
MATLAB EXPO 2018
UNITED KINGDOM
3 October | Silverstone, Northamptonshire
Register at matlabexpo.co.uk

MATLAB EXPO 2018

Image Processing and Computer Vision with MATLAB

Justin Pinkney

Image Processing, Computer Vision, and Deep Learning

Image Processing

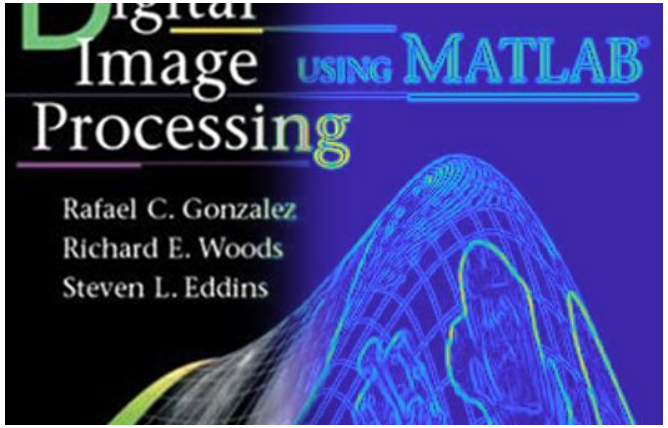

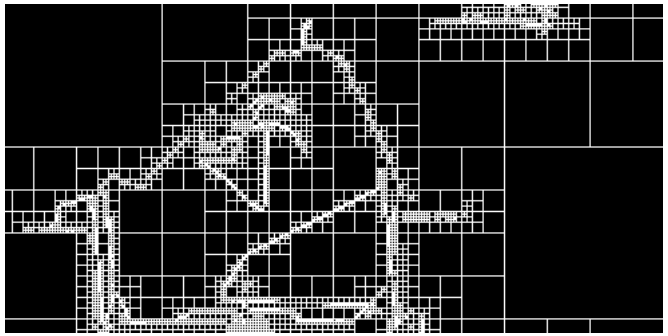

Computer Vision

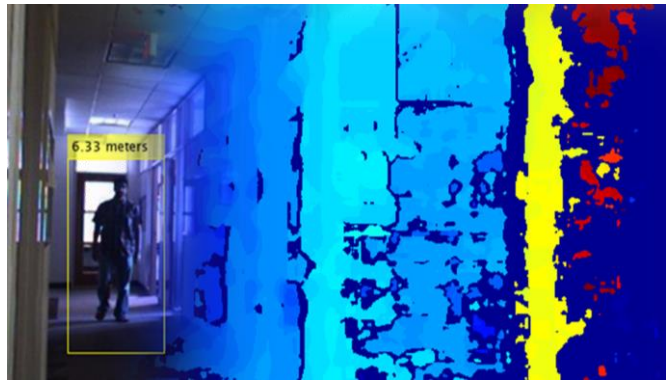

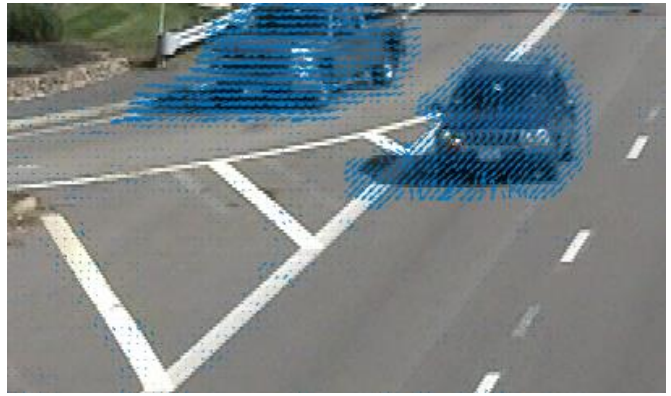

Deep Learning

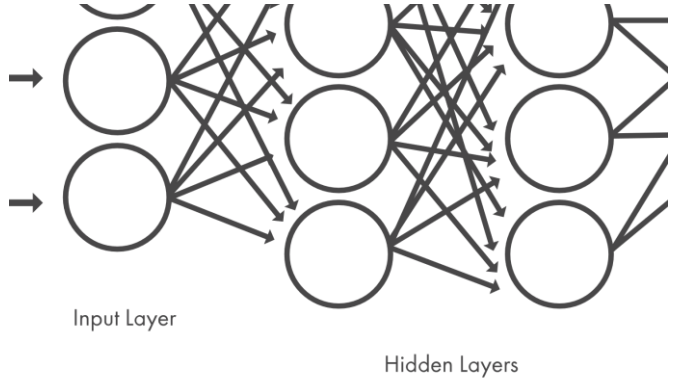

Deep Learning vs Image Processing

The perception:

“Deep learning has made ‘traditional’ image processing obsolete.”

or

“Deep learning needs millions of examples and is only good for classifying pictures of cats anyway.”

Deep Learning and Image Processing

The reality:

Deep learning and image processing are effective tools to solve different problems.

and

Computer Vision tasks are complex, use the right tool for the job.

What are real world problems like?

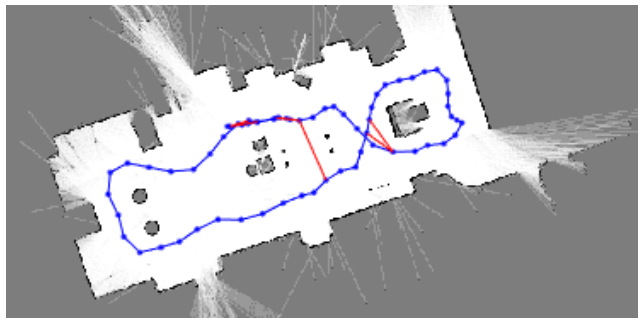

CLEARPATH ROBOTICS™

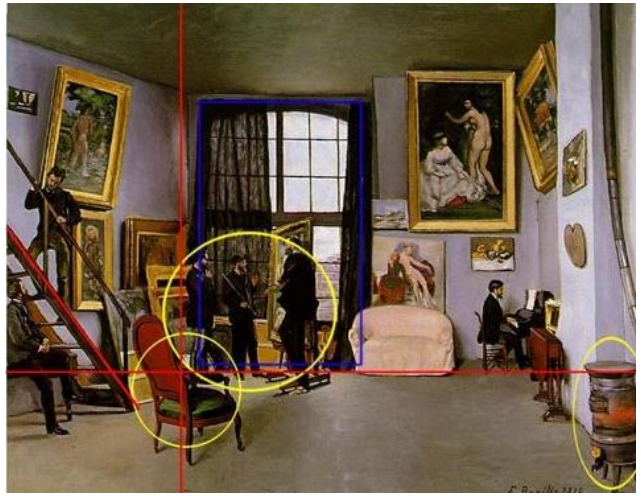

RUTGERS
THE STATE UNIVERSITY OF NEW JERSEY

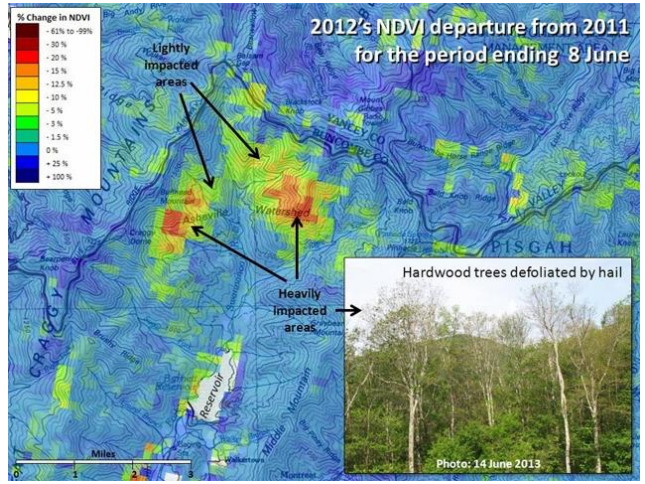

NASA

An Example

5	3			7				
6			1	9	5			
	9	8					6	
8				6				3
4			8		3			1
7				2				6
	6					2	8	
			4	1	9			5
				8			7	9

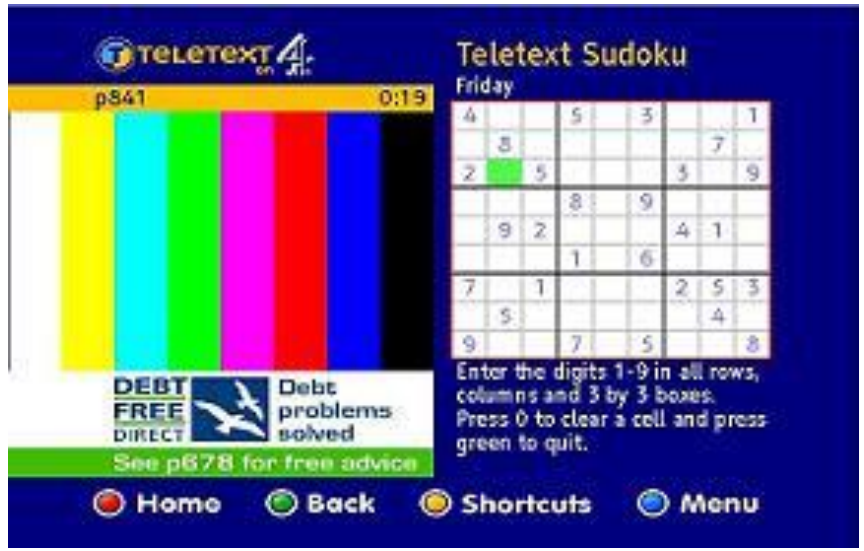

The Sydney Morning Herald

The game's up: jurors playing Sudoku abort trial

By Malcolm Knox
11 June 2008 — 10:00am

AFTER 105 witnesses and three months of evidence, a drug trial costing $1 million was aborted yesterday when it emerged that jurors had been playing Sudoku since the trial's second week

Take a picture → Solve the puzzle

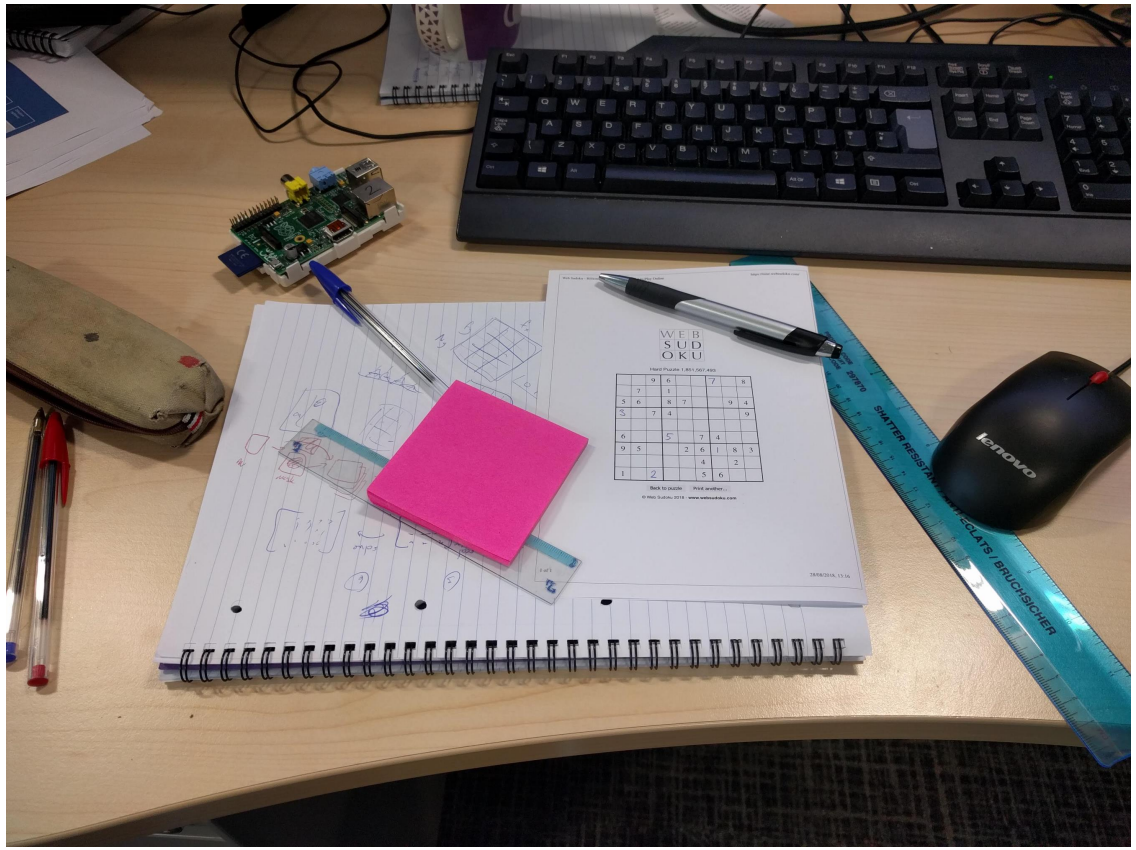

2	4	9	6	5	3	7	1	8
8	7	3	1	4	9	5	6	2
5	6	1	8	7	2	3	9	4
3	1	7	4	6	8	2	5	9
4	9	5	2	3	1	8	7	6
6	2	8	5	9	7	4	3	1
9	5	4	7	2	6	1	8	3
7	8	6	3	1	4	9	2	5
1	3	2	9	8	5	6	4	7

Breaking down the problem

Find the puzzle

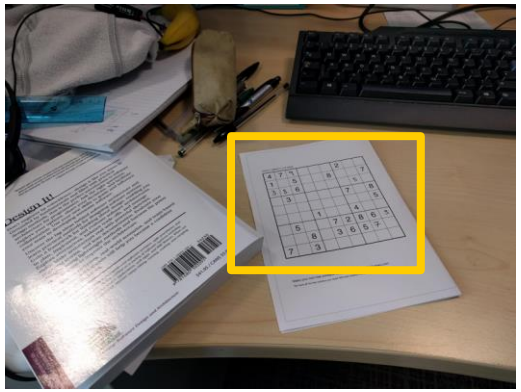

Find the boxes

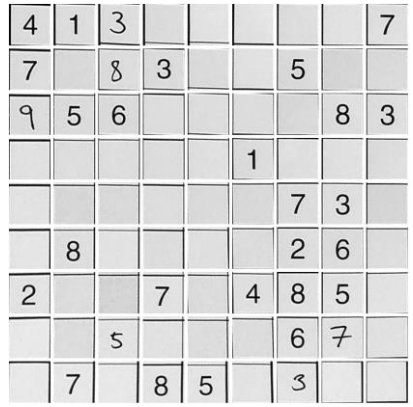

Read the numbers

$$\begin{bmatrix} 4 & 1 & \cdots & ? & ? \\ 7 & ? & \cdots & ? & ? \\ \vdots & \vdots & \cdots & \vdots & \vdots \\ ? & ? & \cdots & 7 & ? \\ ? & 7 & \cdots & ? & ? \end{bmatrix}$$

Solve

$$\begin{bmatrix} 4 & 1 & \cdots & 6 & 8 \\ 7 & 8 & \cdots & 3 & 2 \\ \vdots & \vdots & \cdots & \vdots & \vdots \\ 4 & 2 & \cdots & 7 & 5 \\ 6 & 7 & \cdots & 4 & 1 \end{bmatrix}$$

Deep Learning | **Image Processing** | **Deep Learning** | **Optimisation**

Step 1 – Find the puzzle

Find the puzzle

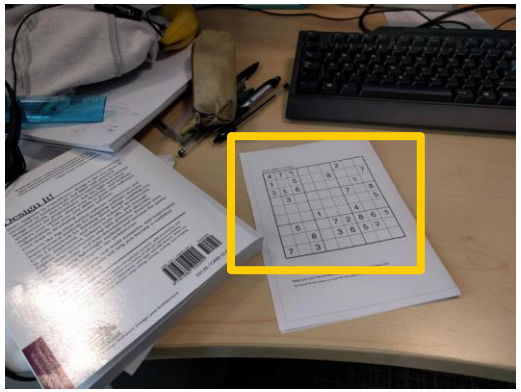

Find the boxes

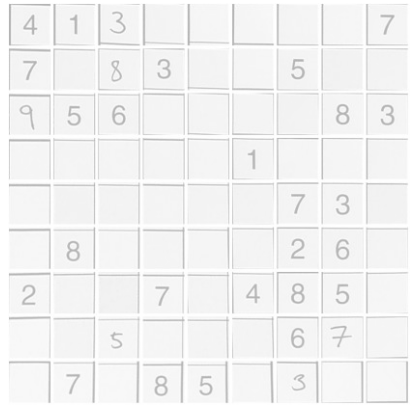

Read the numbers

$$\begin{bmatrix} 4 & 1 & \cdots & ? & ? \\ 7 & ? & \cdots & ? & ? \\ \vdots & \vdots & \cdots & \vdots & \vdots \\ ? & ? & \cdots & 7 & ? \\ ? & 7 & \cdots & ? & ? \end{bmatrix}$$

Solve

$$\begin{bmatrix} 4 & 1 & \cdots & 6 & 8 \\ 7 & 8 & \cdots & 3 & 2 \\ \vdots & \vdots & \cdots & \vdots & \vdots \\ 4 & 2 & \cdots & 7 & 5 \\ 6 & 7 & \cdots & 4 & 1 \end{bmatrix}$$

Step 1 – Find the puzzle

- Varied background and object of interest
- Uncontrolled image acquisition
- Good problem for deep learning

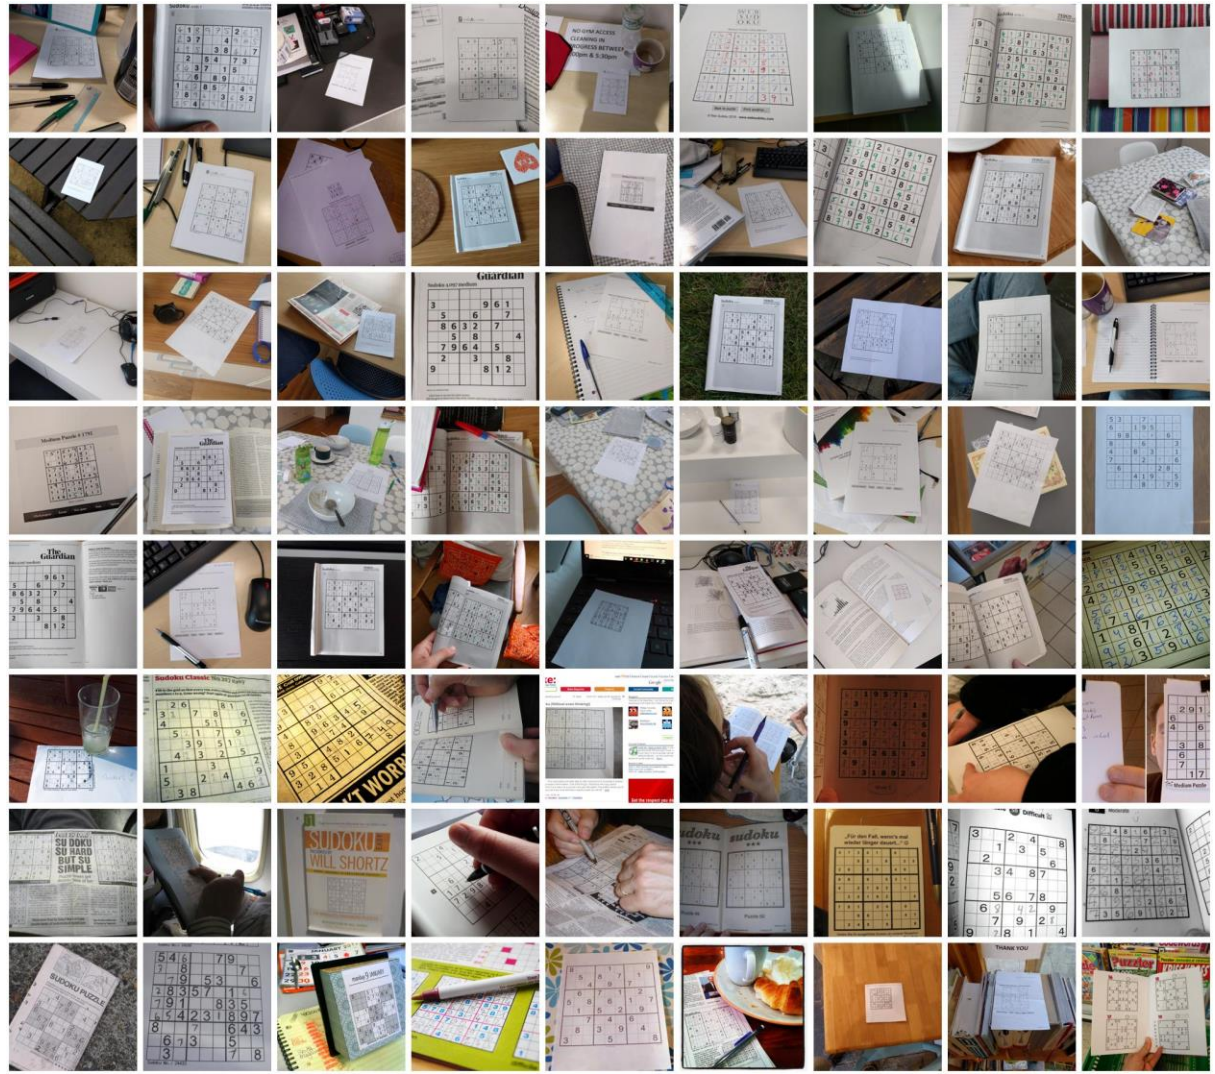

Finding things with deep learning

Object detection

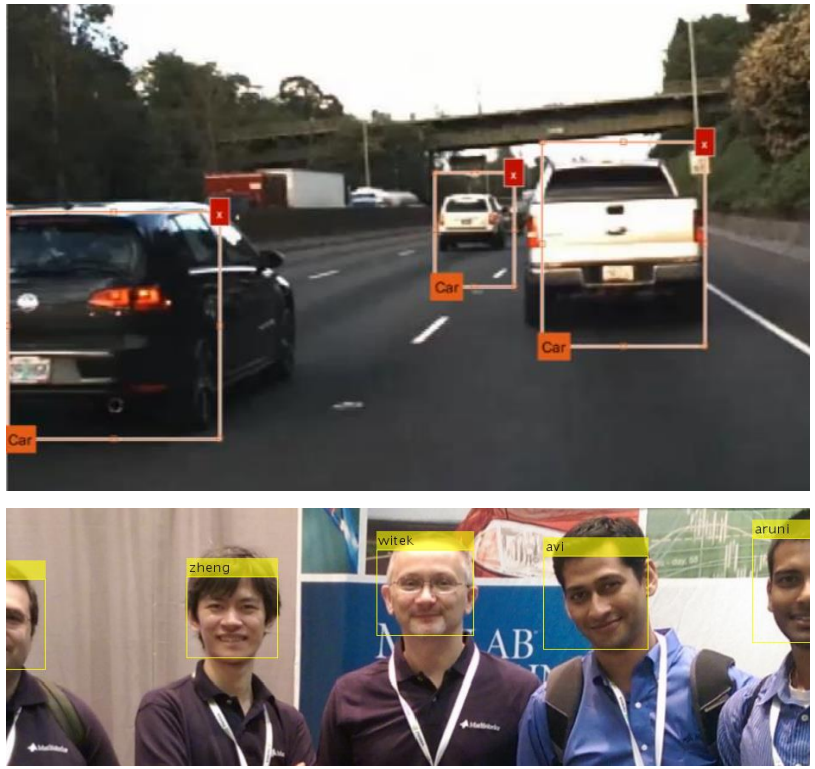

Semantic segmentation

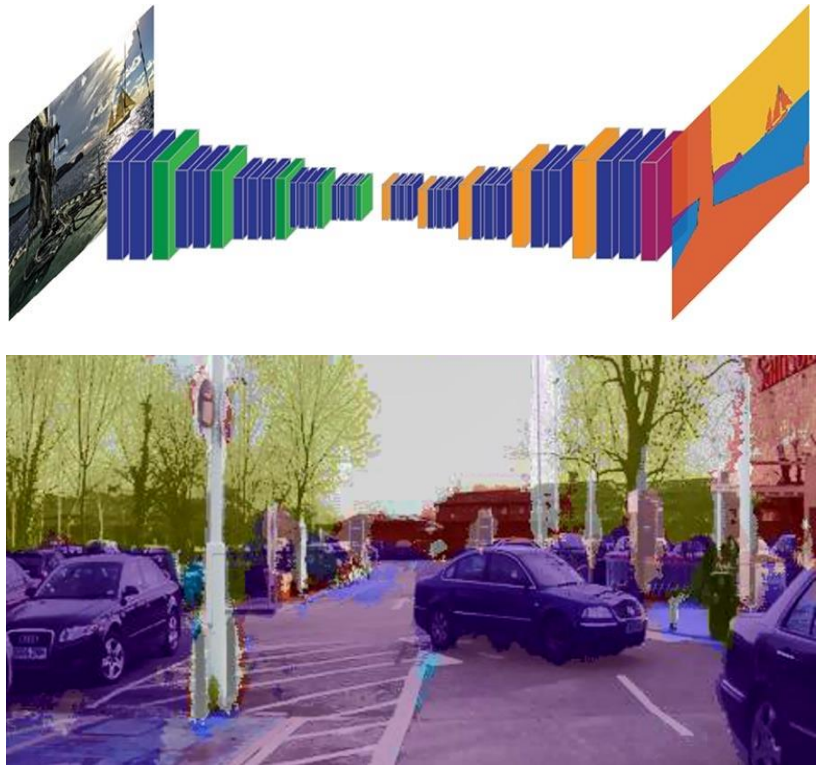

Image Labeler

LABEL

New Session | Load | Save | Import Labels | Label | Zoom In | Zoom Out | Pan | Default Layout | Show Rectangle Labels | Show Scene Labels | Algorithm: Select Algorithm | Automate | View Label Summary | Export Labels

FILE | MODE | VIEW | AUTOMATE LABELING | SUMMARY | EXPORT

ROI Label Definition

Define New ROI Label

To label an ROI, you must first define one or more of the following label types:

- Rectangle label
- Pixel label

Scene Label Definition

Define new scene label

Apply to Image

Remove from Image

To label a scene, you must first define a scene label.

Image

Load images to start labeling.


SUDOKU
WILL SHORTZ
SU DOKU
SU HARD
BUT SU
SIMPLE
SUDOKU PUZZLES

How do we train a semantic segmentation network?

- Manage input data and apply pre-processing

```
%% Input data
inputSize = [512, 512, 3];
train = pixelLabelImageDatastore(imagesTrain, labelsTrain, ...
                                 'OutputSize', inputSize(1:2));
test = pixelLabelImageDatastore(imagesTest, labelsTest, ...
                                'OutputSize', inputSize(1:2));
```

- Use a well established network architecture configured for our problem

```
%% Setup the network
numClasses = 2;
baseNetwork = 'vgg16';
layers = segnetLayers(inputSize, numClasses, baseNetwork);
layers = sudoku.weightLossByFrequency(layers, train);
```

- Set up training options

```
%% Set up the training options
opts = trainingOptions('sgdm', ...
                       'InitialLearnRate', 0.01, ...
                       'ValidationData', test, ...
                       'MaxEpochs', 10, ...
                       'Plots', 'training-progress');
```

- Train the network

```
%% Train
net = trainNetwork(train, layers, opts);
```

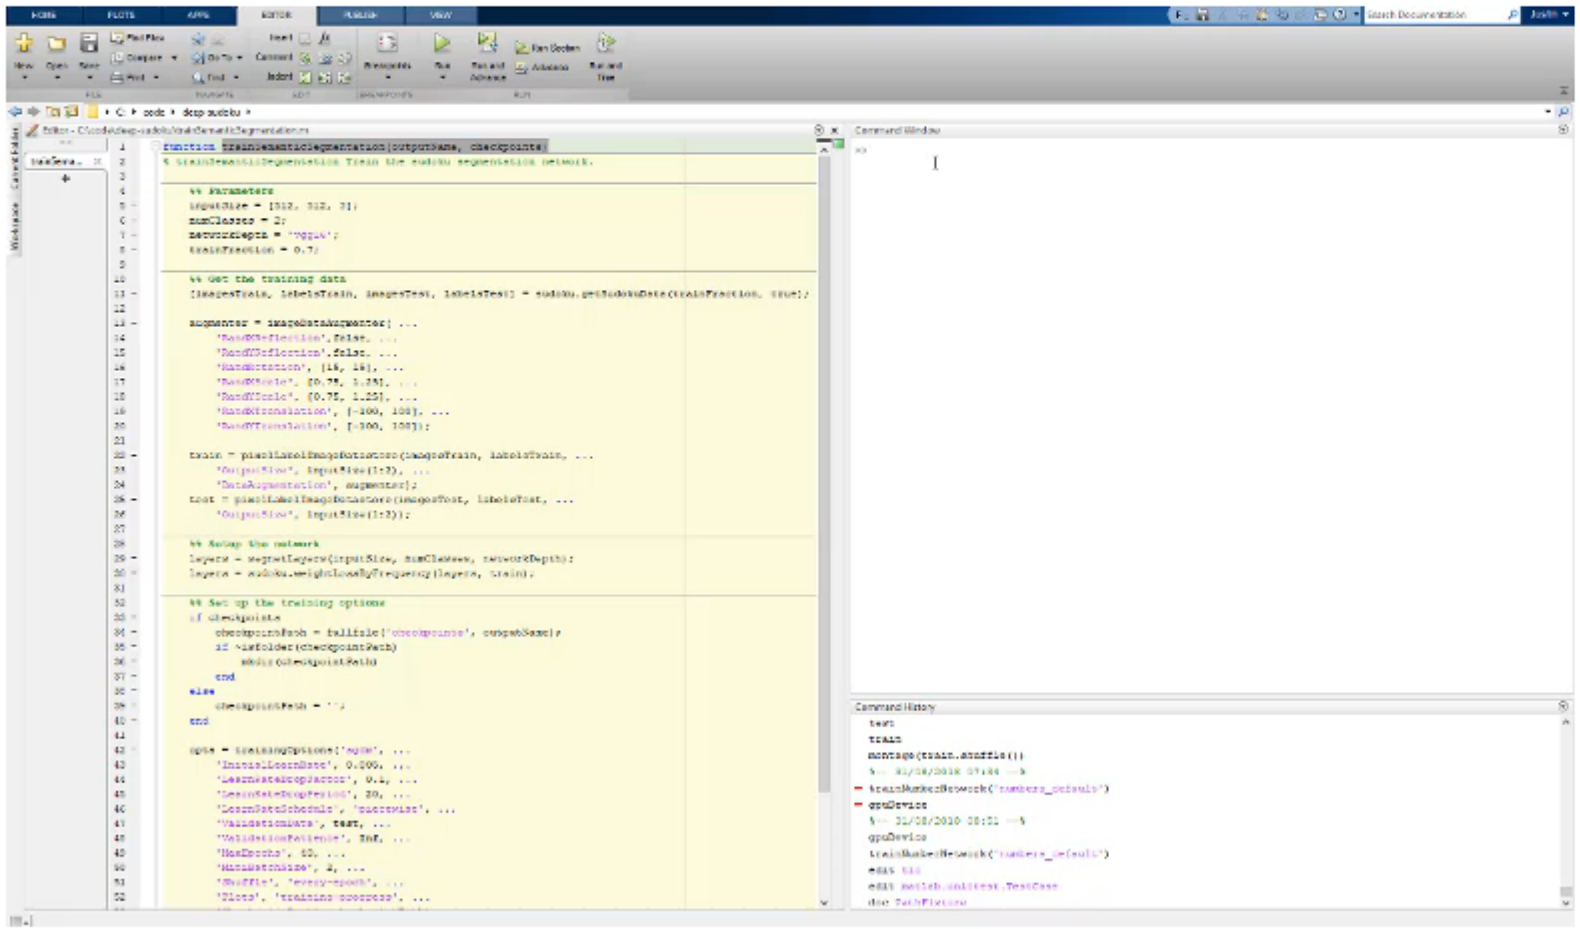
HOME
PLOTS
APPS
EDITOR
PUBLISH
VIEW
Search Documentation
Run Section
Advance
Run and Time
C: ▸ code ▸ deep sudoku ▸
Editor - C:\code\deep-sudoku\trainSemanticSegmentation.m
function trainSemanticSegmentation(outputName, checkpoint)
% trainSemanticSegmentation Train the sudoku segmentation network.
%% Parameters
inputSize = [512, 512, 3];
numClasses = 2;
networkDepth = 'mpp4';
trainFraction = 0.7;
%% Get the training data
[imagesTrain, labelsTrain, imagesTest, labelsTest] = sudoku.getSudokuData(trainFraction, true);
augmenter = imageDataAugmenter(...
'RandXReflection',false, ...
'RandYReflection',false, ...
'RandRotation', [1, 10], ...
'RandXScale', [0.75, 1.25], ...
'RandYScale', [0.75, 1.25], ...
'RandXTranslation', [-100, 100], ...
'RandYTranslation', [-100, 100]);
train = pixelLabelImageDatastore(imagesTrain, labelsTrain, ...
'OutputSize', inputSize(1:2), ...
'DataAugmentation', augmenter);
test = pixelLabelImageDatastore(imagesTest, labelsTest, ...
'OutputSize', inputSize(1:2));
%% Setup the network
layers = segnetLayers(inputSize, numClasses, networkDepth);
layers = sudoku.weightClassByFrequency(layers, train);
%% Set up the training options
if checkpoint
checkpointPath = fullfile('checkpoints', outputName);
if ~isfolder(checkpointPath)
mkdir(checkpointPath)
end
else
checkpointPath = '';
end
opts = trainingOptions('sgdm', ...
'InitialLearnRate', 0.001, ...
'LearnRateDropFactor', 0.1, ...
'LearnRateDropPeriod', 20, ...
'LearnRateSchedule', 'piecewise', ...
'ValidationData', test, ...
'ValidationPatience', Inf, ...
'MaxEpochs', 40, ...
'MiniBatchSize', 2, ...
'Shuffle', 'every-epoch', ...
'Plots', 'training-progress', ...
Command Window
>>
Command History
test
train
montage(train.shuffle())
trainNumberNetwork('numbers_default')
gpuDevice
gpuDevice
trainNumberNetwork('numbers_default')
clc
matlab.unittest.TestCase
FatFixture


The Guardian
Sudoku

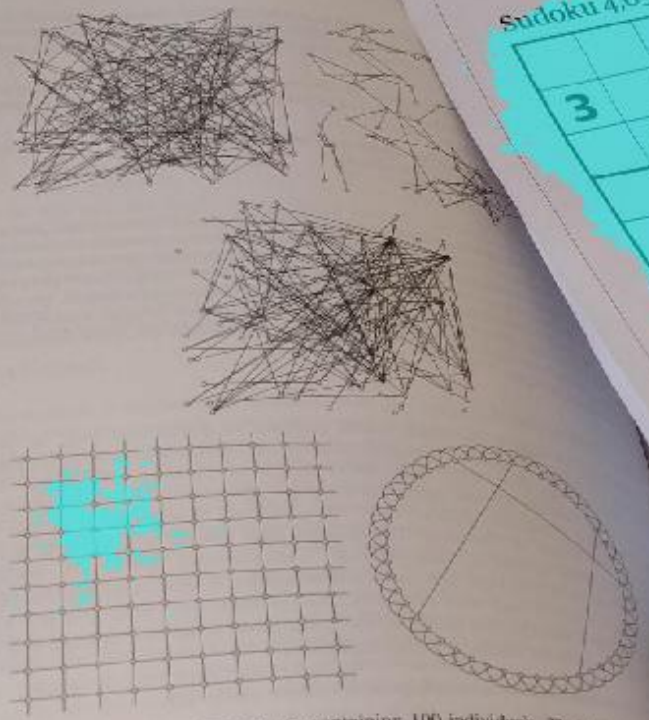

Figure 7.18. Five distinct network types containing 100 individuals. These are from left to r
Random, Spatial (top row), Scale-free (middle row), Lattice, and Small-World (bottom row
Random, Spatial, and Scale-Free networks all use the same position of individuals—although
Random and Scale-Free network, the position of the individuals is irrelevant for forming
In all five graphs, the average number of contacts per individual is approximately 4. For the
network, individuals with high numbers of contacts are shaded gray.

The Guardian

Sudoku 4,097 medium

				9	6	1		
3				6		7		
	5				7			4
	8	6	3	2				
		5		8		5		
	7	9	6	4			8	
		2			3		8	1 2
9								

Click here to access the print version

Fill the grid so that every row, every column and every 3×3 box contains the numbers 1 to 9.

Buy the next issue of the Guardian or subscribe to our Digital Edition to see the completed puzzle.

the vast majority of networks that have been studied, the number of contacts per individual is very heterogeneous, with most individuals having a relatively small number and a few have many contacts (Albert et al. 1999; Barabási and Albert 1999; Jeong et al. 2000; Liljeros et al. 2001). Because the most connected individuals are likely to disproportionately important in disease transmission (see Chapter 3), networks that capture this heterogeneity are therefore vital in understanding the spread of and infections—scale-free networks incorporate these heterogeneities. Scale-free networks are generally created dynamically, adding new individuals to a network one at a time with a preferential mechanism that mimics the natural formation of social contacts. Each new

An Example – Sudoku Solver

Find the puzzle

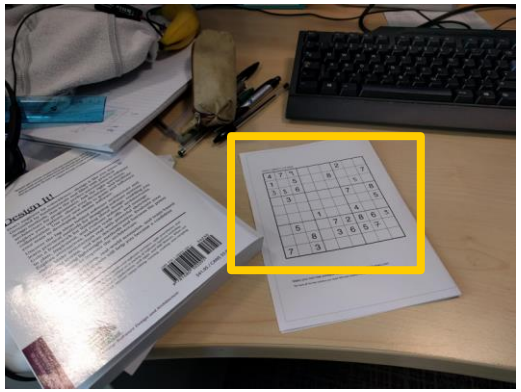

Find the boxes

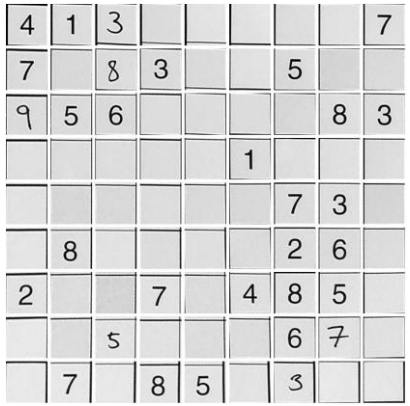

Read the numbers

$$\begin{bmatrix} 4 & 1 & \cdots & ? & ? \\ 7 & ? & \cdots & ? & ? \\ \vdots & \vdots & \cdots & \vdots & \vdots \\ ? & ? & \cdots & 7 & ? \\ ? & 7 & \cdots & ? & ? \end{bmatrix}$$

Solve

$$\begin{bmatrix} 4 & 1 & \cdots & 6 & 8 \\ 7 & 8 & \cdots & 3 & 2 \\ \vdots & \vdots & \cdots & \vdots & \vdots \\ 4 & 2 & \cdots & 7 & 5 \\ 6 & 7 & \cdots & 4 & 1 \end{bmatrix}$$

Deep Learning

Step 2 – Find the boxes

Find the puzzle

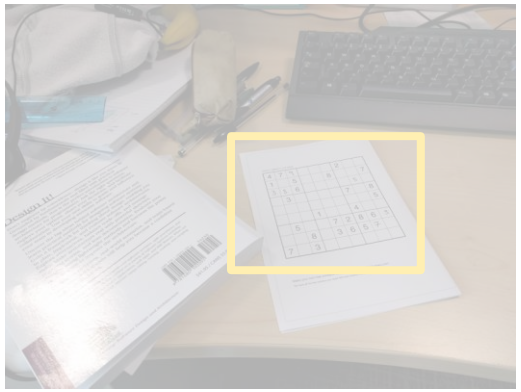

Find the boxes

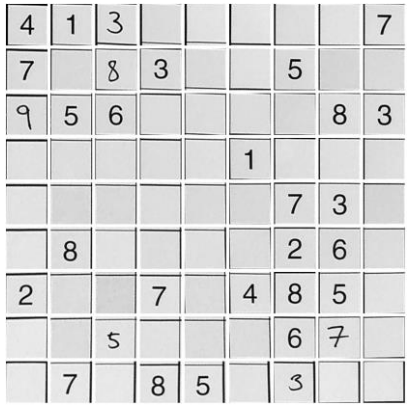

Read the numbers

$$\begin{bmatrix} 4 & 1 & \cdots & ? & ? \\ 7 & ? & \cdots & ? & ? \\ \vdots & \vdots & \cdots & \vdots & \vdots \\ ? & ? & \cdots & 7 & ? \\ ? & 7 & \cdots & ? & ? \end{bmatrix}$$

Solve

$$\begin{bmatrix} 4 & 1 & \cdots & 6 & 8 \\ 7 & 8 & \cdots & 3 & 2 \\ \vdots & \vdots & \cdots & \vdots & \vdots \\ 4 & 2 & \cdots & 7 & 5 \\ 6 & 7 & \cdots & 4 & 1 \end{bmatrix}$$

Step 2 – Find the boxes

- Well defined problem:
  - 4 intersecting straight lines
  - Dark ink on light paper
  - Grid of 9 x 9 equally sized boxes
- Need good precision to localise individual boxes.
- Good image processing problem.

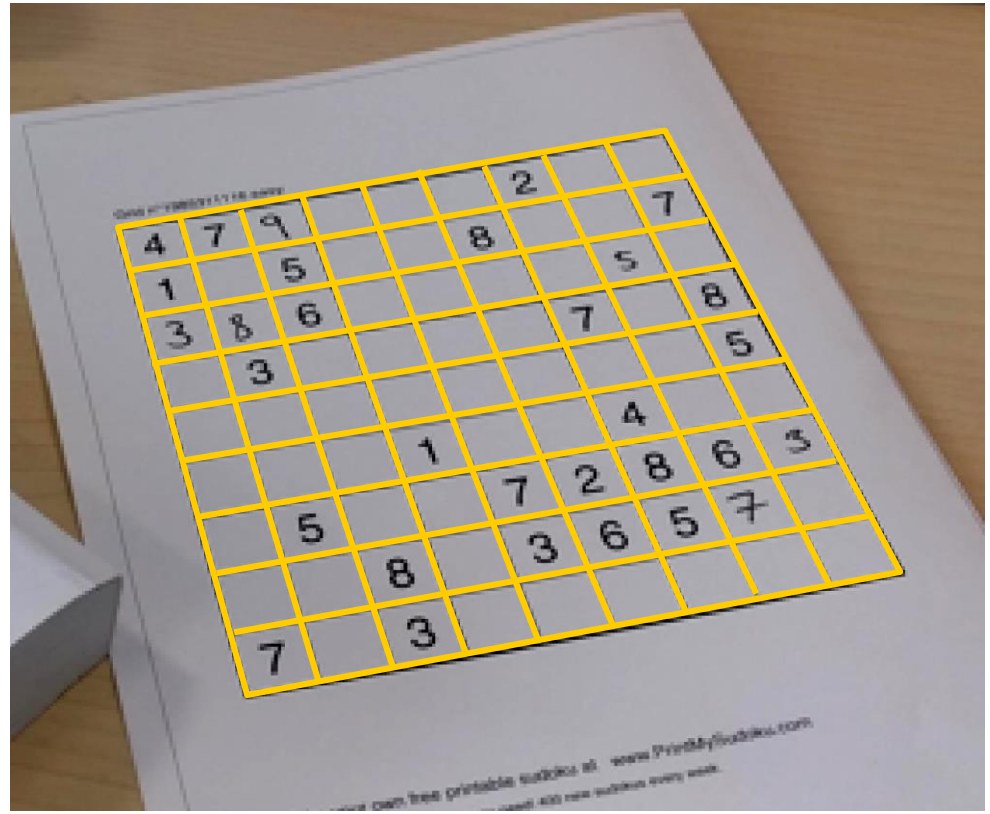

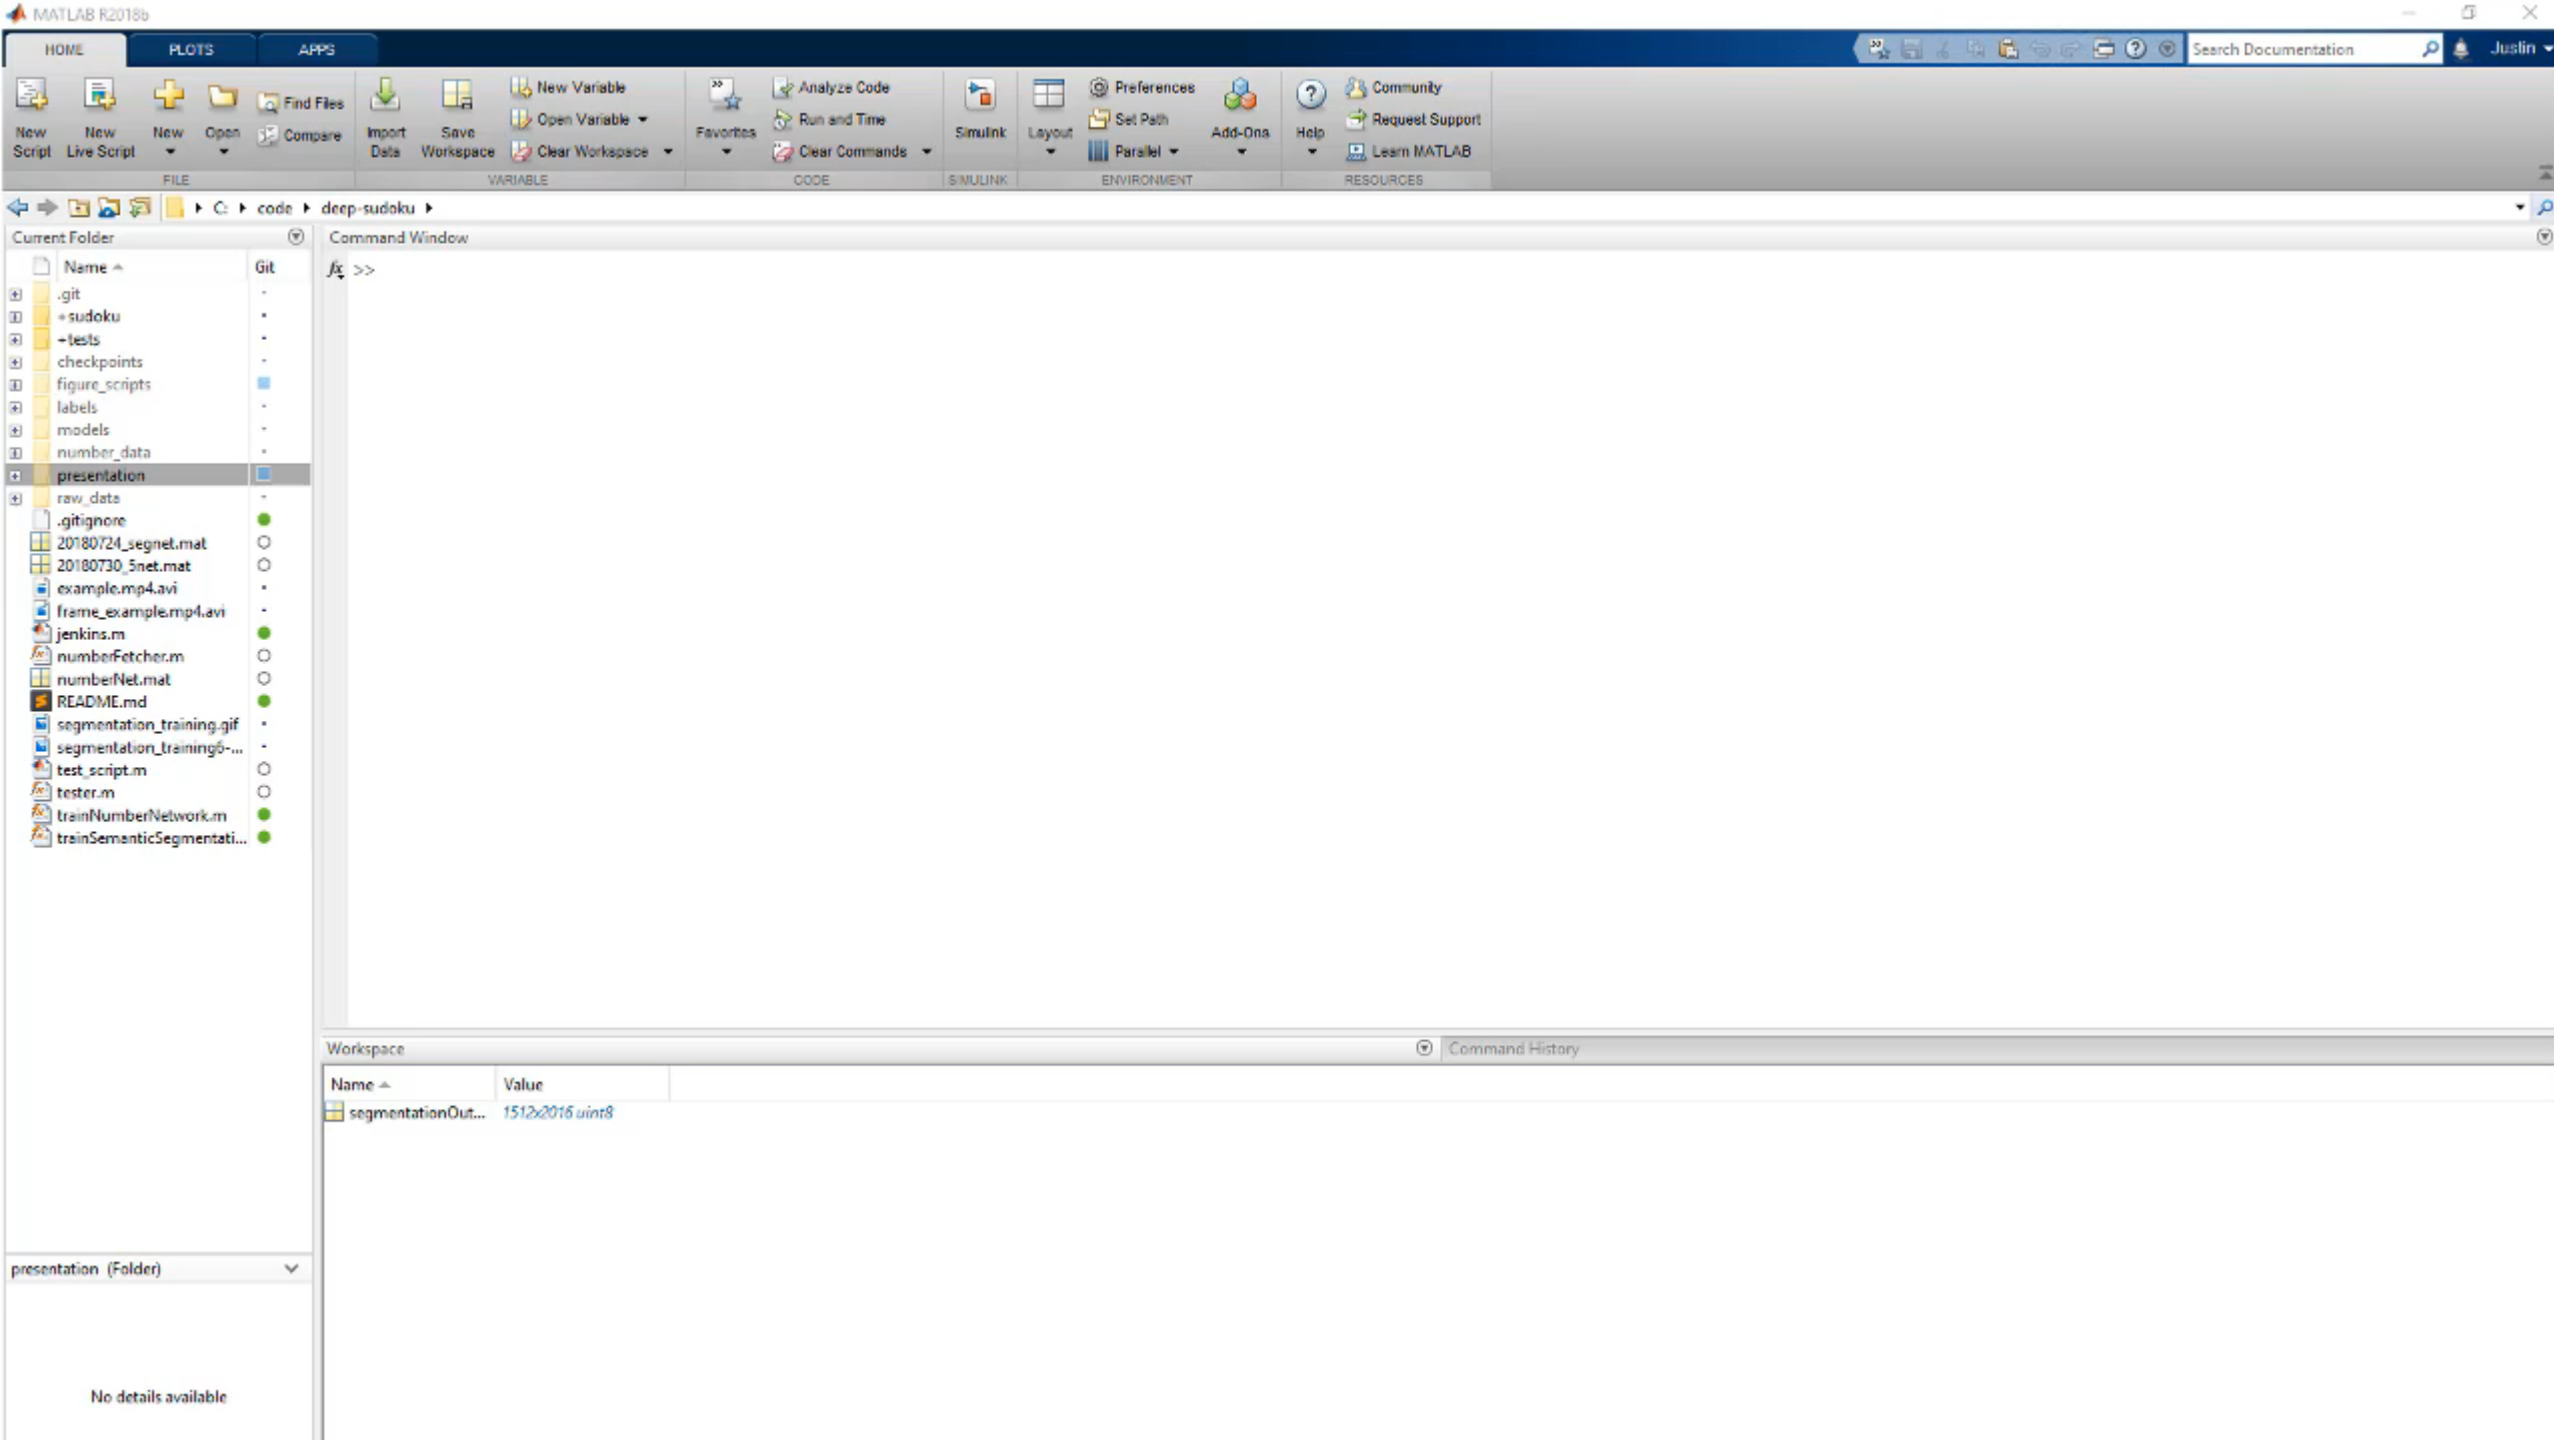
MATLAB R2018b
HOME
PLOTS
APPS
Search Documentation
Justin
New Script
New Live Script
New
Open
Find Files
Compare
Import Data
Save Workspace
New Variable
Open Variable
Clear Workspace
Favorites
Analyze Code
Run and Time
Clear Commands
Simulink
Layout
Preferences
Set Path
Parallel
Add-Ons
Help
Community
Request Support
Learn MATLAB
FILE
VARIABLE
CODE
SIMULINK
ENVIRONMENT
RESOURCES
C: ▸ code ▸ deep-sudoku ▸
Current Folder
Name
Git
.git
+sudoku
+tests
checkpoints
figure_scripts
labels
models
number_data
presentation
raw_data
.gitignore
20180724_segnet.mat
20180730_5net.mat
example.mp4.avi
frame_example.mp4.avi
jenkins.m
numberFetcher.m
numberNet.mat
README.md
segmentation_training.gif
segmentation_training5-...
test_script.m
tester.m
trainNumberNetwork.m
trainSemanticSegmentati...
Command Window
fx >>
Workspace
Command History
Name
Value
segmentationOut...
1512x2016 uint8
presentation (Folder)
No details available

Step 2 – Find the boxes

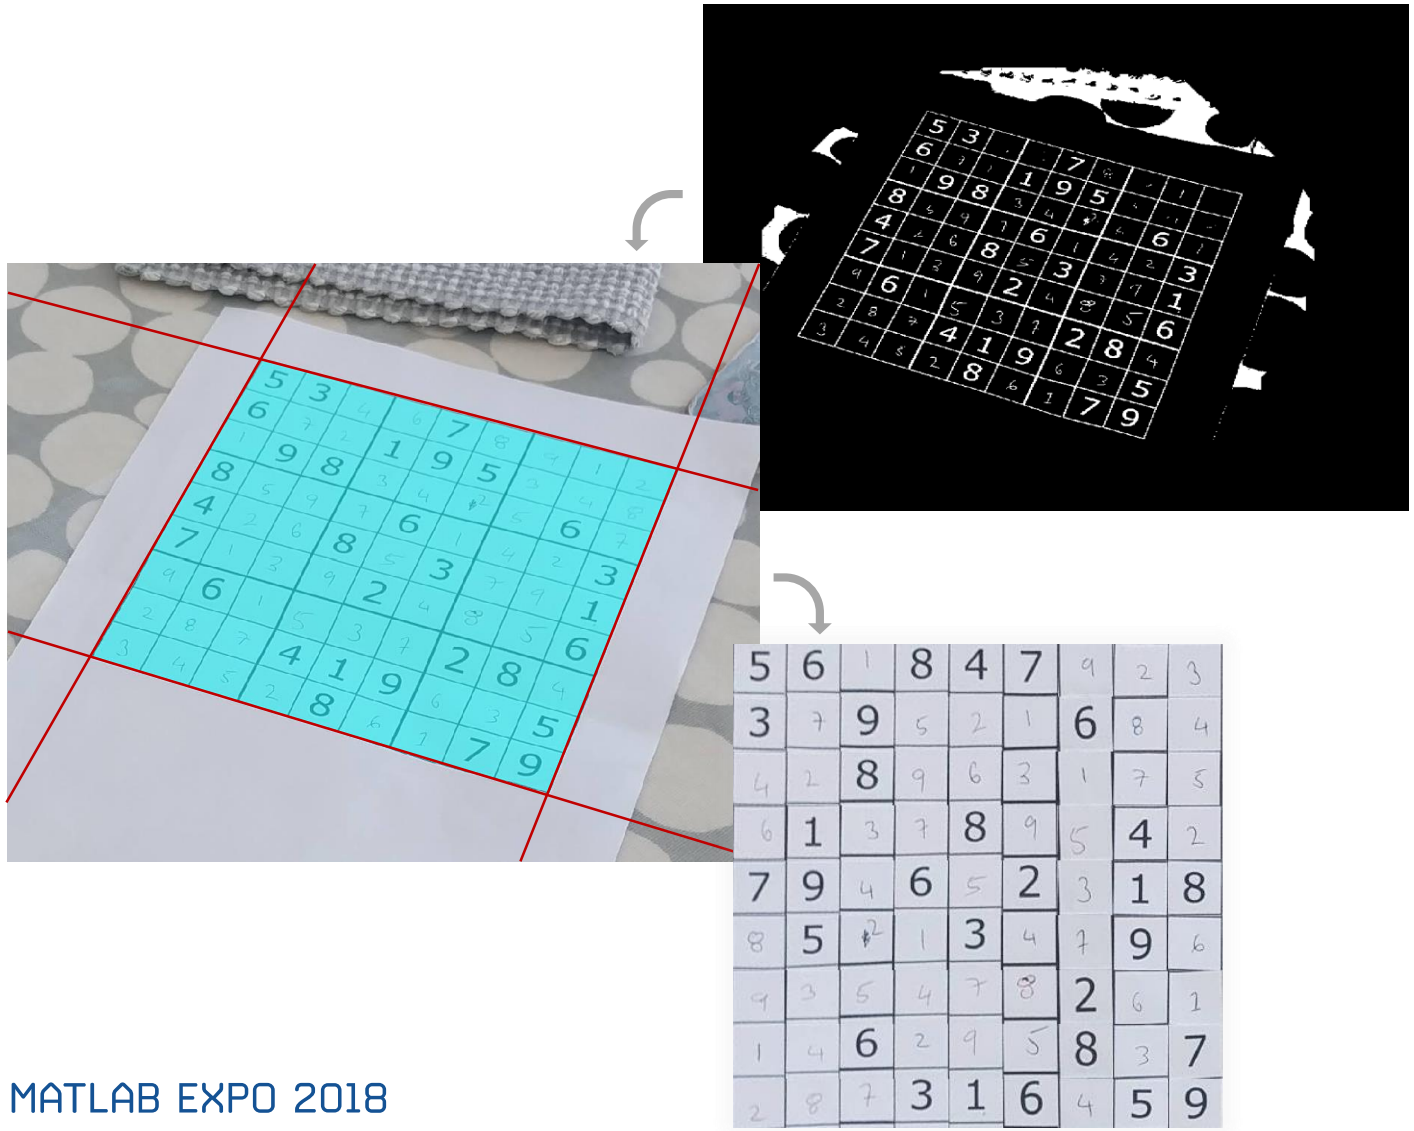

Image processing pipeline

Adaptive thresholding
`imbinarize`
↓
Morphological operations
`imopen/imclose`
↓
Region property analysis
`regionprops`
↓
Robust line detection
`hough/houghpeaks`
↓
Geometric warping
`fitgeotrans/imwarp`

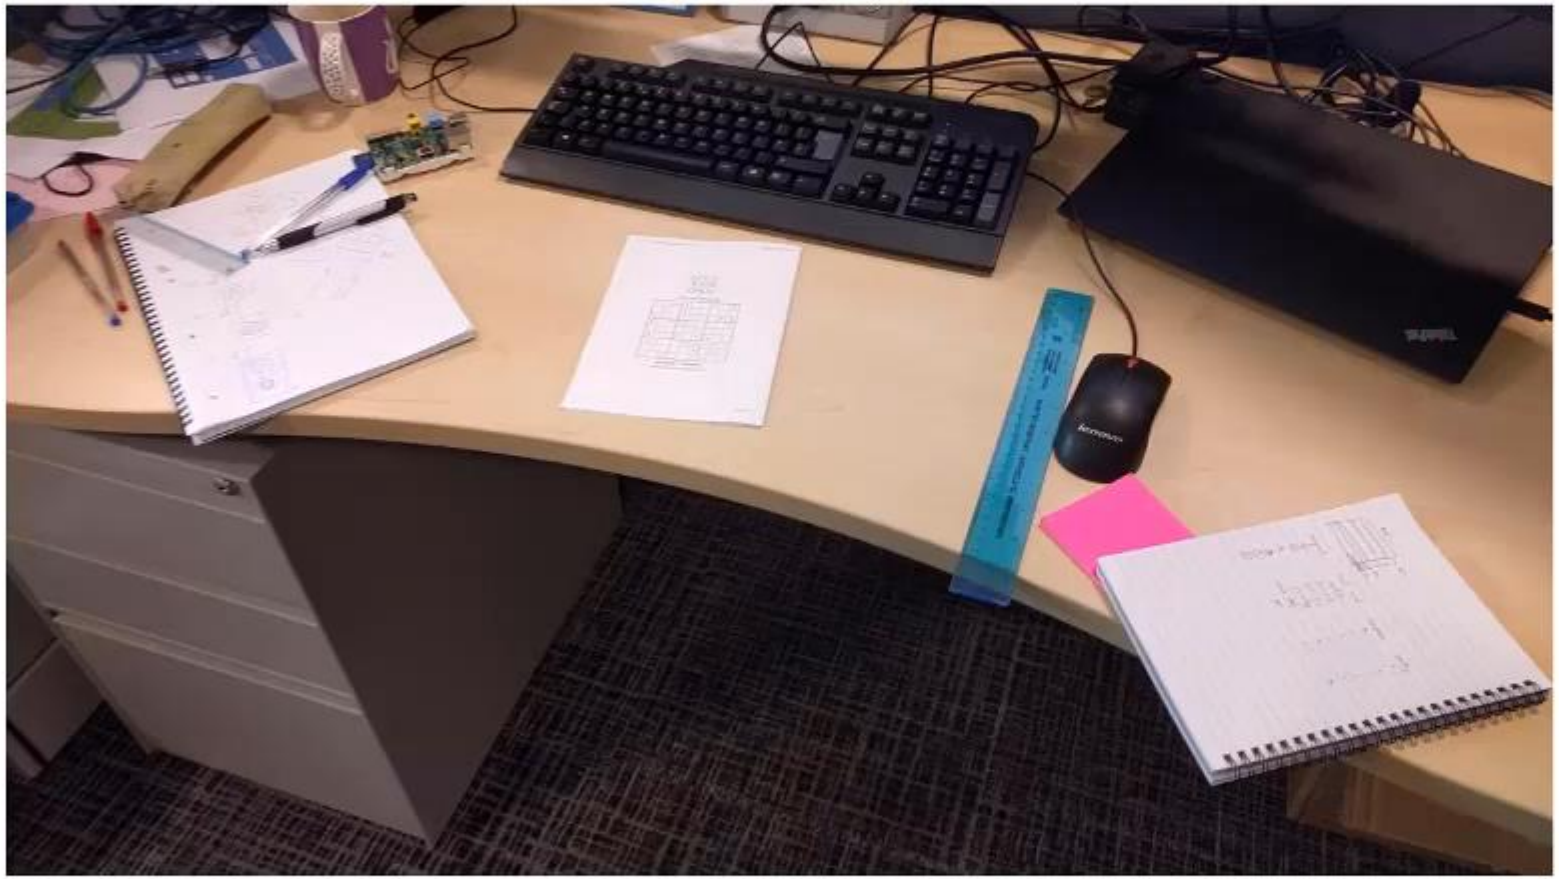

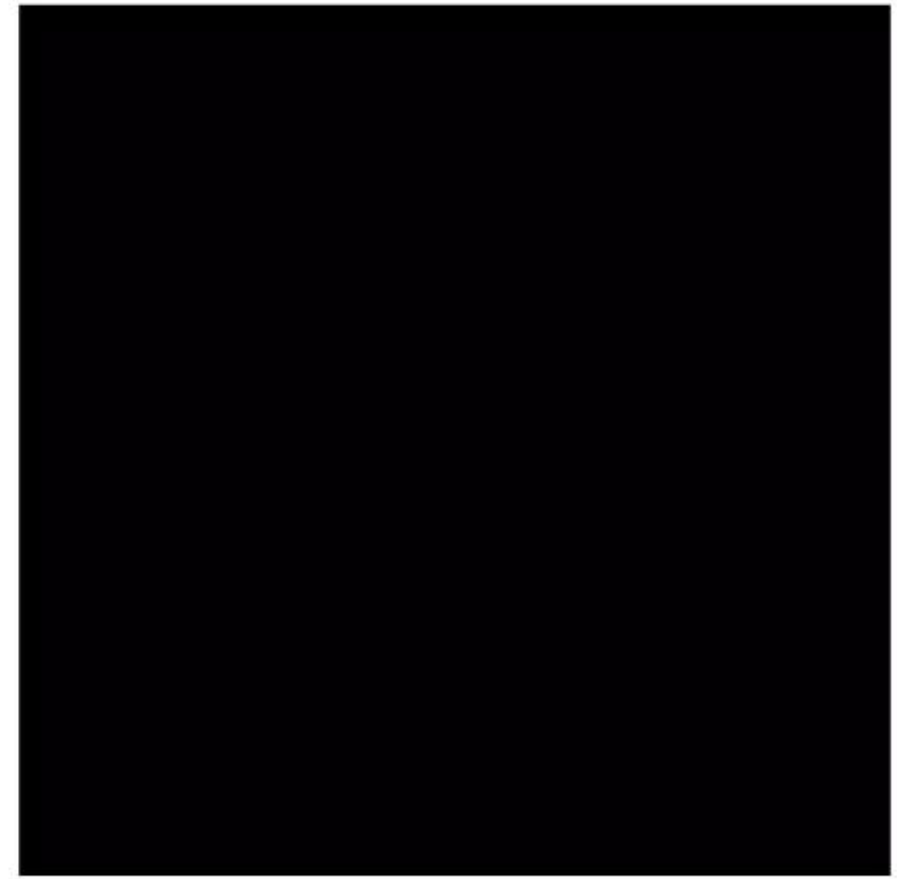

An Example – Sudoku Solver

Find the puzzle

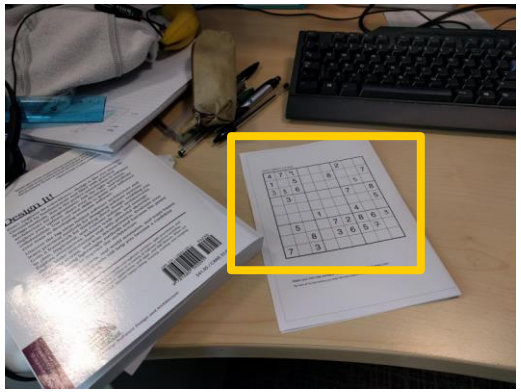

Find the boxes

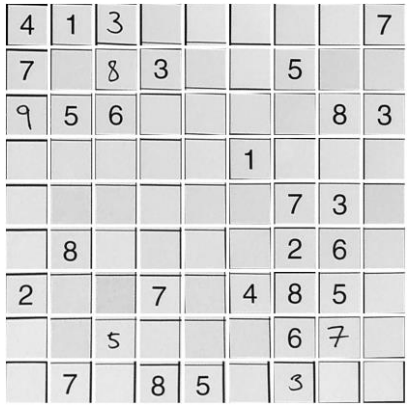

Read the numbers

$$\begin{bmatrix} 4 & 1 & \cdots & ? & ? \\ 7 & ? & \cdots & ? & ? \\ \vdots & \vdots & \cdots & \vdots & \vdots \\ ? & ? & \cdots & 7 & ? \\ ? & 7 & \cdots & ? & ? \end{bmatrix}$$

Solve

$$\begin{bmatrix} 4 & 1 & \cdots & 6 & 8 \\ 7 & 8 & \cdots & 3 & 2 \\ \vdots & \vdots & \cdots & \vdots & \vdots \\ 4 & 2 & \cdots & 7 & 5 \\ 6 & 7 & \cdots & 4 & 1 \end{bmatrix}$$

Deep Learning

Image Processing

Step 3 – Read the numbers

Find the puzzle

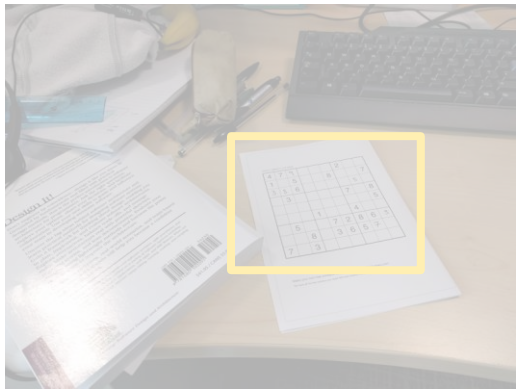

Find the boxes

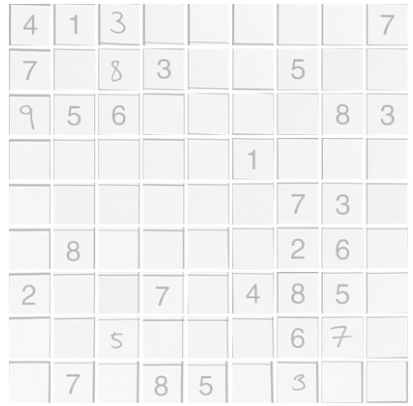

Read the numbers

$$\begin{bmatrix} 4 & 1 & \cdots & ? & ? \\ 7 & ? & \cdots & ? & ? \\ \vdots & \vdots & \cdots & \vdots & \vdots \\ ? & ? & \cdots & 7 & ? \\ ? & 7 & \cdots & ? & ? \end{bmatrix}$$

Solve

$$\begin{bmatrix} 4 & 1 & \cdots & 6 & 8 \\ 7 & 8 & \cdots & 3 & 2 \\ \vdots & \vdots & \cdots & \vdots & \vdots \\ 4 & 2 & \cdots & 7 & 5 \\ 6 & 7 & \cdots & 4 & 1 \end{bmatrix}$$

Step 3 – Read the numbers

How can we read numbers?

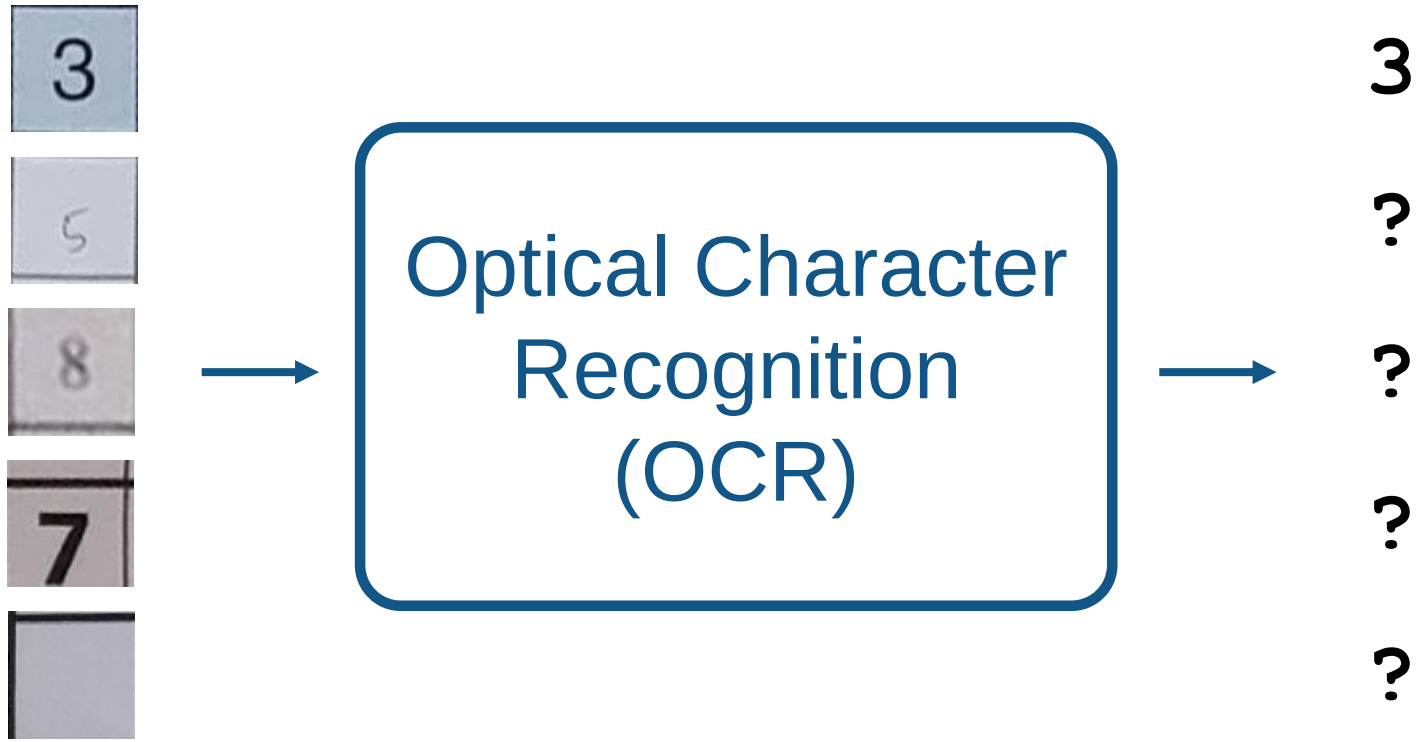

How can we read numbers?

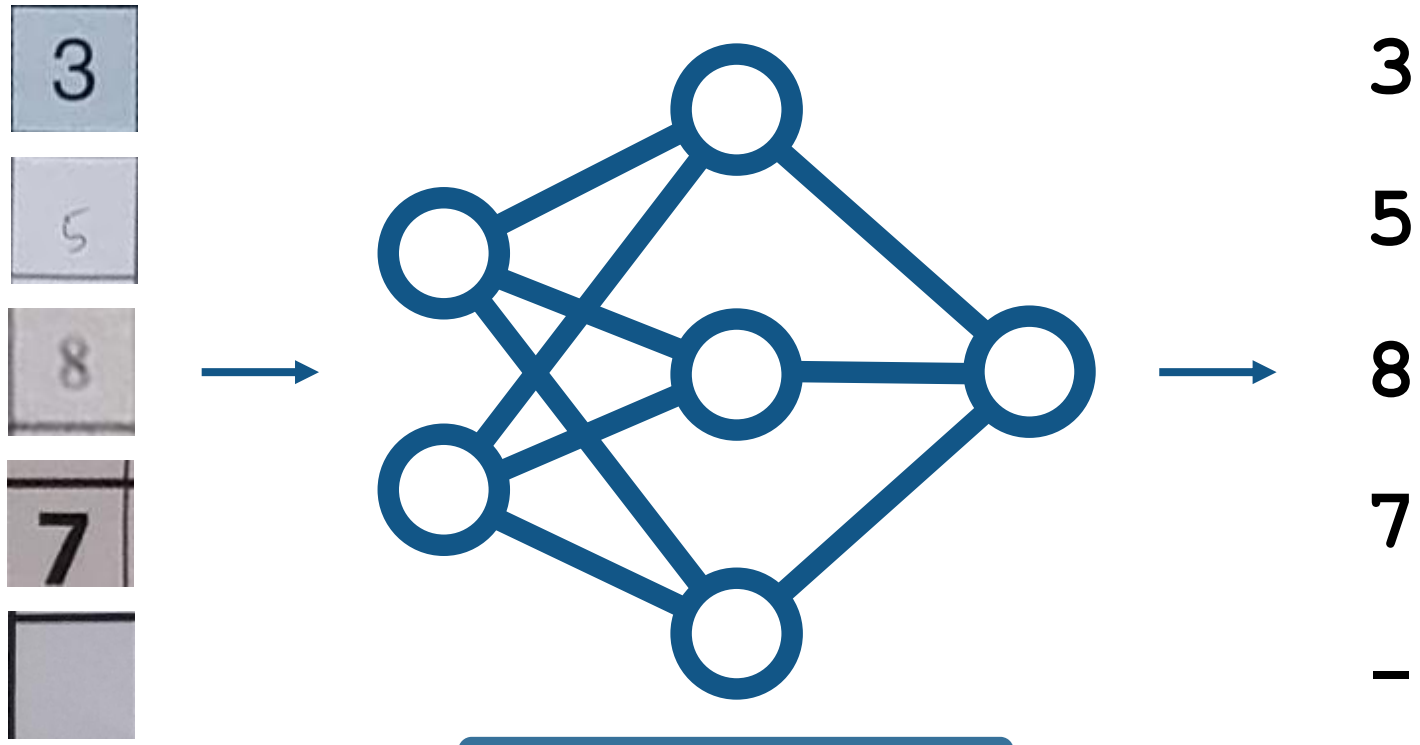

We can make synthetic training data

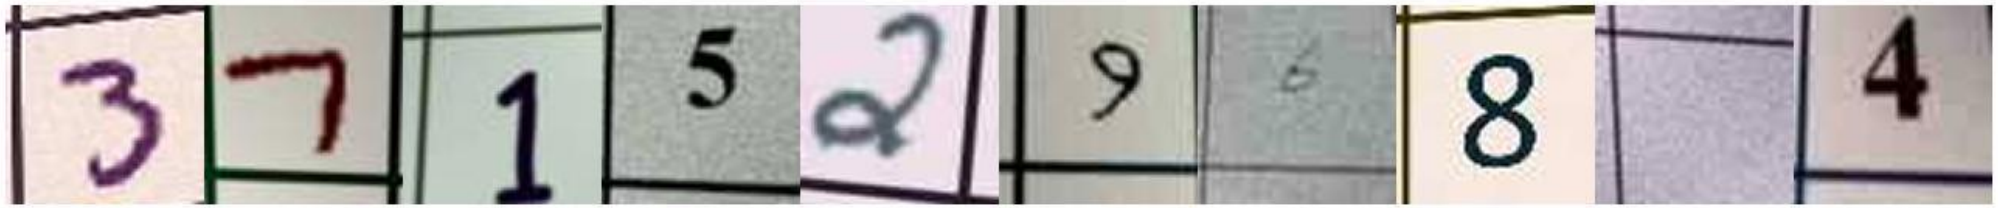

We can make (lots of) synthetic training data

We can make (lots of) synthetic training data

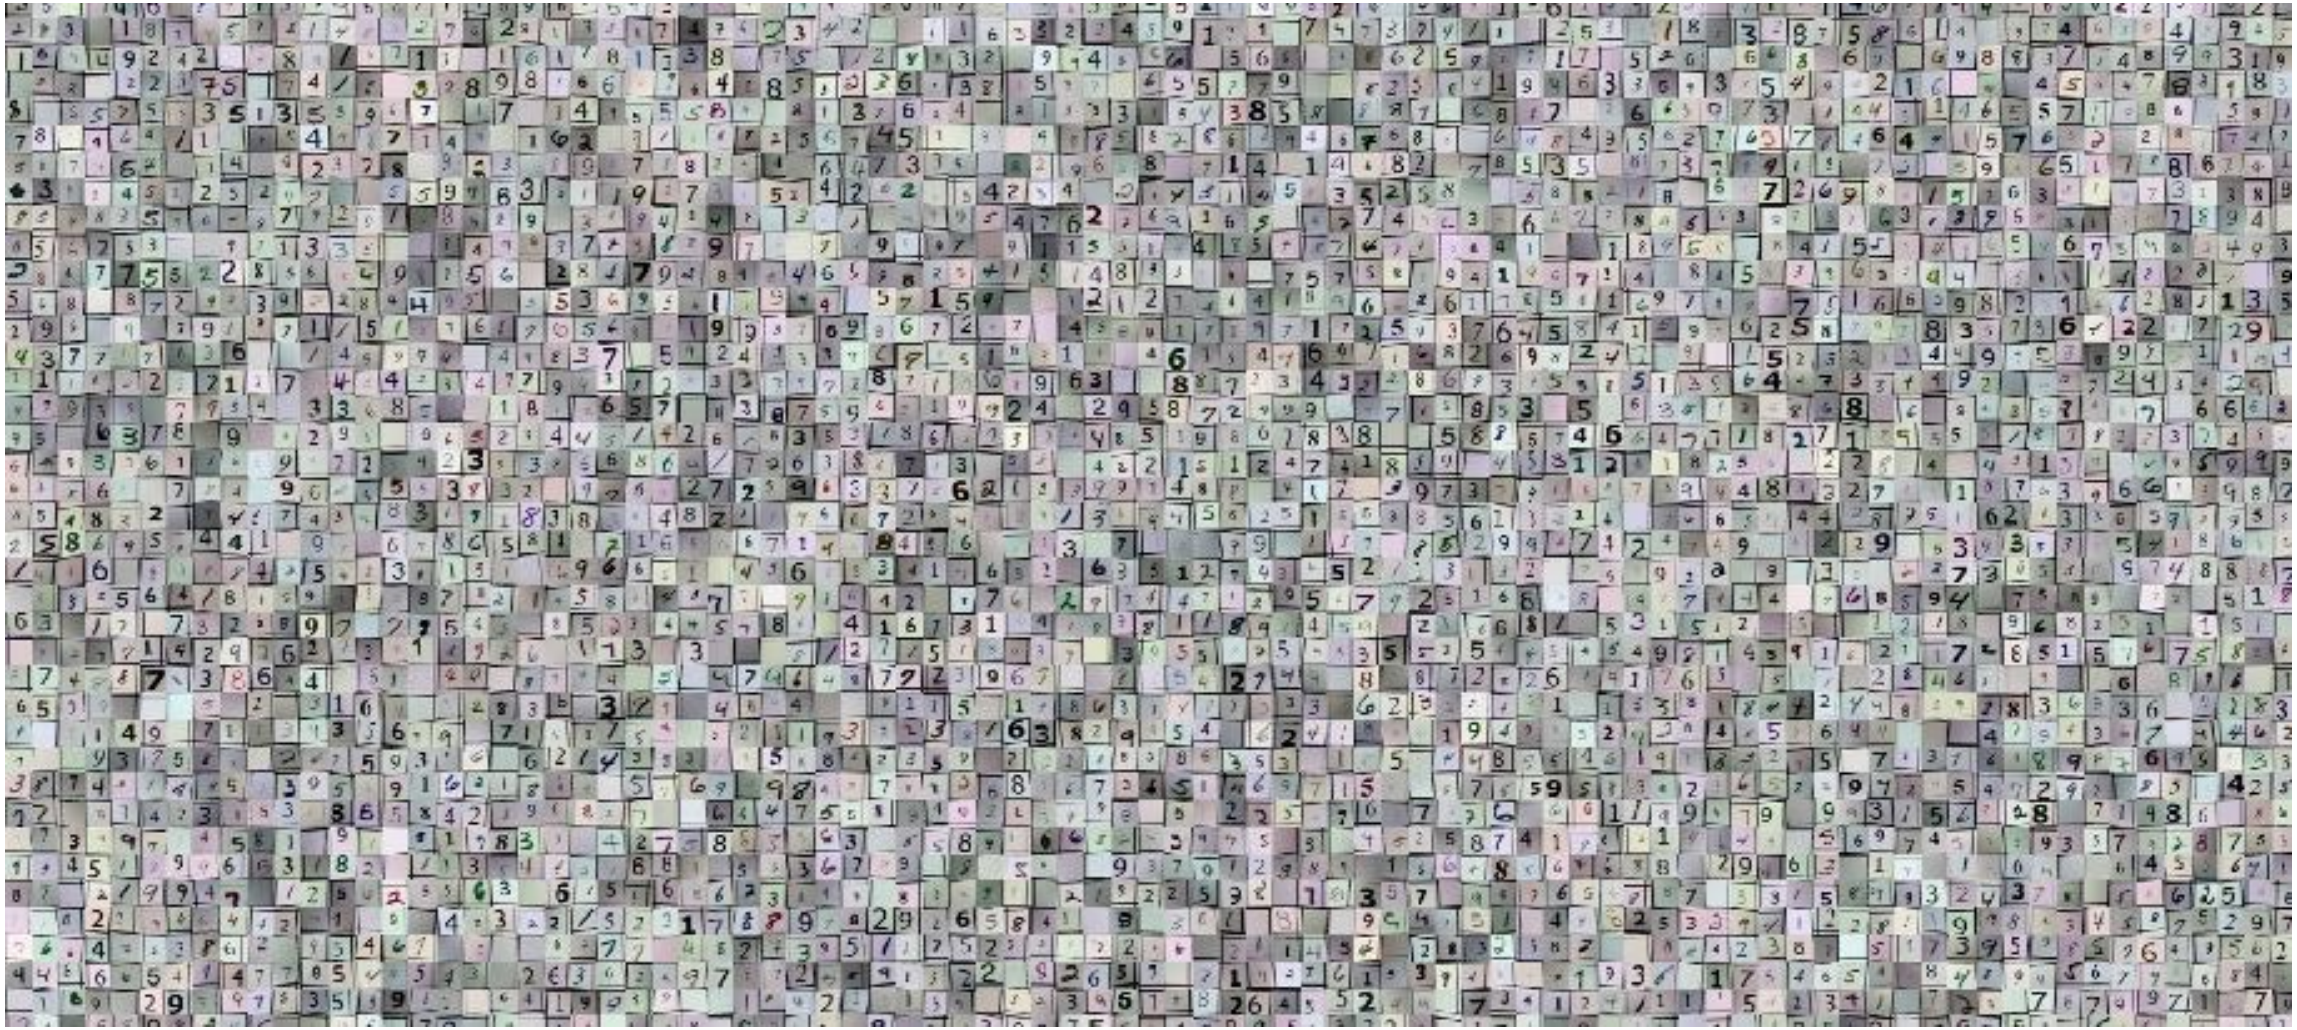

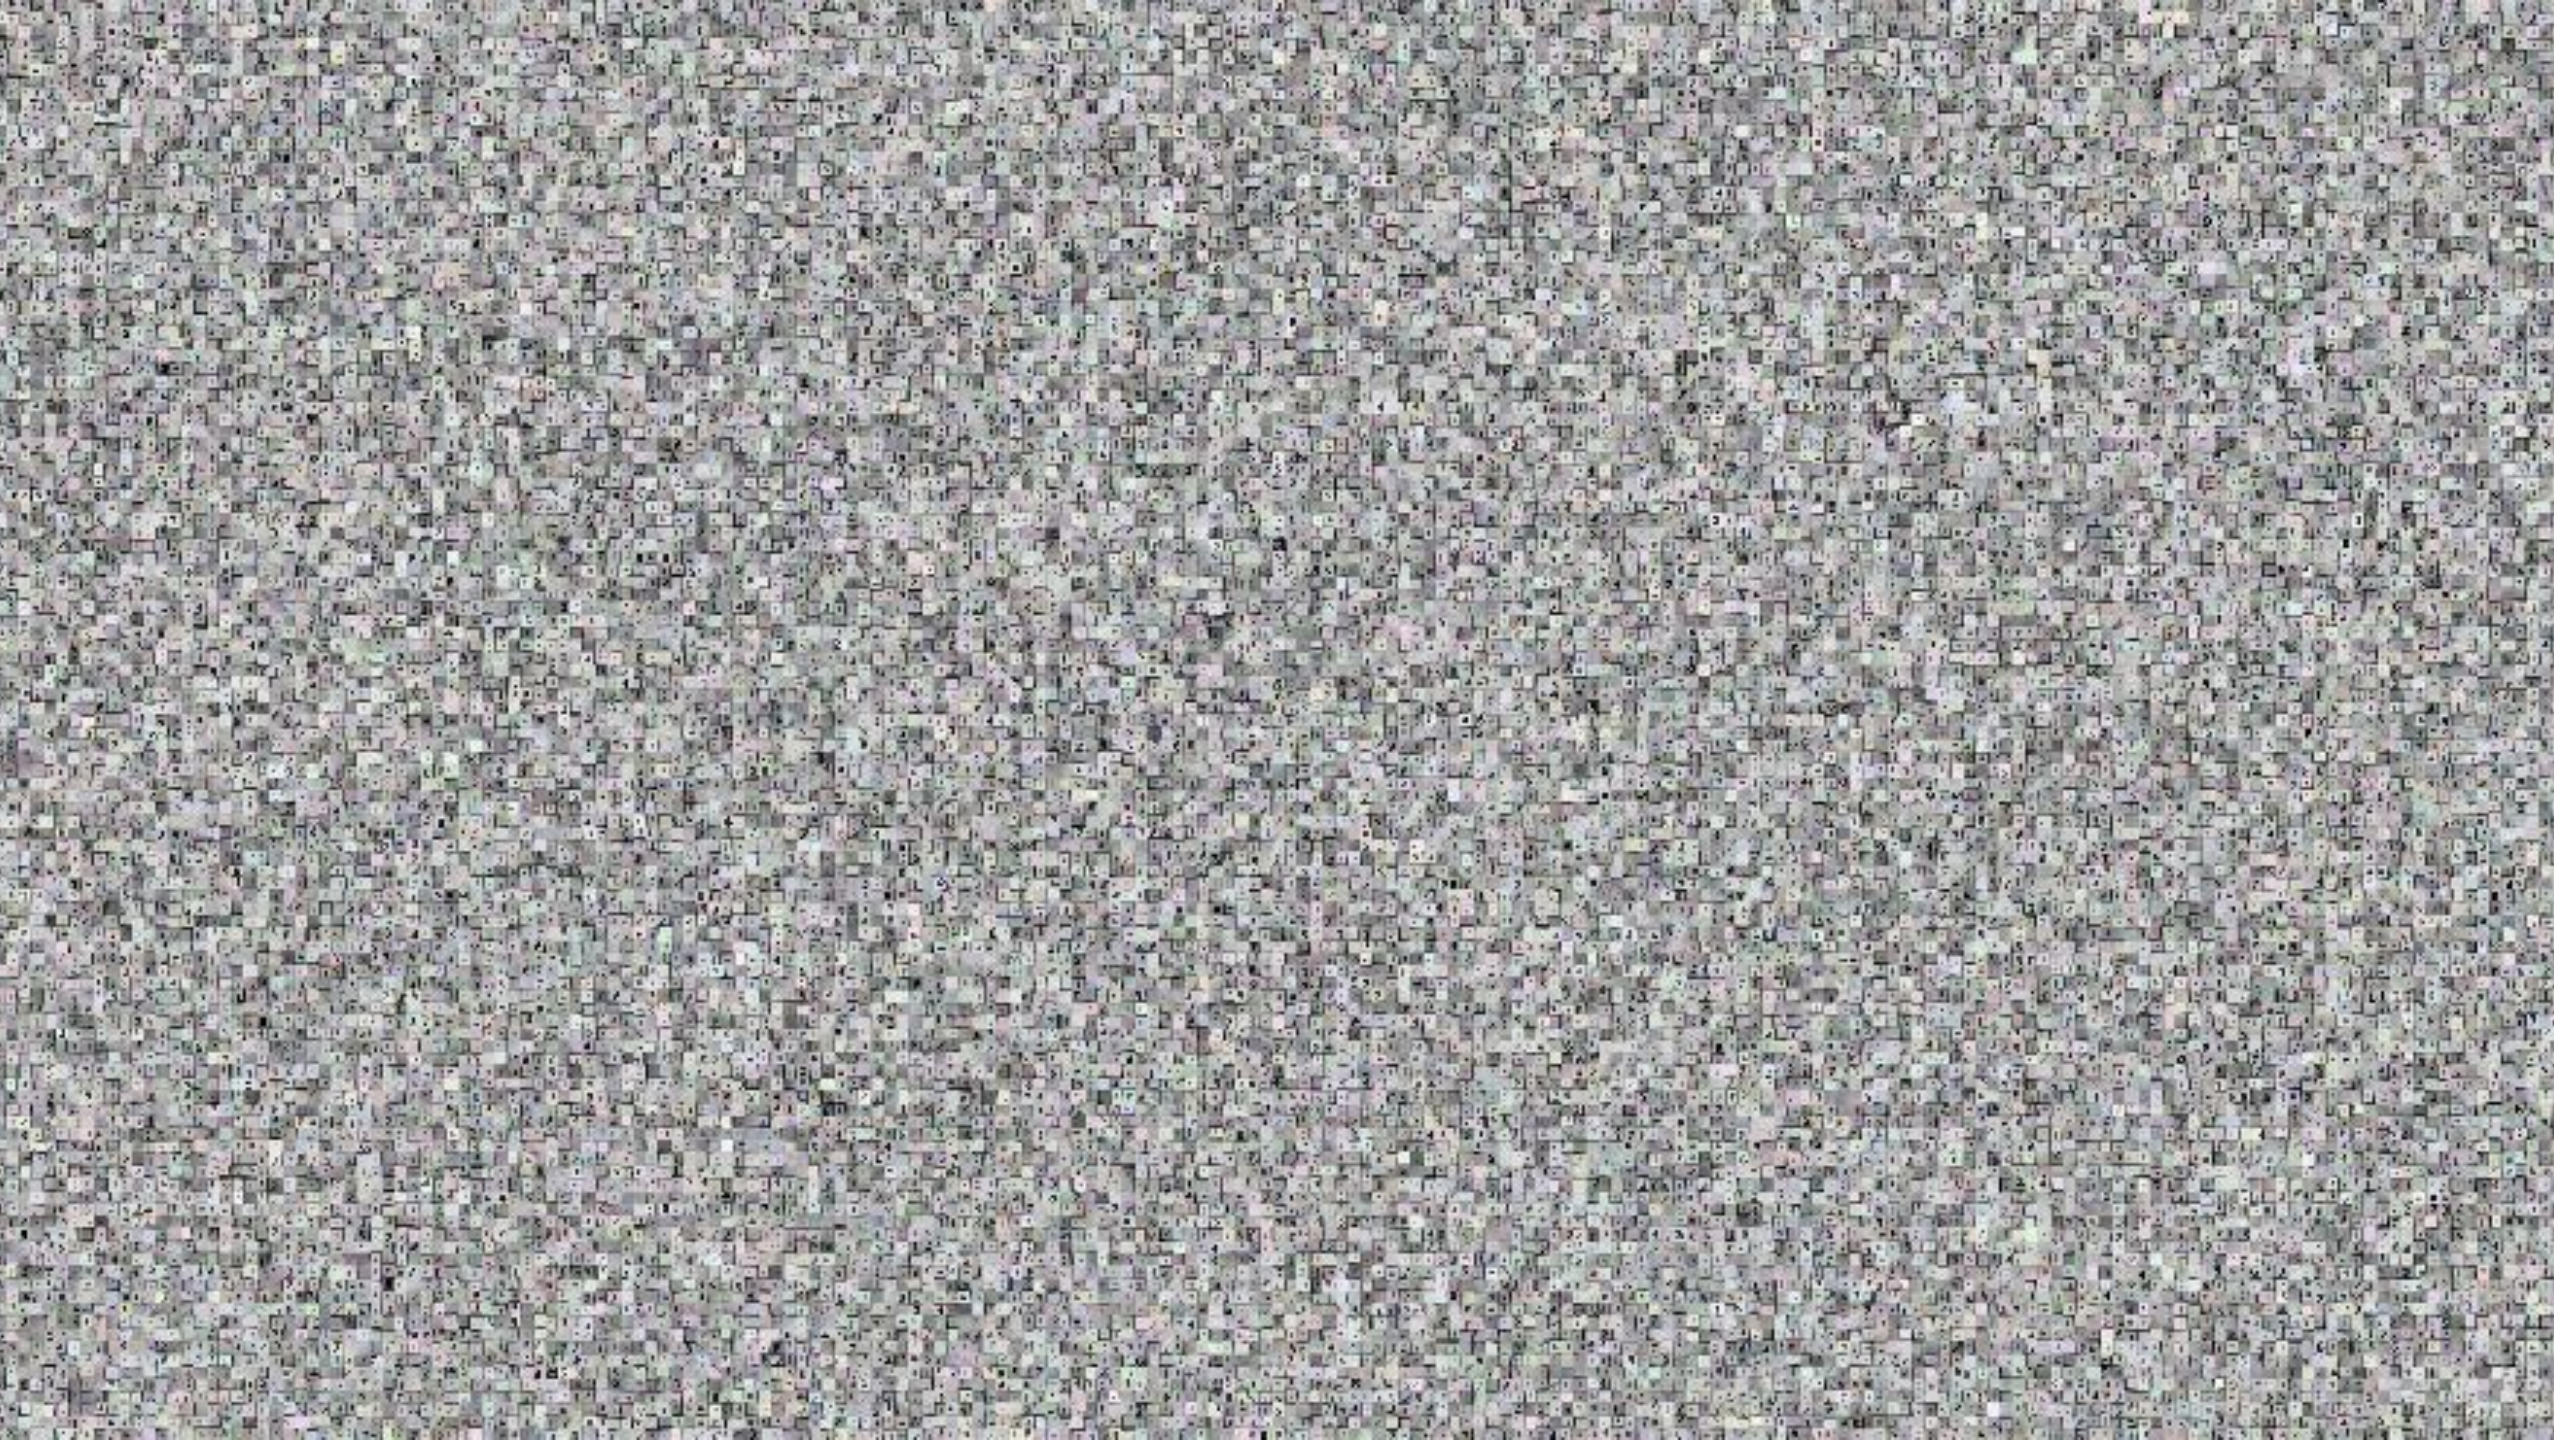

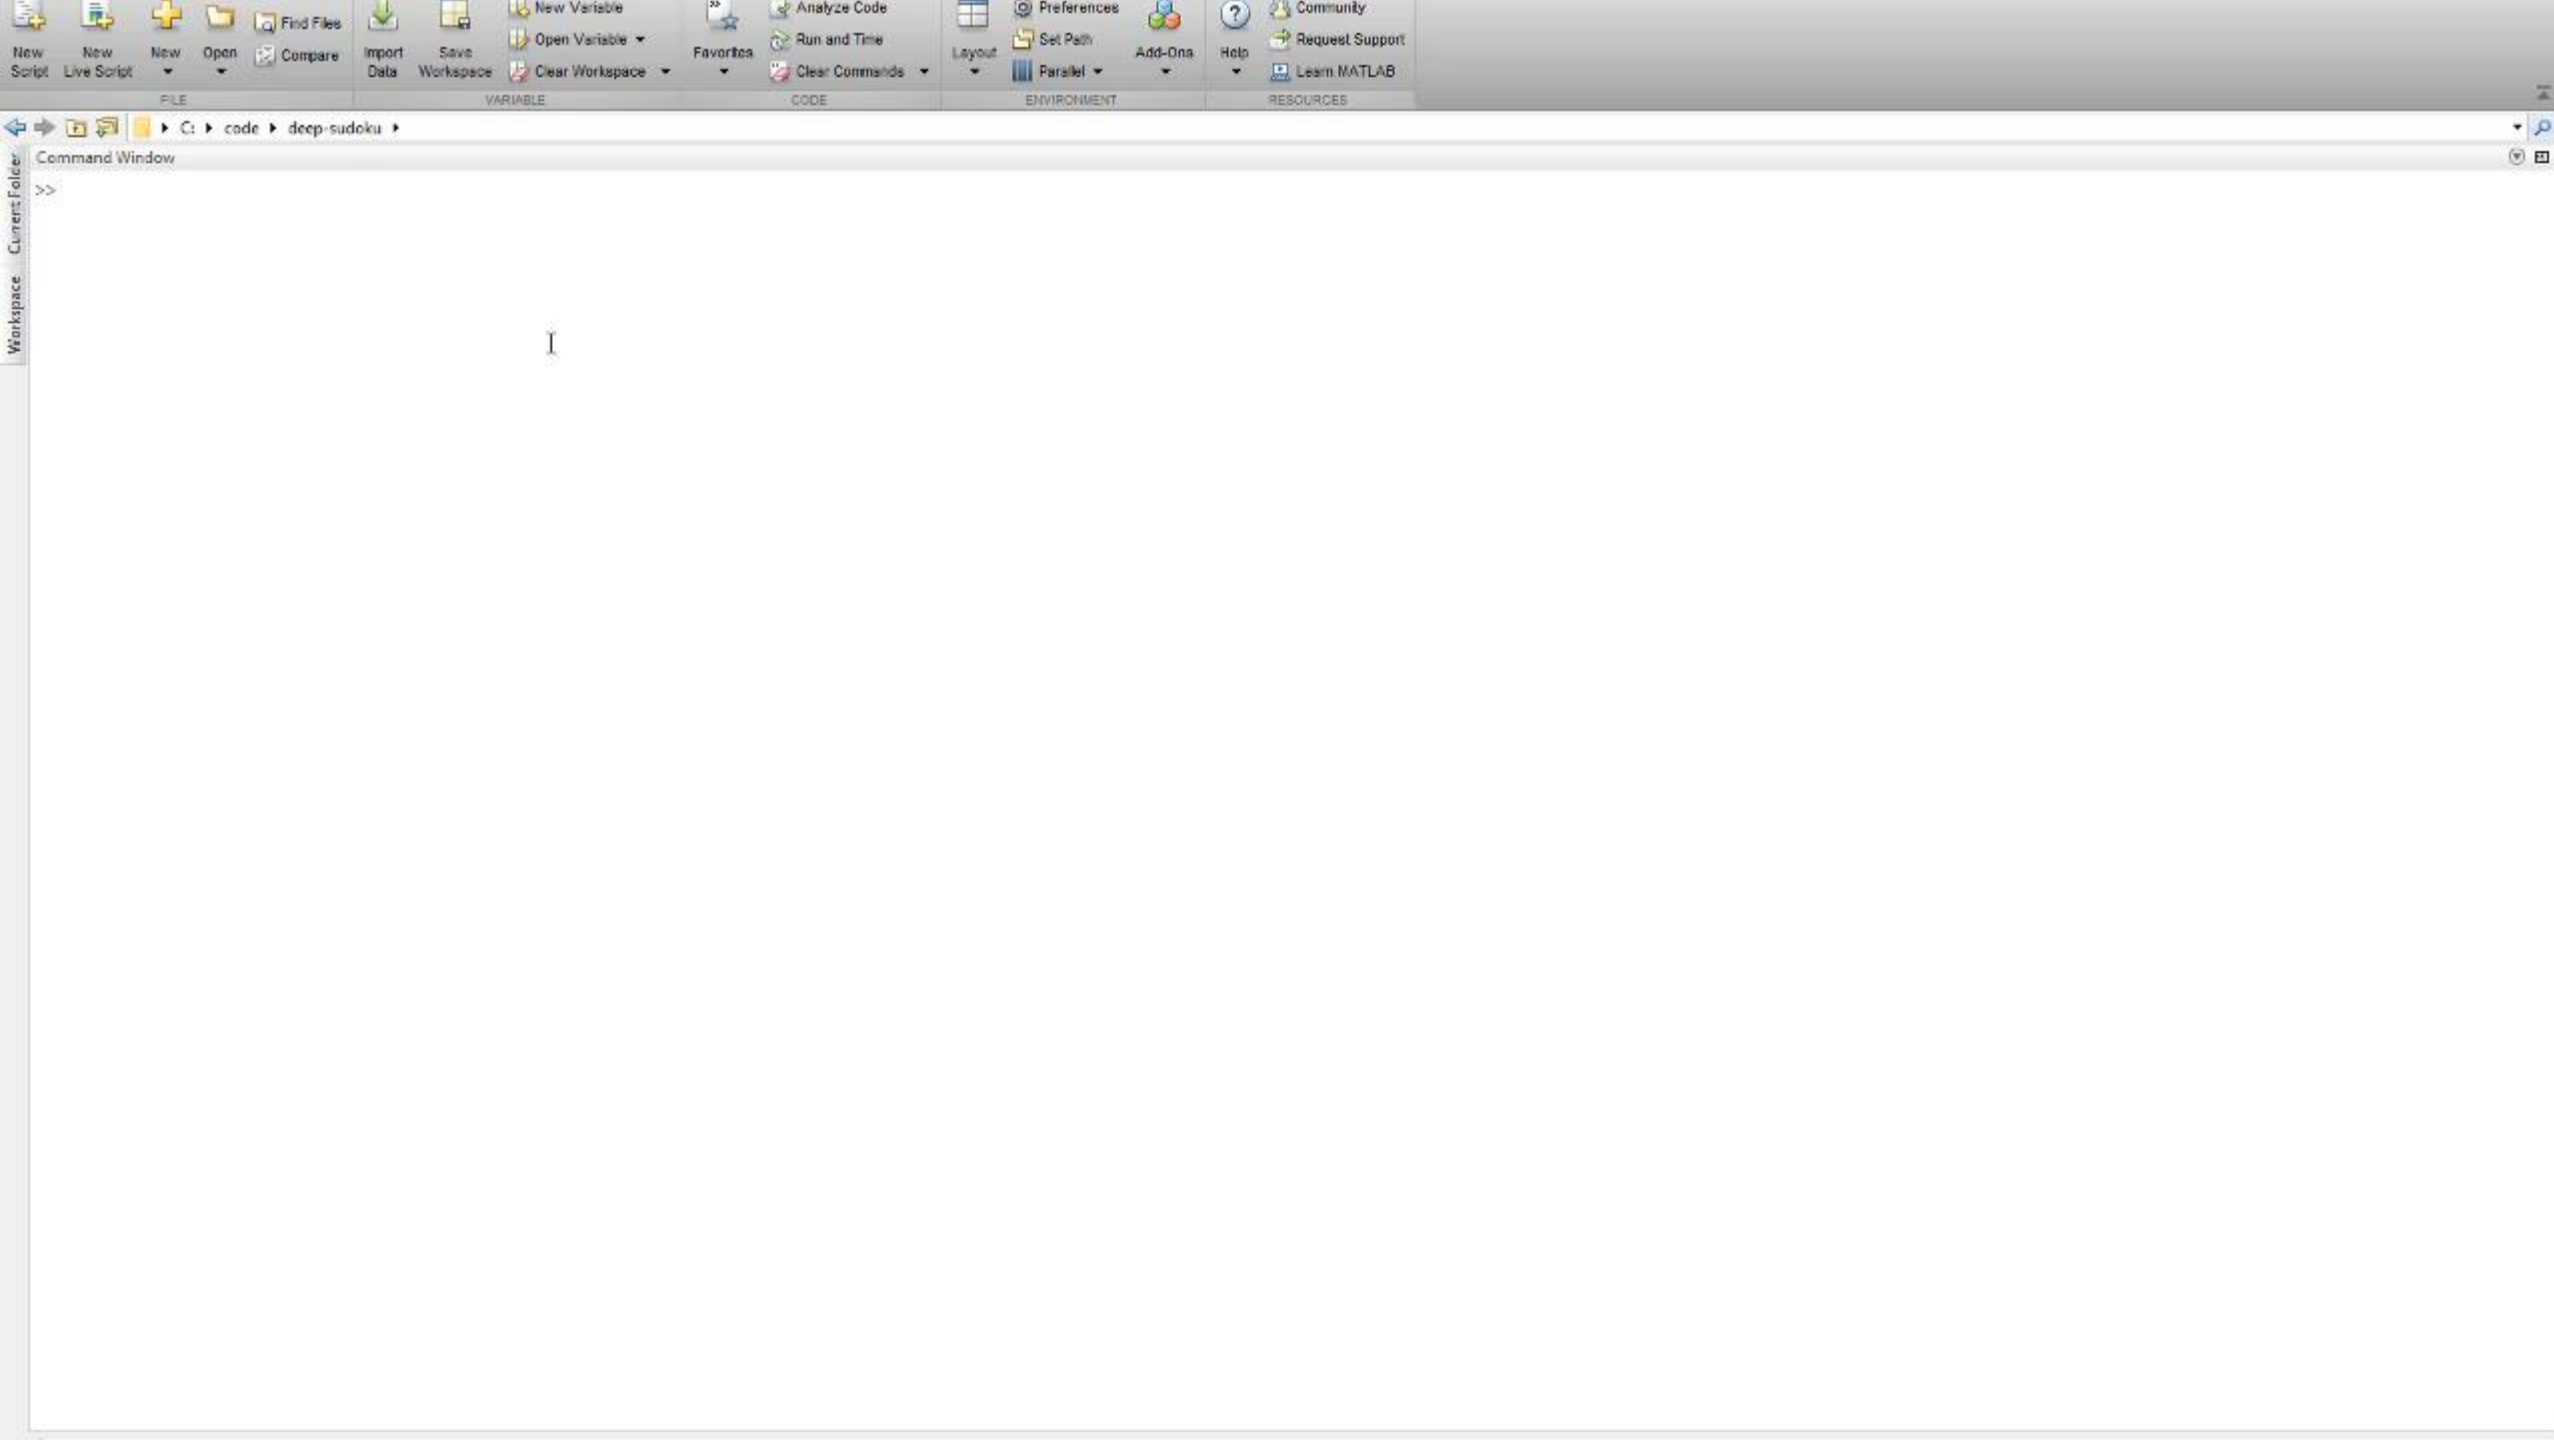
New Script
New Live Script
New
Open
Find Files
Compare
FILE
Import Data
Save Workspace
New Variable
Open Variable
Clear Workspace
VARIABLE
Favorites
Analyze Code
Run and Time
Clear Commands
CODE
Layout
Preferences
Set Path
Parallel
Add-Ons
ENVIRONMENT
Help
Community
Request Support
Learn MATLAB
RESOURCES
C: ▸ code ▸ deep-sudoku ▸
Command Window
>>
Current Folder
Workspace

Debugging a network

97.8% accuracy

Debugging a network

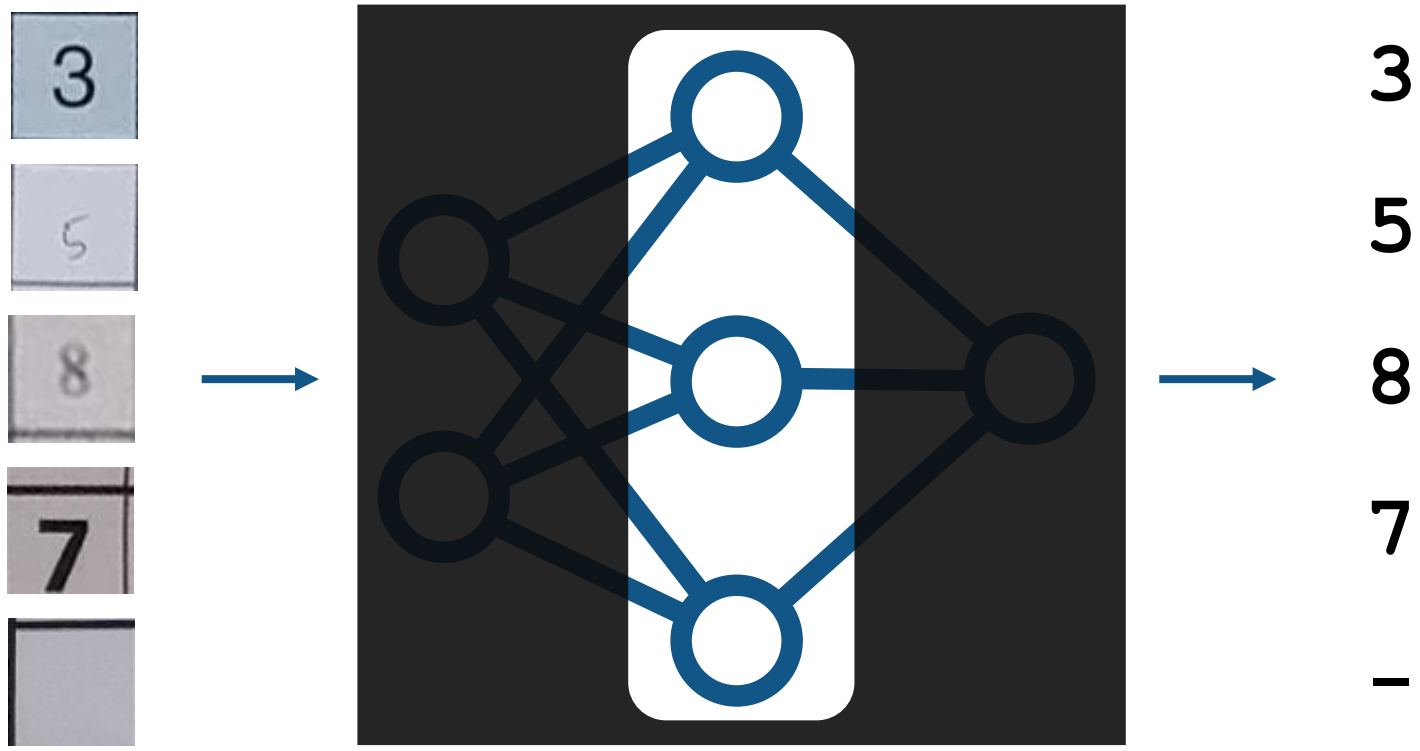

Debugging a network with t-SNE

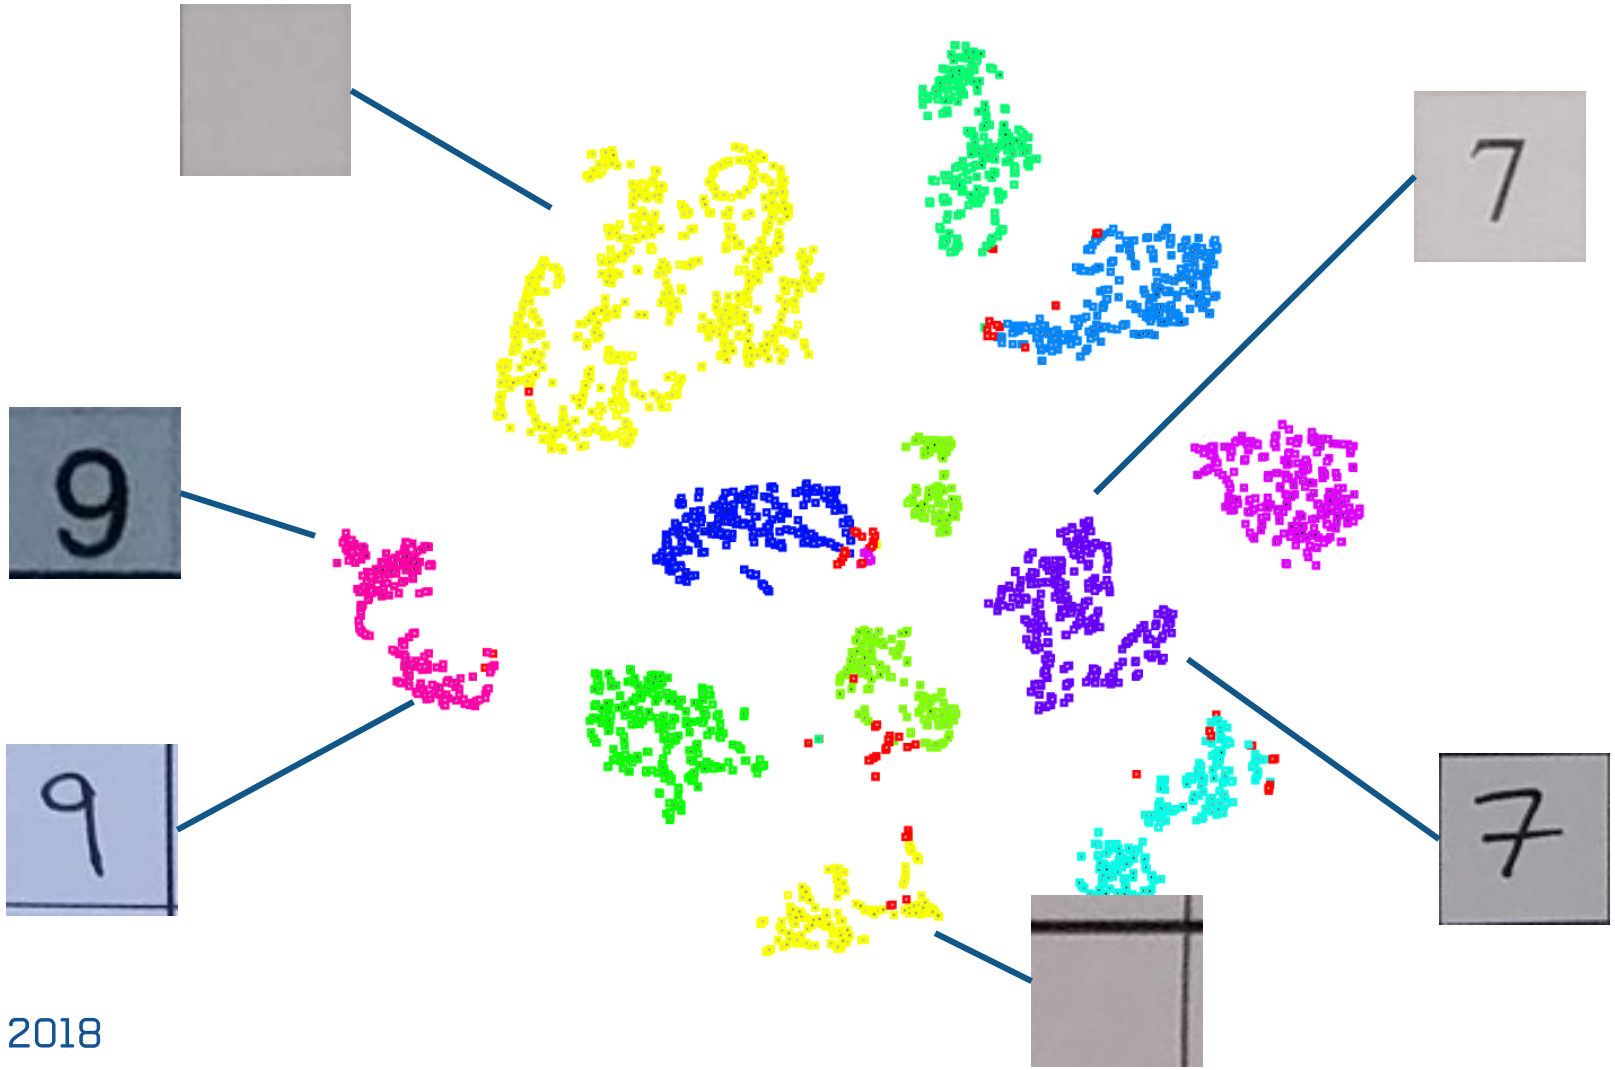

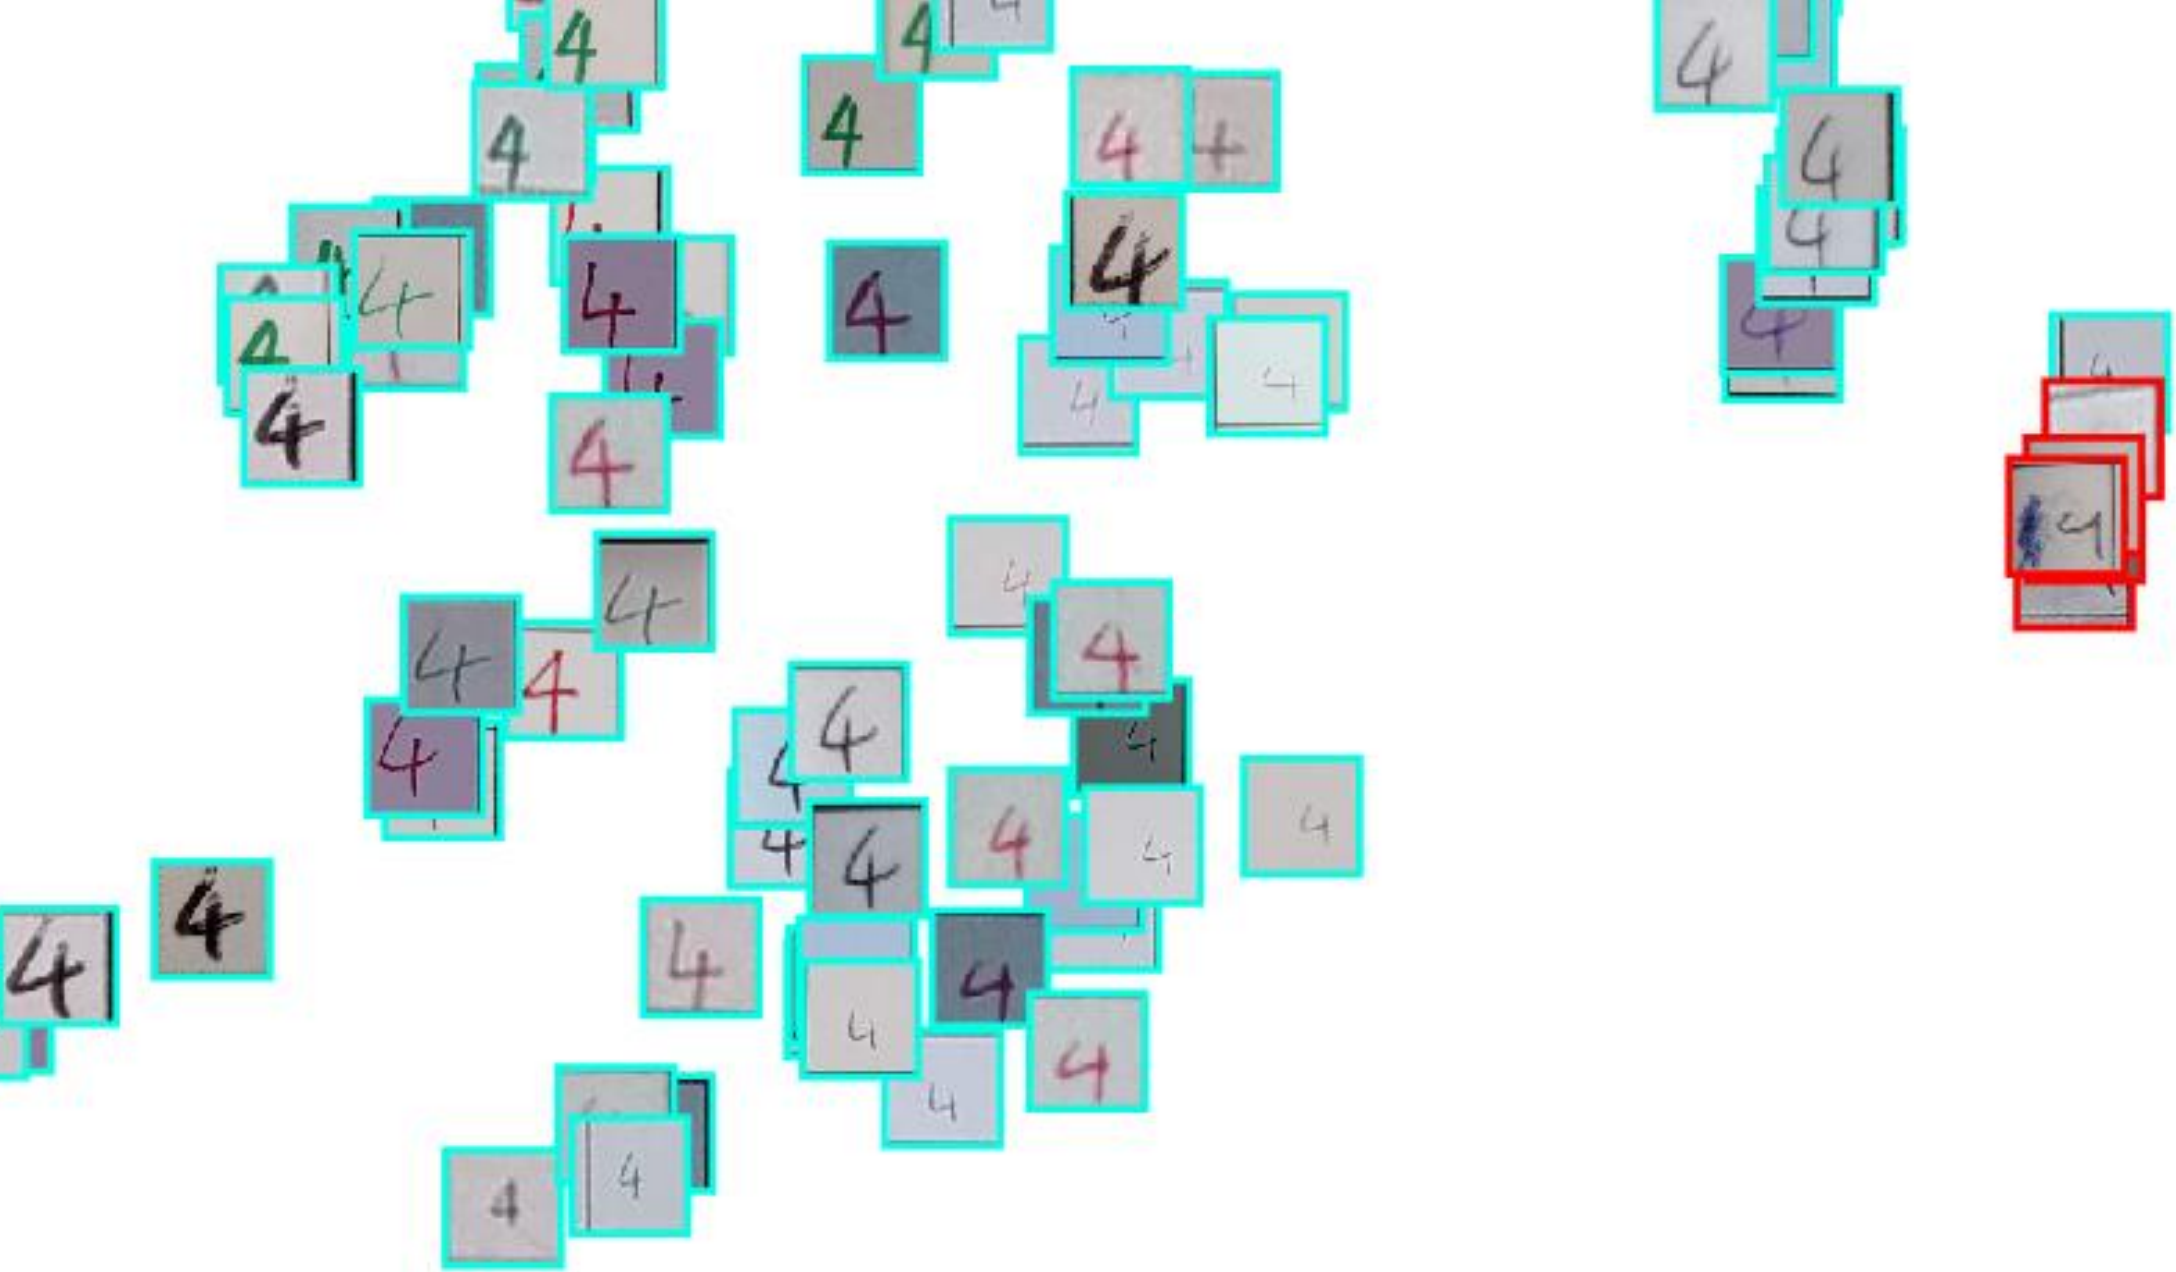

An Example – Sudoku Solver

Find the puzzle

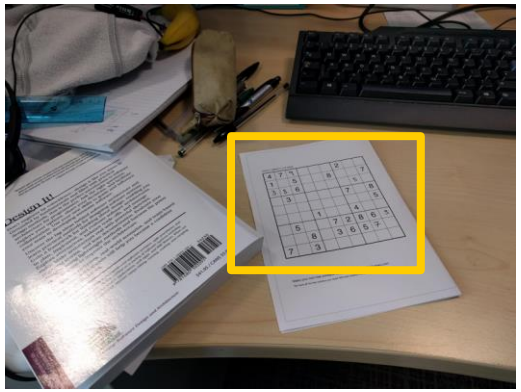

Find the boxes

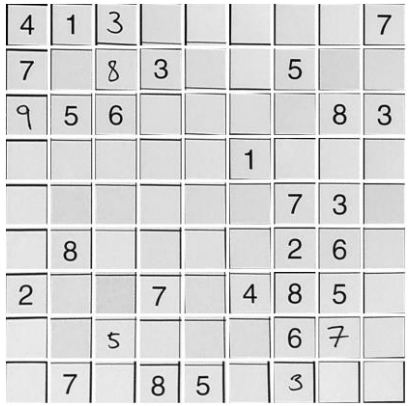

Read the numbers

$$\begin{bmatrix} 4 & 1 & \cdots & ? & ? \\ 7 & ? & \cdots & ? & ? \\ \vdots & \vdots & \cdots & \vdots & \vdots \\ ? & ? & \cdots & 7 & ? \\ ? & 7 & \cdots & ? & ? \end{bmatrix}$$

Solve

$$\begin{bmatrix} 4 & 1 & \cdots & 6 & 8 \\ 7 & 8 & \cdots & 3 & 2 \\ \vdots & \vdots & \cdots & \vdots & \vdots \\ 4 & 2 & \cdots & 7 & 5 \\ 6 & 7 & \cdots & 4 & 1 \end{bmatrix}$$

Deep Learning | **Image Processing** | **Deep Learning**

Step 4 – Solve

Find the puzzle

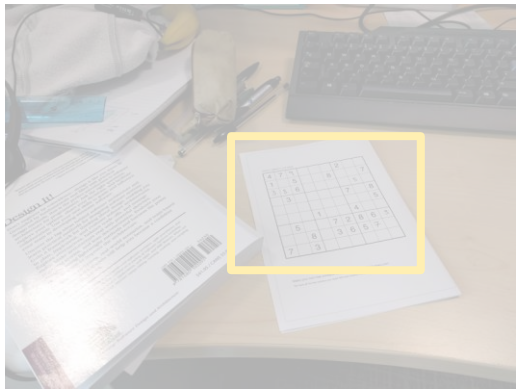

Find the boxes

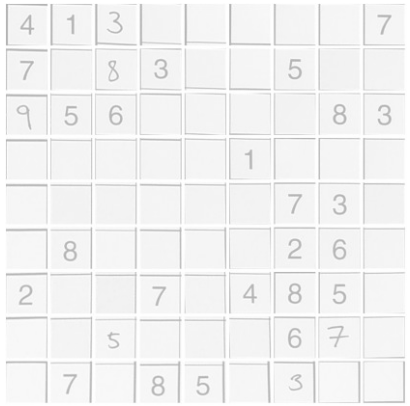

Read the numbers

$$\begin{bmatrix} 4 & 1 & \cdots & ? & ? \\ 7 & ? & \cdots & ? & ? \\ \vdots & \vdots & \cdots & \vdots & \vdots \\ ? & ? & \cdots & 7 & ? \\ ? & 7 & \cdots & ? & ? \end{bmatrix}$$

Solve

$$\begin{bmatrix} 4 & 1 & \cdots & 6 & 8 \\ 7 & 8 & \cdots & 3 & 2 \\ \vdots & \vdots & \cdots & \vdots & \vdots \\ 4 & 2 & \cdots & 7 & 5 \\ 6 & 7 & \cdots & 4 & 1 \end{bmatrix}$$

How can we solve a sudoku?

MathWorks® Products Solutions Academia Support Community Events

Technical Articles and Newsletters

Search Technical Art

Overview | Search Technical Articles | Newsletters | Cleve's Corner Collection | Sign Up

Solving Sudoku with MATLAB

By Cleve Moler, MathWorks

Human puzzle-solvers and computer programs use very different Sudoku-solving techniques. The fascination with solving Sudoku by hand derives from the discovery and mastery of a myriad of subtle combinations and patterns that provide hints about the final solution. It is not easy to program a computer to duplicate these human pattern-recognition capabilities. For th reason, most Sudoku-solving programs take a very different approach, relying on the computer's almost limitless capacity to carry out brute-force trial and error. That is the approach that I used for the MATLAB® program.

The Sudoku Challenge

As you probably know, solving a Sudoku involves filling in a 9-by-9 grid so that each row, column, and major 3-by-3 block contains all the digits 1 through 9. The initial grid is populated with a few digits, known as *clues*. In contrast to magic squares and other numeric puzzles, no arithmetic is involved; the elements in a Sudoku grid could just as well be letters of the alphat or any other symbols.

Figure 1 shows an initial grid. I especially like the symmetry in this example, which is due to Gordon Royle of the University c Western Australia. Figure 2 shows the solution.

	2			3			4	
6								3
		4				5		
			8		6			

MathWorks® Products Solutions Academia Support Community Events

Documentation

CONTENTS Close

‹ Examples Home

‹ Optimization Toolbox

‹ Mixed-Integer Linear Programming

‹ Optimization Toolbox

‹ Linear Programming and Mixed-Integer Linear Programming

‹ Problem-Based Optimization

Solve Sudoku Puzzles Via Integer Programming: Problem-Based

ON THIS PAGE

Initial Puzzle

Binary Integer Programming Approach

Express the Rules for Sudoku as Constraints

Express Clues

Sudoku in Optimization Problem Form

Function to Draw the Sudoku Puzzle

Solve Sudoku Puzzles Via Integer Programming: Problem-Based

This example shows how to solve a Sudoku puzzle using binary integer programming. For the solver-based approach, see S

You probably have seen Sudoku puzzles. A puzzle is to fill a 9-by-9 grid with integers from 1 through 9 so that each integer a partially populated with clues, and your task is to fill in the rest of the grid.

Initial Puzzle

Here is a data matrix `B` of clues. The first row, `B(1,2,2)`, means row 1, column 2 has a clue 2. The second row, `B(1,5,3)`, m

```
B = [1,2,2;
     1,5,3;
     1,8,4;
     2,1,6;
     2,9,3;
     3,3,4;
     3,7,5;
     4,4,8;
     4,6,6;
     5,1,8;
     5,5,1;
     5,9,6;
     6,4,7;
     6,6,5;
     7,3,7;
     7,7,6;
     8,1,4;
     8,9,8;
     9,2,3;
     9,5,4;
     9,8,2];

drawSudoku(B) % For the listing of this program, see the end of this example.
```

	2			3			4	
6								3
		4				5		

Step 4 – Solve

2	4	9	6	5	3	7	1	8
8	7	3	1	4	9	5	6	2
5	6	1	8	7	2	3	9	4
3	1	7	4	6	8	2	5	9
4	9	5	2	3	1	8	7	6
6	2	8	5	9	7	4	3	1
9	5	4	7	2	6	1	8	3
7	8	6	3	1	4	9	2	5
1	3	2	9	8	5	6	4	7

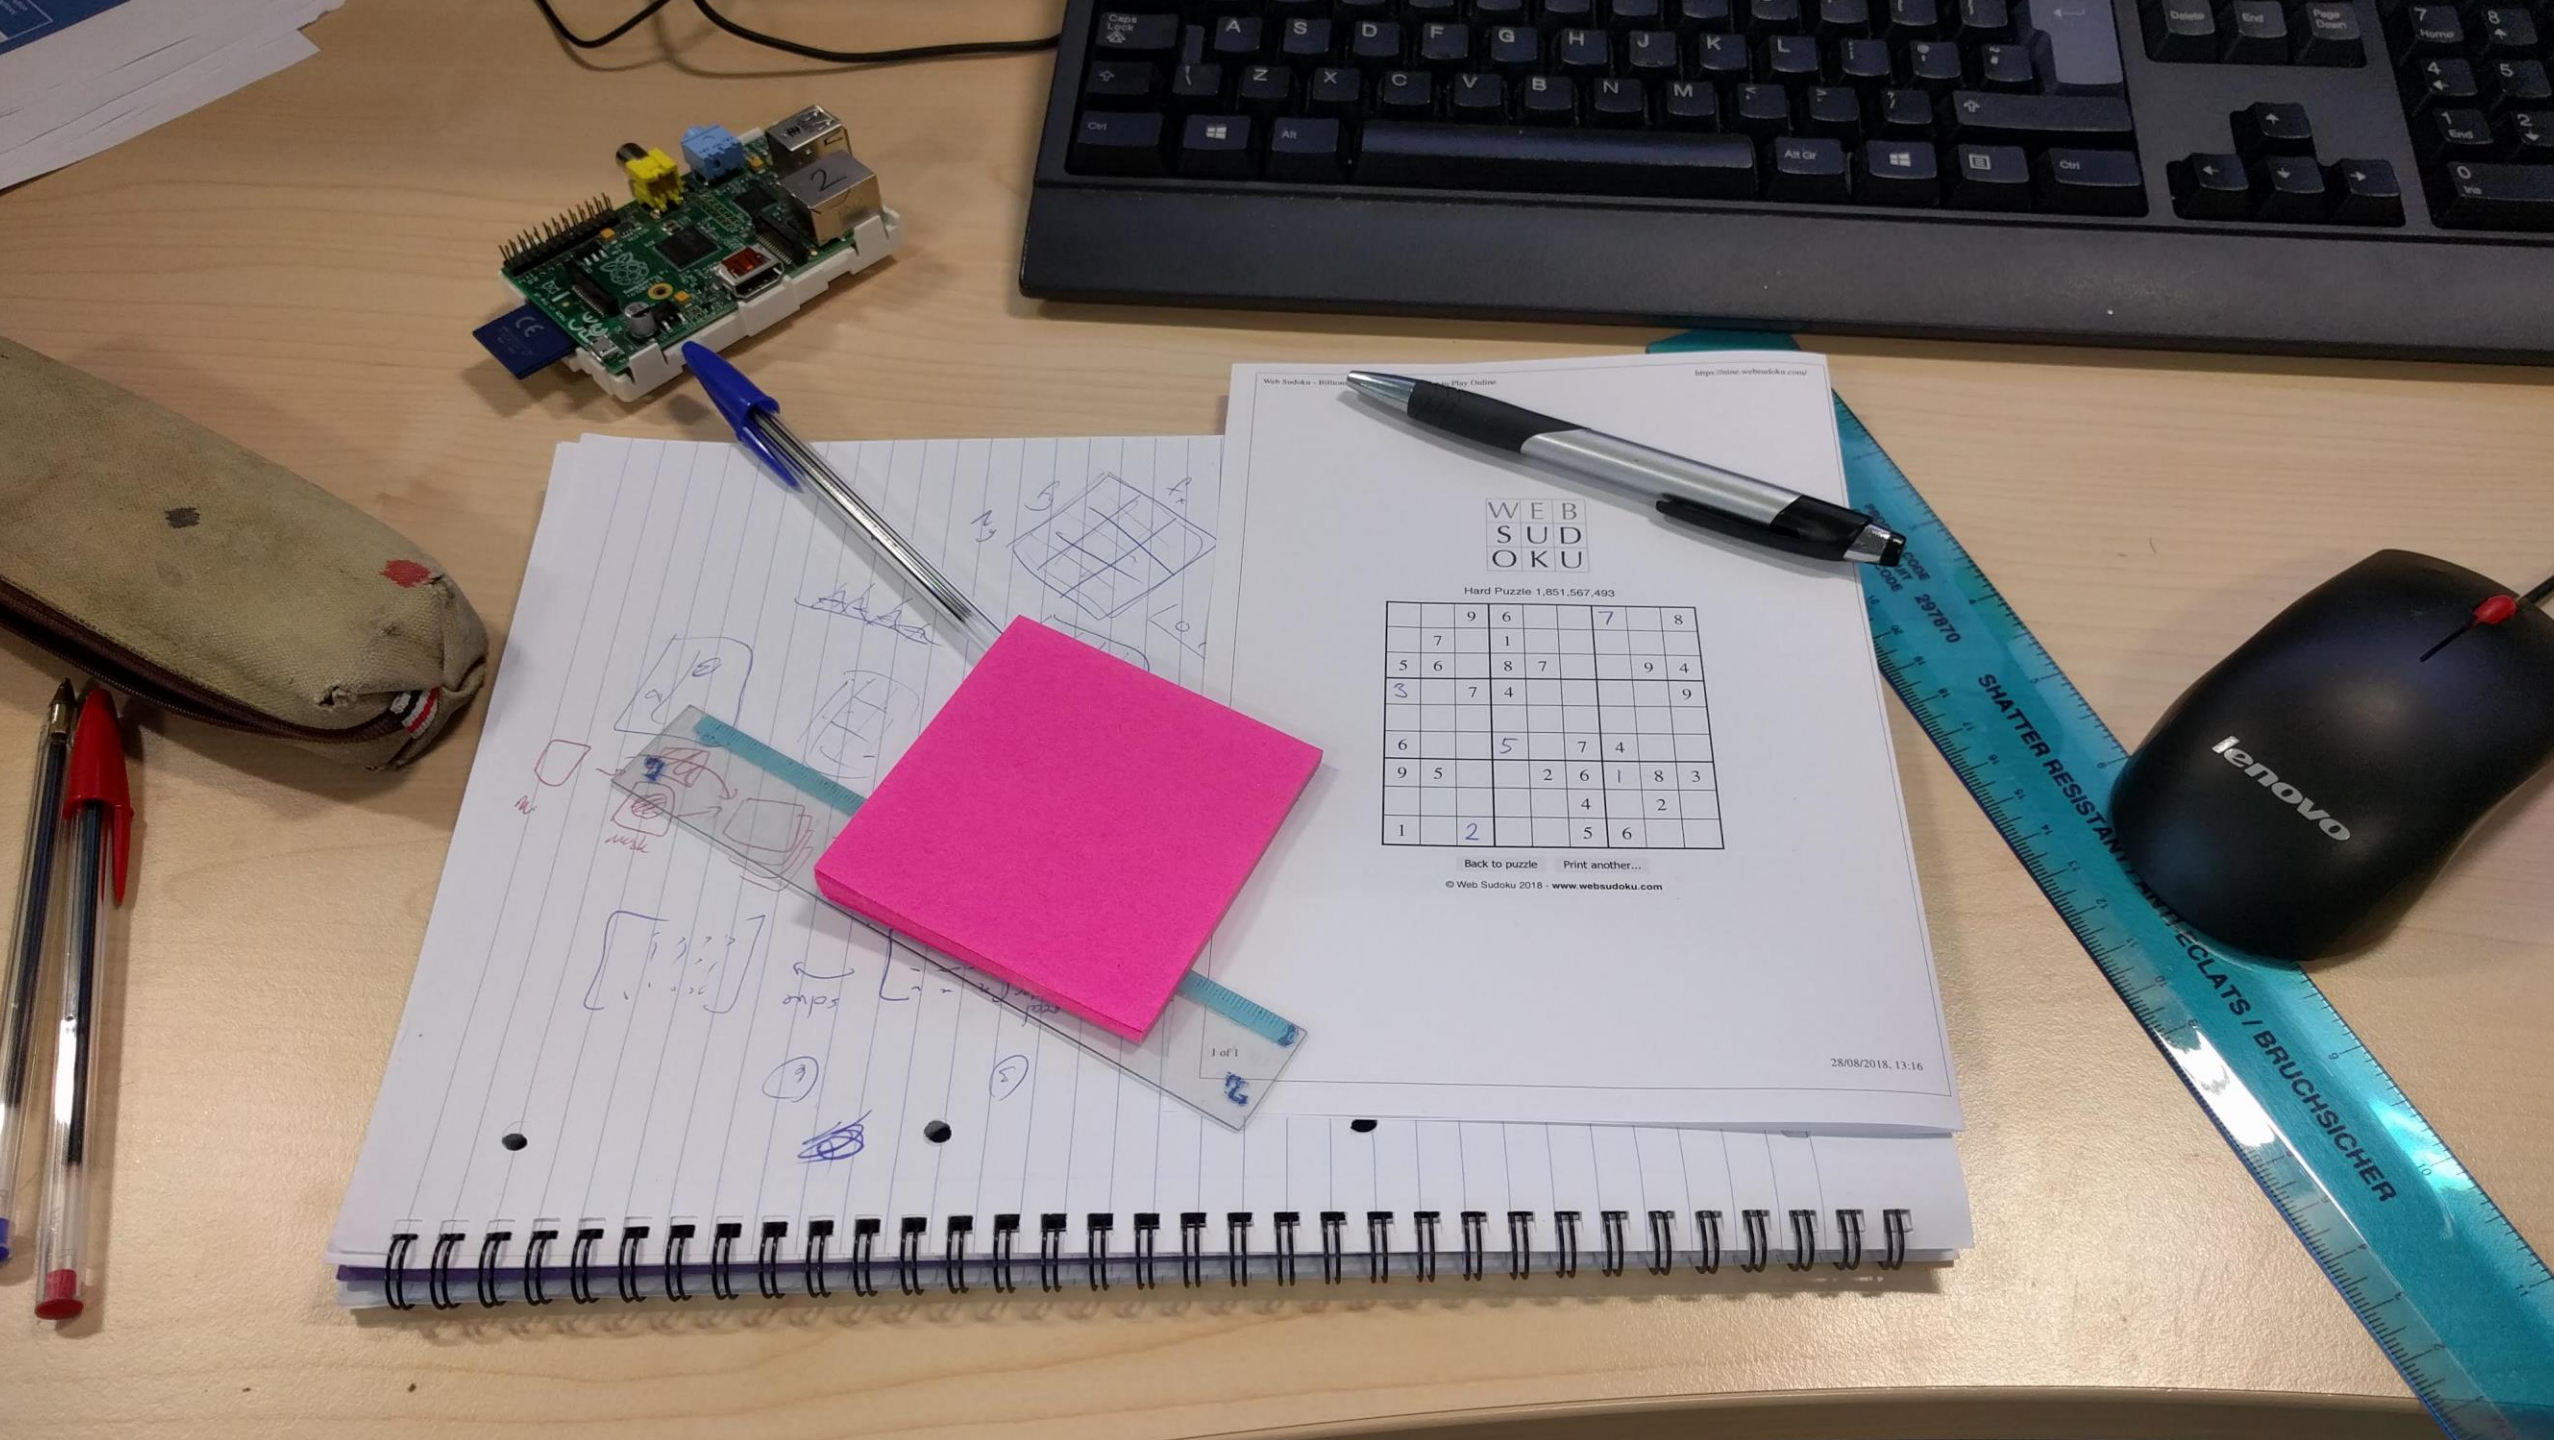
WEB
SUD
OKU
Hard Puzzle 1,851,567,493
Back to puzzle
Print another...
© Web Sudoku 2018 - www.websudoku.com
28/08/2018, 13:16
lenovo
SHATTER RESISTANT / ECLATS / BRUCHSICHER

Deep Learning

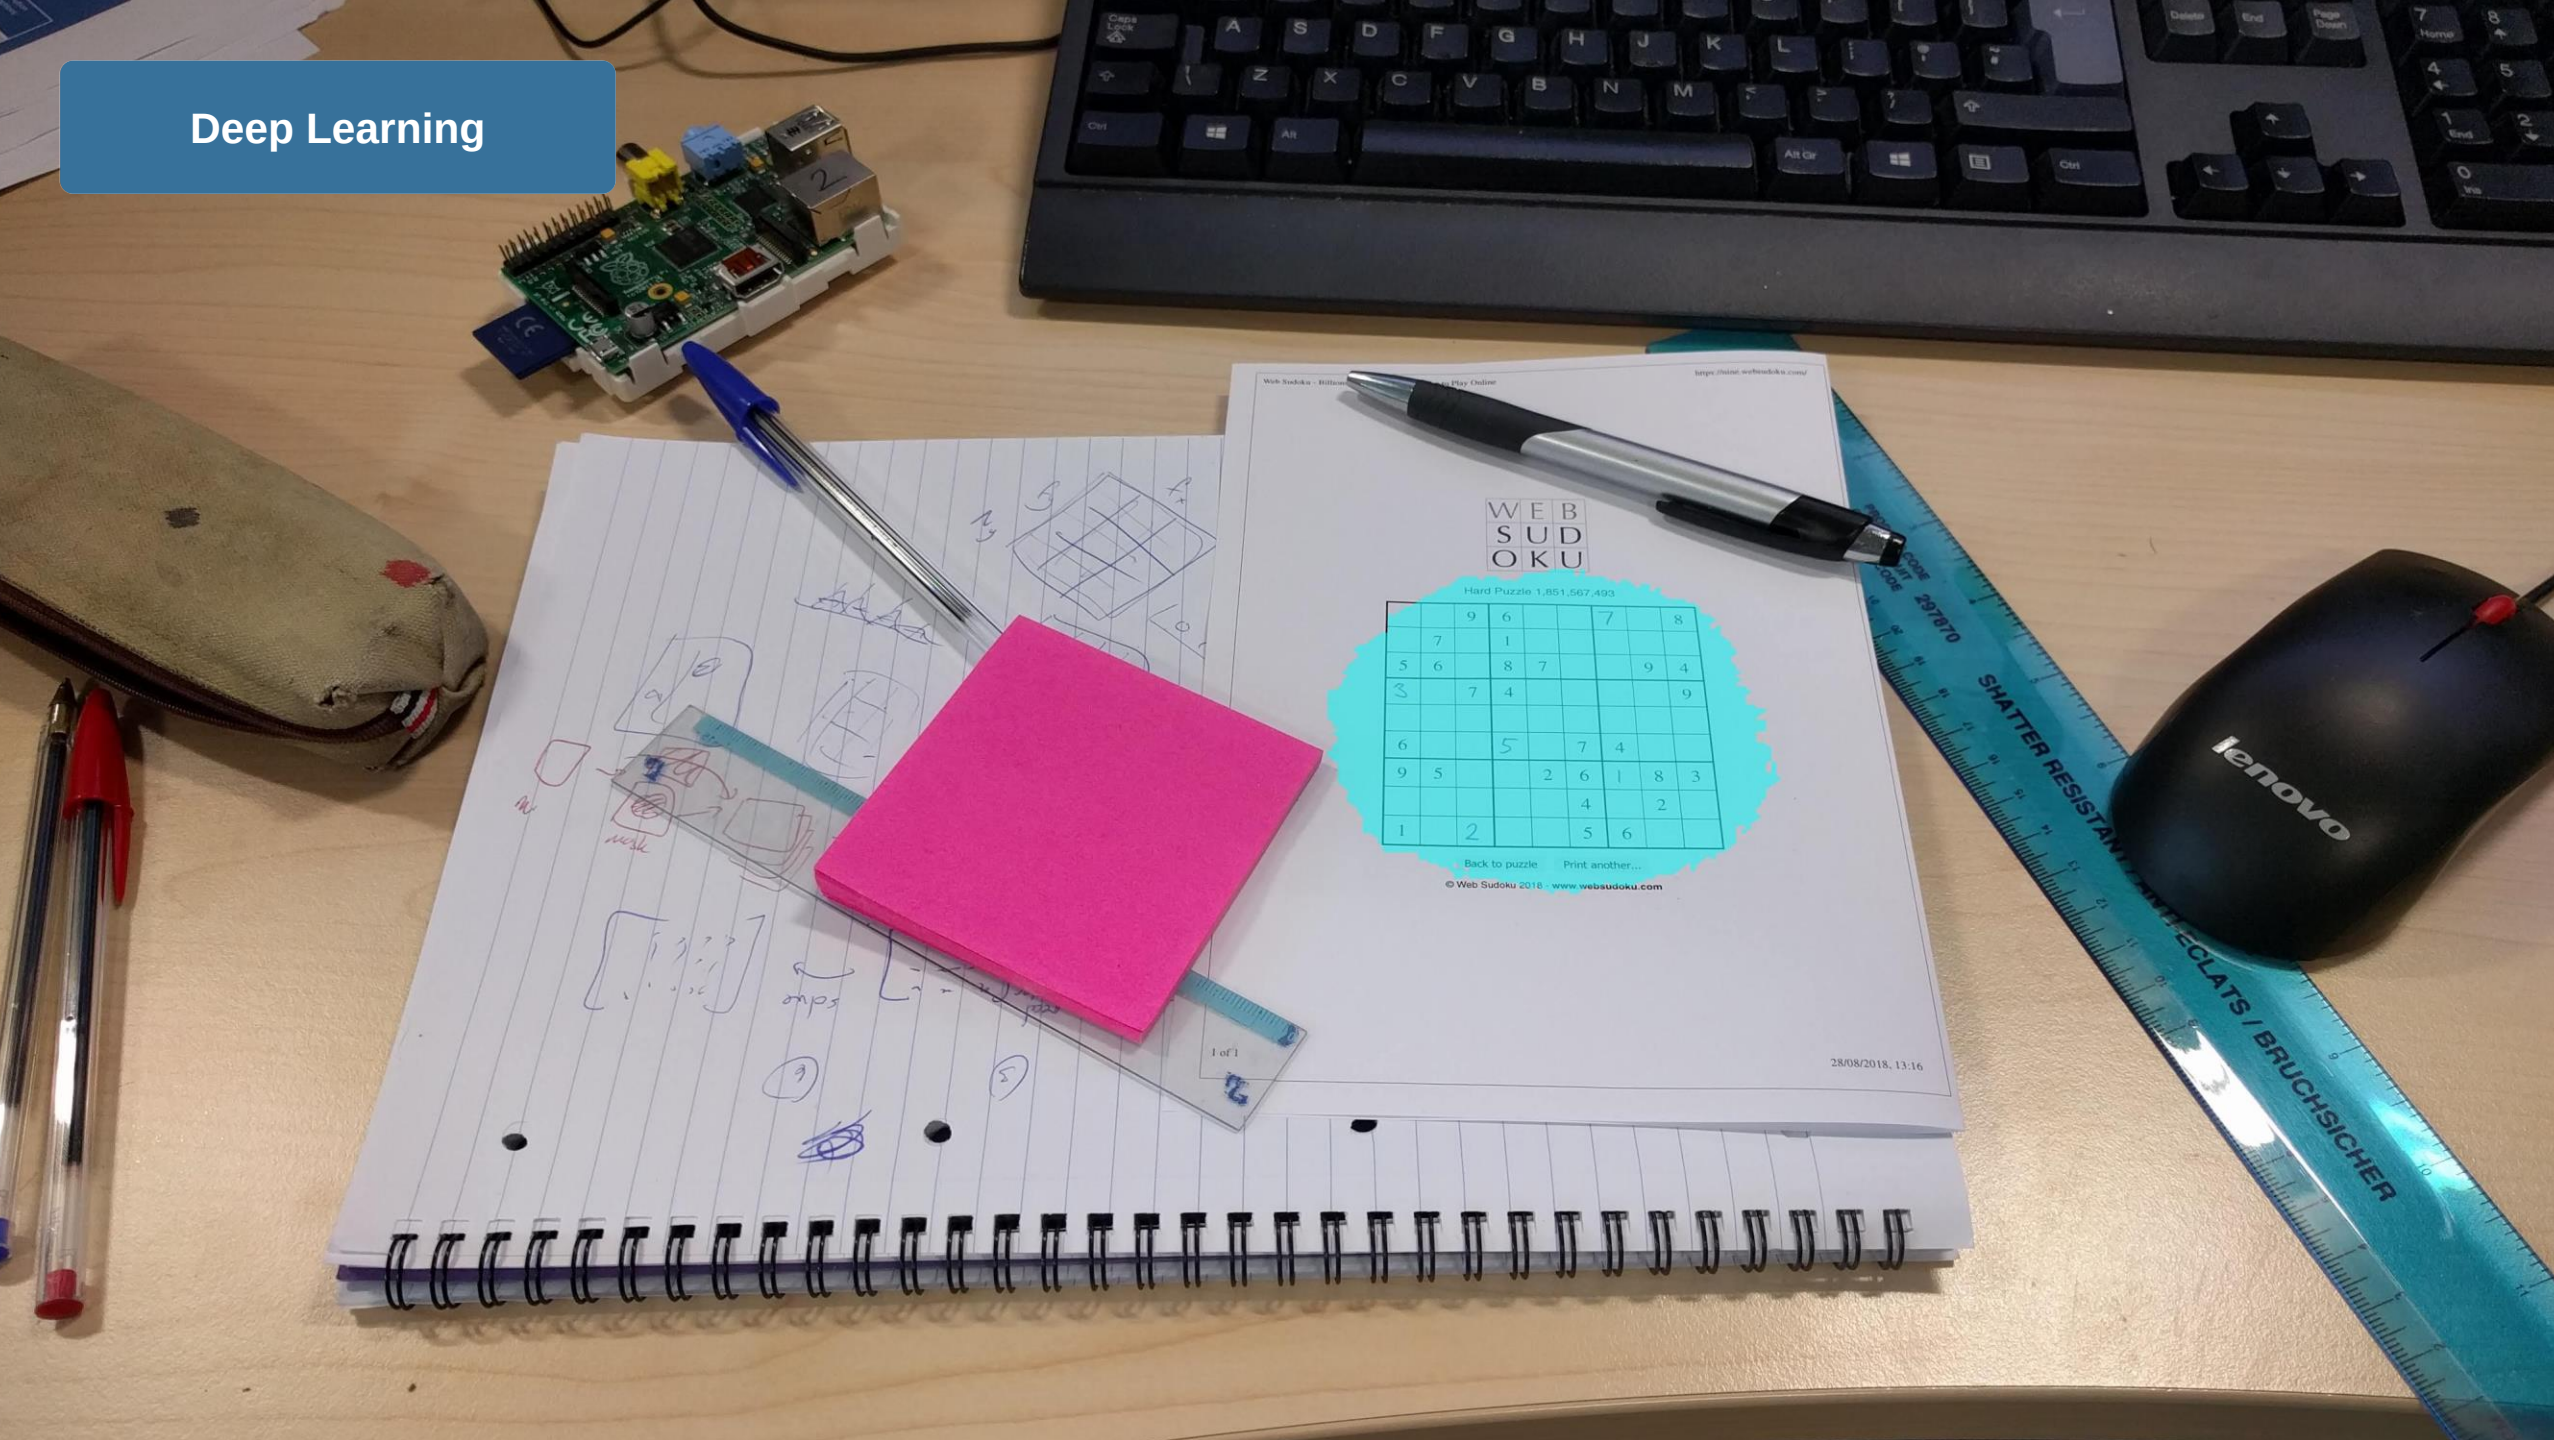

Image Processing

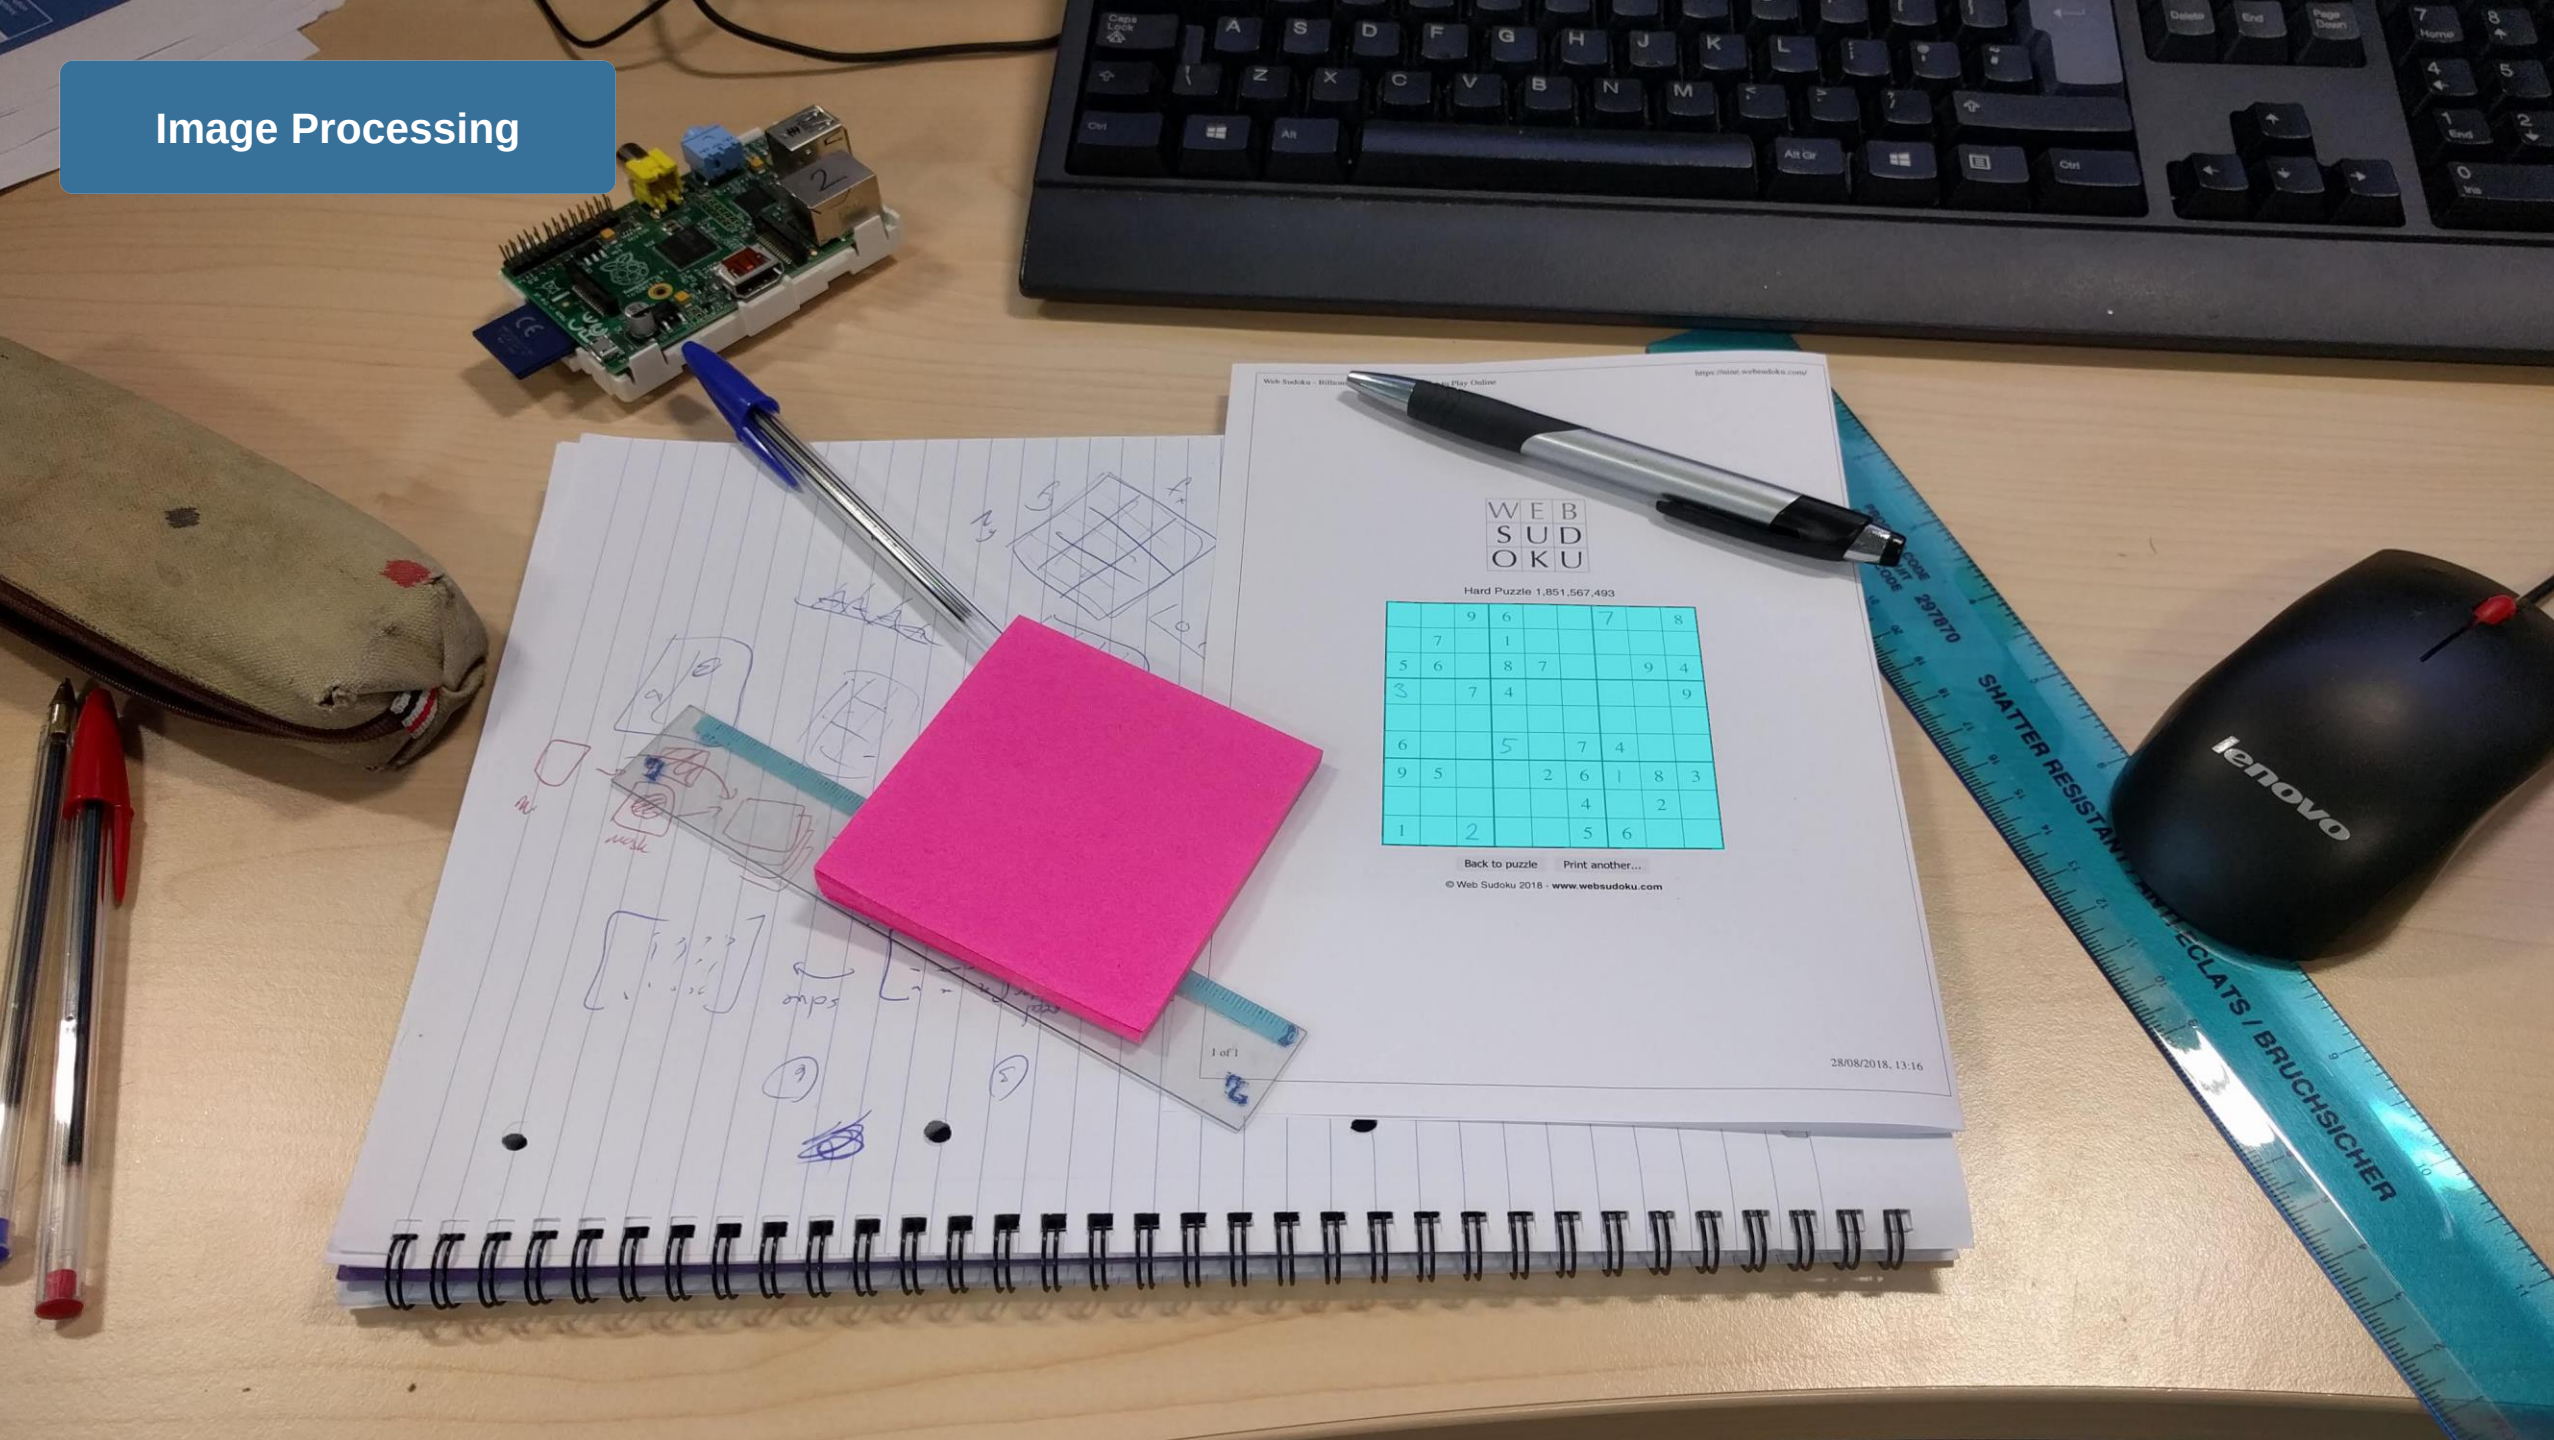

Deep Learning

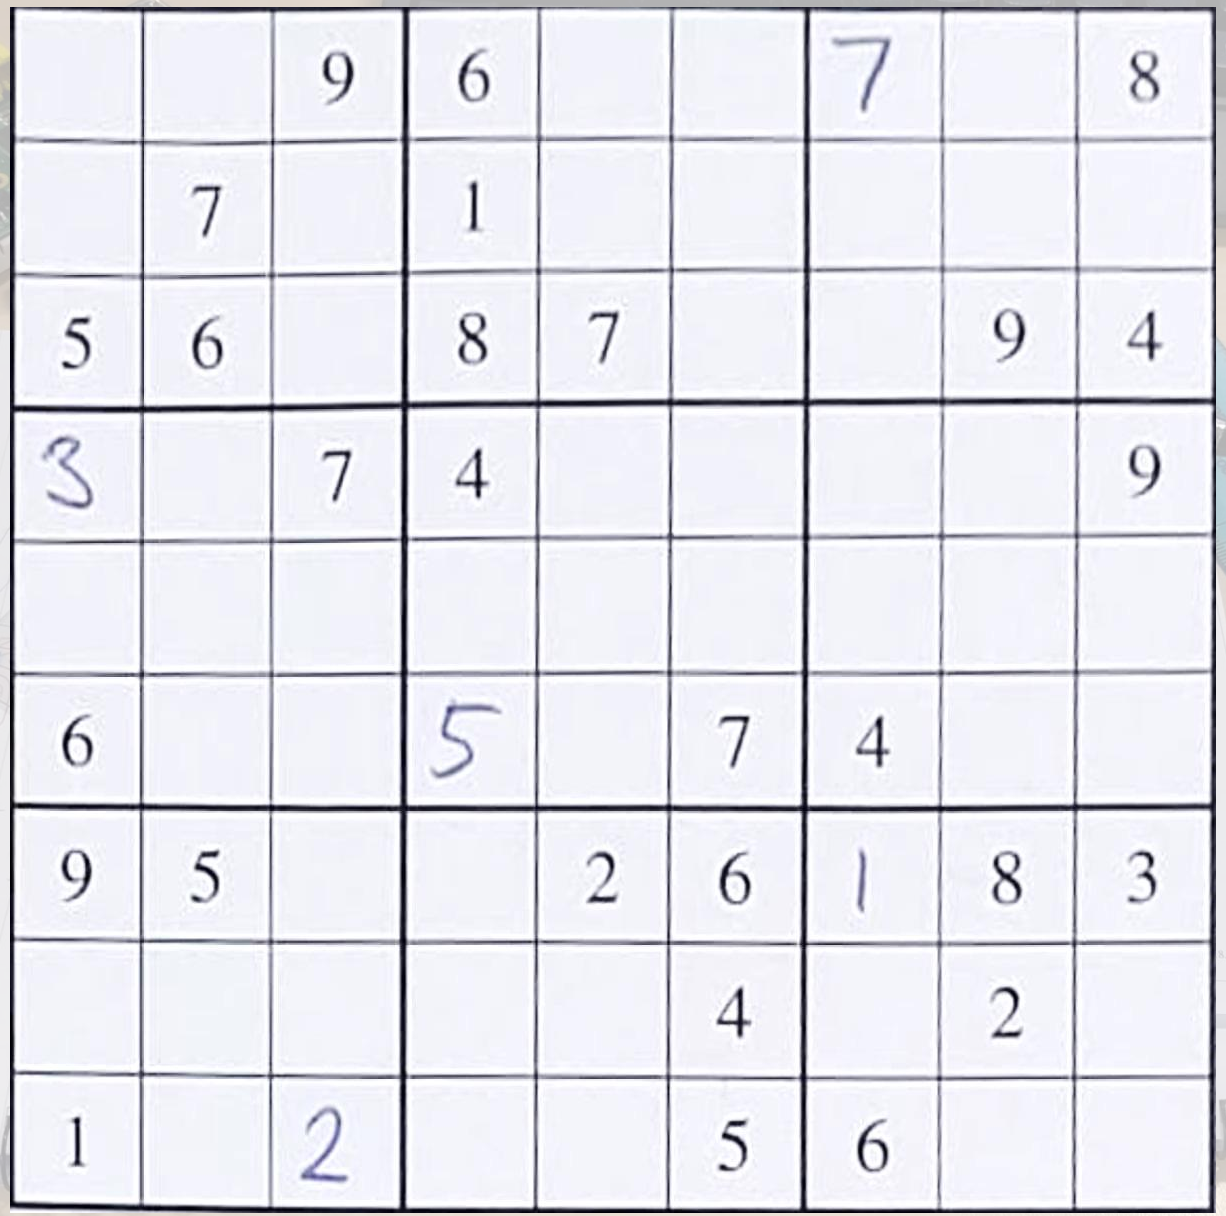

Optimisation

2	4	9	6	5	3	7	1	8
8	7	3	1	4	9	5	6	2
5	6	1	8	7	2	3	9	4
3	1	7	4	6	8	2	5	9
4	9	5	2	3	1	8	7	6
6	2	8	5	9	7	4	3	1
9	5	4	7	2	6	1	8	3
7	8	6	3	1	4	9	2	5
1	3	2	9	8	5	6	4	7

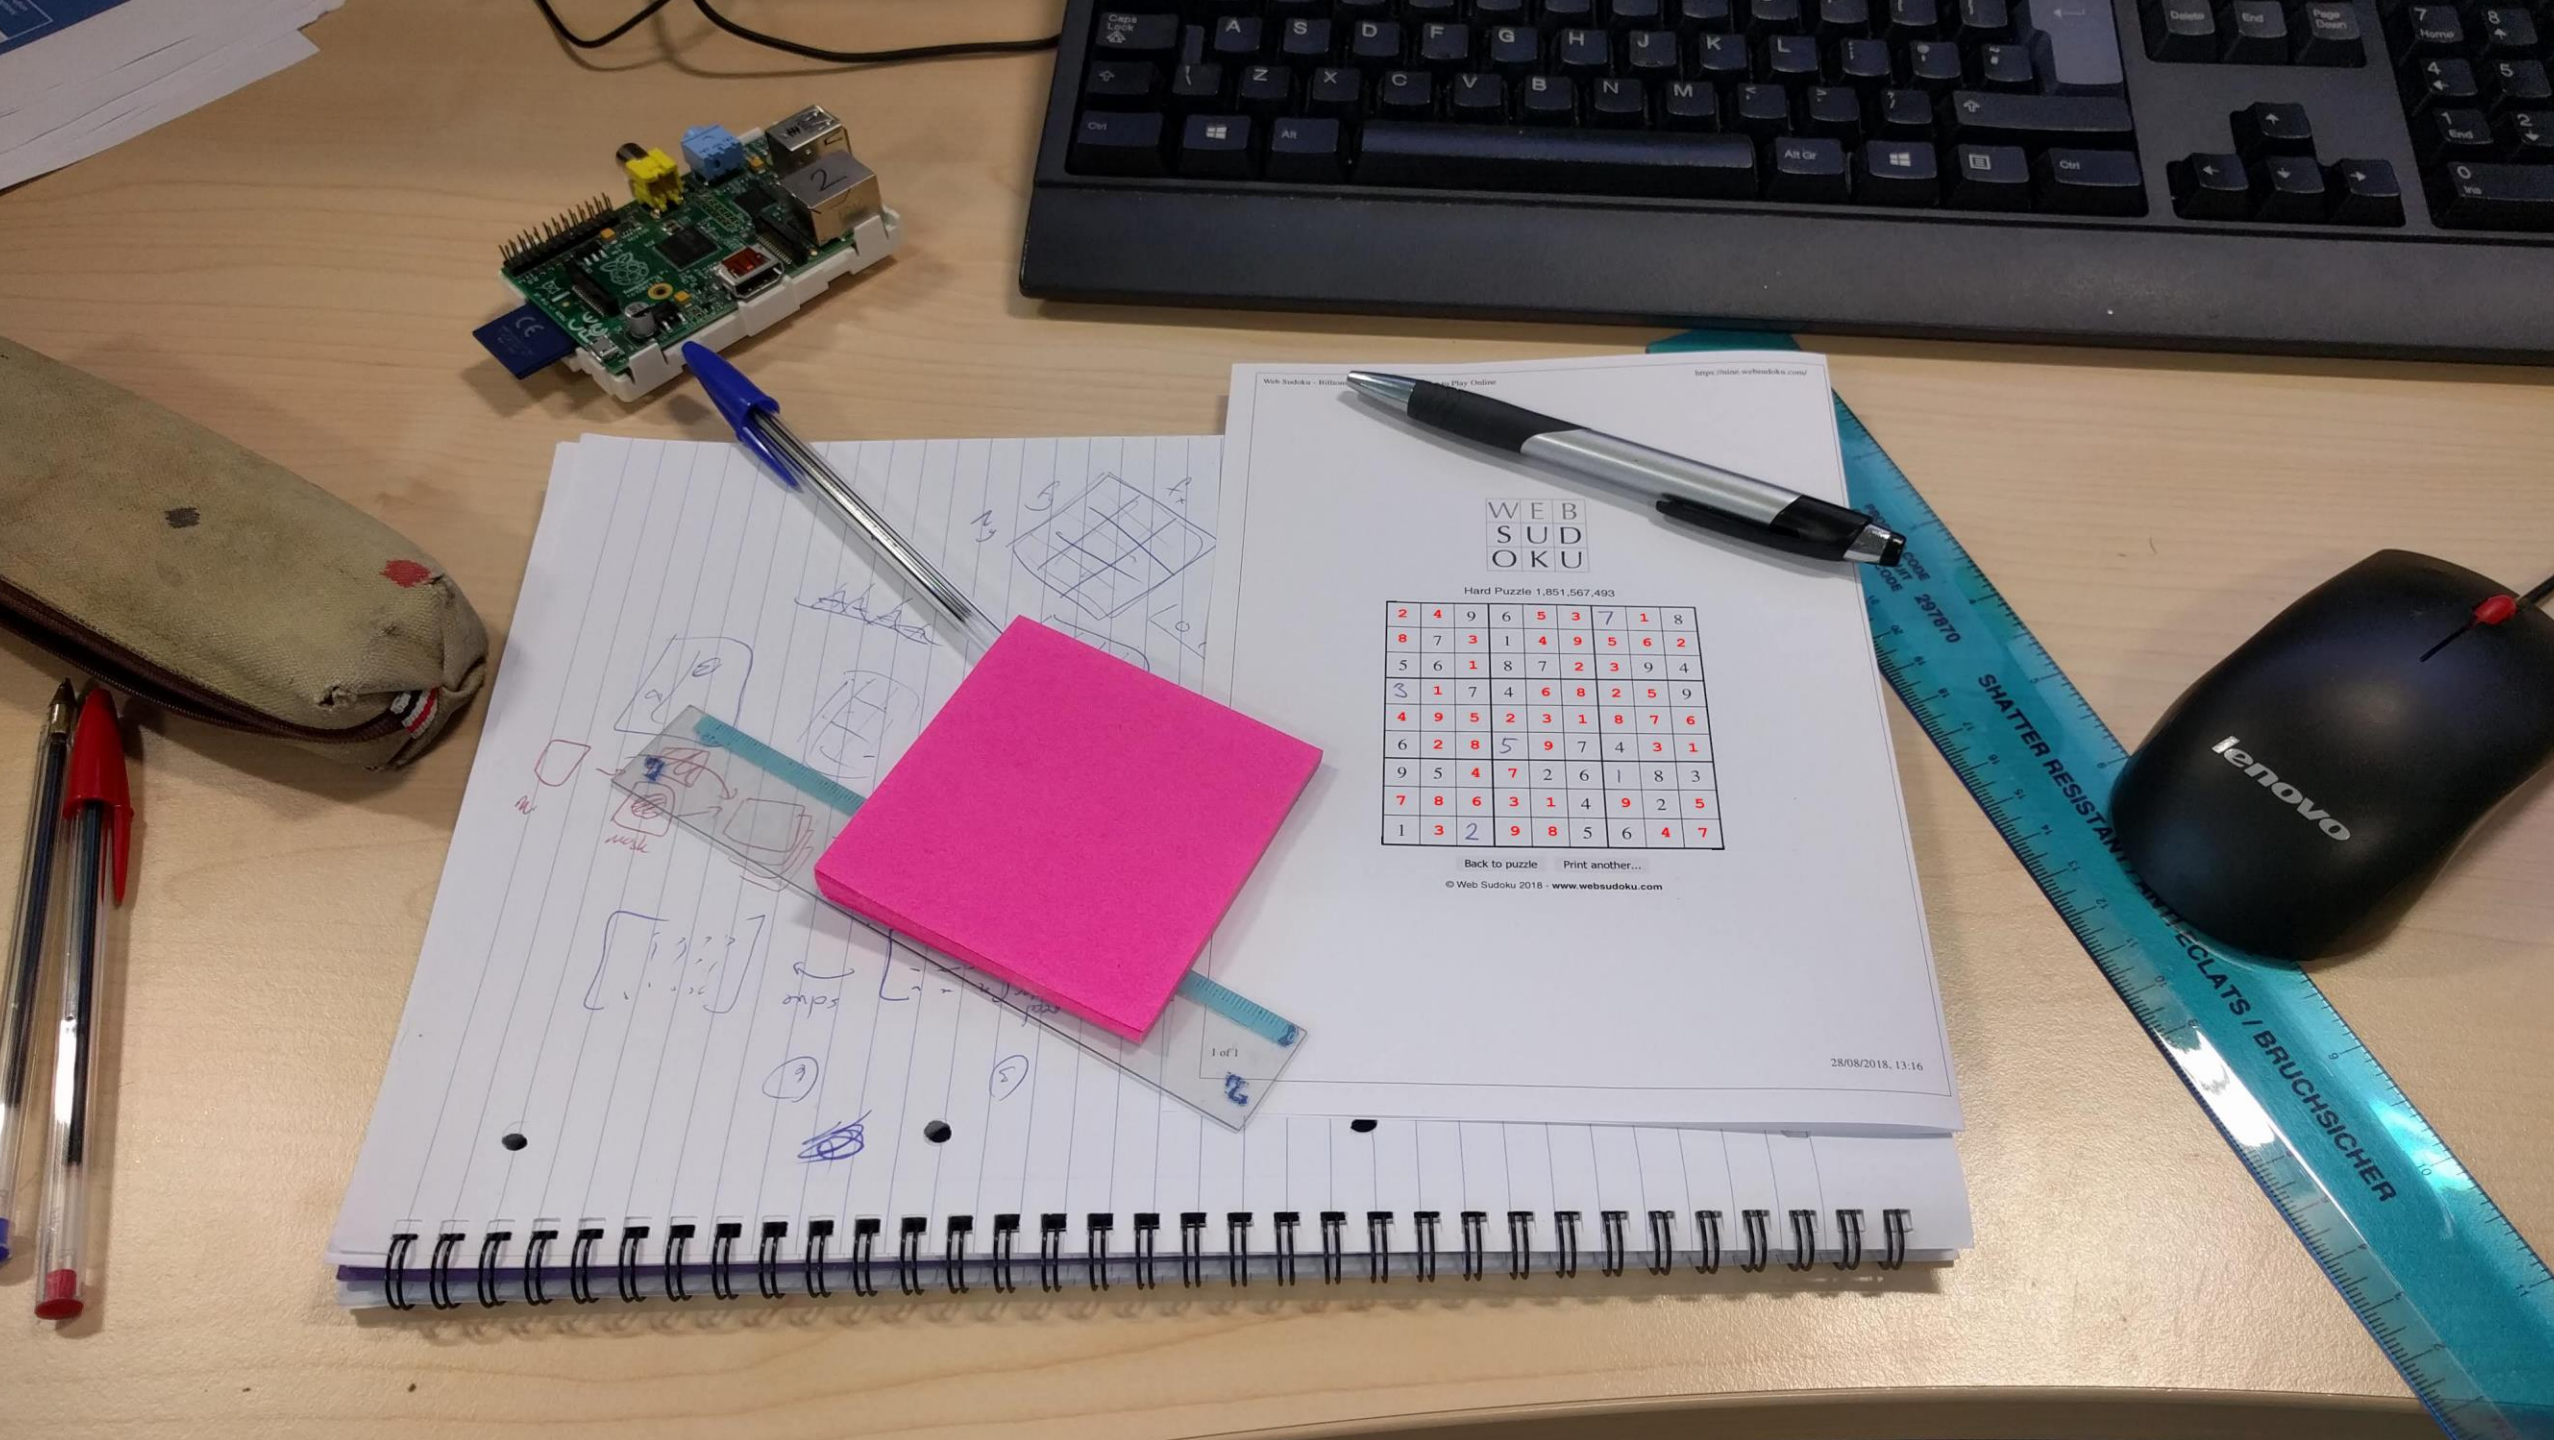
WEB
SUD
OKU
Hard Puzzle 1,851,567,493
Back to puzzle
Print another...
© Web Sudoku 2018 - www.websudoku.com
28/08/2018, 13:16
lenovo
SHATTER RESISTANT
ECLATS / BRUCHSICHER

An Example – Sudoku Solver

Find the puzzle

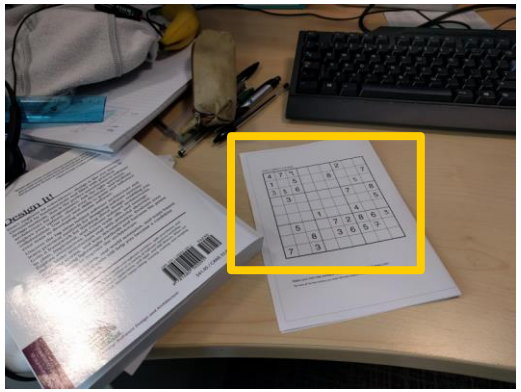

Find the boxes

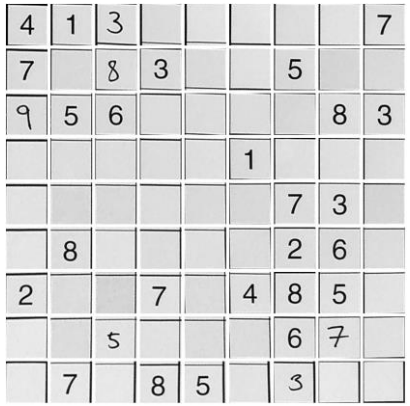

Read the numbers

$$\begin{bmatrix} 4 & 1 & \cdots & ? & ? \\ 7 & ? & \cdots & ? & ? \\ \vdots & \vdots & \cdots & \vdots & \vdots \\ ? & ? & \cdots & 7 & ? \\ ? & 7 & \cdots & ? & ? \end{bmatrix}$$

Solve

$$\begin{bmatrix} 4 & 1 & \cdots & 6 & 8 \\ 7 & 8 & \cdots & 3 & 2 \\ \vdots & \vdots & \cdots & \vdots & \vdots \\ 4 & 2 & \cdots & 7 & 5 \\ 6 & 7 & \cdots & 4 & 1 \end{bmatrix}$$

Deep Learning | **Image Processing** | **Deep Learning** | **Optimisation**

Image Processing and Computer Vision with MATLAB

- Both deep learning and image processing are great tools for computer vision.
- Use the right combination of tools to get the job done.
- MATLAB makes it easy and efficient to do both image processing and deep learning together.
- MATLAB helps you integrate a computer vision algorithm into the rest of your workflow.

Try out the example code
(Search for: Deep Sudoku Solver)

Try what's new in Image Processing and Computer Vision toolboxes

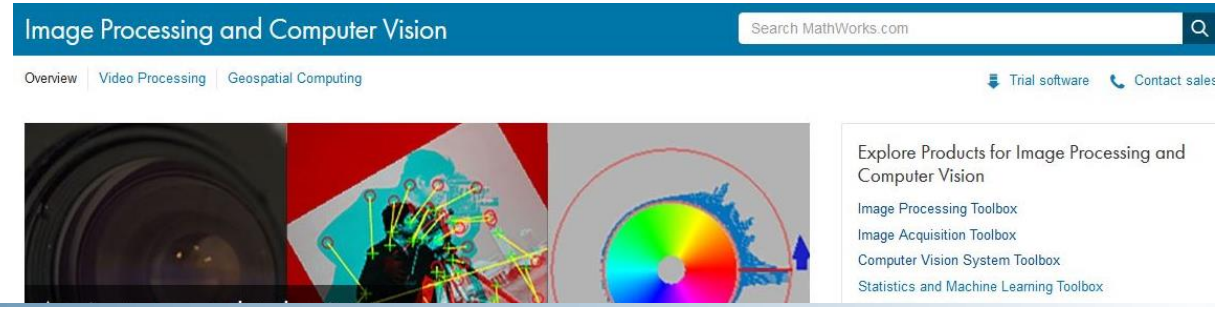

Talk to us about your (computer vision) problems