8.9.2016 Brno

TCC 2016

Deep Learning with MATLAB

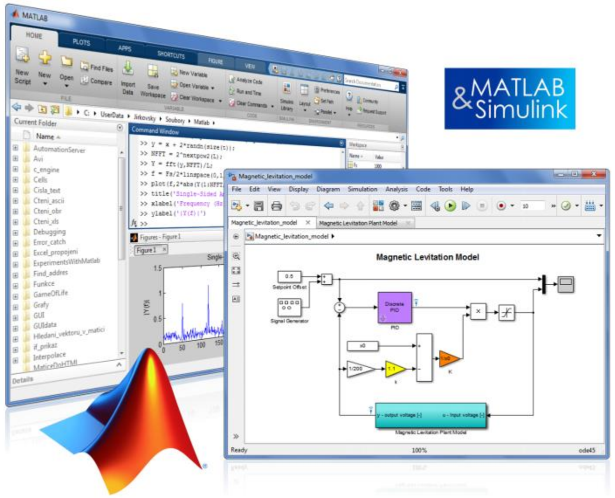

Jan Studnička
studnicka@humusoft.cz

www.humusoft.cz
info@humusoft.cz

www.mathworks.com

HUMUSOFT

Computer Vision Applications

Computer Vision

- **Pedestrian and traffic sign detection**
- **Landmark identification**
- **Scene recognition**
- **Medical diagnosis and drug discovery**
- **Public Safety / Surveillance**
- **Automotive**
- **Robotics**

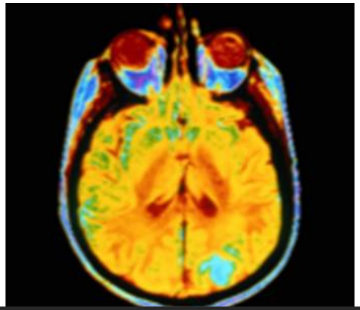

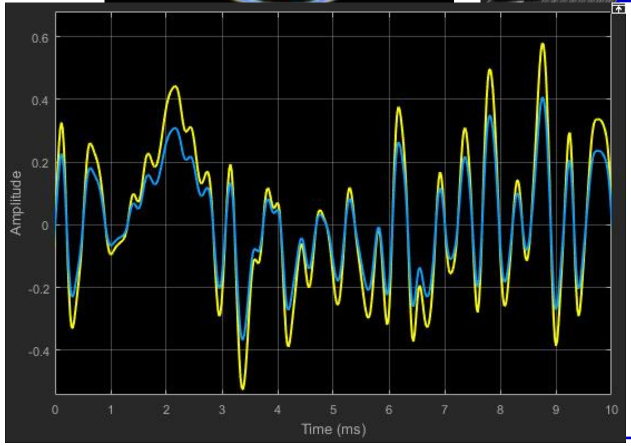

and many more…

What is Deep Learning ?

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

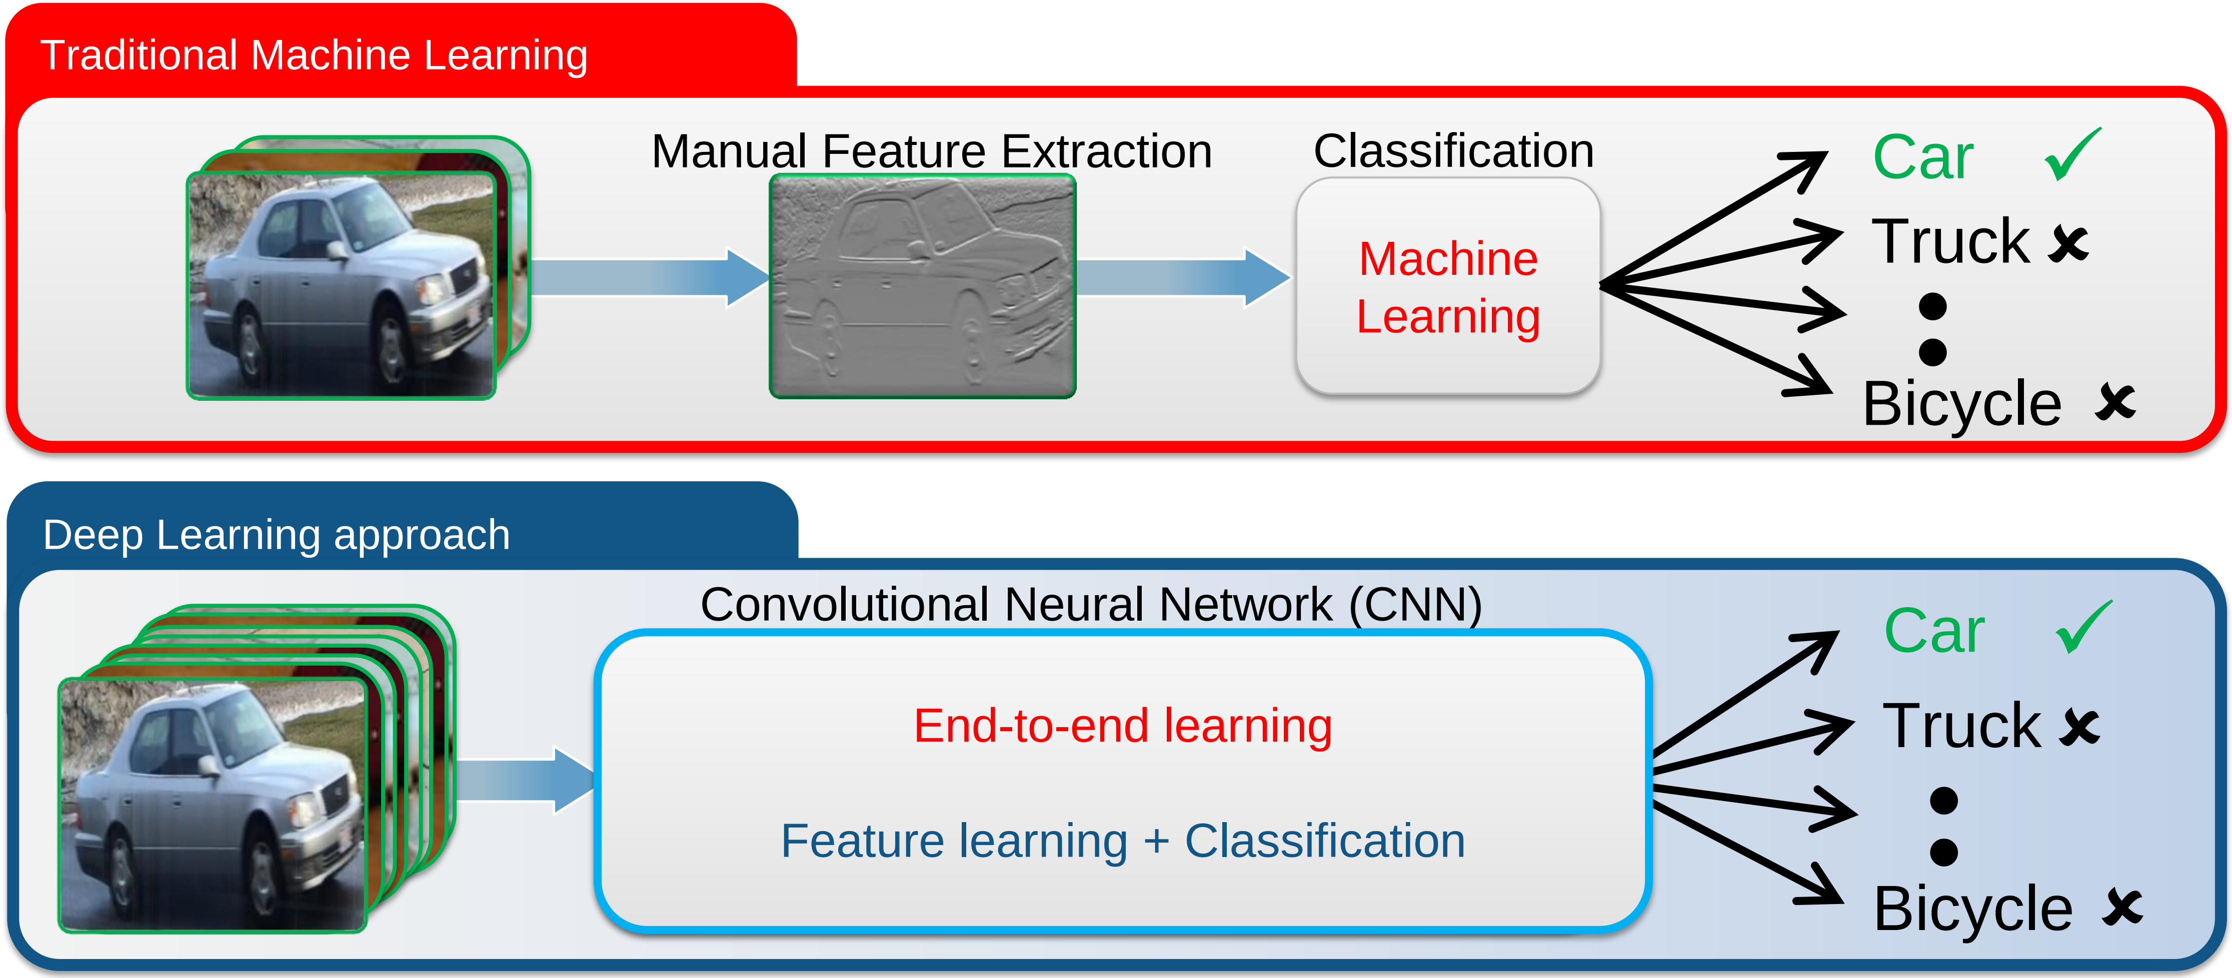

Deep Learning with MATLAB for Computer Vision

- **Autoencoders**
  - Example: Classify digits in images

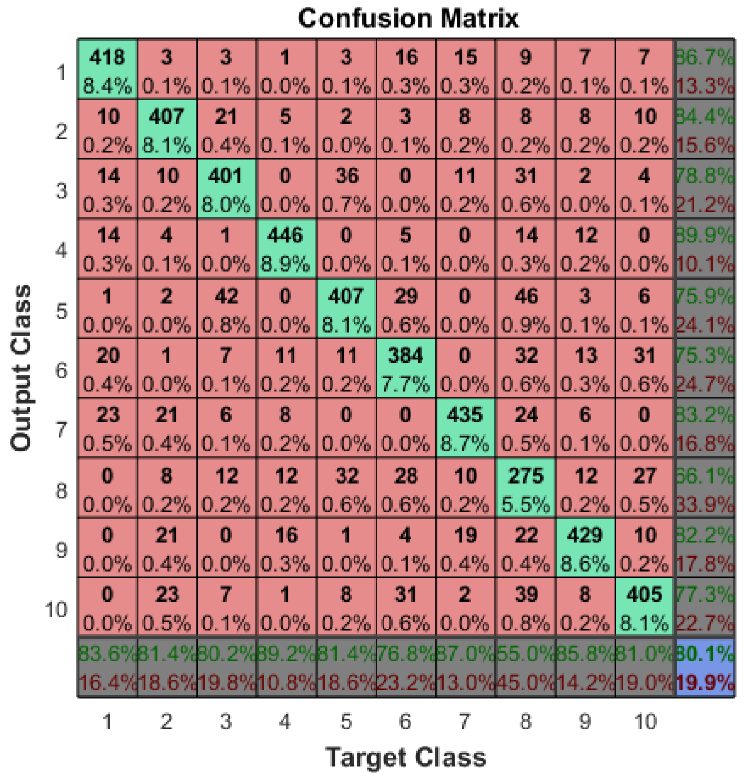

- **Convolutional Neural Networks (CNN)**
  - Trained on massive sets of data
  - High accuracy

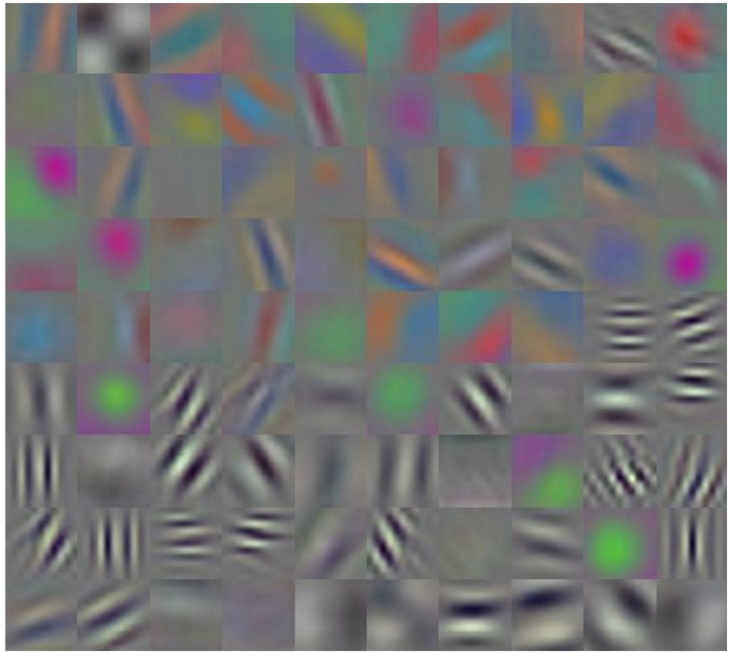

A visualization of learned weights of the first layer of a CNN.

Neural Network

- **Single neuron**

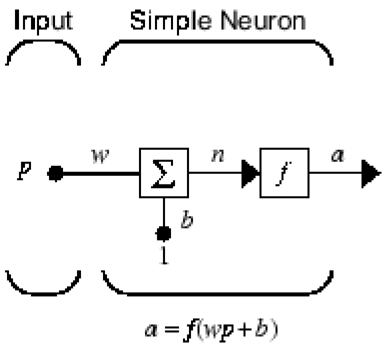

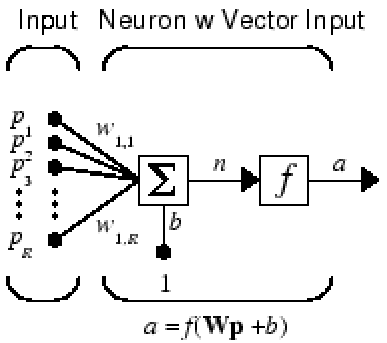

Where

R = number of elements in input vector

- **Layer of Neurons**

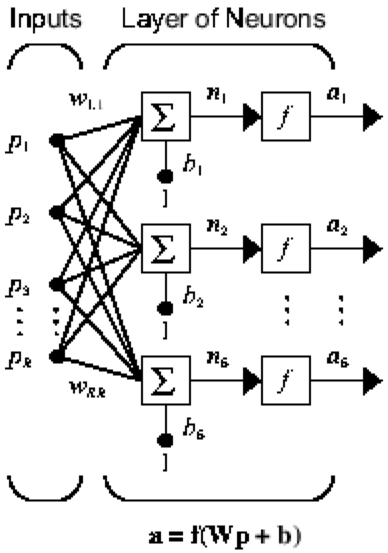

Where

R = number of elements in input vector

S = number of neurons in layer

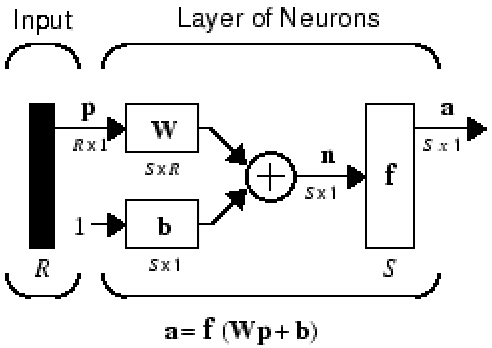

Where...

R = number of elements in input vector

S = number of neurons in layer 1

Autoencoders

- **Unsupervised Learning**
  - Hidden layer – Encoder

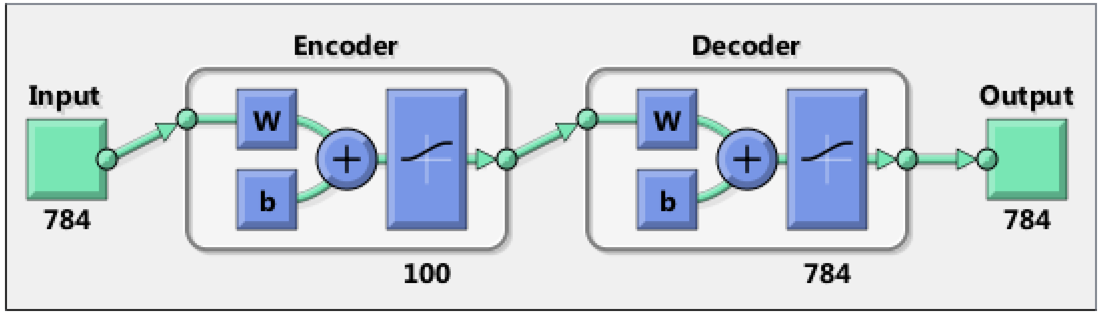

- **Pretrain Deep Neural Network**
  - Hidden layers – Encoders of pretrained Autoencoders

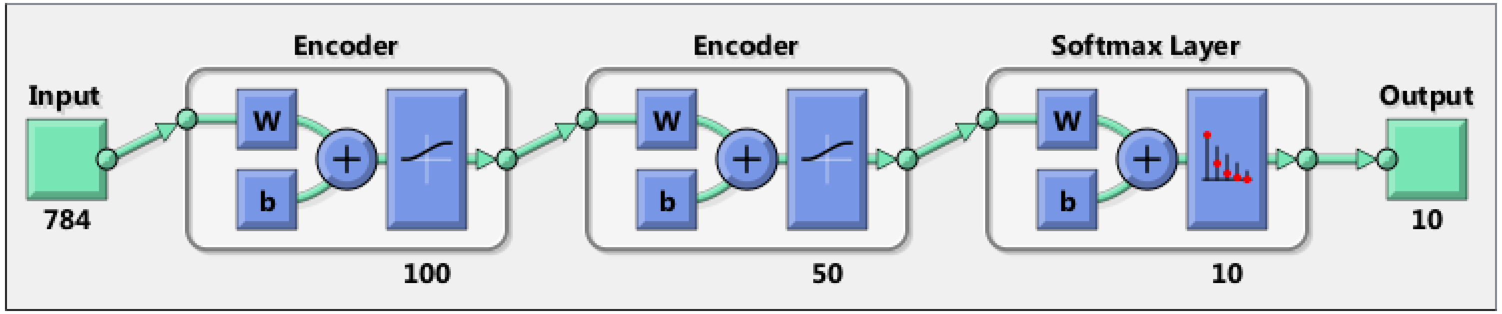

Digit Classification

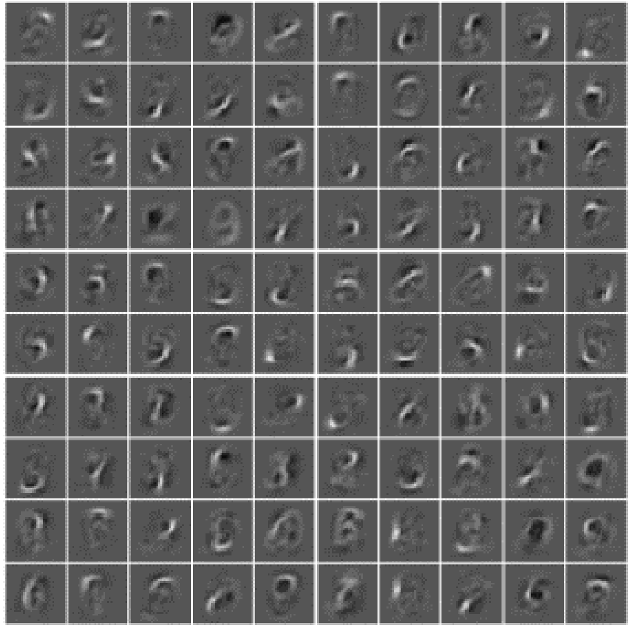

- **Classify digits in images**
- **Data:**
  - 28 x 28 pixels
  - 10 digit classes
  - 5000 samples
- **Solution:**

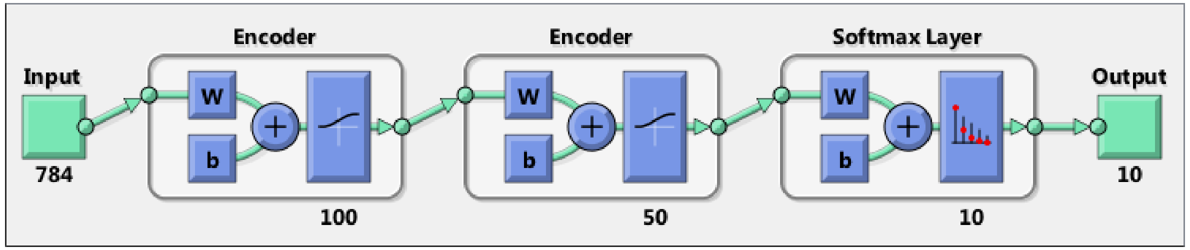

  - 2 hidden layers – autoencoders
  - Classification – Softmax layer
  - Stack the Encoders with the Softmax layer to form a Deep Network
  - Fine-tune the entire Deep Network – Classification

Convolutional Neural Networks: Live Object Recognition with Webcam

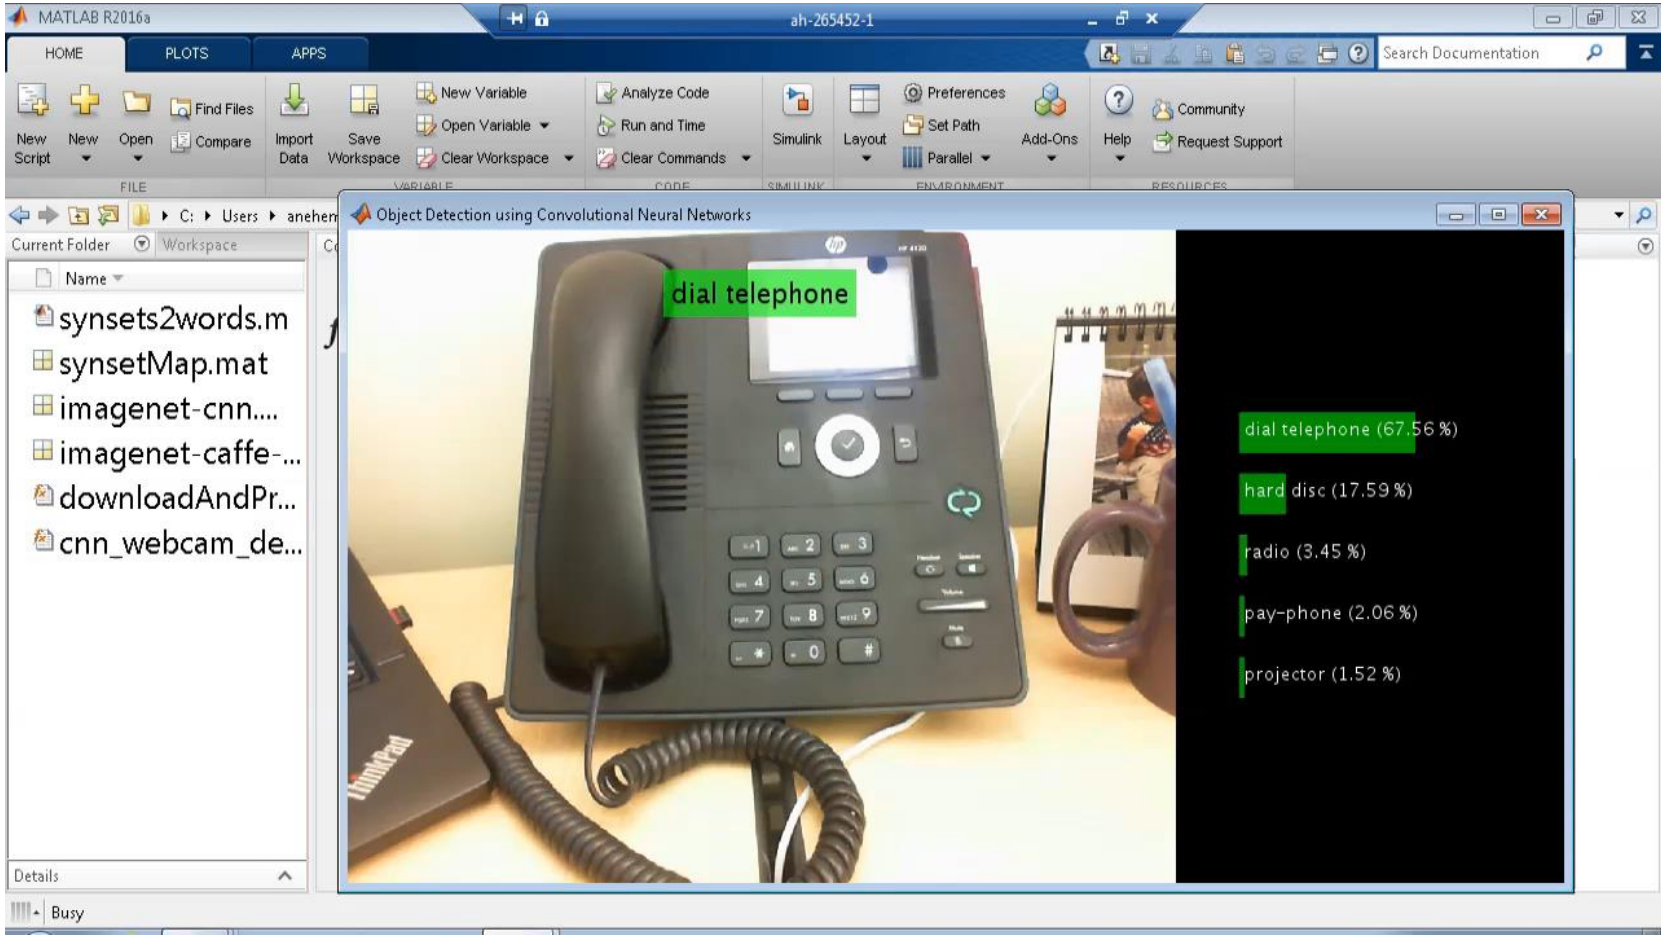

Why is Deep Learning so Popular ?

- **Results: Achieved substantially better results on ImageNet large scale recognition challenge**
  - **95% + accuracy on ImageNet 1000 class challenge**

Year	Error Rate
Pre-2012 (traditional computer vision and machine learning techniques)	> 25%
2012 (Deep Learning)	~ 15%
2015 (Deep Learning)	<5 %

- **Computing Power: GPU's and advances to processor technologies have enabled us to train networks on massive sets of data.**

- **Data: Availability of storage and access to large sets of labeled data**
  - **E.g. ImageNet , PASCAL VoC , Kaggle**

Two Approaches for Deep Learning

1. Train a Deep Neural Network from Scratch

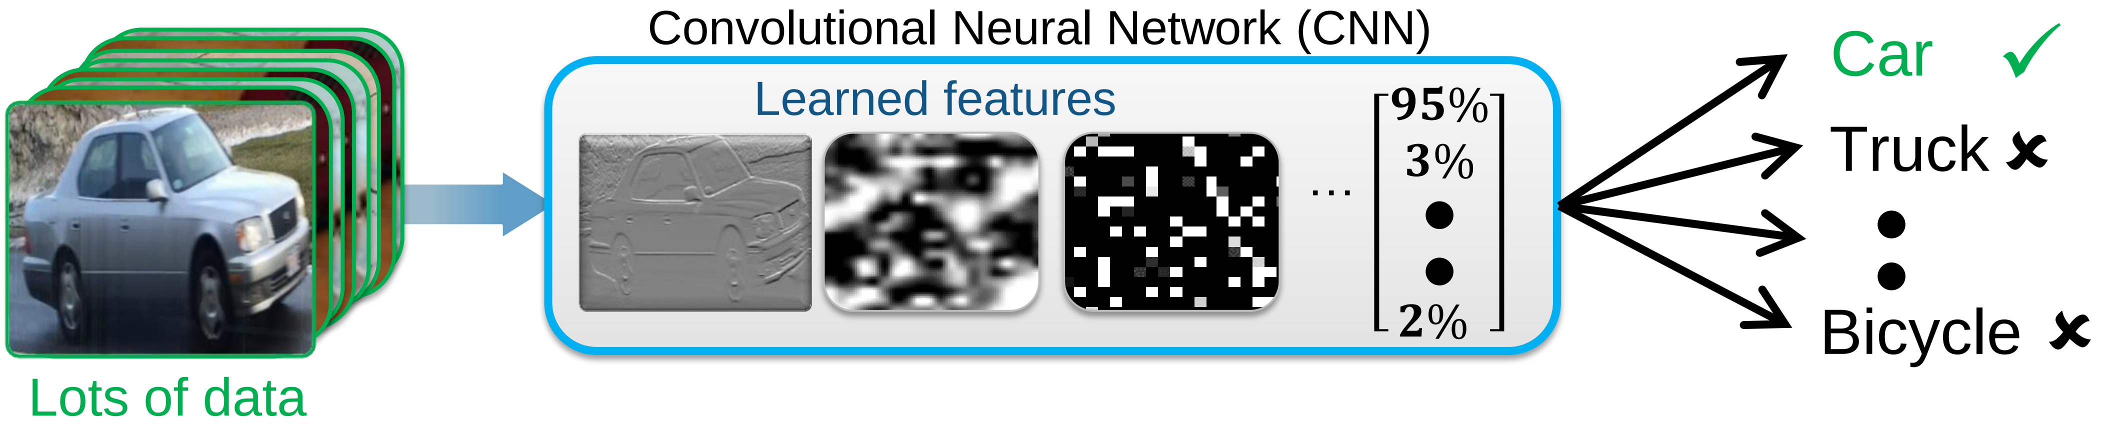

2. Fine-tune a pre-trained model (transfer learning)

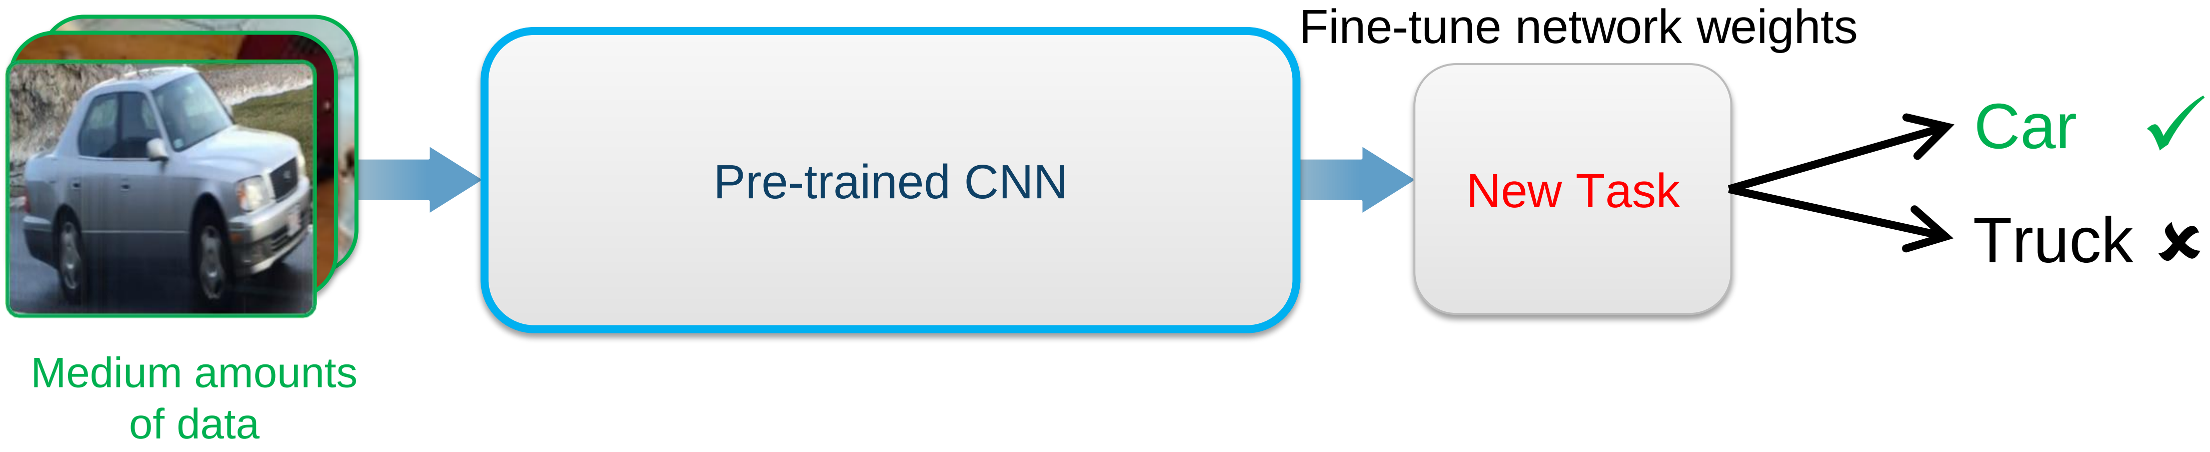

Two Deep Learning Approaches

Approach 1: Train a Deep Neural Network from Scratch

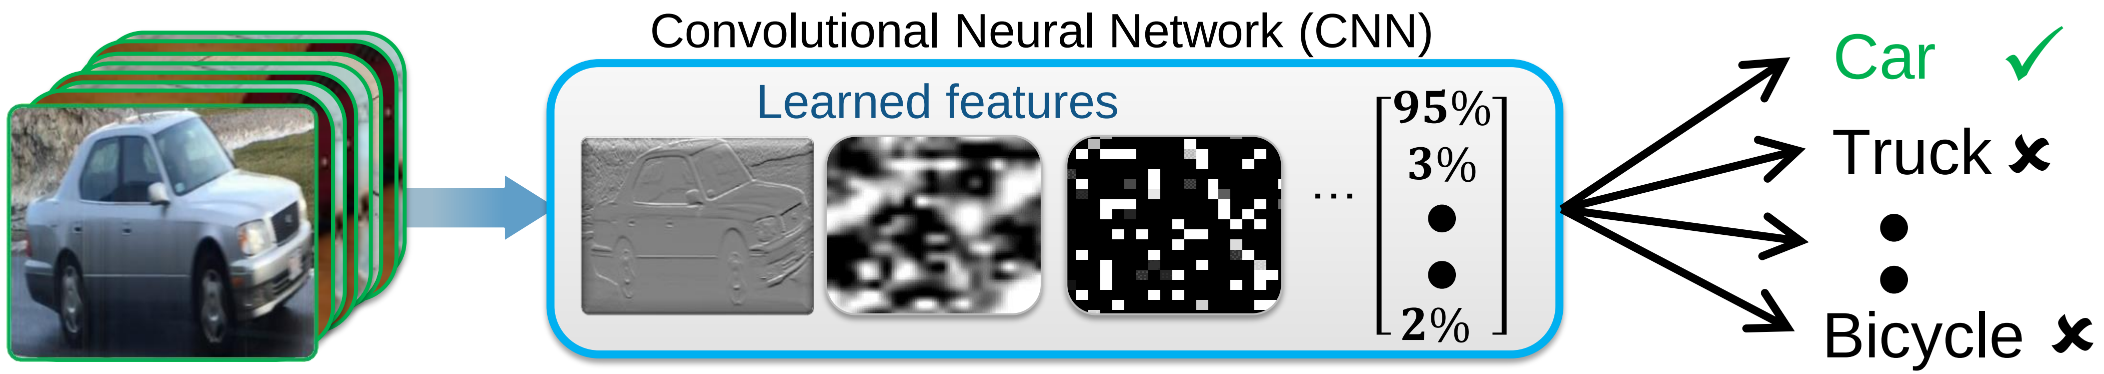

Recommended only when:

Training data	1000s to millions of labeled images
Computation	Compute intensive (requires GPU)
Training Time	Days to Weeks for real problems
Model accuracy	High (can overfit to small datasets)

Two Deep Learning Approaches

Approach 2: Fine-tune a pre-trained model (transfer learning)

CNN trained on massive sets of data

- Learned robust representations of images from larger data set
- Can be fine-tuned for use with ***new data or task*** with small – medium size datasets

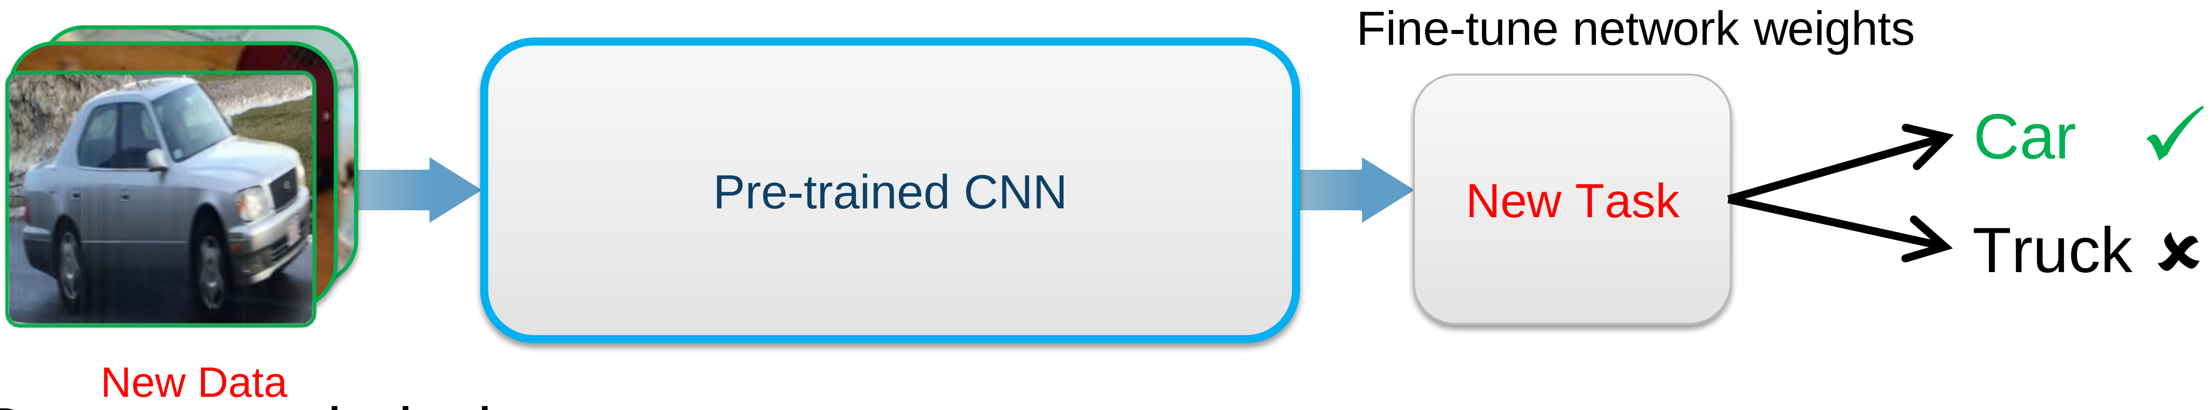

Recommended when:

Training data	100s to 1000s of labeled images (small)
Computation	Moderate computation (GPU optional)
Training Time	Seconds to minutes
Model accuracy	Good, depends on the pre-trained CNN model

Convolutional Neural Networks

- **Train "deep" neural networks on structured data (e.g. images, signals, text)**
- **Implements Feature Learning: Eliminates need for "hand crafted" features**
- **Trained using GPUs for performance**

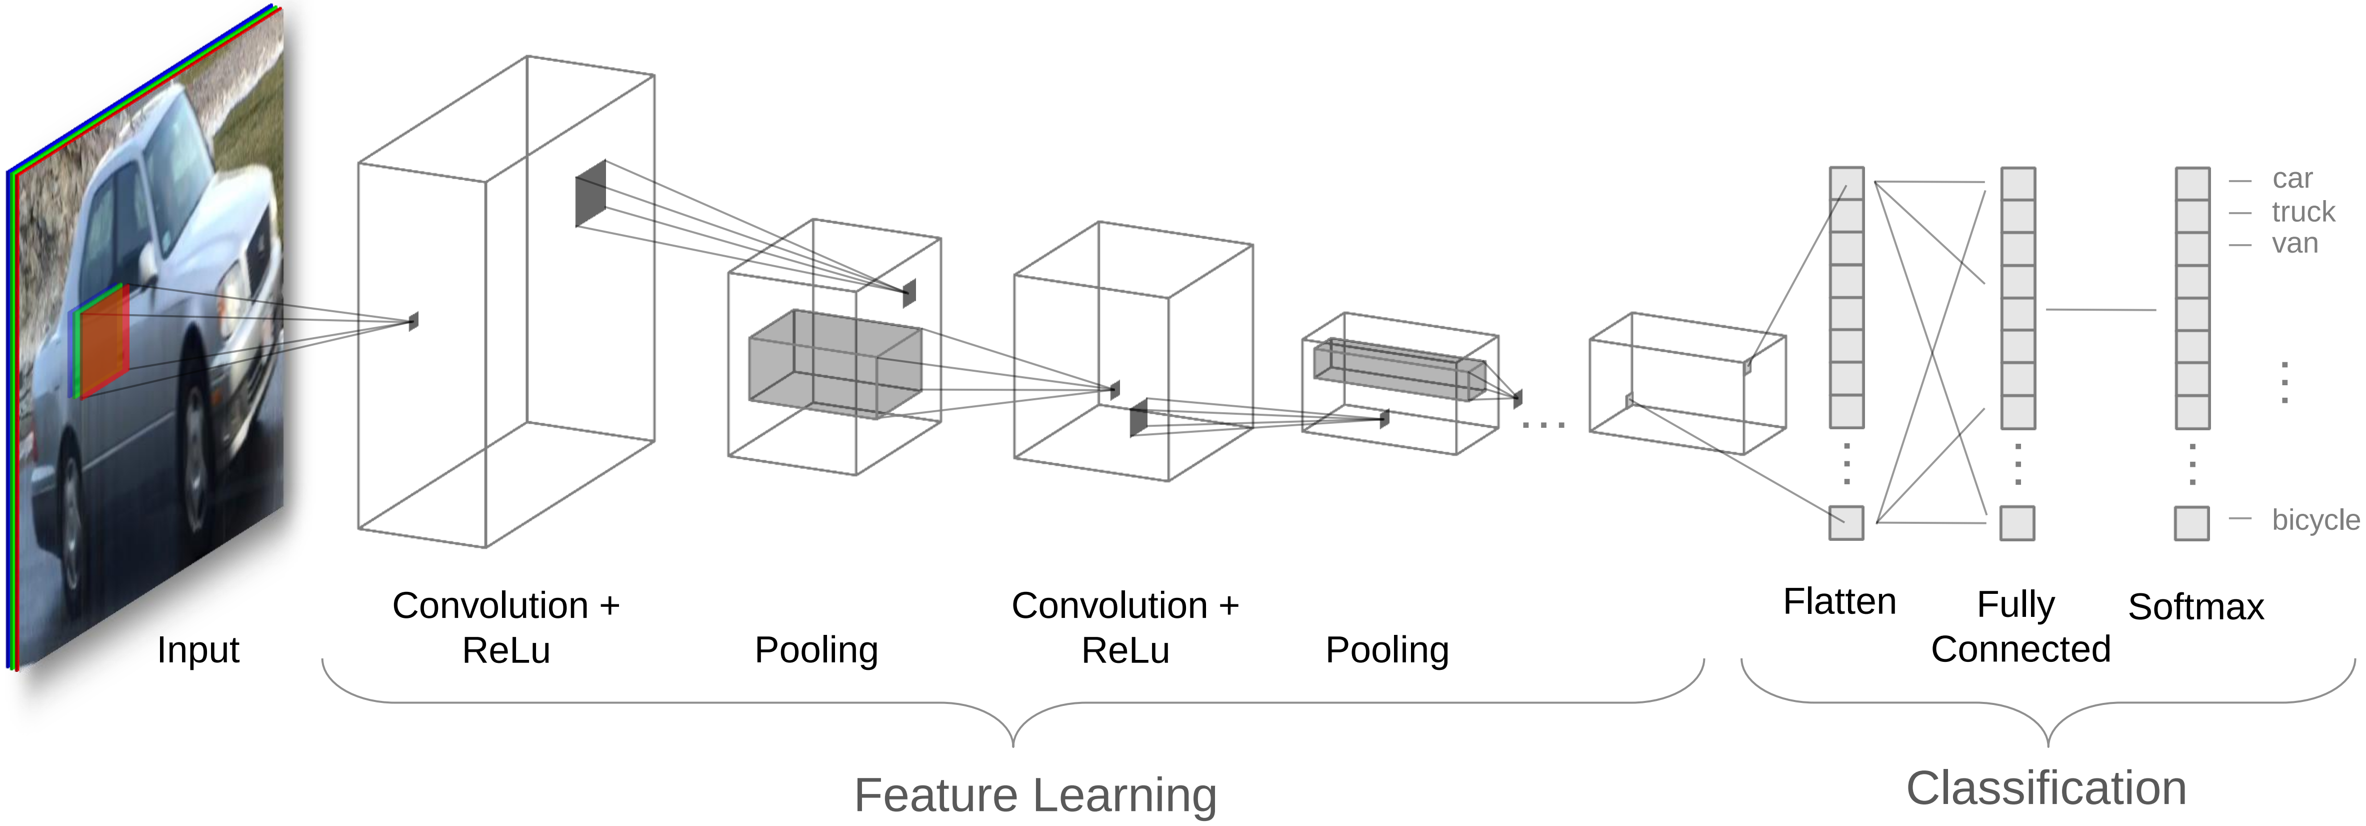

Convolutional Neural Networks

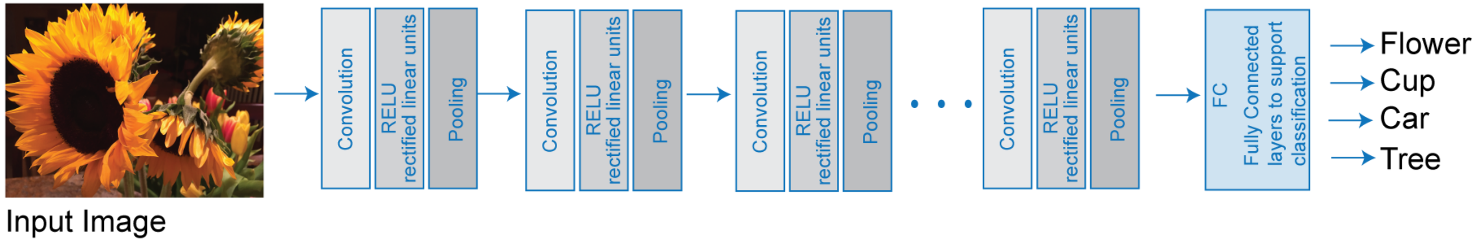

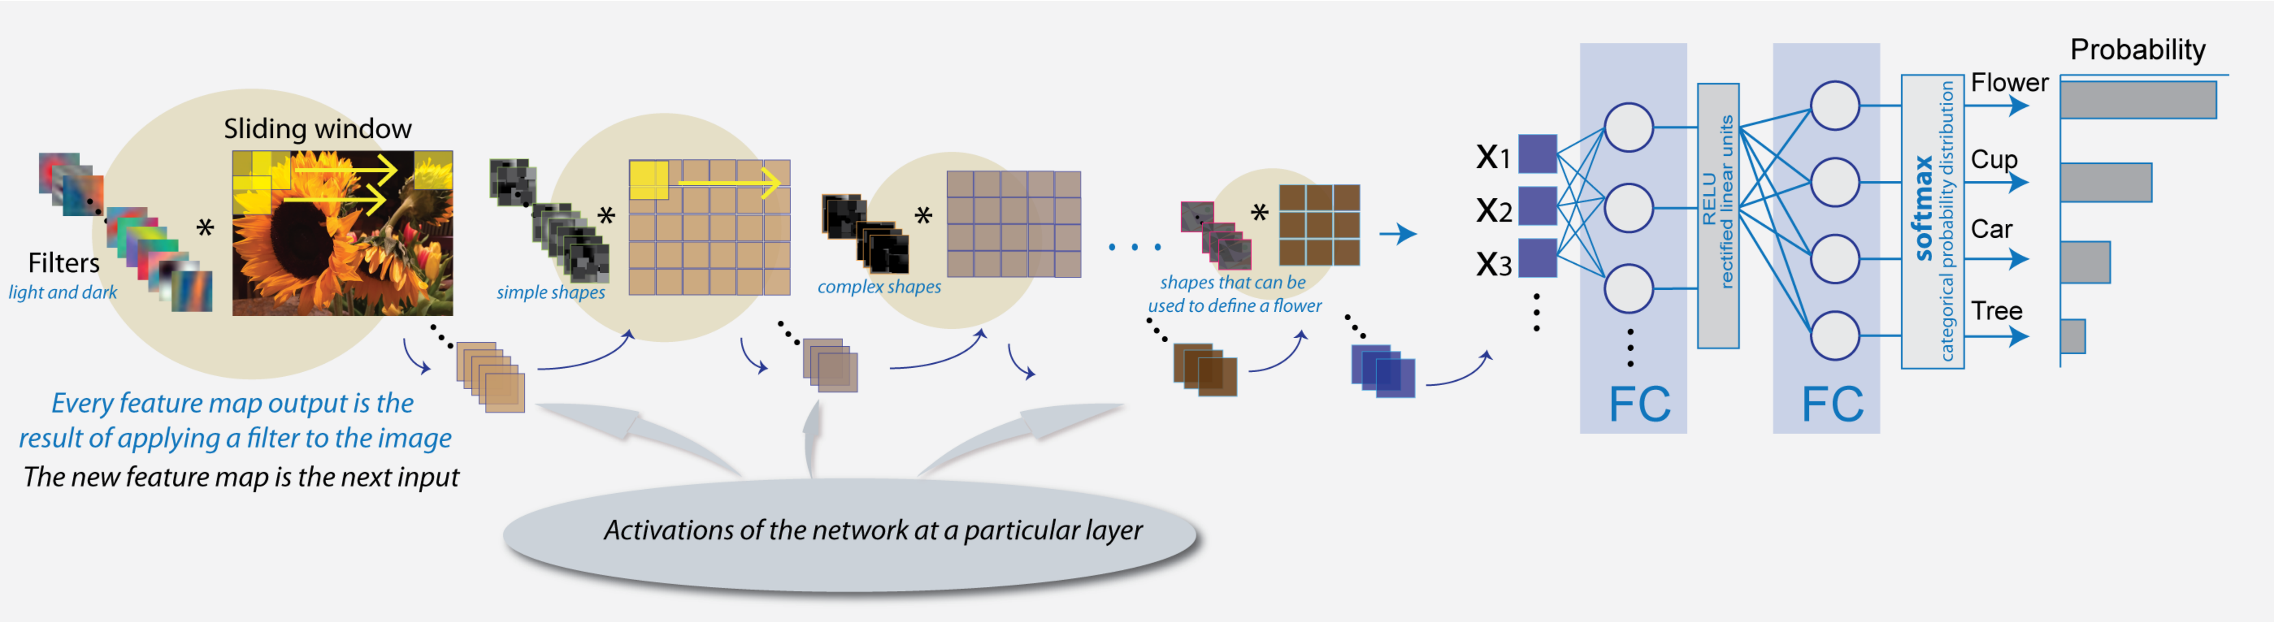

Demo

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

New Data

Pre-trained CNN
(AlexNet – 1000 Classes)

Car

SUV

New Task – 2 Class Classification

Demo

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

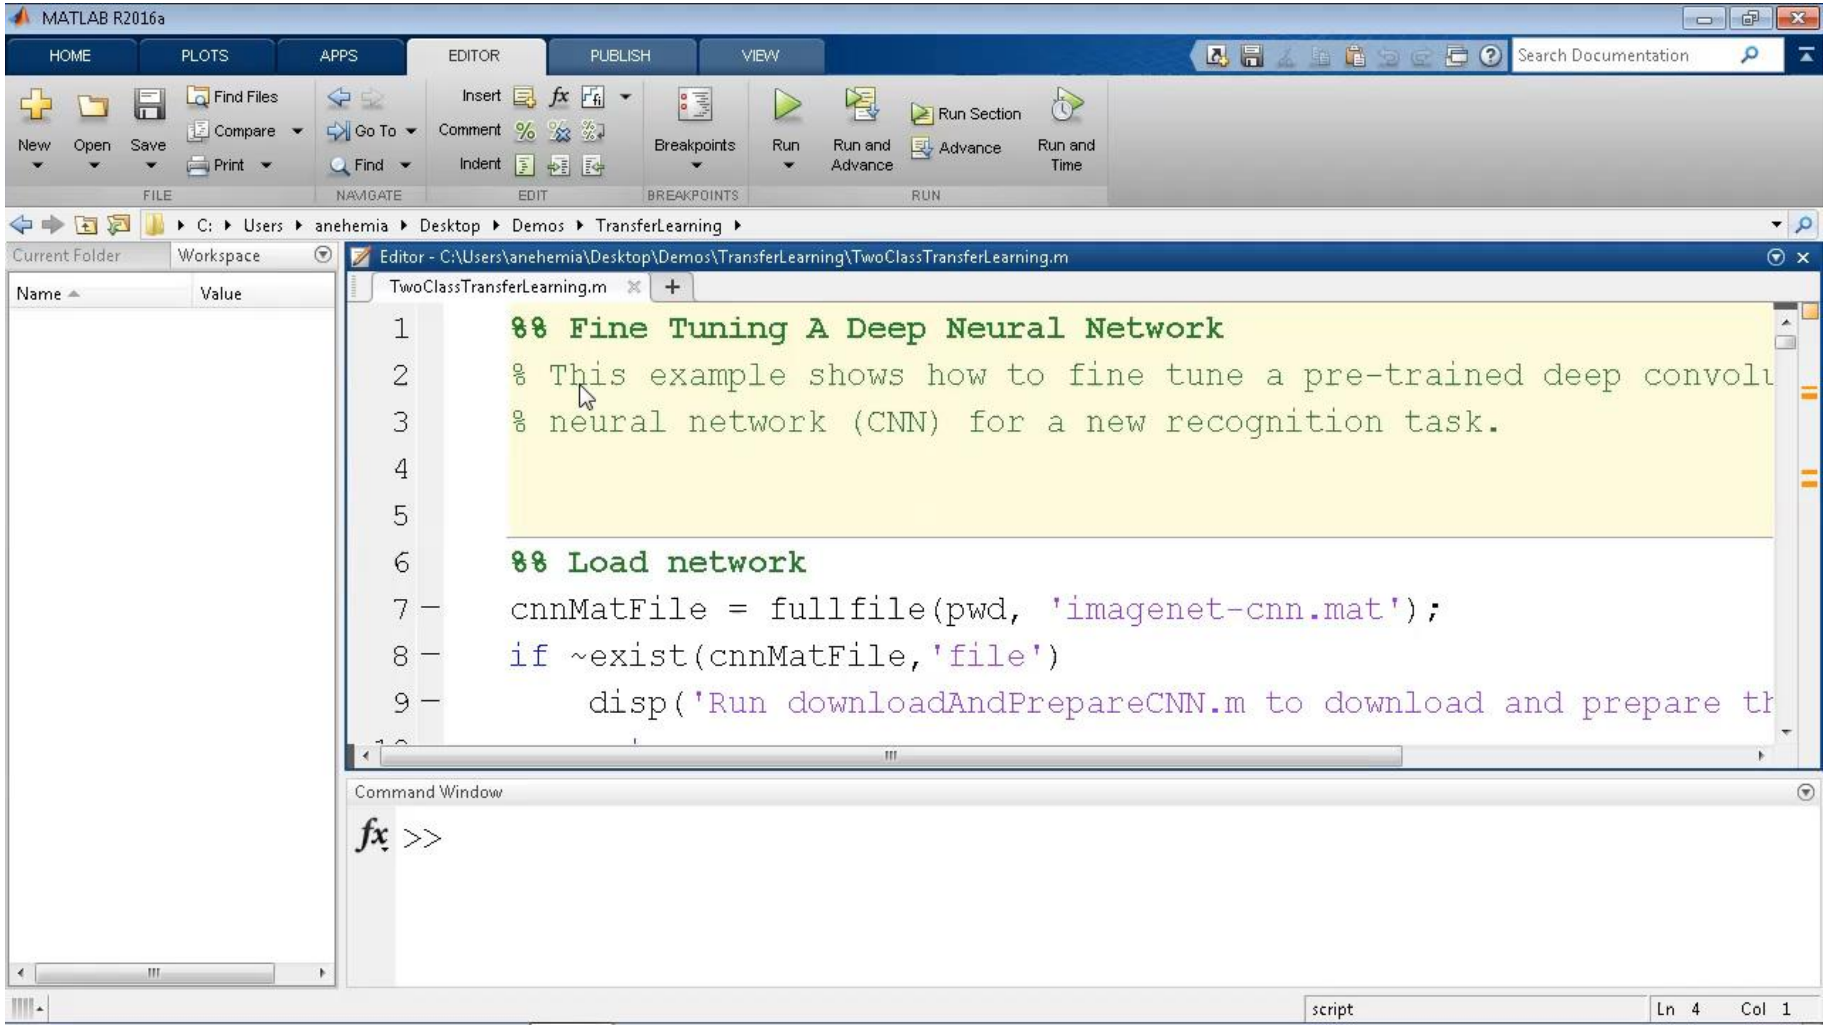

Addressing Challenges in Deep Learning for CV

Challenge	Solution
Managing large sets of labeled images	**`imageSet`** or **`imageDataStore`** to handle large sets of images
Resizing, Data augmentation	**`imresize`**, **`imcrop`**, **`imadjust`**, **`imageInputLayer`**, etc.
Background in neural networks (deep learning)	Intuitive interfaces, well-documented architectures and examples
Computation intensive task (requires GPU)	Training supported on GPUs No GPU expertise is required
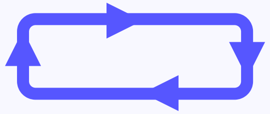	Automate. Offload computations to a cluster and test multiple architectures

International reseller of MathWorks for the Czech Republic and Slovakia

Pobřežní 20

186 00 Praha 8

Česká republika

Email: info@humusoft.cz

www.humusoft.cz

Tel.: +420 284 011 720

www.facebook.com/humusoft

www.youtube.com/humusoft

www.twitter.com/humusoft