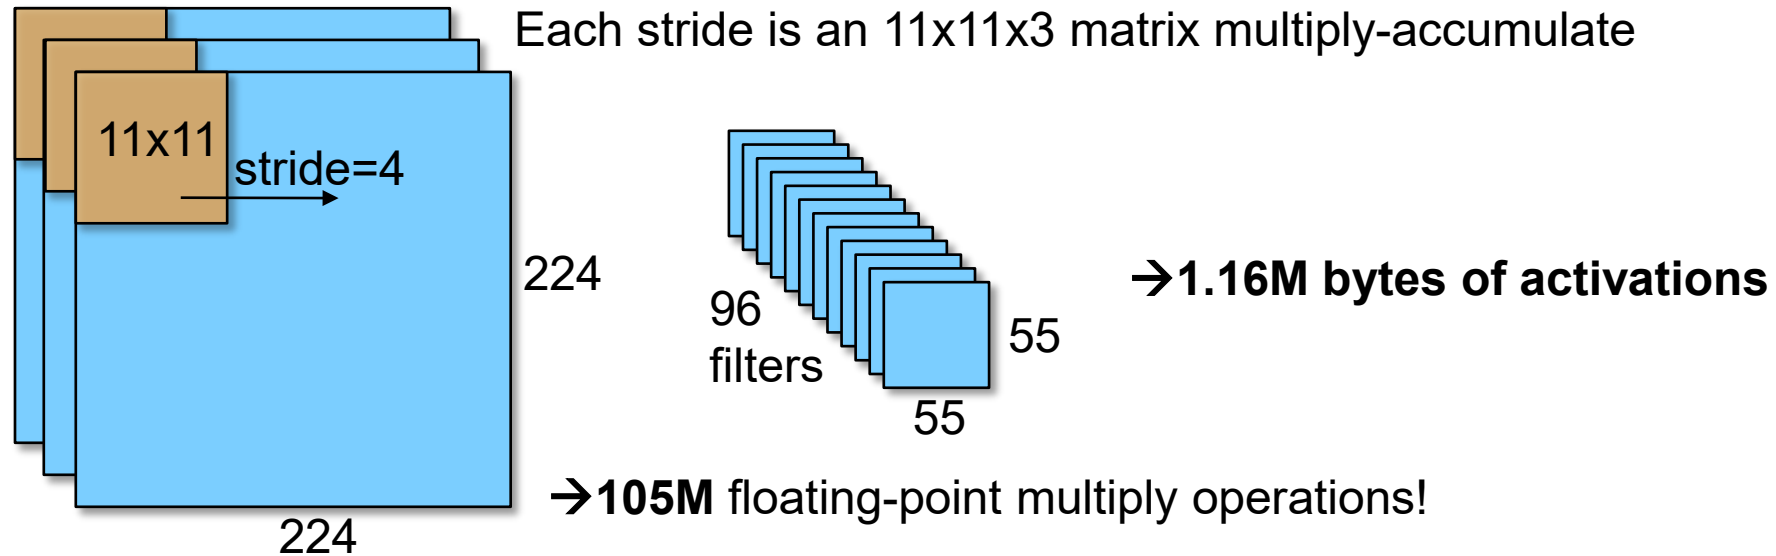
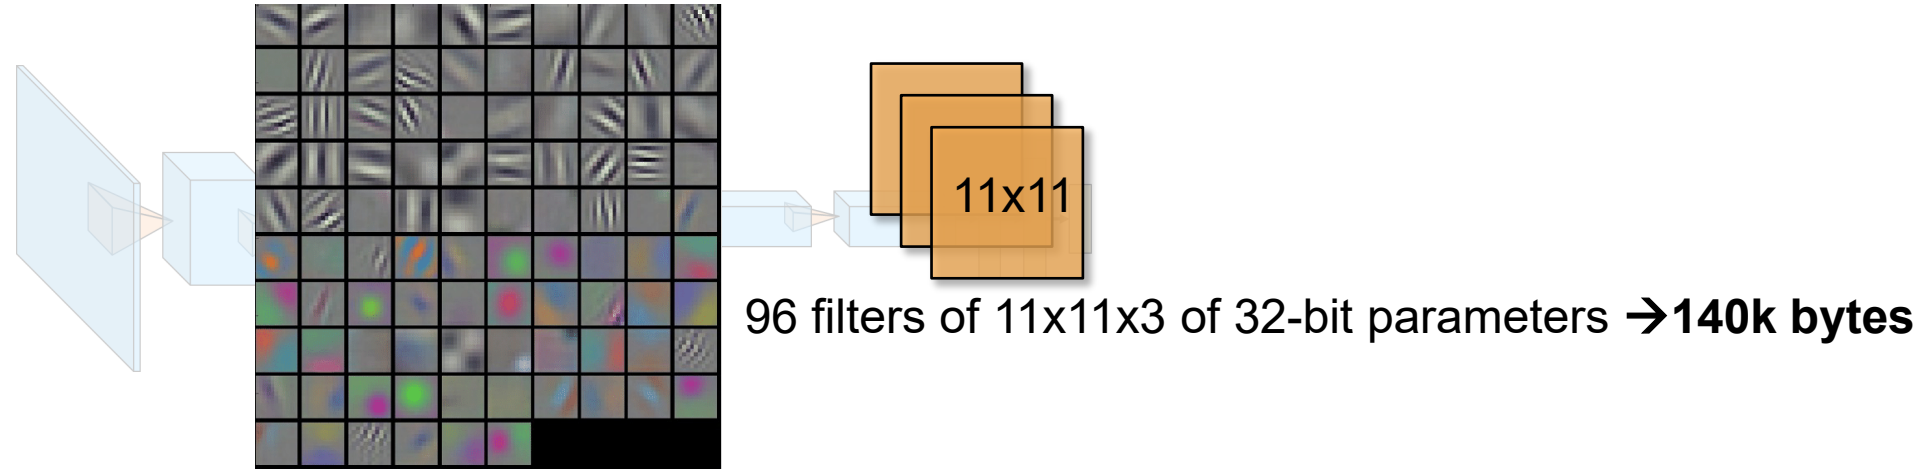
System Requirements Drive Network Design



Camera specs
Accuracy
Latency
Cost
Power

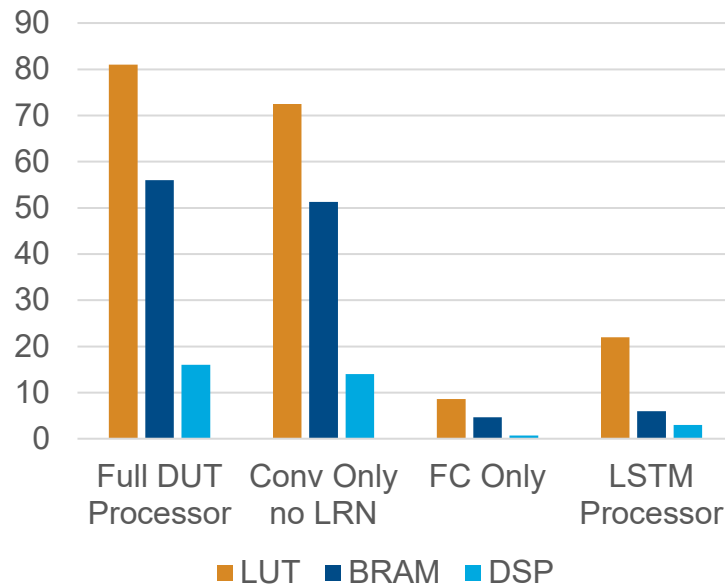


Challenges of Deploying Deep Learning to FPGA Hardware:

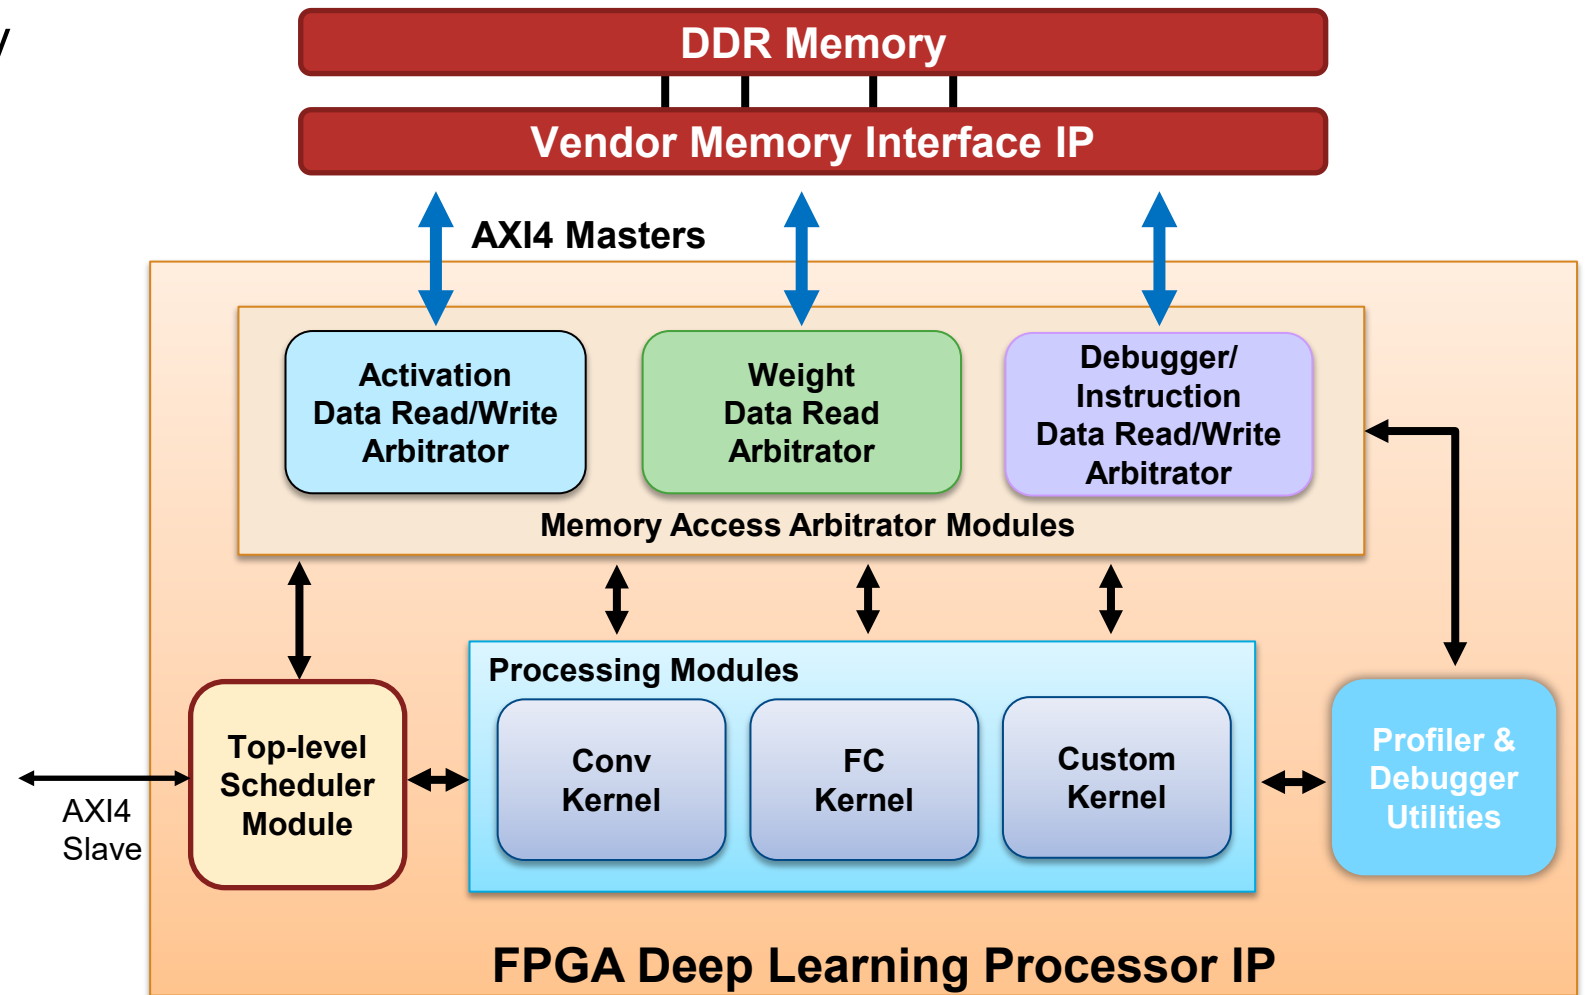


Customizable Deep Learning Processor

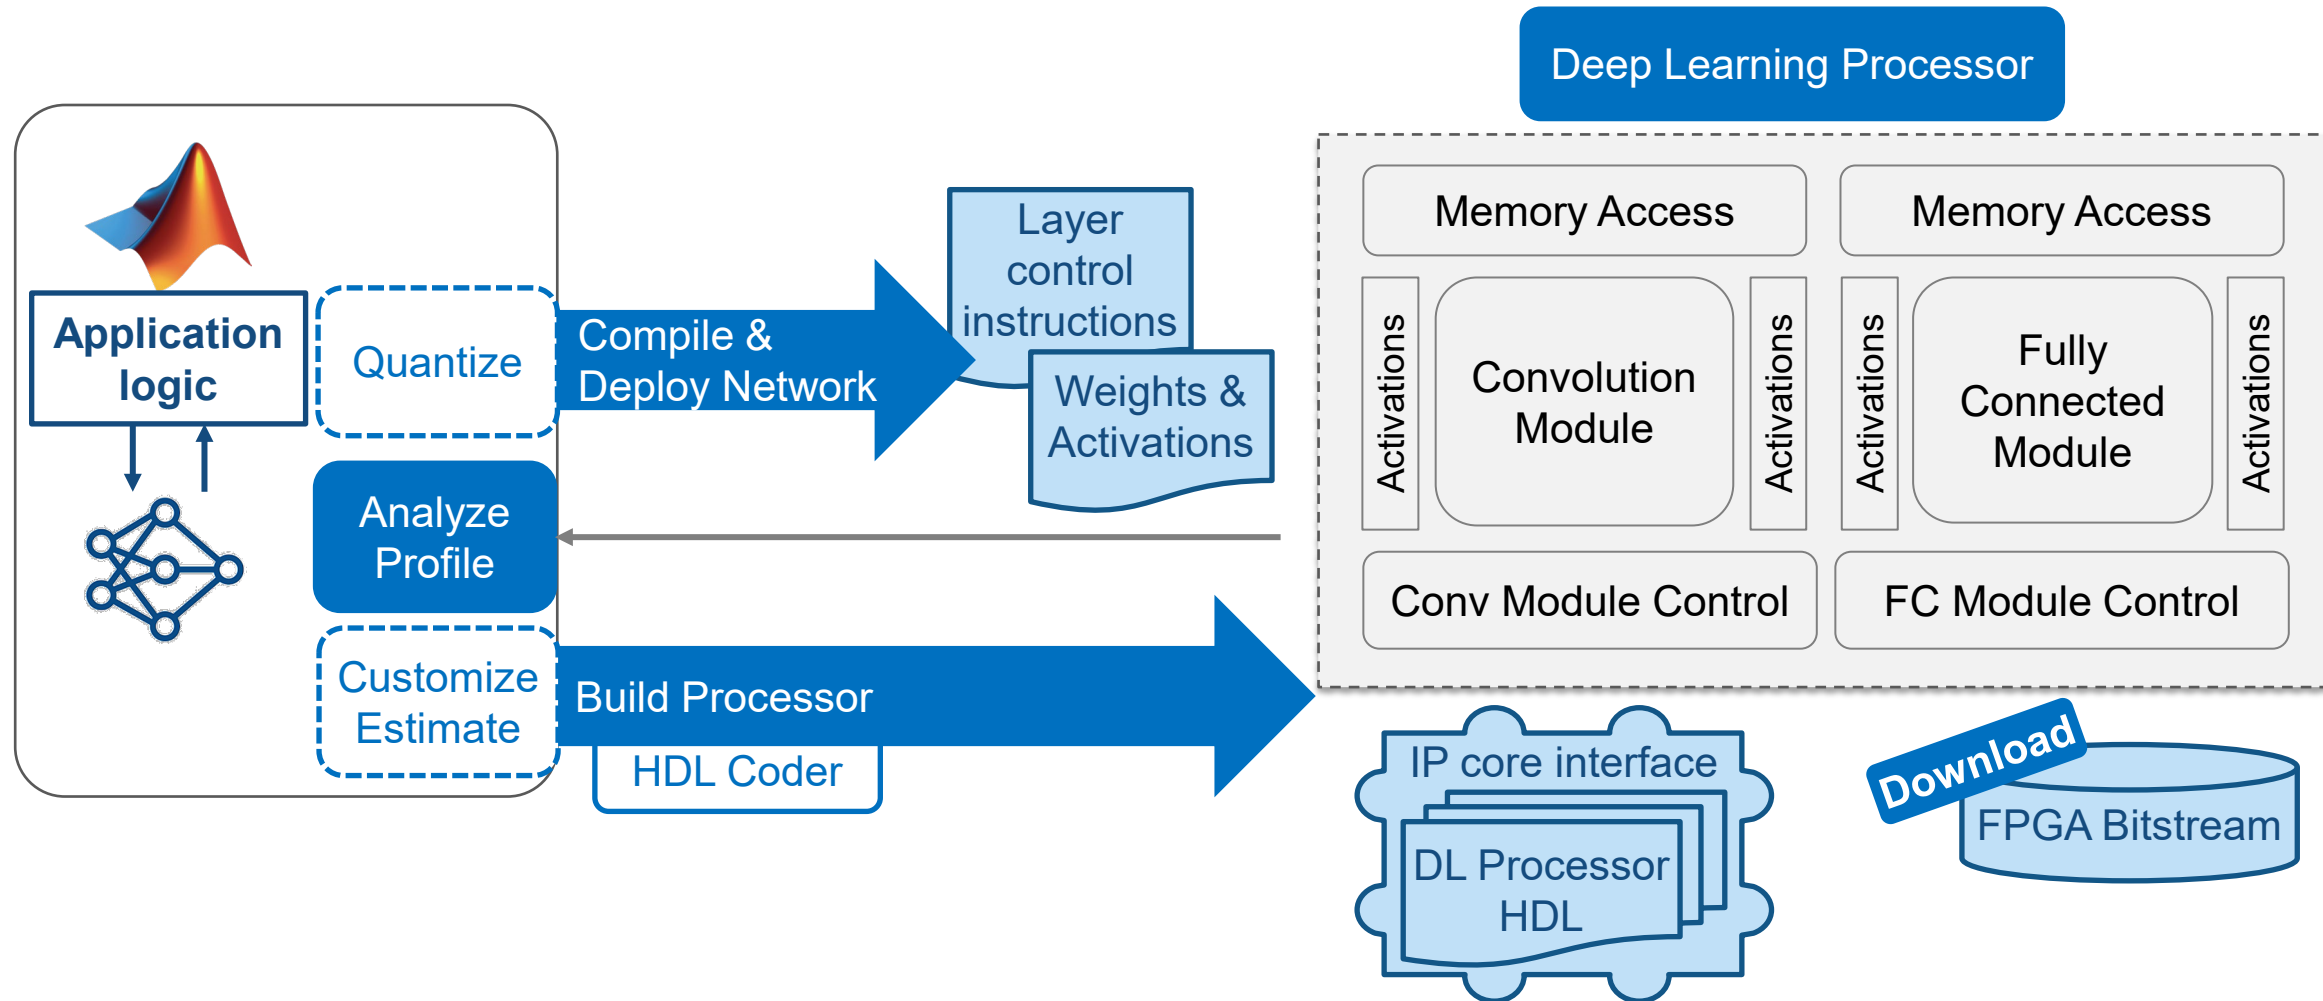
- Spend FPGA resource for only the layer kernels used in your network



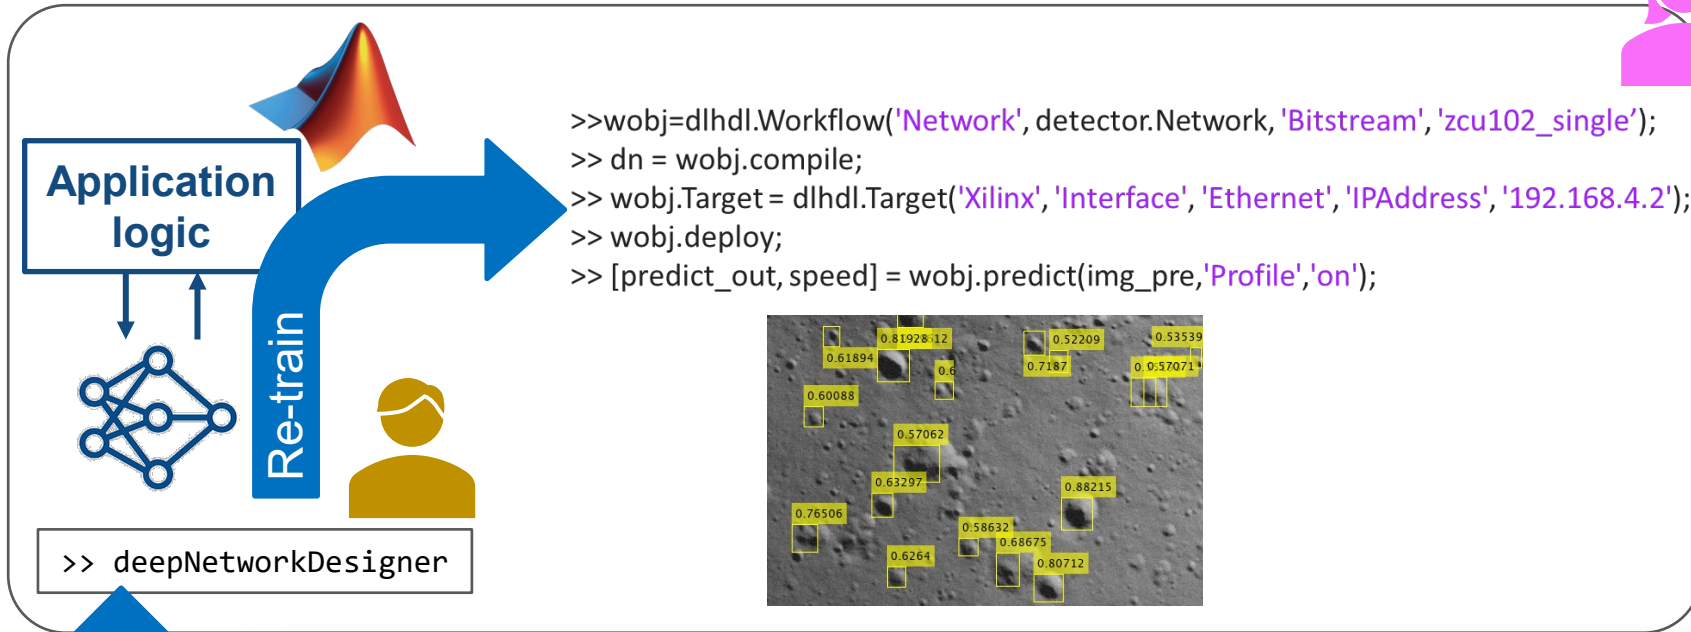
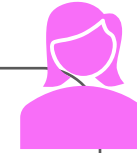
Percentage resource usage on ZCU102 board



Deep Learning HDL Processor steps

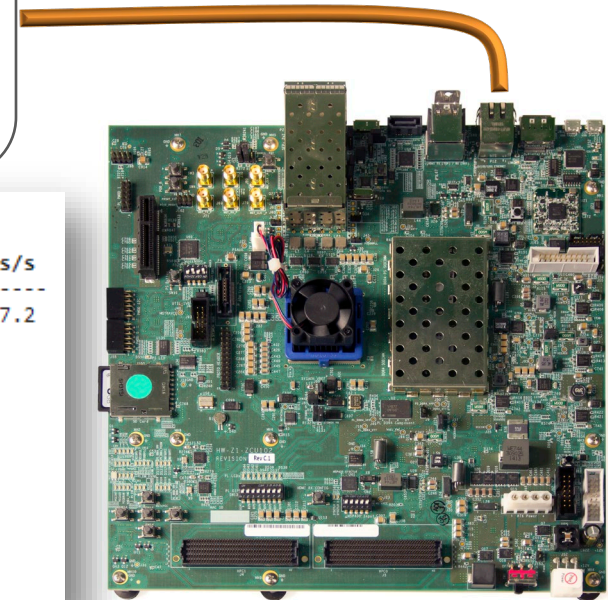


Profile FPGA Prototype and Iterate in MATLAB



Layer control instructions

Weights & Activations

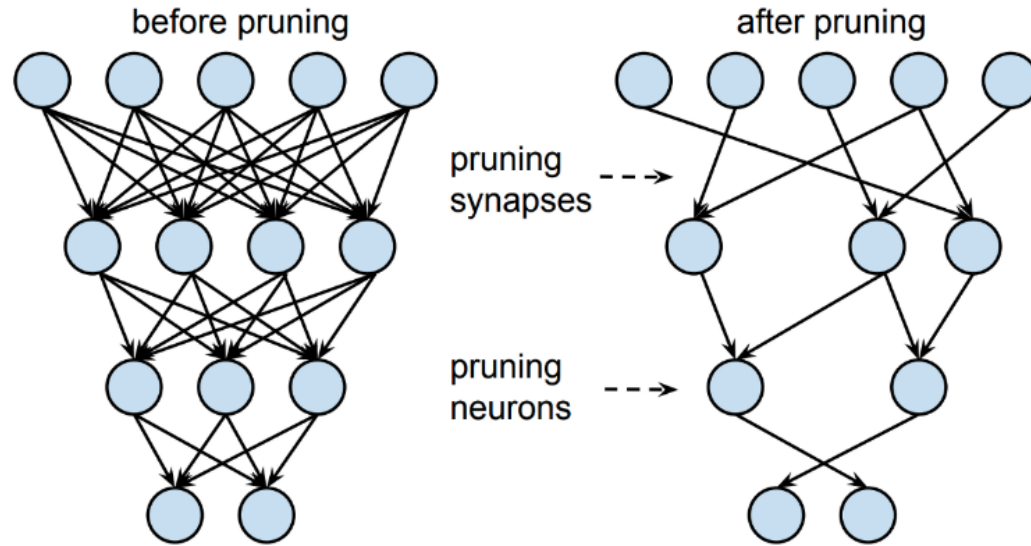


Deep Learning Processor Profiler Performance Results

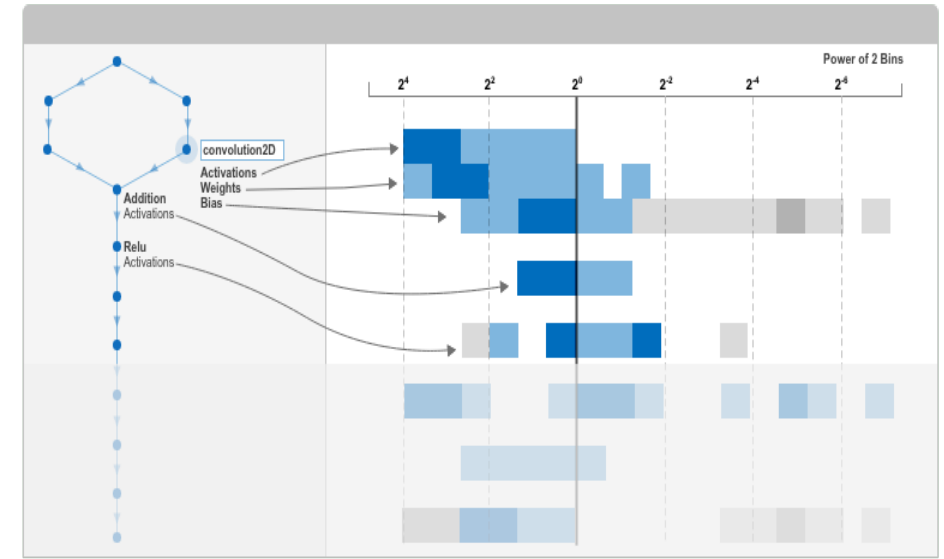
	LastFrameLatency(cycles)	LastFrameLatency(seconds)	FramesNum	Total Latency	Frames/s
Network	1729236	0.00786	1	1729753	127.2
conv_1	205293	0.00093			
maxpool1	161098	0.00073			
conv_2	212825	0.00097			
maxpool2	79776	0.00036			
conv_3	178397	0.00081			
maxpool3	44220	0.00020			
conv_4	161974	0.00074			
yolov2Conv1	305766	0.00139			
yolov2Conv2	306061	0.00139			
yolov2ClassConv	73795	0.00034			

* The clock frequency of the DL processor is: 220MHz

Two Compression Techniques

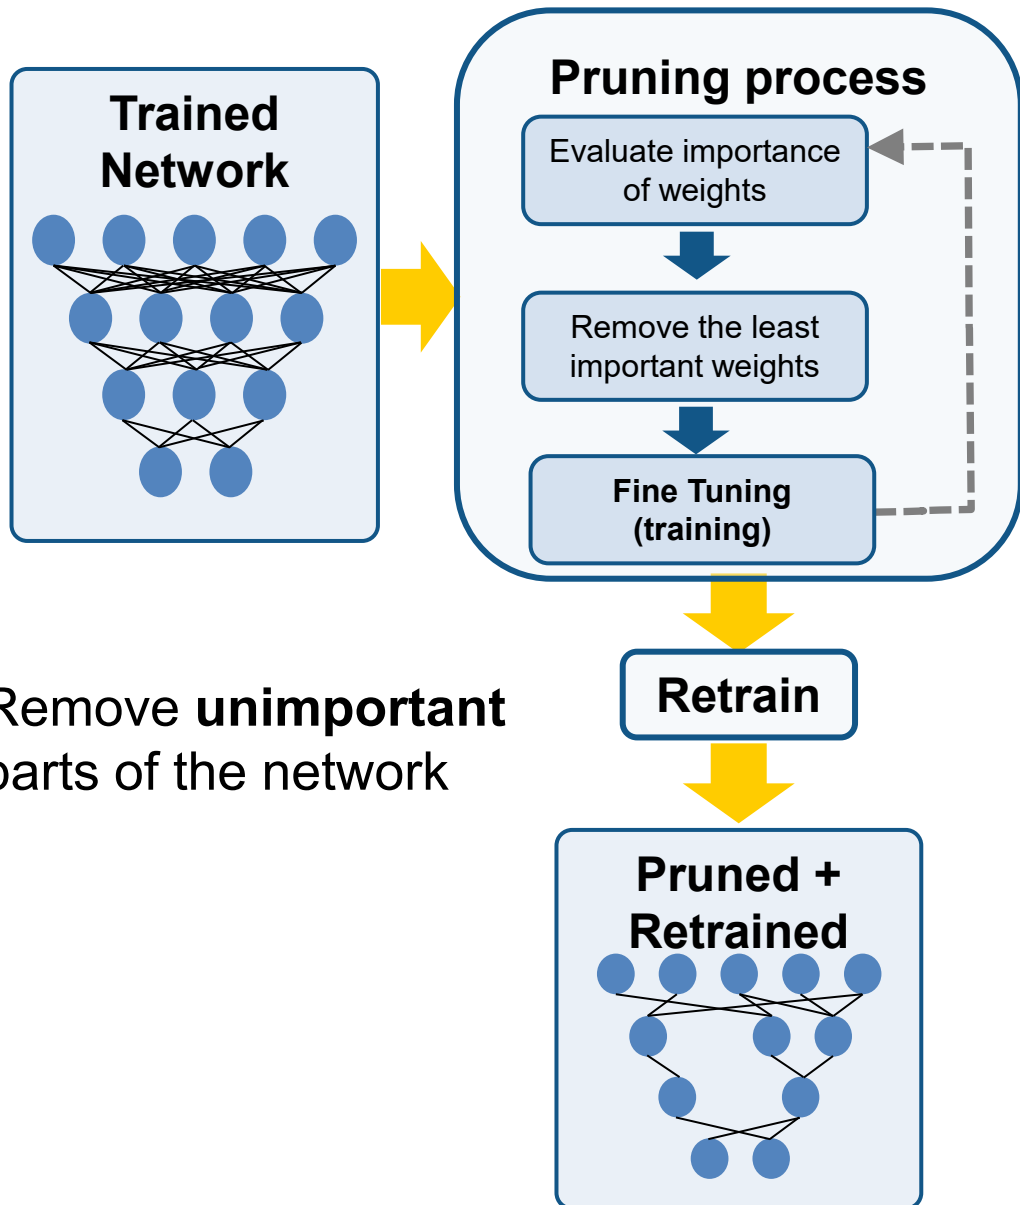


Pruning
deep neural networks



Quantization of
deep neural networks

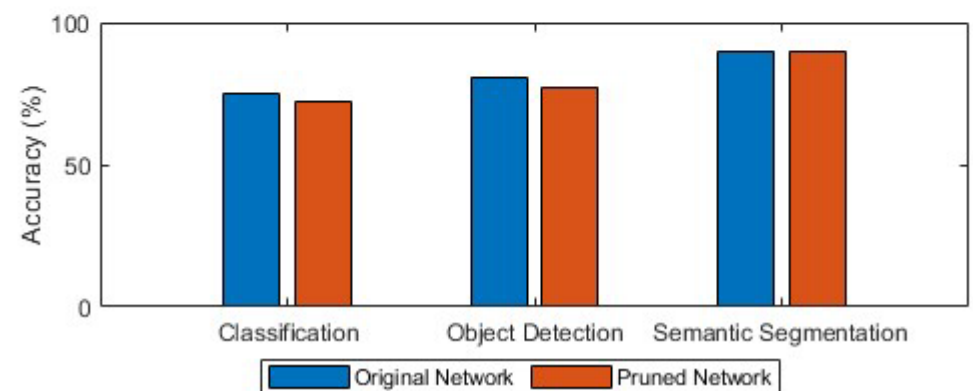
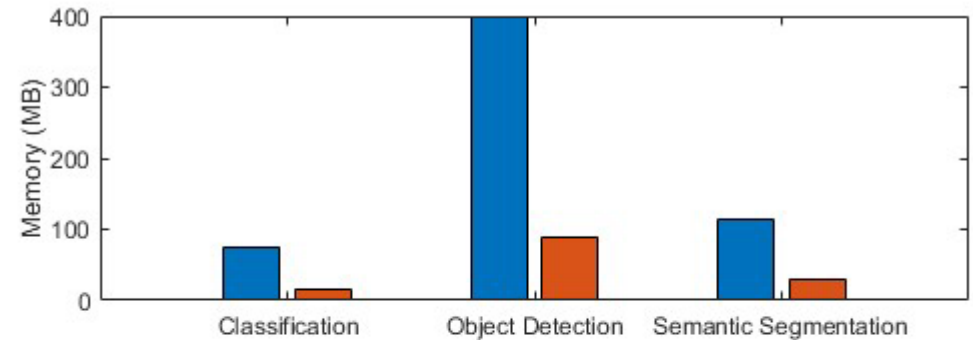
Taylor Approximation Pruning



Remove **unimportant** parts of the network

```
prunableNetwork = taylorPrunableNetwork(dlnet)
```

```
prunableNetwork = TaylorNetworkPruner with properties ...
```



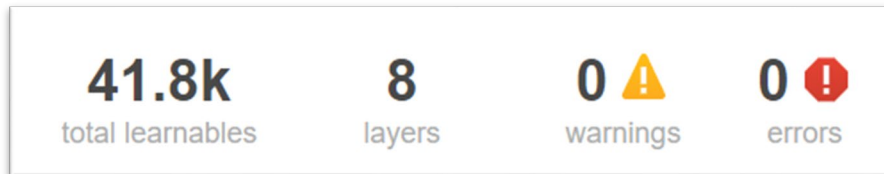
Recurrent Neural Networks (RNN) can be simplified by compressing LSTM Layers



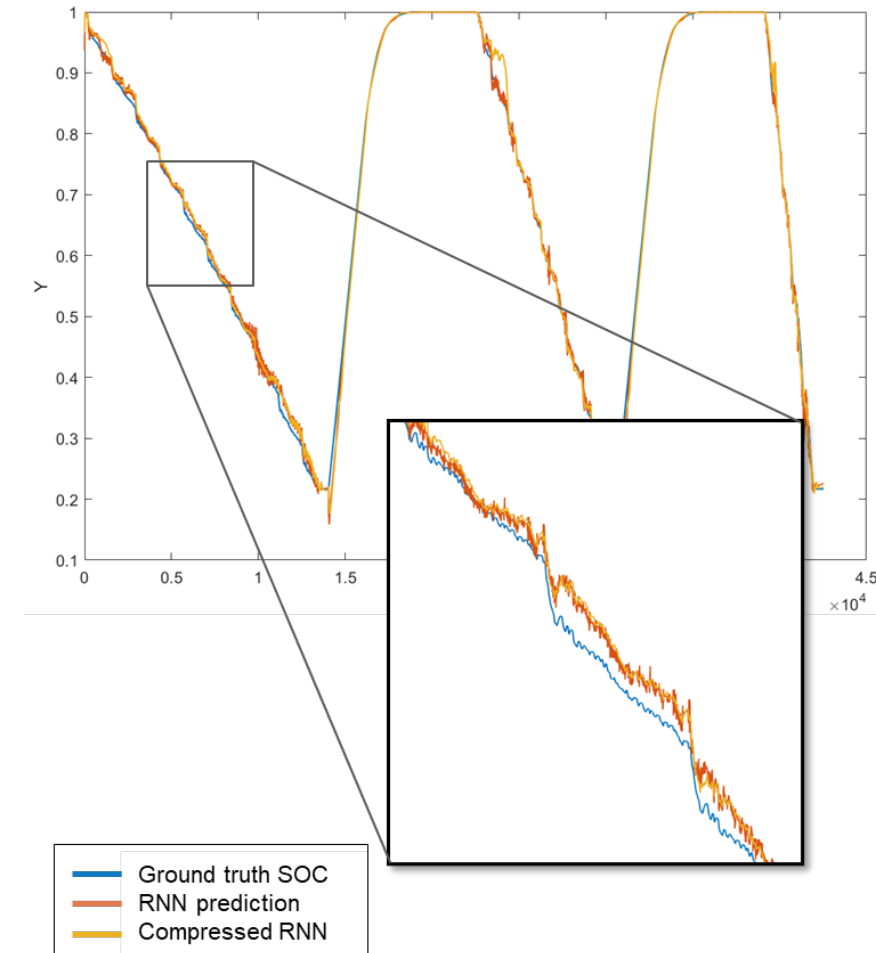
- **Data:** time series of temps/currents/voltages
→ predict **state of charge** at each time step
- **Model:** Recurrent regression NN with two LSTMs



- **New in R2022b - Projected Layer Pruning**



- Memory Footprint: **1.85 MB → 167 KB (91% reduction)**
- Inference Latency: **3.3 sec → 1.7 sec (48% reduction)**



Deep Network Quantizer - Int8 Quantization

The screenshot displays the Deep Network Quantizer interface with four numbered callouts:

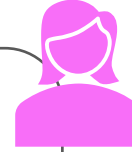
- 1 Import Network:** Points to the 'New' button in the top-left corner.
- 2 Calibrate:** Points to the 'Calibrate' button in the top toolbar.
- 3 Quantize and Validate:** Points to the 'Quantize and Validate' button in the top toolbar.
- 4 Export quantized network:** Points to the 'Export' button in the top toolbar.

The interface is divided into three main sections:

- Layer Graph:** A vertical list of network layers on the left, with 'conv_3' highlighted.
- Calibration Statistics Table:** A table in the center showing the minimum and maximum values for various layers. The 'conv_3' Bias row is highlighted in blue.
- Dynamic Range of Calibrated Layers:** A heatmap on the right showing the distribution of values for each layer across different quantization ranges. The title 'Dynamic Range of Calibrated Layers' is circled in orange.

Layer Name	Min Value	Max Value	Quantization Status
input			
Activations	0.0000	1.0000	
conv_1			<input checked="" type="checkbox"/>
Weights	-1.7574	1.7767	
Bias	-2.5100	2.5245	
Activations	-8.3809	6.8080	
relu_1			<input checked="" type="checkbox"/>
Activations	0.0000	6.8080	
maxpool1			<input checked="" type="checkbox"/>
Activations	0.0000	6.8080	
conv_2			<input checked="" type="checkbox"/>
Weights	-0.2267	0.2915	
Bias	-1.5915	1.7443	
Activations	-6.6541	8.7700	
relu_2			<input checked="" type="checkbox"/>
Activations	0.0000	8.7700	
maxpool2			<input checked="" type="checkbox"/>
Activations	0.0000	8.7700	
conv_3			<input checked="" type="checkbox"/>
Weights	-0.2329	0.1731	
Bias	-1.8097	1.9424	
Activations	-5.8875	6.7268	
relu_3			<input checked="" type="checkbox"/>
Activations	0.0000	6.7268	
maxpool3			<input checked="" type="checkbox"/>

Quantize Deep Learning Network and Processor in MATLAB



Application logic

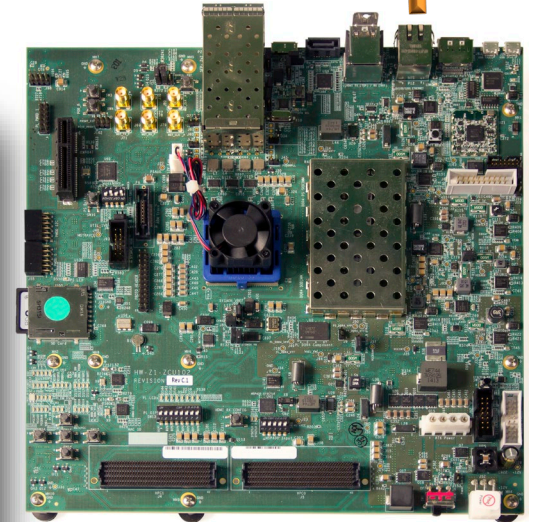
```

>> wobj=dlhdl.Workflow('Network', detector.Network, 'Bitstream','zcu102_int8');
>> dn = wobj.compile;
>> wobj.Target = dlhdl.Target('Xilinx', 'Interface', 'Ethernet', 'IPAddress', '192.168.4.2');
>> wobj.deploy;
>> [predict_out, speed] = wobj.predict(img_pre, 'Profile', 'on');
                    
```

>> deepNetworkQuantizer

Layer control instructions

Weights & Activations

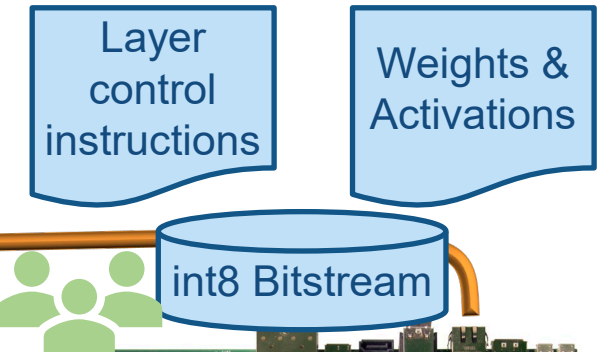
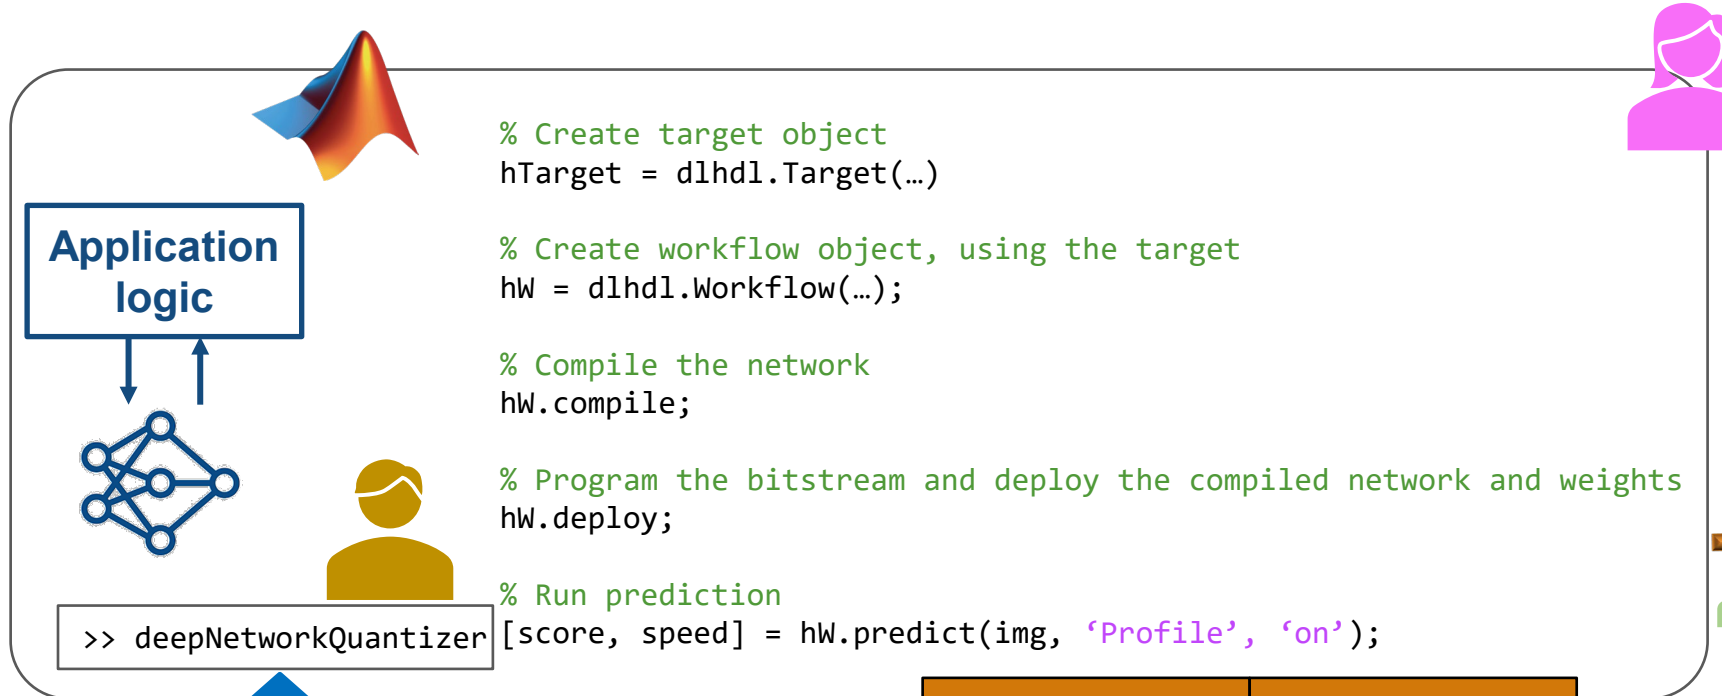


Deep Learning Processor Profiler Performance Results

Network	LastFrameLatency(cycles)	LastFrameLatency(seconds)	FramesNum	Total Latency	Frames/s
Network	575217	0.00230	1	575799	434.2
conv_1	99836	0.00040			
maxpool1	65019	0.00026			
conv_2	66616	0.00027			
maxpool2	31568	0.00013			
conv_3	53503	0.00021			
maxpool3	18154	0.00007			
conv_4	46168	0.00018			
yolov2Conv1	84962	0.00034			
yolov2Conv2	85134	0.00034			
yolov2ClassConv	24226	0.00010			

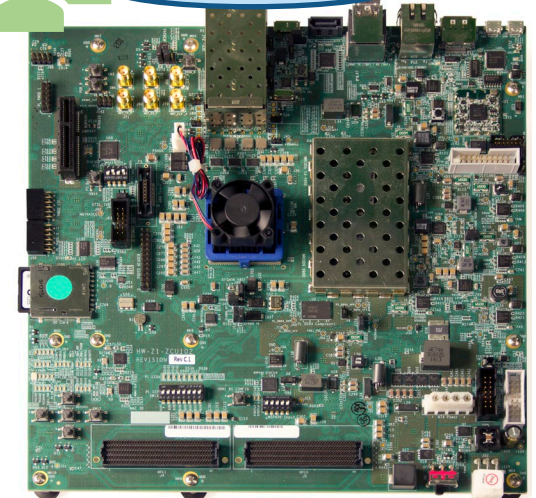
* The clock frequency of the DL processor is: 250MHz

Converge on an FPGA-Optimized Deep Learning Network



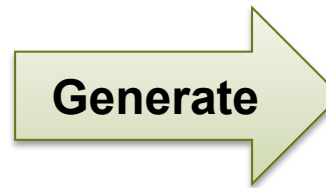
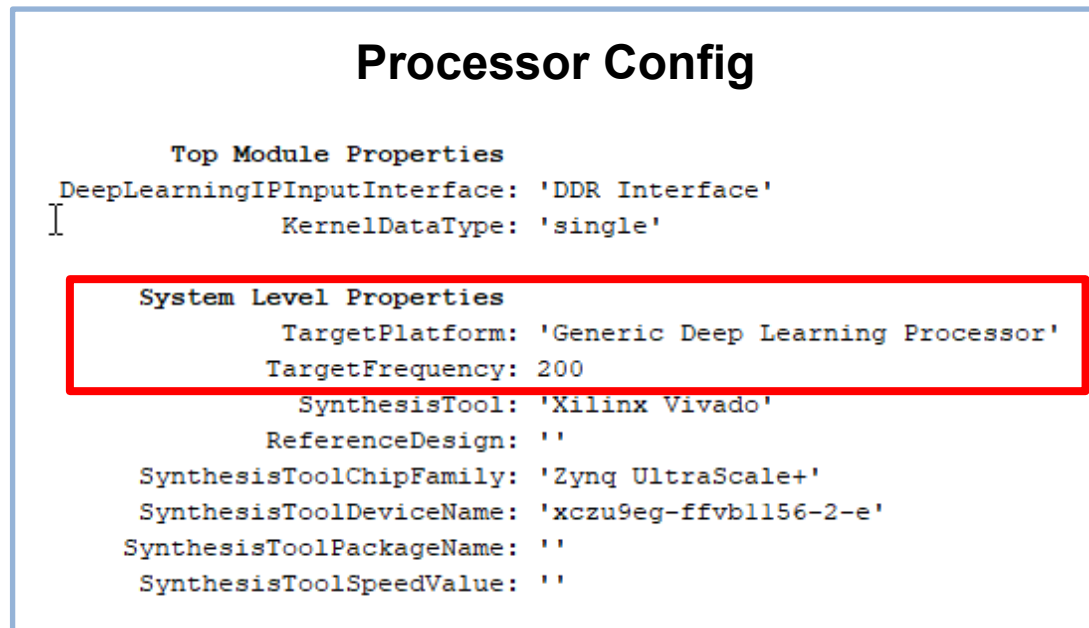
Parameters	Speed
48 MB	127.2 fps
44 MB	433.8 fps

Bitstream Name	ConvThreadNumber	FCThreadNumber	Lookup Table(LUT) Utilization(%)	Block RAM (BRAM) Utilization(%)	DSP Utilization (%)
zcu102_single	16	4	90	63.7	15
zcu102_int8	64	16	62	49	31

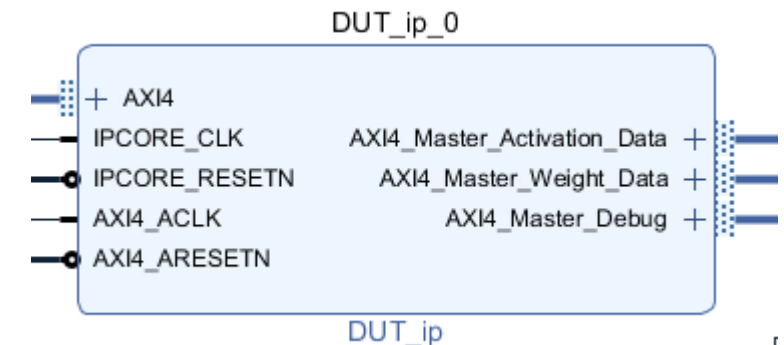


Integrate the Deep Learning Processor into your bigger system

- Generate Generic Deep Learning Processor IP core
- Define clean input/output frame hand-shaking protocol
- Drop the generated Deep Learning IP core into your bigger system



```
>> dlhdl.buildProcessor(hPC)
```

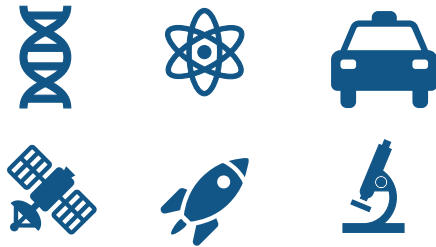


Network Examples

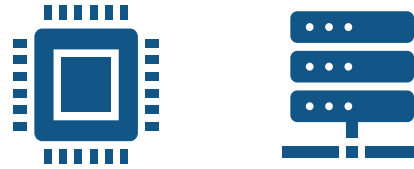
Network Examples	Application Area	Type	Release
VGG16/VGG19	Classification	CNN	R2021b
Darknet19	Classification	CNN	
ResNet18/ResNet50	Classification/Detection	CNN	
YOLO v2	Object detection	CNN	
MobileNet v2	Classification/Detection	CNN	
1-Dimensional CNN networks	Classification/Detection	CNN	R2022a
Segmentation networks	Segmentation	CNN	R2022b
LSTM networks	Signal processing	RNN	
YOLO v3	Object detection	CNN, MIMO	
GRU network	Signal processing	RNN	R2023a
YAMnet (Audio toolbox)	Classification/Detection	CNN	

Why MATLAB & MathWorks for AI?

Domain-specialized workflows for engineering and science



Multi-platform deployment of full applications and systems



SIMULINK[®]



Platform productivity



Interoperability with Python and DL Python-based frameworks



People





ANY
QUESTIONS
?

Examples

Deep Learning HDL Toolbox

- Get Started with Deep Learning HDL Toolbox 5
- Prototype Deep Learning Networks on FPGA 14
- Deep Learning Processor Customization and IP Generation 5
- System Integration of Deep Learning Processor IP Core 3
- Deep Learning INT8 Quantization 5

Networks on FPGA



Deploy Transfer Learning Network for Lane Detection

Create, compile, and deploy a dihdl.Workflow object that has a convolutional neural network. The network can detect and output lane

[Open Live Script](#)

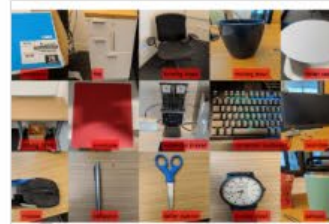


Image Category Classification by Using Deep Learning

Create, compile, and deploy a dihdl.Workflow object with alexnet as the network object by using the Deep Learning HDL Toolbox™

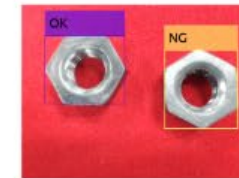
[Open Live Script](#)



Image Classification Using DAG Network Deployed to FPGA

Train, compile, and deploy a dihdl.Workflow object that has ResNet-18 as the network object by using the Deep Learning HDL

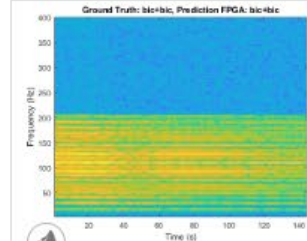
[Open Live Script](#)



Defect Detection

Deploy a custom trained series network to detect defects in objects such as hexagon nuts. The custom networks were trained by using

[Open Live Script](#)



Bicyclist and Pedestrian Classification by Using FPGA

Deploy a custom trained series network to detect pedestrians and bicyclists based on their micro-Doppler signatures. This network is

[Open Live Script](#)



Visualize Activations of a Deep Learning Network by Using LogoNet

Feed an image to a convolutional neural network and display the activations of the different layers of the network. Examine the activations

[Open Live Script](#)



Running Convolution-Only Networks by Using FPGA Deployment

Typical series classification networks include a sequence of convolution layers followed by one or more fully connected layers.

[Open Live Script](#)



Vehicle Detection Using YOLO v2 Deployed to FPGA

Deep learning is a powerful machine learning technique that you can use to train robust object detectors. Several techniques for object

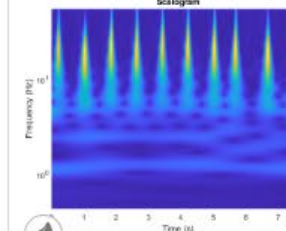
[Open Live Script](#)



Vehicle Detection Using DAG Network Based YOLO v2 Deployed to FPGA

Train and deploy a you look only once (YOLO) v2 object detector.

[Open Live Script](#)



Classify ECG Signals Using DAG Network Deployed To FPGA

Classify human electrocardiogram (ECG) signals by deploying a trained directed acyclic graph (DAG) network.

[Open Live Script](#)



Prototype and Verify Deep Learning Networks Without Target Hardware

Rapidly prototype your custom deep learning network and bitstream by visualizing intermediate layer activation results and verifying

[Open Live Script](#)

Training Resources



Machine Learning Onramp

6 modules | 2 hours | Languages

Learn the basics of practical machine learning methods for classification problems.

Free



Machine Learning with MATLAB

7 modules | 12 hours | Languages

Explore data and build predictive models.



Deep Learning Onramp

5 modules | 2 hours | Languages

Get started quickly using deep learning methods to perform image recognition.

Free



Deep Learning with MATLAB

13 modules | 8 hours | Languages

Learn the theory and practice of building deep neural networks with real-life image and sequence data.



Reinforcement Learning Onramp

5 modules | 3 hours | Languages

Master the basics of creating intelligent controllers that learn from experience.

Free

<https://matlabacademy.mathworks.com/>

Deep Learning Onramp

0%

Start course

Share | Certificate | Settings

Course Description

Get started quickly using deep learning methods to perform image recognition.



Course Author
Renee Bach

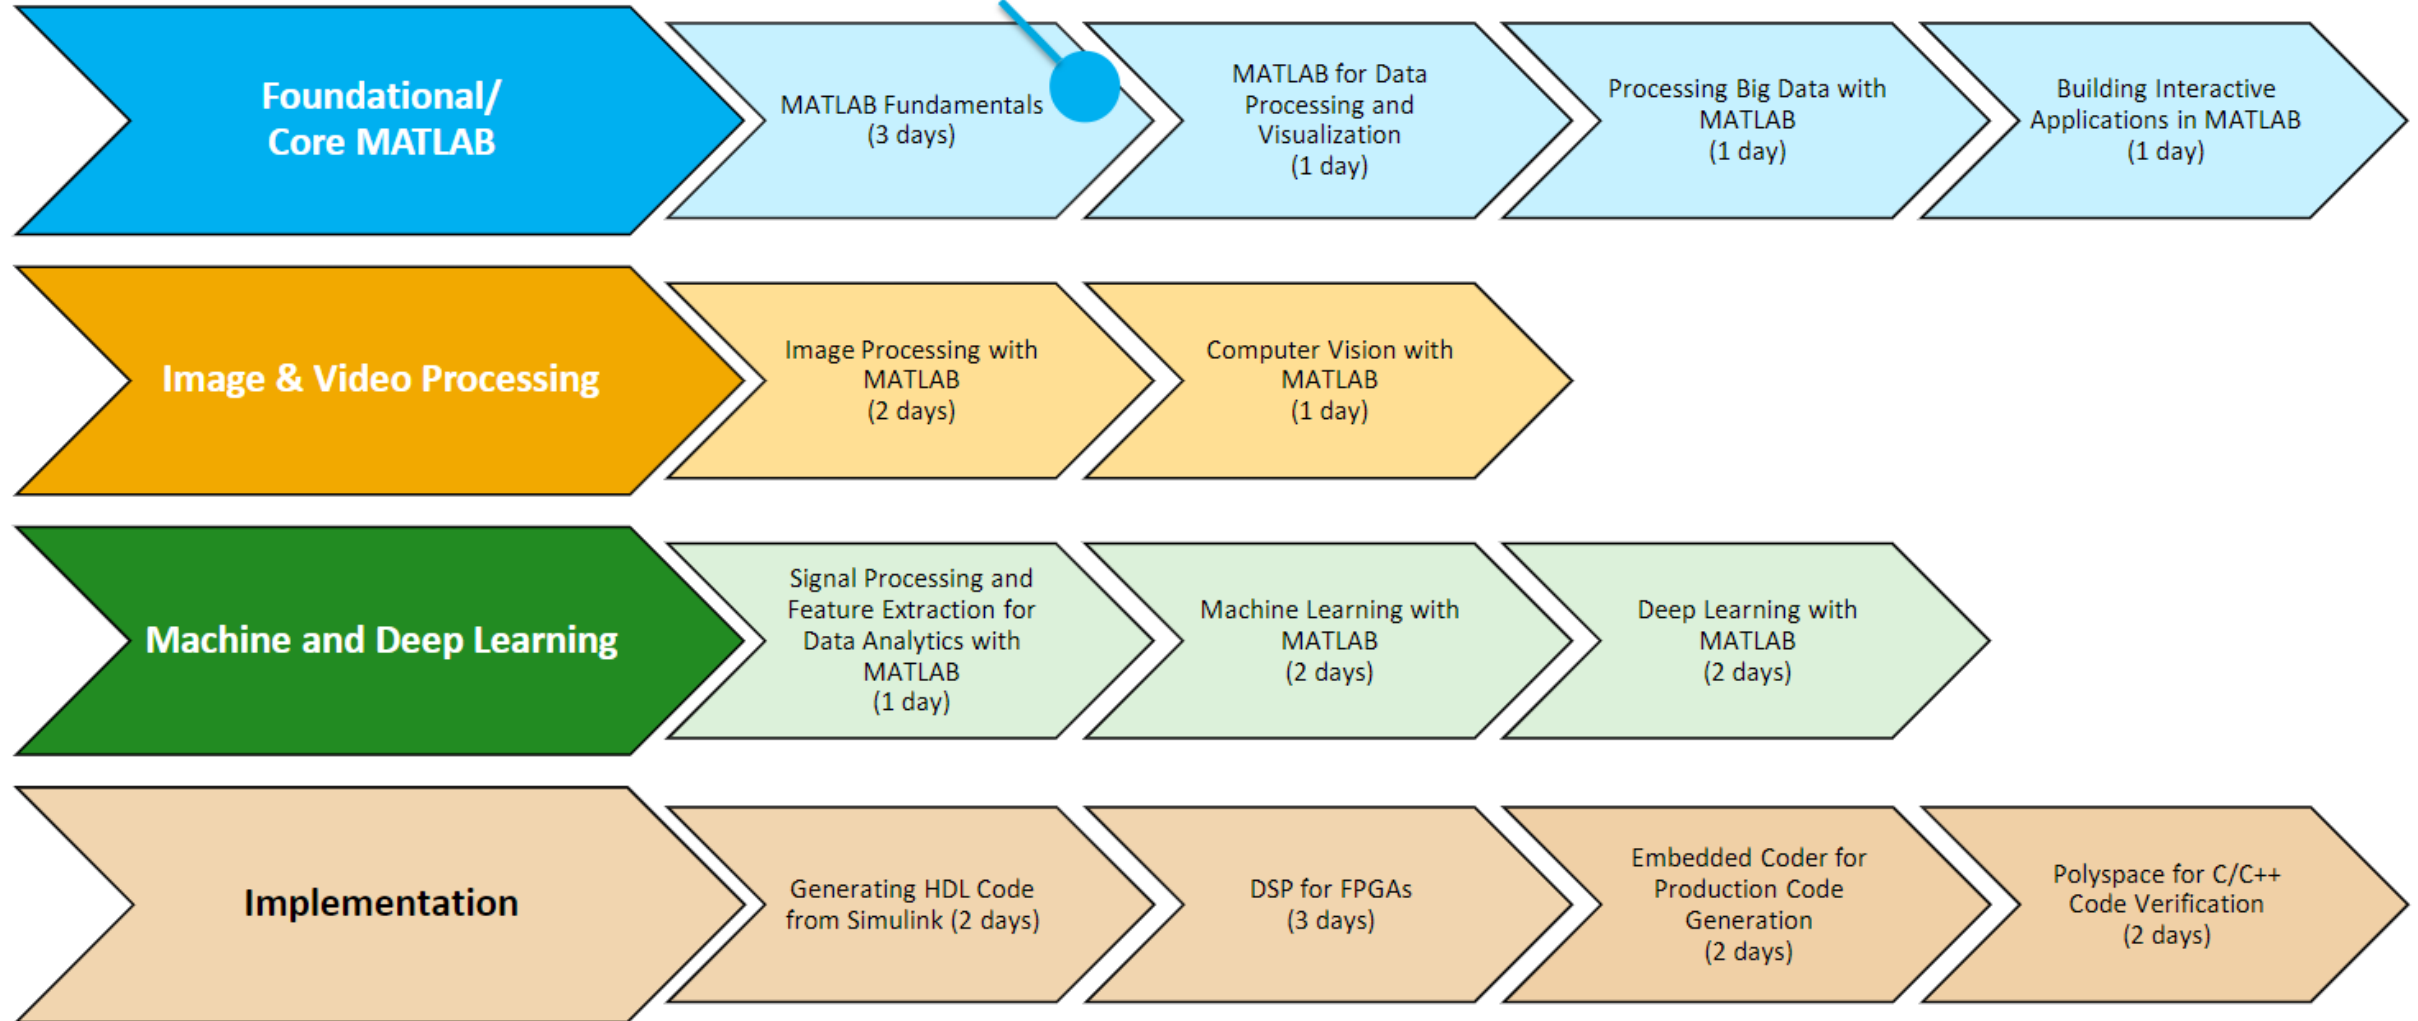
Format Self-paced online
Duration 2 hours
Language English (set language)

Modules

- > Introduction 5 min
- > Using Pretrained Networks 20 min
- > Managing Collections of Image Data 30 min
- > Performing Transfer Learning 60 min
- > Conclusion 10 min

MathWorks training options for AI topics

MATLAB Skills Assessment

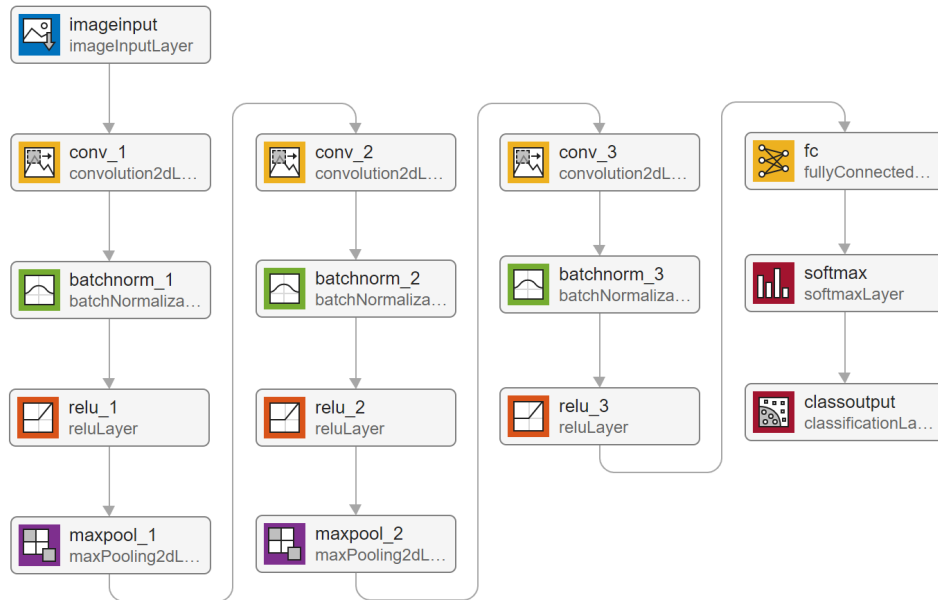


Resources for Further Learning

- Crater Detection - Deep Learning
 - [Deep Learning Solutions in MATLAB](#)
 - [Deep Learning Verification Library](#)
 - [Deep Learning Models](#)
 - [MATLAB with TensorFlow and PyTorch](#)
 - [Importing Models from TensorFlow, PyTorch, and ONNX](#)
 - [TensorFlow-Keras Layers Supported for Conversion into Built-In MATLAB Layers](#)
 - [What's New in Interoperability with TensorFlow and PyTorch](#)
- Crater Detection - Deep Learning → FPGA
 - [Deep Learning HDL Toolbox](#)
 - [Deep Learning HDL Toolbox Supported Networks, Layers, Boards and Tools](#)

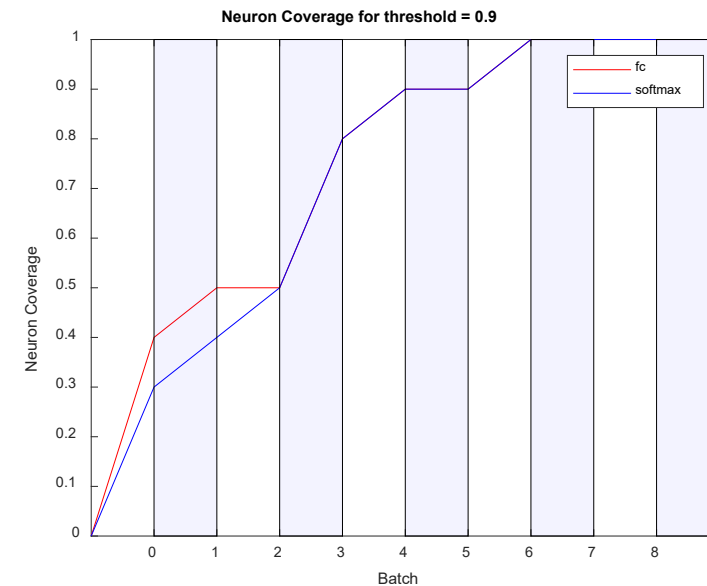
Neuron Coverage for Deep Learning

<https://github.com/matlab-deep-learning/neuron-coverage-for-deep-learning>



		LayerCoverage
1	fc	0.2000
2	softmax	0.1000

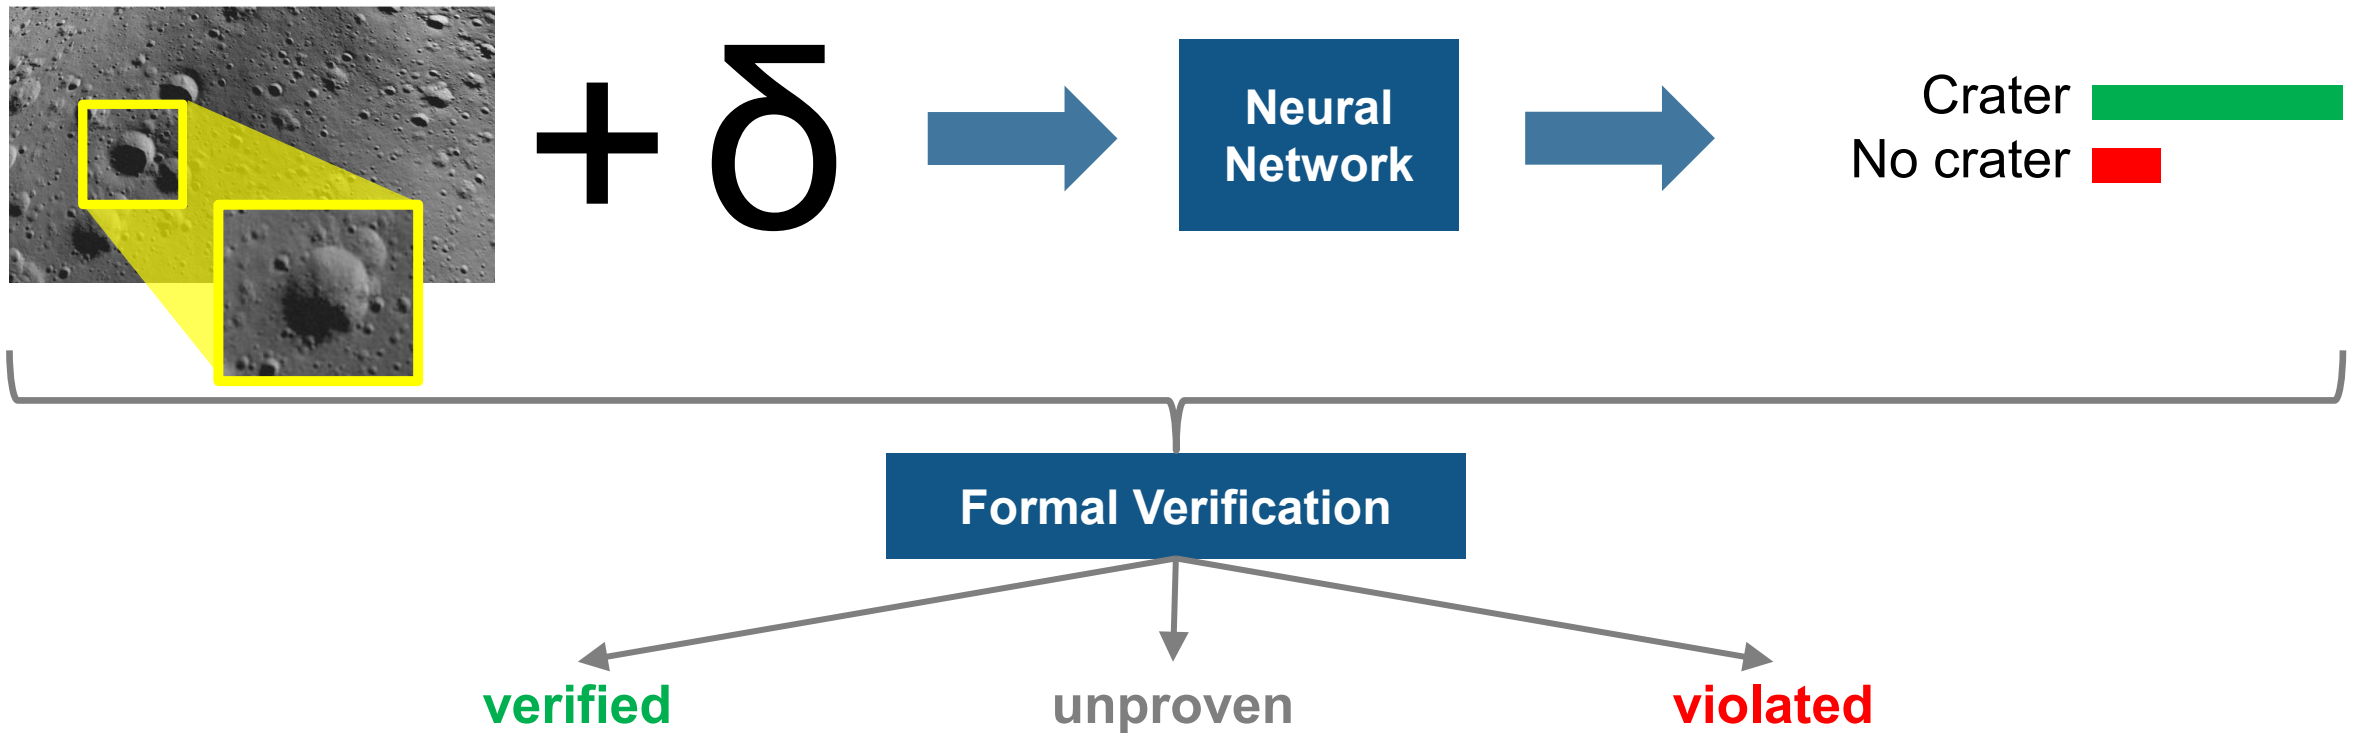
		LayerCoverage
1	fc	0.3000
2	softmax	0.2000



Deep Learning Toolbox Verification Library

R2022b

Verify deep learning network robustness against adversarial examples and to compute the output bounds for a set of input bounds.



<https://www.mathworks.com/help/deeplearning/deep-learning-verification.html>

<https://www.mathworks.com/matlabcentral/fileexchange/118735-deep-learning-toolbox-verification-library>