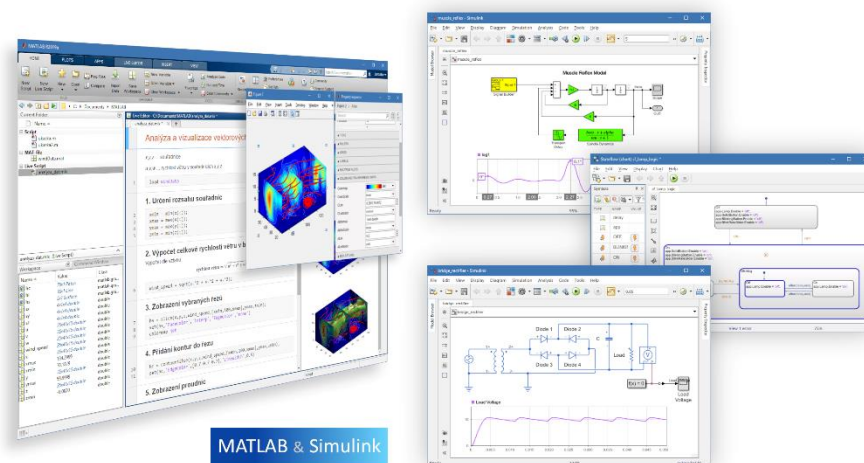


10.9.2020 Brno

TCC 2020

Deep learning

nové možnosti pro začátečníky i pokročilé uživatele



Jaroslav Jirkovský
jirkovsky@humusoft.cz

www.humusoft.cz
info@humusoft.cz

www.mathworks.com

AI, Machine Learning, and Deep Learning

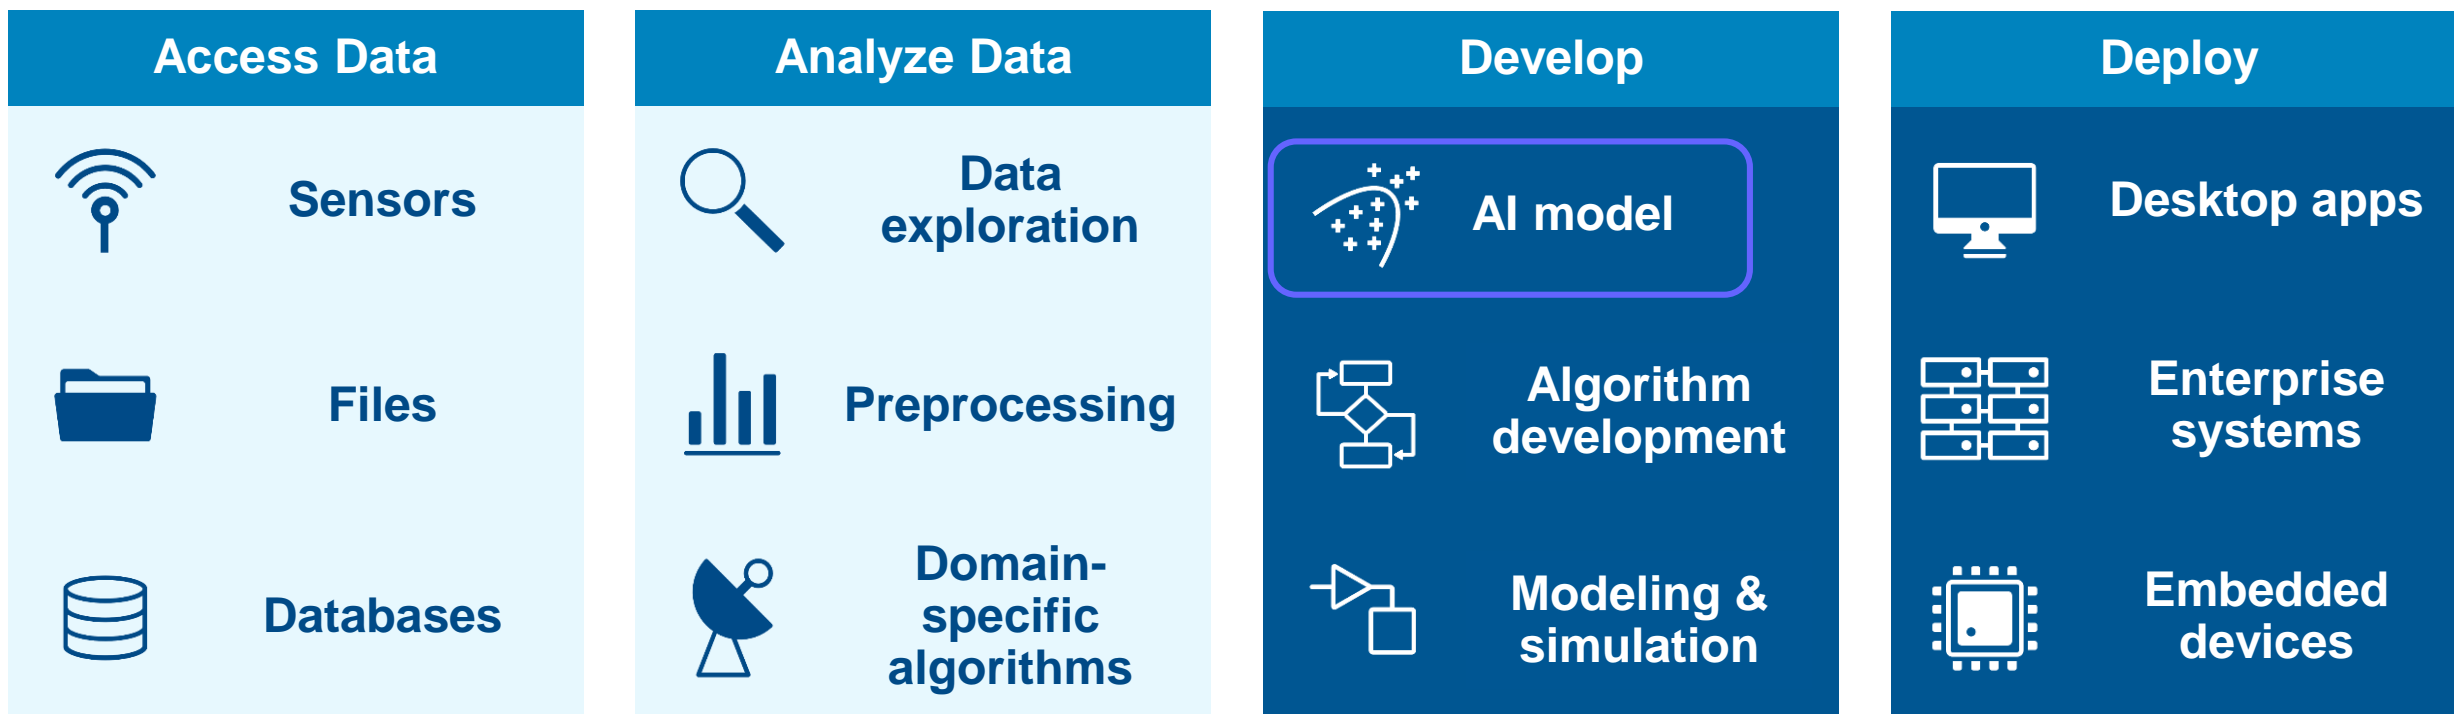
AI

Machine Learning

Deep Learning

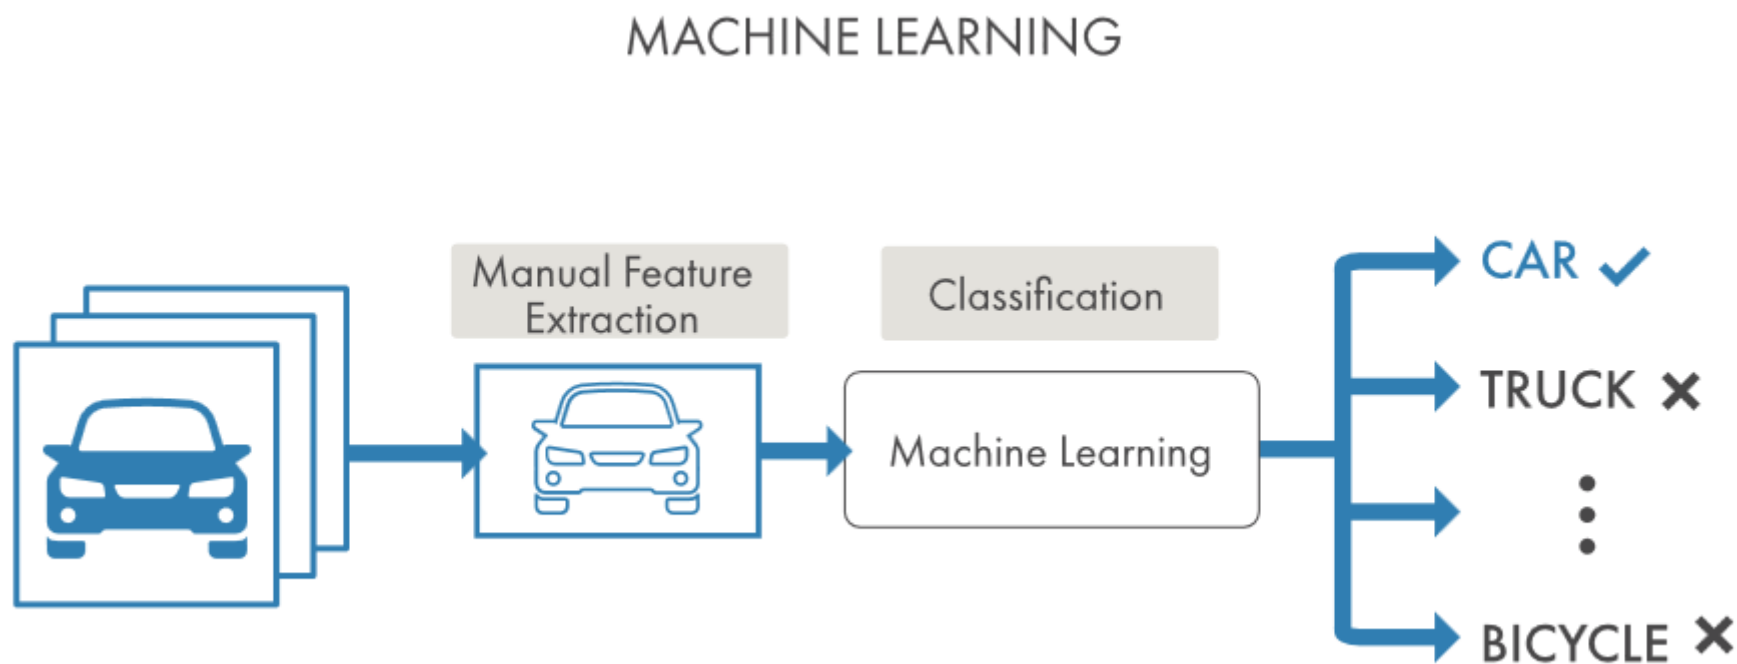
AI is Just One Part of System Development Workflow

with MATLAB and Simulink



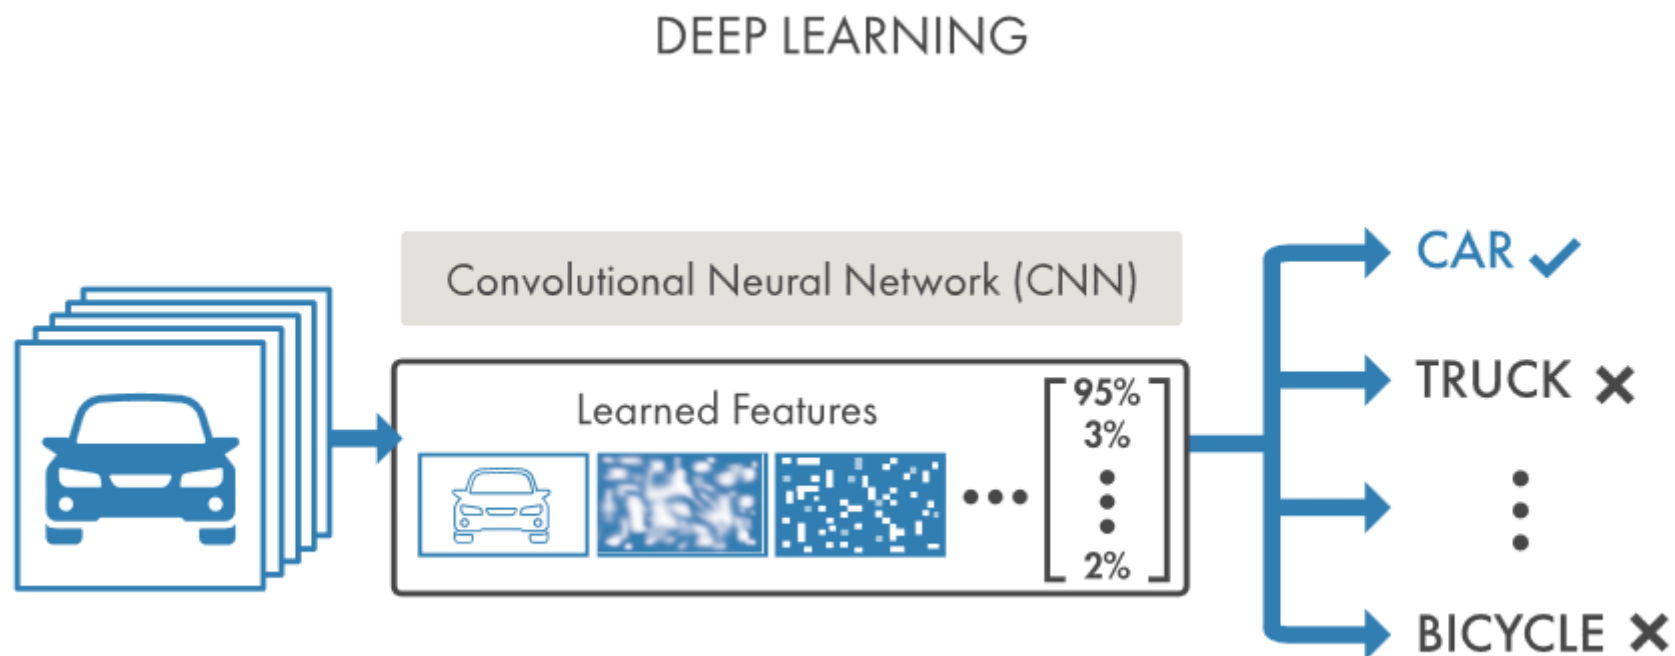
What is Machine Learning ?

Machine learning uses **data** and produces a **program** to perform a **task**






What is Deep Learning ?

Deep learning performs **end-end learning** by learning **features, representations and tasks** directly from images, text and sound



Deep Learning is Ubiquitous

- Computer Vision
- Signal Processing
- Robotics & Controls
- ...





MATLAB for Deep Learning

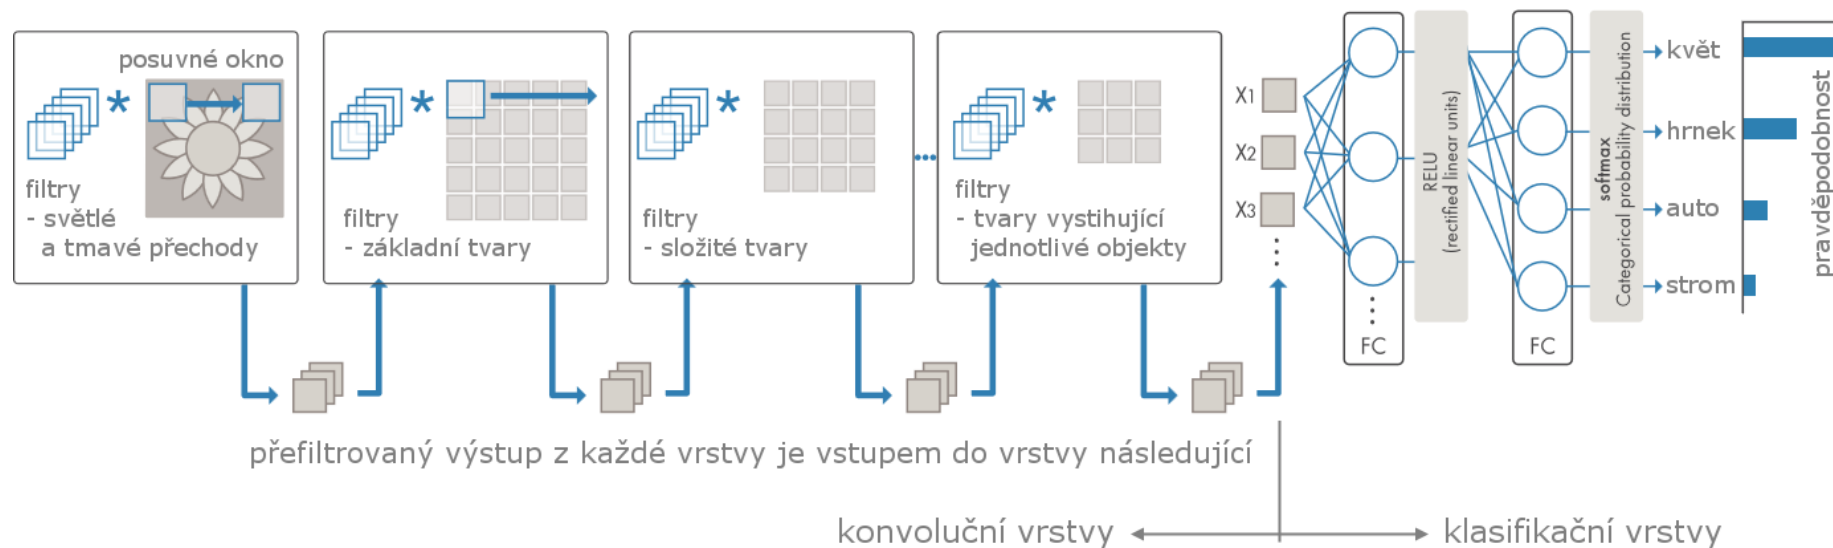
- **Network Architectures and Algorithms**
- **Training and Visualization**
- **Access the Latest Pretrained Models**
- **Scaling and Acceleration**
- **Handling Large Sets of Images**

- **Classification and Regression**
- **Object Detection**
- **Semantic Segmentation**


- **Embedded Deployment**



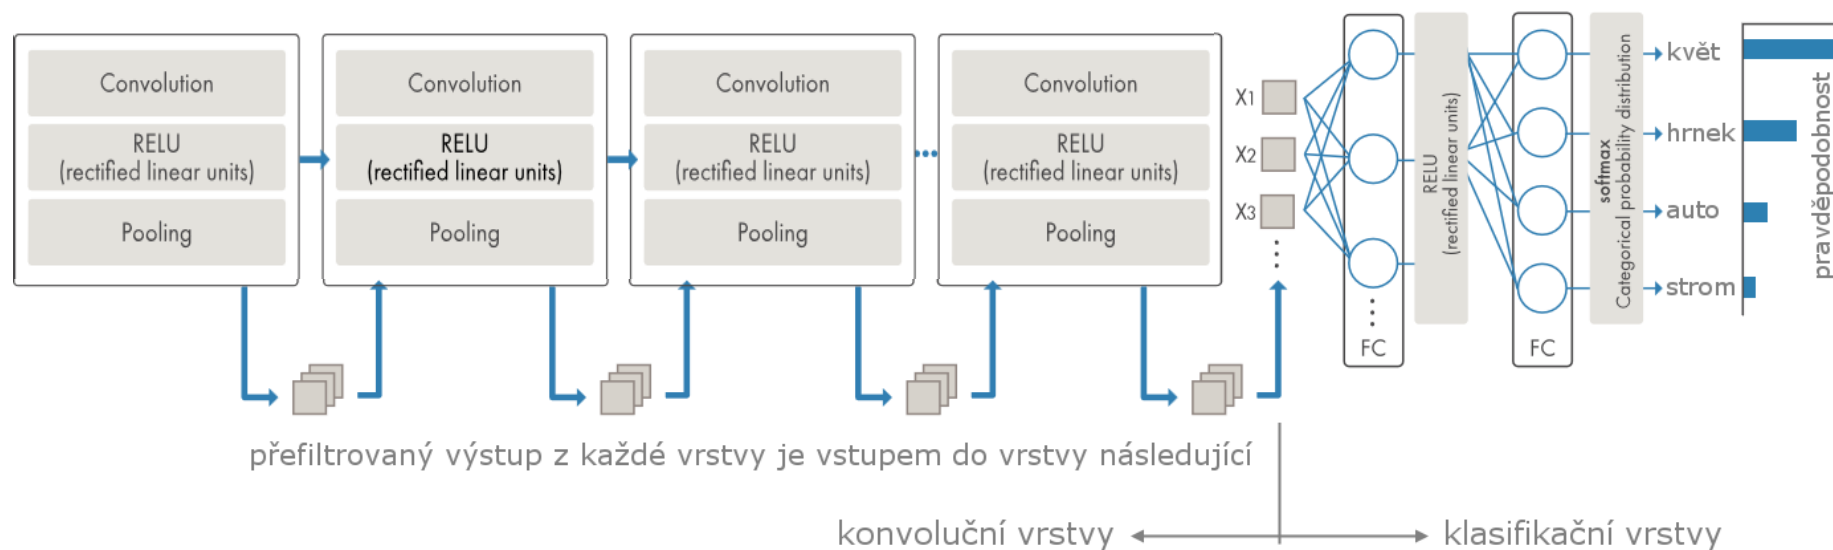

Convolutional Neural Networks (CNN)



What do filters do?



Convolutional Neural Networks (CNN)




What do filters do?




Great for classification:

- Convolution Layer
- ReLU Layer
- Max Pooling Layer

Deep Neural Networks in MATLAB – 3 Approaches



Deep Network Designer



training
using
APP

```
options = trainingOptions('sgdm');
convnet = trainNetwork(data, layers, opts);
```

```
results = classify(convnet, newData);
```

- for most deep learning tasks

Standard Framework

```
layers = [imageInputLayer([28 28 1])
convolution2dLayer(5, 20)
reluLayer()
maxPooling2dLayer(2, 'Stride', 2)
fullyConnectedLayer(10)
softmaxLayer()
classificationLayer()];
```

Extended Framework

- custom training loops
- automatic differentiation
- shared weights
- custom loss functions
- ...
- GANs, Siamese networks, ...

>30 Layers


imageInputLayer	Image input layer		
image3dInputLayer	3-D image input layer		
convolution2dLayer	2-D convolutional layer		
convolution3dLayer	3-D convolutional layer		
groupedConvolution2dLayer	leakyReluLayer	Leaky Rectified Linear Unit (ReLU) layer	
transposedConv2dLayer	clippedReluLayer	Clipped Rectified Linear Unit (ReLU) layer	
transposedConv3dLayer	eluLayer	Exponential linear unit (ELU) layer	
fullyConnectedLayer	tanhLayer	Hyperbolic tangent (tanh) layer	
reluLayer	batchNormalizationLayer	maxPooling2dLayer	Max pooling layer
	crossChannelNormalizationLayer	maxPooling3dLayer	3-D max pooling layer
	dropoutLayer	maxUnpooling2dLayer	Max unpooling layer
	averagePooling2dLayer	additionLayer	Addition layer
	averagePooling3dLayer	concatenationLayer	Concatenation layer
		depthConcatenationLayer	Depth concatenation layer
		softmaxLayer	Softmax layer
		classificationLayer	Classification output layer
		regressionLayer	Create a regression output layer

- Author custom layers in MATLAB using the Custom Layer API
 - including automatic differentiation

Transfer Learning using Pre-Trained Networks

- **Pre-Trained Networks**

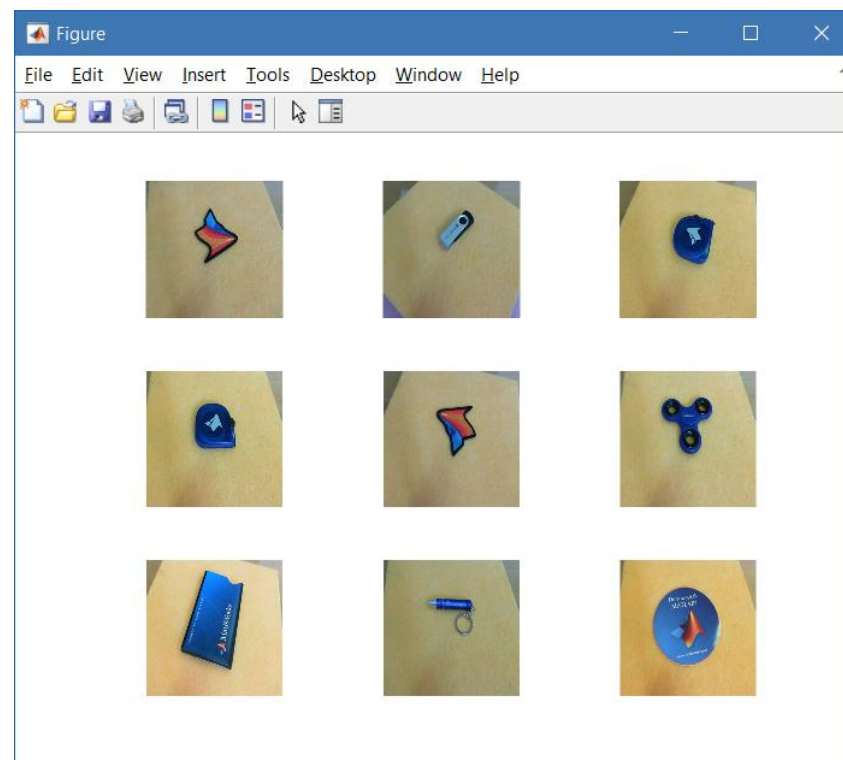
- AlexNet
- VGG-16 and VGG-19
- GoogLeNet
- ResNet-50 and ResNet-101
- Inception-v3
- Inception-ResNet-v2
- SqueezeNet
- and more ...



- **ONNX Model Converter**

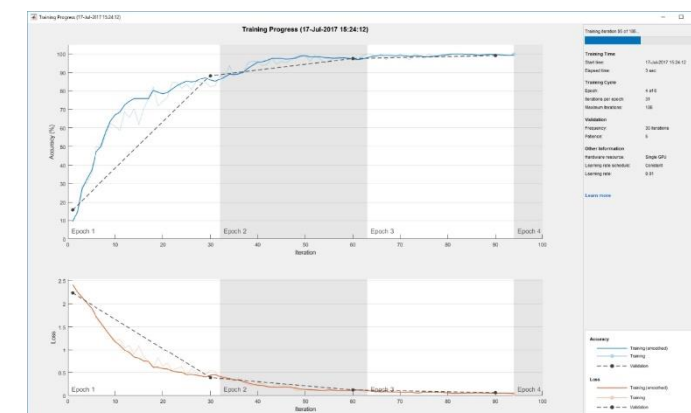
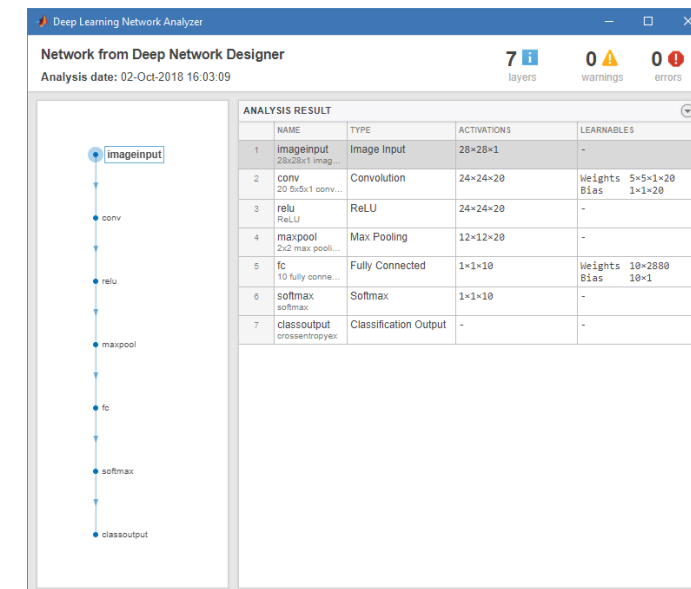
Example: Fine-tune a pre-trained model (transfer learning)

- <https://www.mathworks.com/help/deeplearning/gs/get-started-with-deep-network-designer.html>




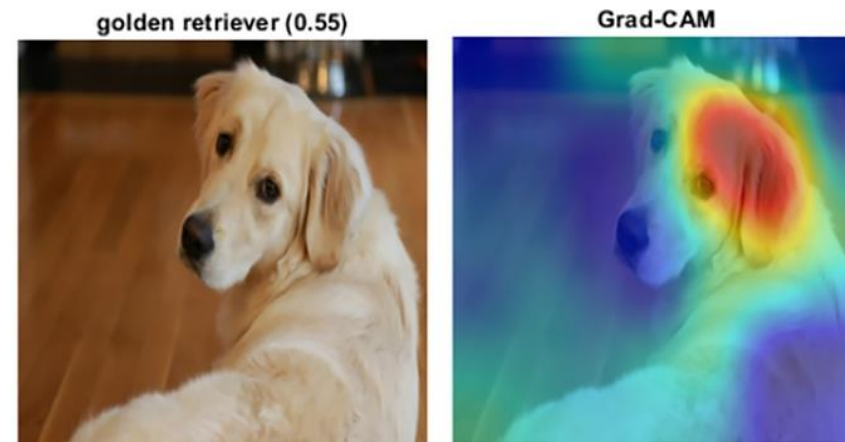
Training, Validation and Visualization

- **Network Analyzer (analyzeNetwork)**
 - find problems in network architectures before training
- **Monitor training progress**
 - plots for accuracy, loss, validation metrics, and more
- **Automatically validate network performance**
 - stop training when the validation metrics stop improving
- **Perform hyperparameter tuning**
 - using Bayesian optimization



Debugging and Visualization

- Visualize activations and filters from intermediate layers
- CAM (Class Activation Mapping)
- Grad-CAM
- Occlusion sensitivity maps
- Deep Dream visualization



Handling Large Sets of Images

- **Use imageDataStore**
 - easily read and process large sets of images
- **Access data stored in**
 - local files
 - networked storage
 - databases
 - big data file systems
- **Efficiently resize and augment image data**
 - increase the size of training datasets
 - `imageDataAugmenter`, `augmentedImageDatastore`



Customizations (Extended Framework)

- **Define and train complex networks using**
 - custom training loops
 - automatic differentiation
 - shared weights
 - custom loss functions
- **Custom layers support**
 - define new layers, including layers with multiple inputs and outputs
- **Multi-Input, Multi-Output Networks**
 - create and train networks with multiple inputs and multiple outputs
- **Build advanced network architectures**
 - GANs, Siamese networks, attention networks, ...



Using Custom Training Loops

1. Define your network
 - *lgraph* object, same as standard approach, without classification layer
2. Convert network object to *dlnetwork* object
3. Define your custom training options
 - as a set of variables, not options-object
4. Define you custom training loop
 - for loops over number of epochs and iterations
 - read data, convert to *dlarray* (and *gpuArray* for GPU computing)
 - calculate model gradients and loss (use automatic differentiation)
 - run solver for network update
5. Use the trained network
 - *predict* function, convert to class selection



Example: Extended Framework

- <https://www.mathworks.com/help/deeplearning/ug/train-network-using-custom-training-loop.html>

```
% Loop over epochs.
for epoch = 1:numEpochs
    % Shuffle data.
    idx = randperm(numel(YTrain));
    XTrain = XTrain(:,:,,idx);
    YTrain = YTrain(idx);

    % Loop over mini-batches.
    for i = 1:numIterationsPerEpoch
        iteration = iteration + 1;

        % Read mini-batch of data and convert the labels to dummy
        % variables.
        idx = (i-1)*miniBatchSize+1:i*miniBatchSize;
        X = XTrain(:,:,,idx);

        Y = zeros(numClasses, miniBatchSize, 'single');
        for c = 1:numClasses
            Y(c,YTrain(idx)==classes(c)) = 1;
        end

        % Convert mini-batch of data to dlarray.
        dlX = dlarray(single(X),'SSCB');

        % If training on a GPU, then convert data to gpuArray.
```

Generative Adversarial Network (GAN)

- **Generate data with similar characteristics as the input real data**
- **Two networks that train together**

- **Generator**
 - input vector of random values
 - generates data with the same structure as the training data

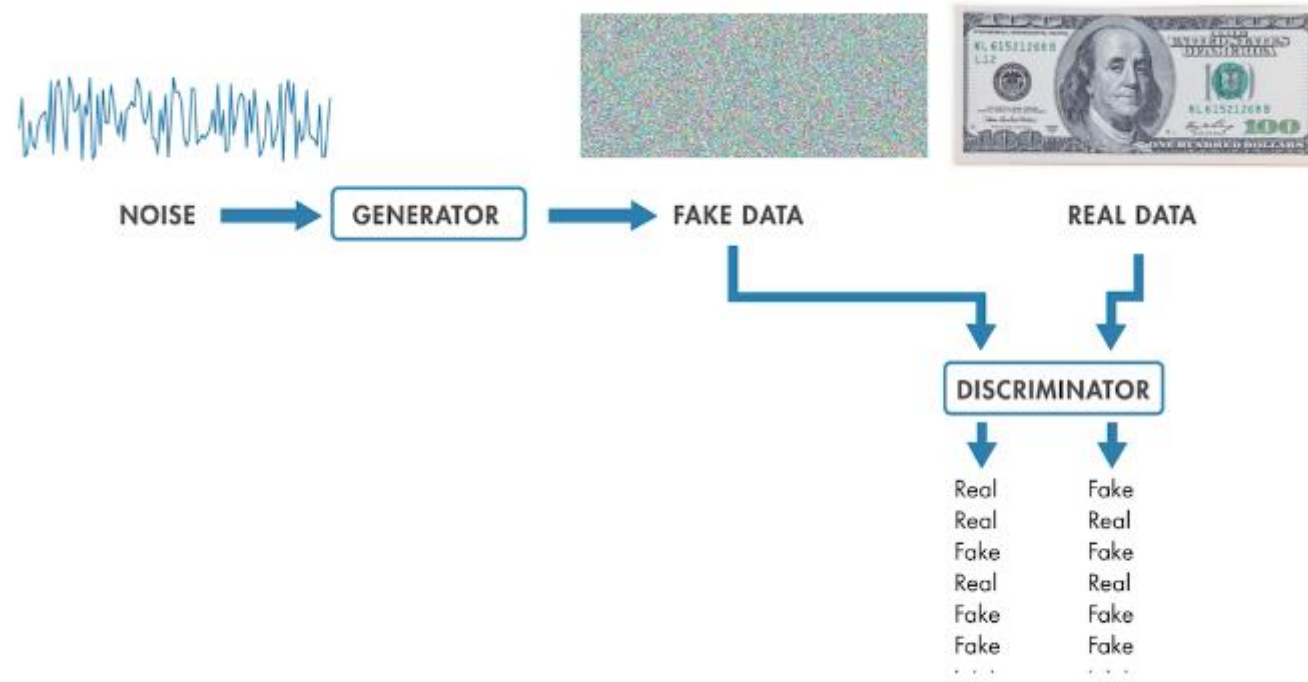
- **Discriminator**
 - observations from the training data, and generated data
 - classify the observations as "real" or "generated".

Generative Adversarial Network (GAN)

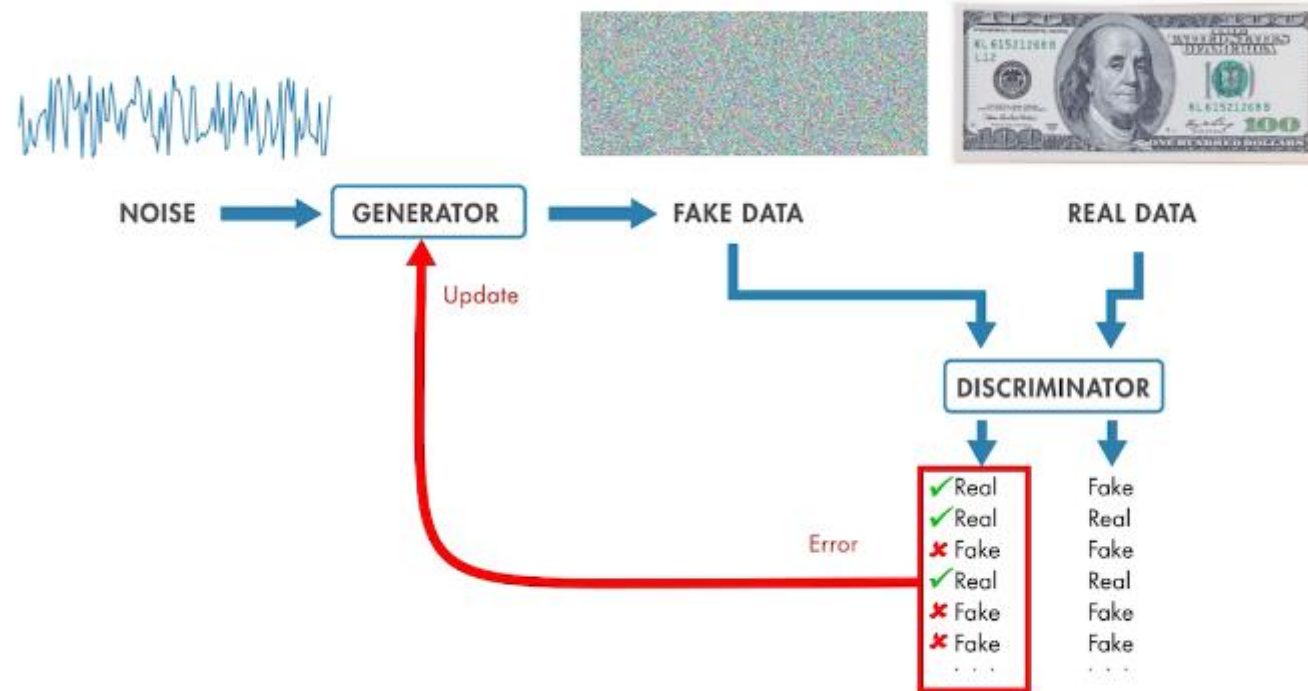
GENERATOR

DISCRIMINATOR

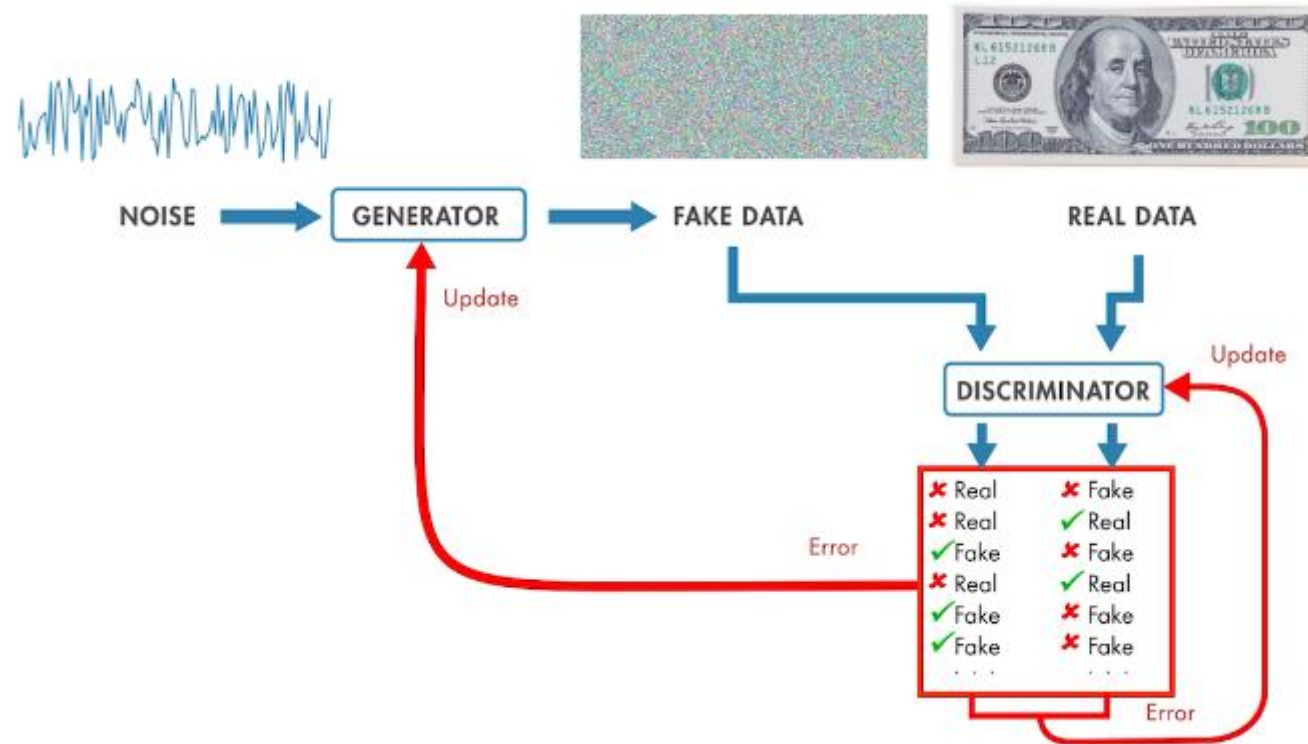
Generative Adversarial Network (GAN)



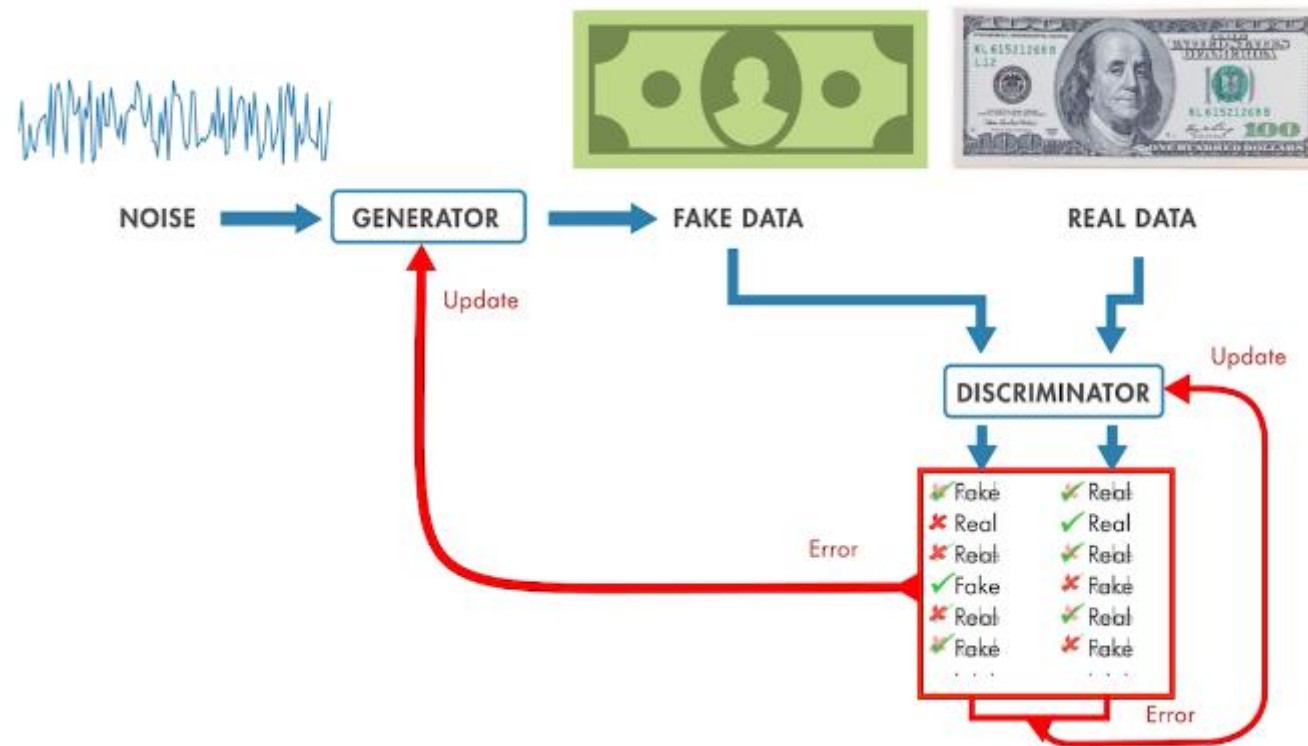
Generative Adversarial Network (GAN)



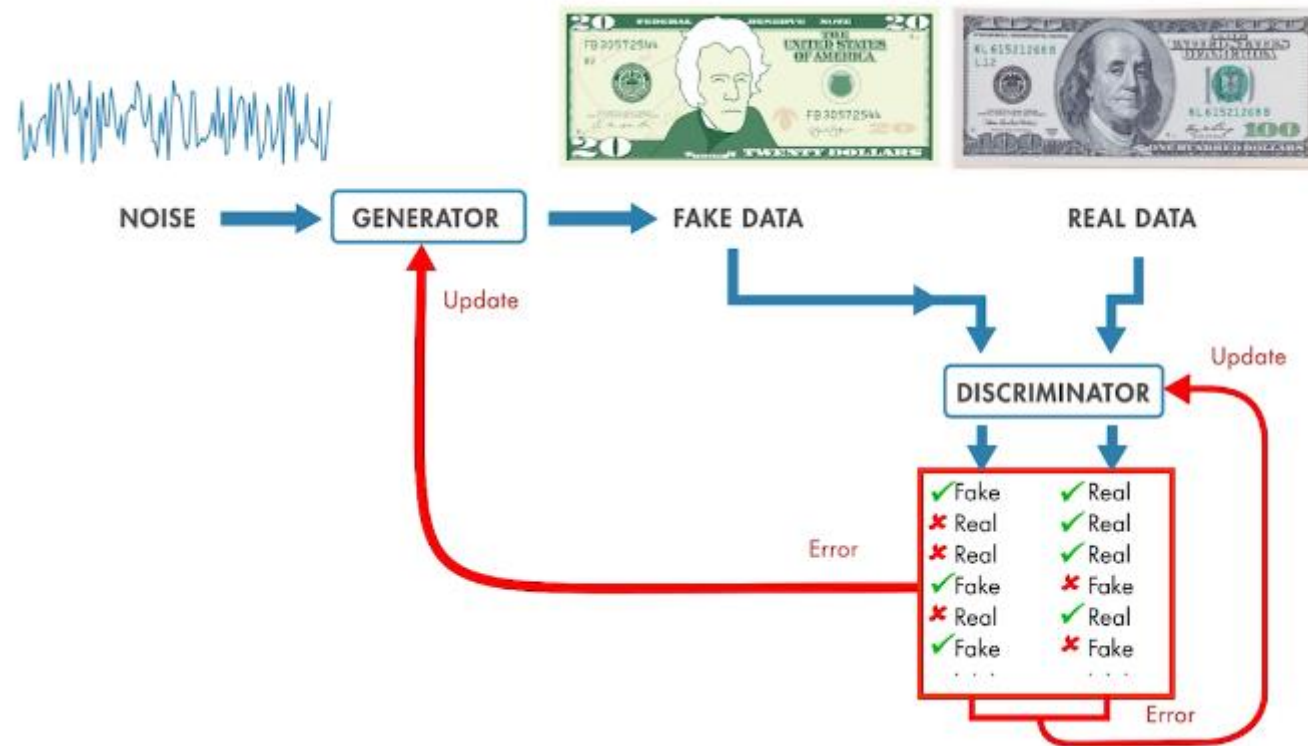
Generative Adversarial Network (GAN)



Generative Adversarial Network (GAN)



Generative Adversarial Network (GAN)



Example: Generative Adversarial Network

- <https://www.mathworks.com/help/deeplearning/ug/train-generative-adversarial-network.html>

Generated Images



Other deep learning tasks with images

- **Regression**


- predict continuous variable from the image

- **Object Detection**

- recognizing and locating the object in a scene
- multiple objects in one image


- **Semantic Segmentation**

- classify individual pixels



Deep learning for signal processing

- **Leverage CNNs with signals**
 - „convert“ signal into image using time-frequency representations
 - how spectral content of signal evolves over time
 - many time-frequency representations available
 - spectrogram, cwt, stft
 - doc "[Time-Frequency Gallery](#)"
- **Special network layers for signals – LSTM networks**
 - classification and prediction



Multi-Platform Deployment

- **Deploy deep learning models anywhere**
 - CUDA
 - C code
 - enterprise systems
 - or the cloud
 - **Generate code that leverages optimized libraries**
 - Intel® (MKL-DNN)
 - NVIDIA (TensorRT, cuDNN)
 - ARM® (ARM Compute Library)
- ⇒ **deployable models with high-performance inference speed.**




Latest Features

- **What's New in MATLAB for Deep Learning?**
 - <https://www.mathworks.com/solutions/deep-learning/features.html>



Jak začít s prostředím MATLAB?

- On-line kurzy zdarma
 - MATLAB Onramp, Simulink Onramp, Stateflow Onramp
 - [Deep Learning Onramp](#), [Machine Learning Onramp](#)
 - časová náročnost: 2 hodiny
 - <https://matlabacademy.mathworks.com/>



Děkuji za pozornost