Vision for Robotics

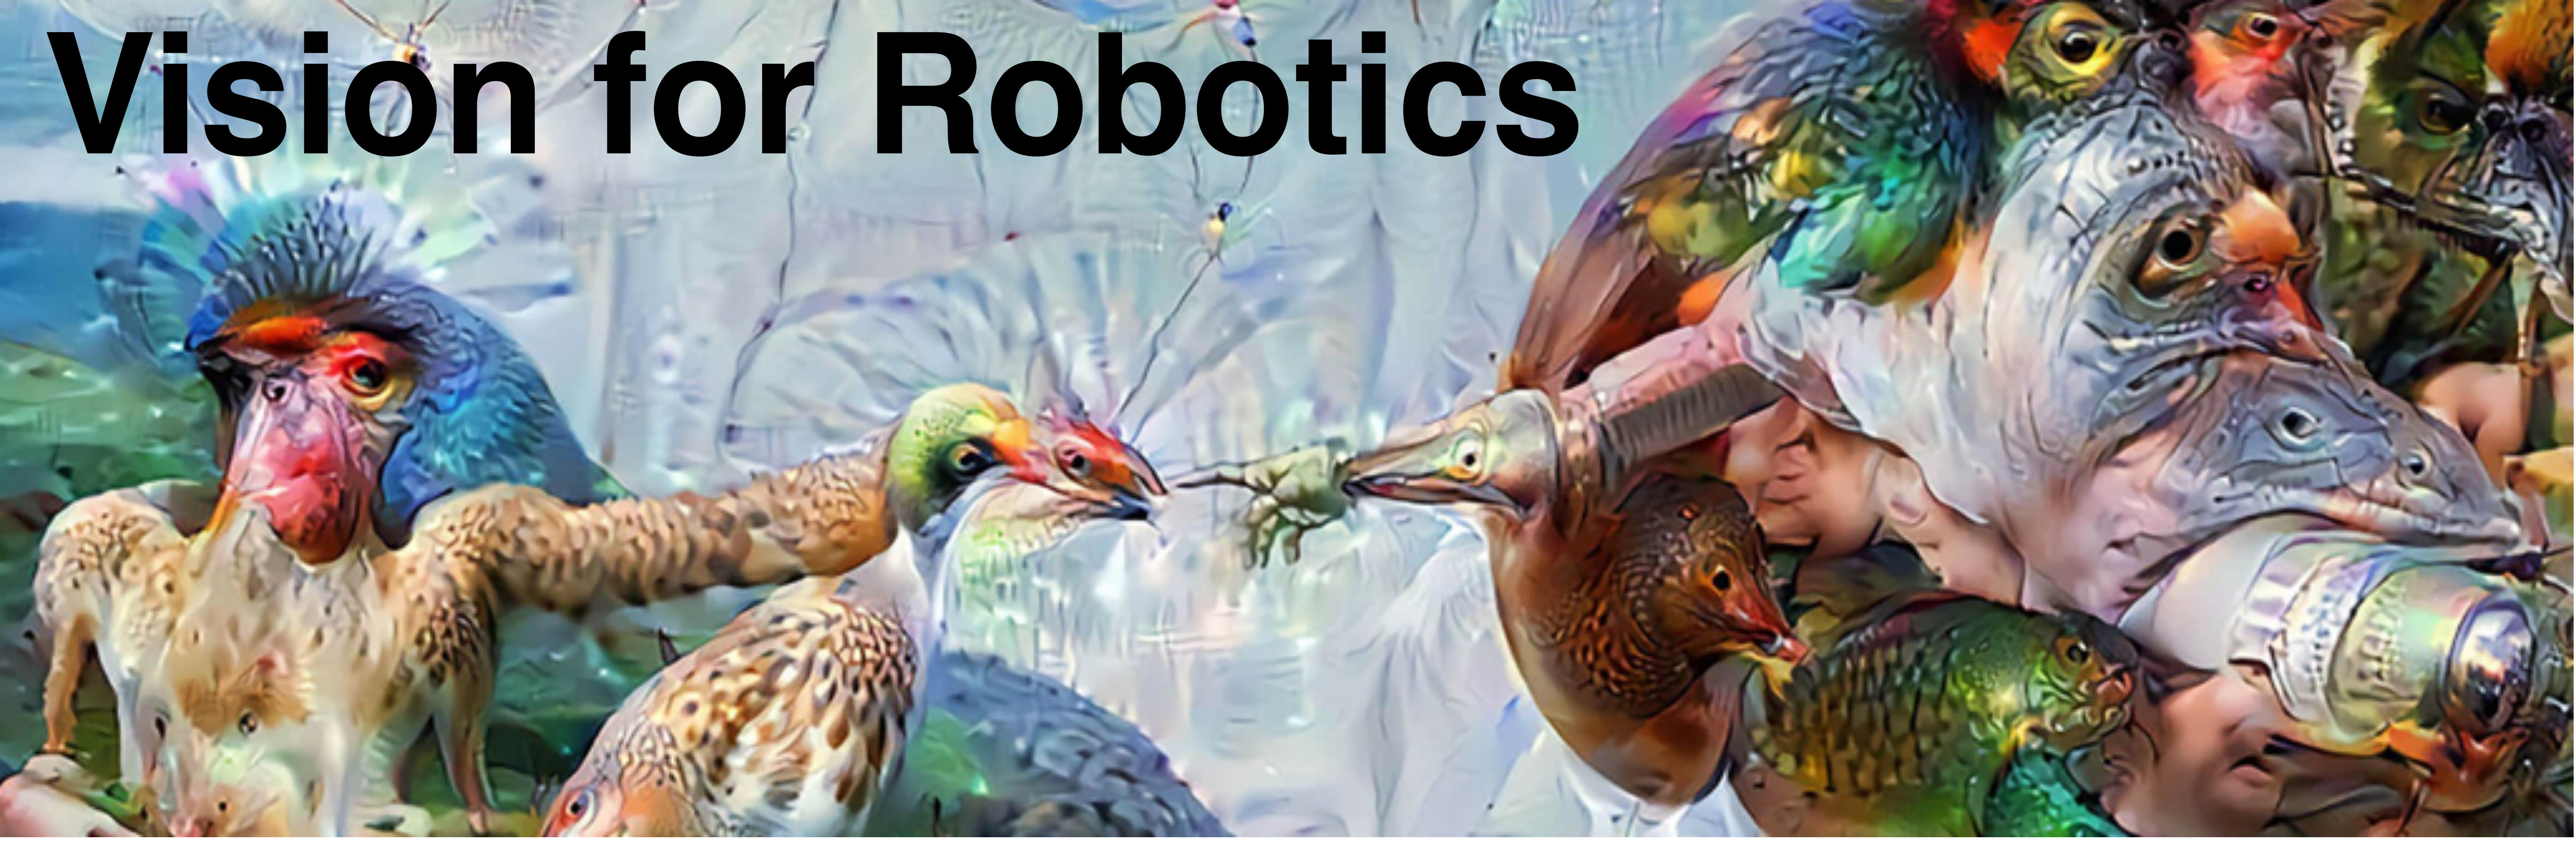

https://cw.fel.cvut.cz/b181/courses/b3b33vir

Karel Zimmermann

Czech Technical University in Prague
Faculty of Electrical Engineering, Department of Cybernetics

Outline

- Introduction of the ViR-team
- Outline of the course
- Organization (homework, tests, semestral work)
- Test0

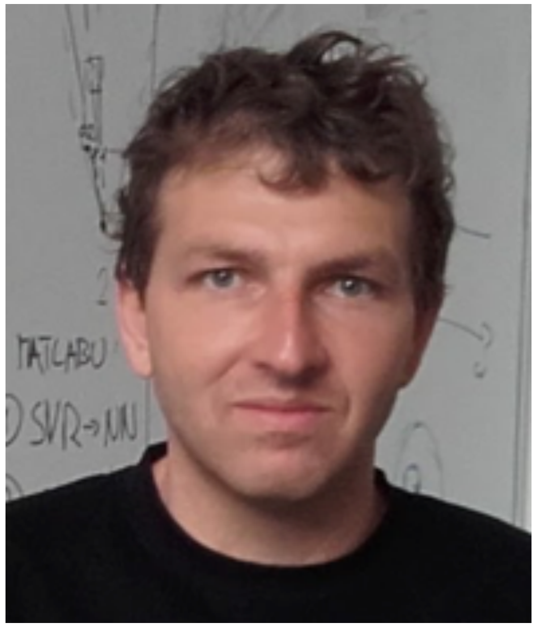

- **Karel Zimmermann**
  (associate professor at CTU)
- main lecturer

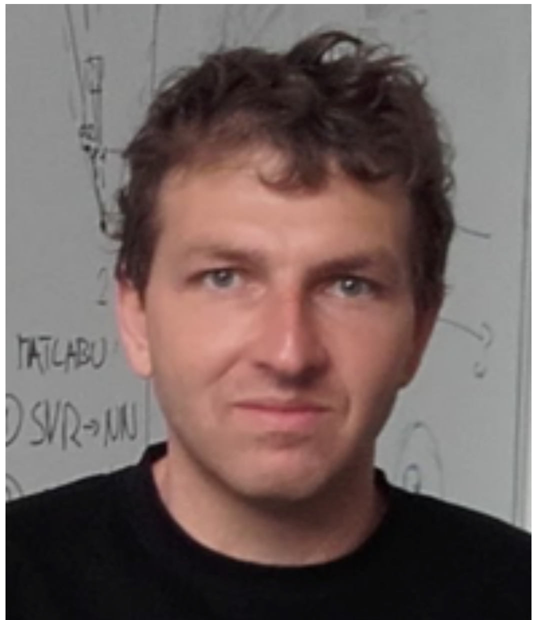

- **Karel Zimmermann**
  (associate professor at CTU)
- main lecturer

- **Tomáš Petříček**
  (postdoctoral researcher in Toyota Research Lab at CTU)
- head of the labs
- 3D mapping & deep learning

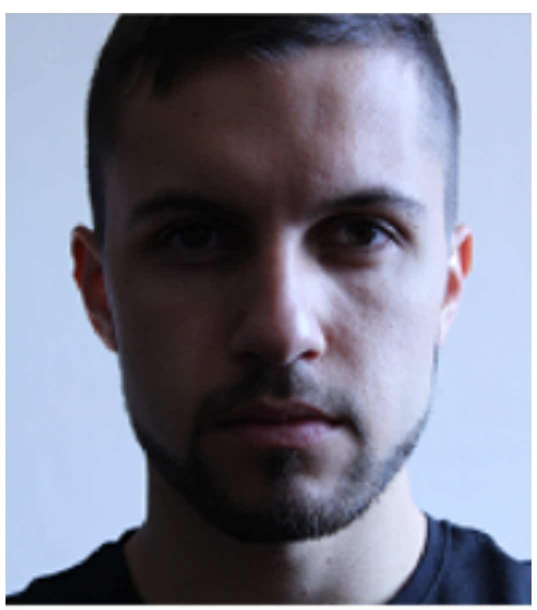

- **Teymur Azayev**
  (PhD student since 2018)
- lab teacher
- motion control & deep learning

- Introduction of the ViR-team
- Outline of the course
- Organization (homework, tests, semestral work)
- Test0

Motivation for unusual organization

- What I did not like as a student:
  - lectures are boring and make me sleepy
  - weak connection between (i) theory (math, statistics, algebra) and (ii) applications (robotics)
  - non-interactive lectures
  - semestral work:
    - second half of lectures cannot be used
    - no space for own creativity
- What I do not like as teacher:
  - lectures are boring and make me sleepy
  - weak motivation of students for continuous studying
  - weak motivation of students for interactive discussions
  - weak motivation of students for originality
  - strong motivation of students for plagiarism

Outline

intense interactive lectures

date	wee	due	Topic
01.10.20	1	T0	Outline, organization of the course.
01.10.20	1		Linear regression and classification + priors
08.10.20	2	T1	Neural nets + backprop
08.10.20	2		Convolutional network + backprop
15.10.20	3	T2	Learning (gradients, adaptive methods) + layers (activation functions,
15.10.20	3		Frameworks for learning
22.10.20	4	HW1	Recognition, detection, segmentation
22.10.20	4		Camera
29.10.20	5	T3	Depth sensors (stereo, kinect, time-of-flight)
29.10.20	5		Mapping I (basic SLAM from camera and depth sensor)
05.11.20	6	HW2	Mapping II (recurrent nets, learnable SLAM, stereo and monodepth)
05.11.20	6		Motion control from images
12.11.20	7		Domain transfer (GANs, time-constrastive nets, style transfer nets,
12.11.20	7		Semestral work assignment
19.11.20	8		— supervisor consultations —
19.11.20	8		— supervisor consultations —
26.11.20	9		— supervisor consultations —
26.11.20	9		— supervisor consultations —
03.12.20	10		— supervisor consultations —
03.12.20	10		— supervisor consultations —
10.12.20	11		— supervisor consultations —
10.12.20	11		— supervisor consultations —
17.12.20	12		— supervisor consultations —
17.12.20	12		— supervisor consultations —
07.01.20	15	SW	Semestral work presentations
07.01.20	15	T4	Exam test

Outline

strong motivation for continuous studying

date	wee	due	Topic
01.10.20	1	T0	Outline, organization of the course.
01.10.20	1		Linear regression and classification + priors
08.10.20	2	T1	Neural nets + backprop
08.10.20	2		Convolutional network + backprop
15.10.20	3	T2	Learning (gradients, adaptive methods) + layers (activation functions,
15.10.20	3		Frameworks for learning
22.10.20	4	HW1	Recognition, detection, segmentation
22.10.20	4		Camera
29.10.20	5	T3	Depth sensors (stereo, kinect, time-of-flight)
29.10.20	5		Mapping I (basic SLAM from camera and depth sensor)
05.11.20	6	HW2	Mapping II (recurrent nets, learnable SLAM, stereo and monodepth)
05.11.20	6		Motion control from images
12.11.20	7		Domain transfer (GANs, time-constrastive nets, style transfer nets,
12.11.20	7		Semestral work assignment
19.11.20	8		— supervisor consultations —
19.11.20	8		— supervisor consultations —
26.11.20	9		— supervisor consultations —
26.11.20	9		— supervisor consultations —
03.12.20	10		— supervisor consultations —
03.12.20	10		— supervisor consultations —
10.12.20	11		— supervisor consultations —
10.12.20	11		— supervisor consultations —
17.12.20	12		— supervisor consultations —
17.12.20	12		— supervisor consultations —
07.01.20	15	SW	Semestral work presentations
07.01.20	15	T4	Exam test

Outline

date	wee	due	Topic
01.10.20	1	T0	Outline, organization of the course.
01.10.20	1		Linear regression and classification + priors
08.10.20	2	T1	Neural nets + backprop
08.10.20	2		Convolutional network + backprop
15.10.20	3	T2	Learning (gradients, adaptive methods) + layers (activation functions,
15.10.20	3		Frameworks for learning
22.10.20	4	HW1	Recognition, detection, segmentation
22.10.20	4		Camera
29.10.20	5	T3	Depth sensors (stereo, kinect, time-of-flight)
29.10.20	5		Mapping I (basic SLAM from camera and depth sensor)
05.11.20	6	HW2	Mapping II (recurrent nets, learnable SLAM, stereo and monodepth)
05.11.20	6		Motion control from images
12.11.20	7		Domain transfer (GANs, time-constrastive nets, style transfer nets,
12.11.20	7		Semestral work assignment
19.11.20	8		— supervisor consultations —
19.11.20	8		— supervisor consultations —
26.11.20	9		— supervisor consultations —
26.11.20	9		— supervisor consultations —
03.12.20	10		— supervisor consultations —
03.12.20	10		— supervisor consultations —
10.12.20	11		— supervisor consultations —
10.12.20	11		— supervisor consultations —
17.12.20	12		— supervisor consultations —
17.12.20	12		— supervisor consultations —
07.01.20	15	SW	Semestral work presentations
07.01.20	15	T4	Exam test

creative original work

Outline

date	wee	due	Topic
01.10.20	1	T0	Outline, organization of the course.
01.10.20	1		Linear regression and classification + priors
08.10.20	2	T1	Neural nets + backprop
08.10.20	2		Convolutional network + backprop
15.10.20	3	T2	Learning (gradients, adaptive methods) + layers (activation functions,
15.10.20	3		Frameworks for learning
22.10.20	4	HW1	Recognition, detection, segmentation
22.10.20	4		Camera
29.10.20	5	T3	Depth sensors (stereo, kinect, time-of-flight)
29.10.20	5		Mapping I (basic SLAM from camera and depth sensor)
05.11.20	6	HW2	Mapping II (recurrent nets, learnable SLAM, stereo and monodepth)
05.11.20	6		Motion control from images
12.11.20	7		Domain transfer (GANs, time-constrastive nets, style transfer nets,
12.11.20	7		Semestral work assignment
19.11.20	8		— supervisor consultations —
19.11.20	8		— supervisor consultations —
26.11.20	9		— supervisor consultations —
26.11.20	9		— supervisor consultations —
03.12.20	10		— supervisor consultations —
03.12.20	10		— supervisor consultations —
10.12.20	11		— supervisor consultations —
10.12.20	11		— supervisor consultations —
17.12.20	12		— supervisor consultations —
17.12.20	12		— supervisor consultations —
07.01.20	15	SW	Semestral work presentations
07.01.20	15	T4	Exam test

Outline

date	wee	due	Topic
01.10.20	1	T0	Outline, organization of the course.
01.10.20	1		Linear regression and classification + priors
08.10.20	2	T1	Neural nets + backprop
08.10.20	2		Convolutional network + backprop
15.10.20	3	T2	Learning (gradients, adaptive methods) + layers (activation functions,
15.10.20	3		Frameworks for learning
22.10.20	4	HW1	Recognition, detection, segmentation
22.10.20	4		Camera
29.10.20	5	T3	Depth sensors (stereo, kinect, time-of-flight)
29.10.20	5		Mapping I (basic SLAM from camera and depth sensor)
05.11.20	6	HW2	Mapping II (recurrent nets, learnable SLAM, stereo and monodepth)
05.11.20	6		Motion control from images
12.11.20	7		Domain transfer (GANs, time-constrastive nets, style transfer nets,
12.11.20	7		Semestral work assignment
19.11.20	8		— supervisor consultations —
19.11.20	8		— supervisor consultations —
26.11.20	9		— supervisor consultations —
26.11.20	9		— supervisor consultations —
03.12.20	10		— supervisor consultations —
03.12.20	10		— supervisor consultations —
10.12.20	11		— supervisor consultations —
10.12.20	11		— supervisor consultations —
17.12.20	12		— supervisor consultations —
17.12.20	12		— supervisor consultations —
07.01.20	15	SW	Semestral work presentations
07.01.20	15	T4	Exam test

We will formulate classification and regression problems as Bayesian parameter estimation of a probability distribution

- Pre-requisites:
  - ALG (basic linear algebra)
  - PSI (probability, ML and Bayes rule)

$$\mathbf{w}^* = \arg\min_{\mathbf{w}} \left(\sum_i -\log(p(y_i|\mathbf{x}_i, \mathbf{w})) \right) + (-\log p(\mathbf{w}))$$

loss function prior/regulariser

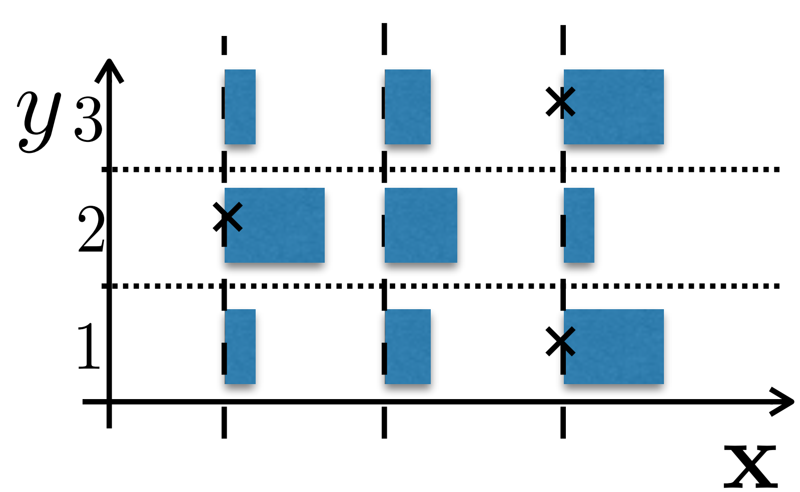

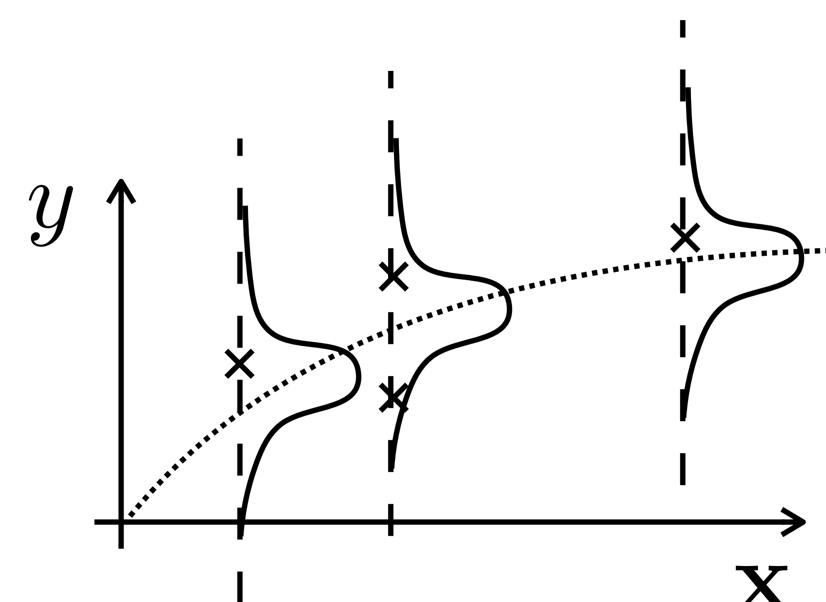

Czech Technical University in Prague
Faculty of Electrical Engineering, Department of Cybernetics

Learning of single neuron

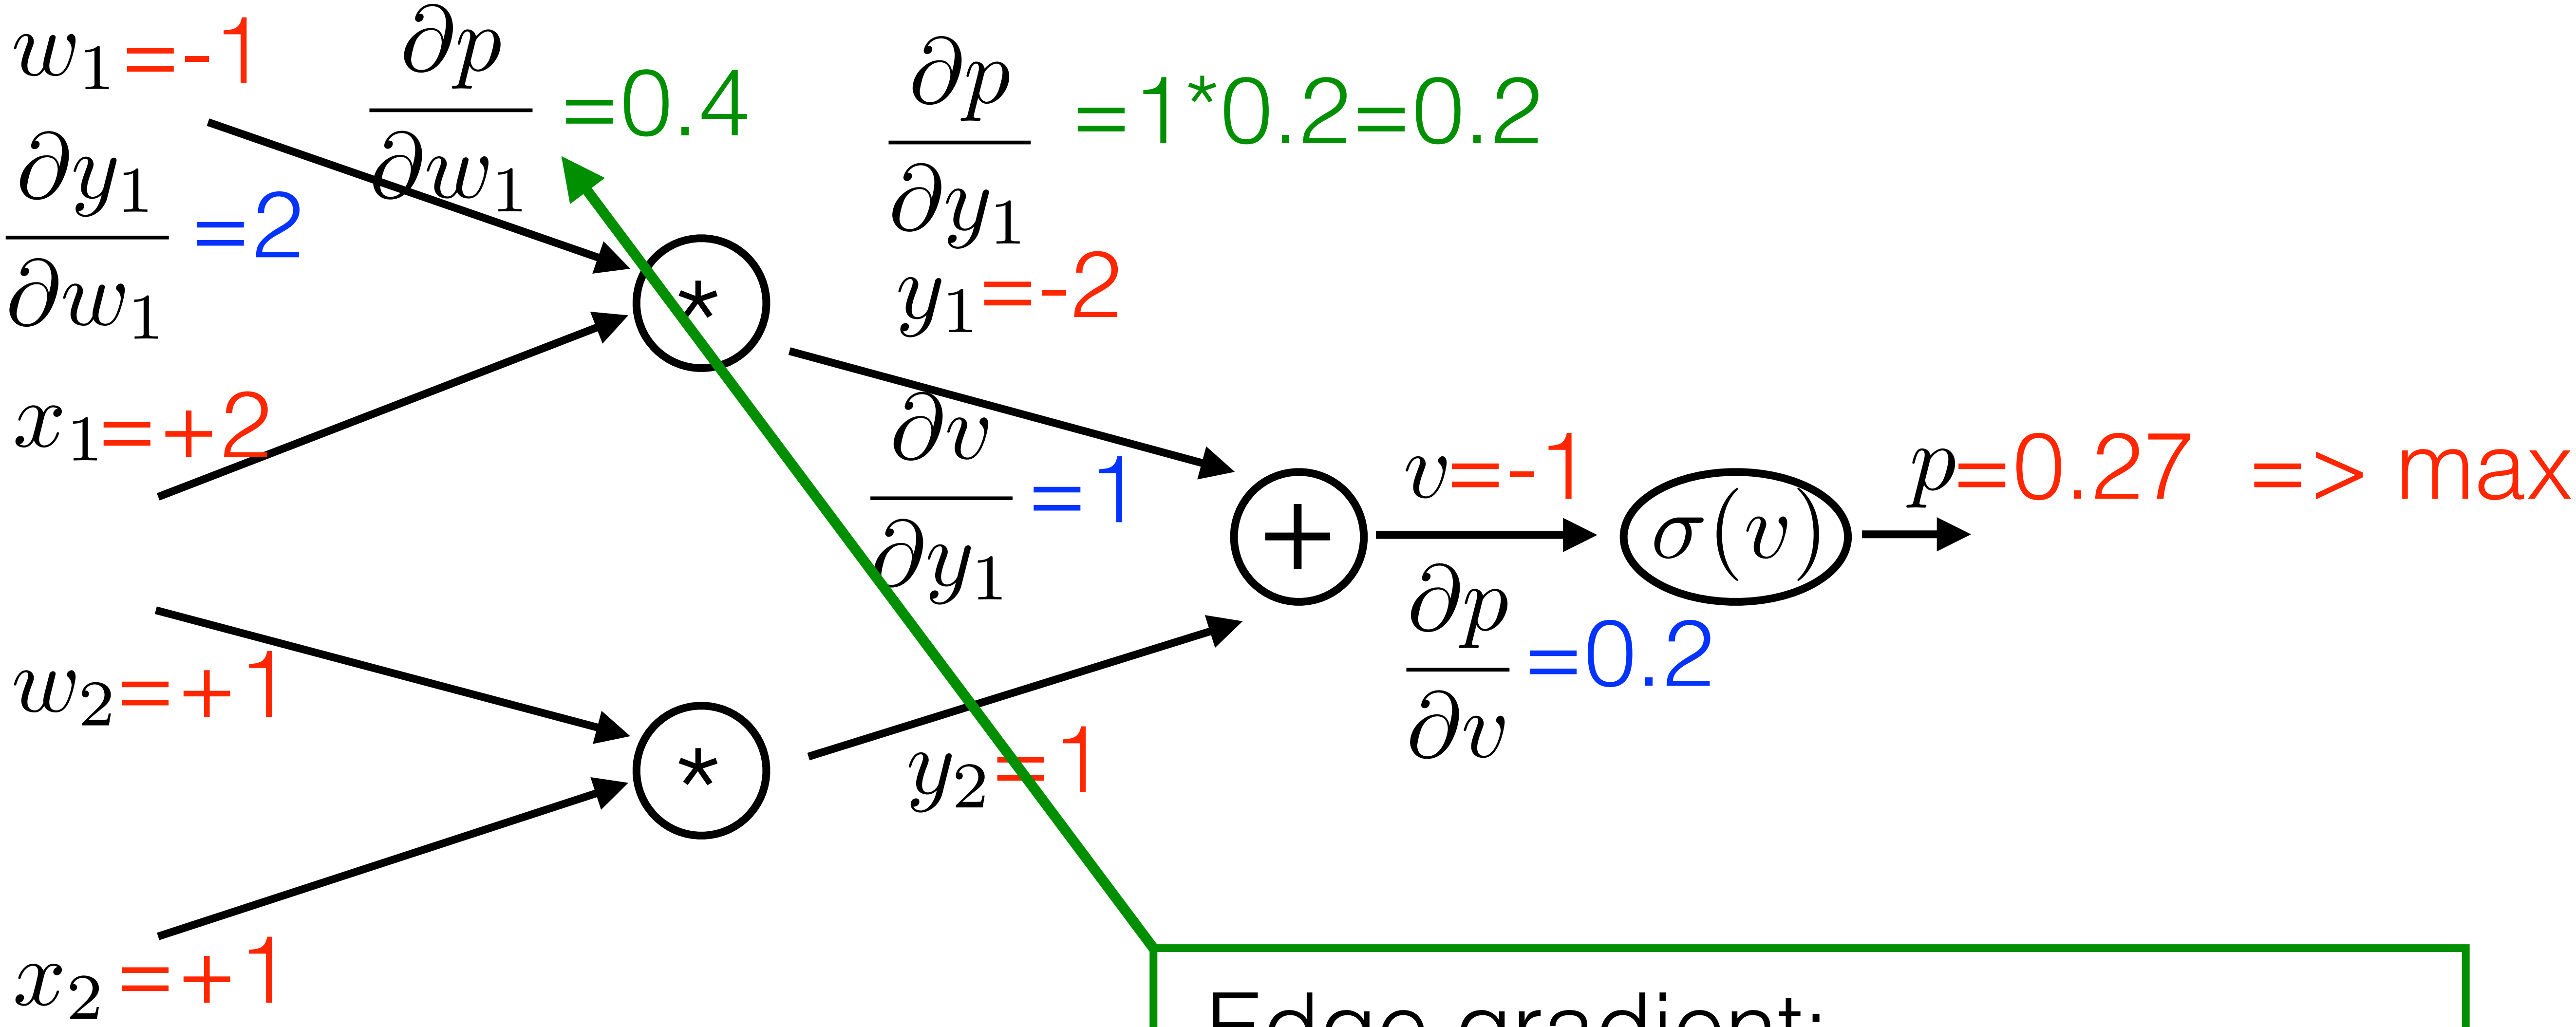

Edge gradient:

$$\frac{\partial p}{\partial w_1} = \frac{\partial p}{\partial y_1}\frac{\partial y_1}{\partial w_1}$$

$$w_1 = w_1 + \frac{\partial p}{\partial w1}$$

Chain-rule in computational graph $\frac{\partial p}{\partial w_1} = \frac{\partial p}{\partial v}\frac{\partial v}{\partial y_1}\frac{\partial y_1}{\partial w_1}$

Learning of fully connected neural network

- Pre-requisites:
  - Math II (partial derivatives, chain-rule)

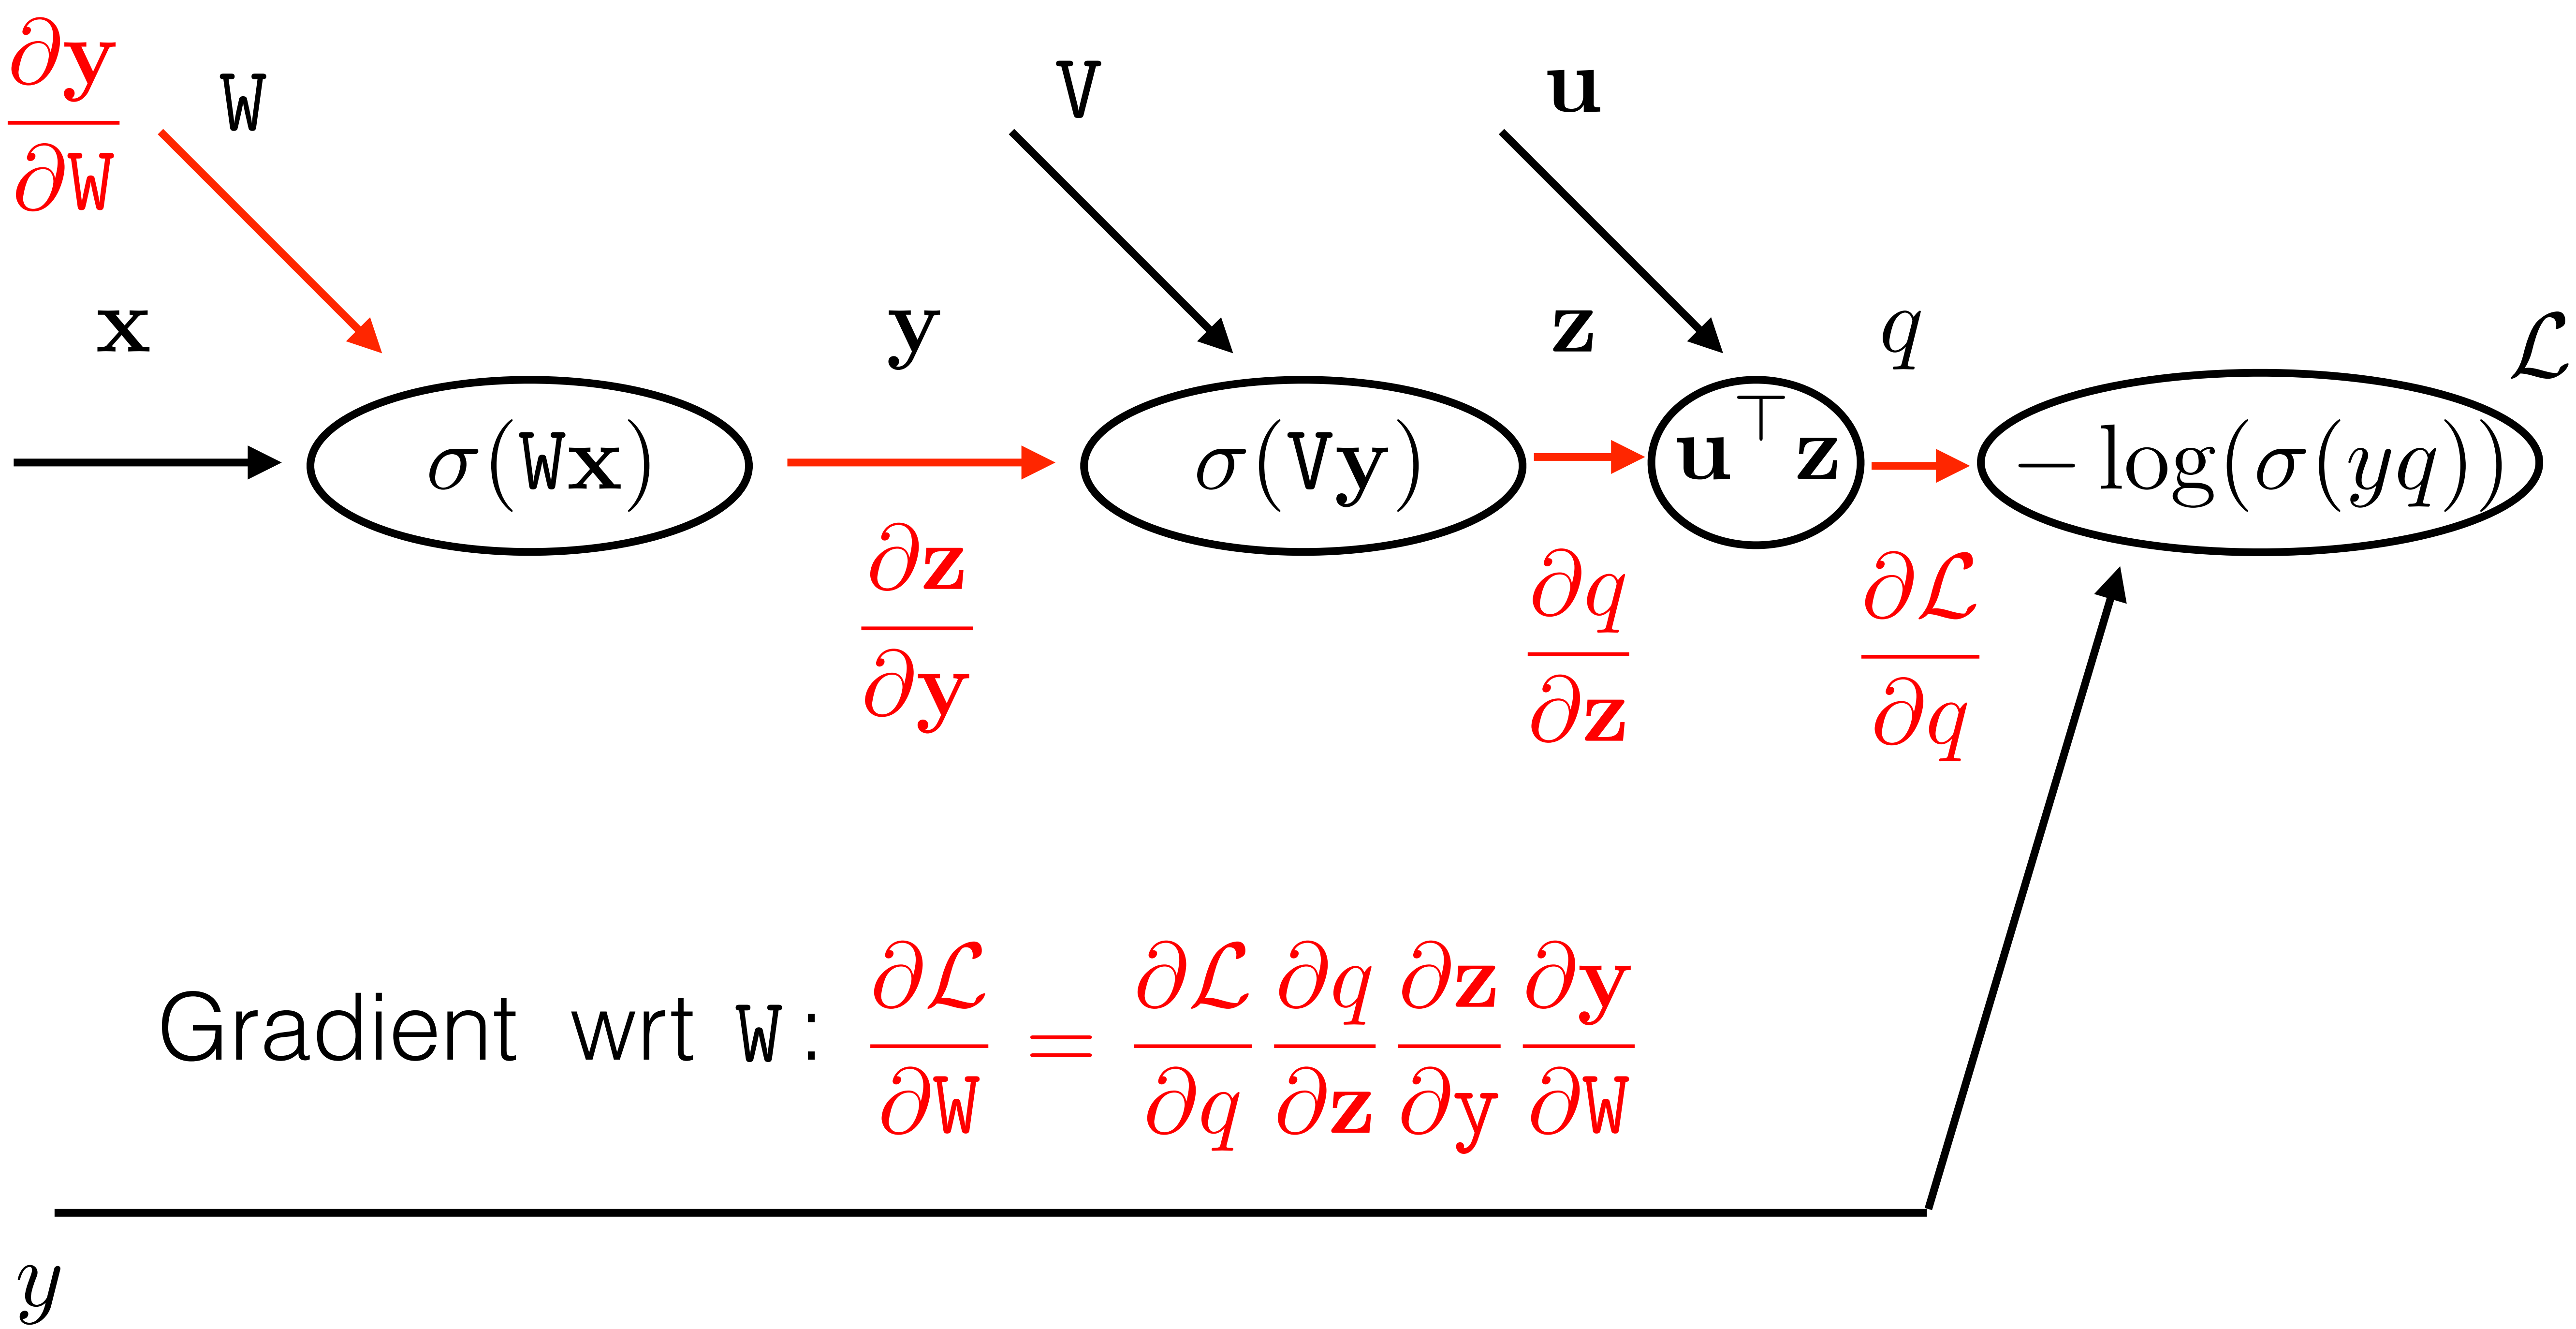

Gradient wrt W: $\frac{\partial \mathcal{L}}{\partial \mathtt{W}} = \frac{\partial \mathcal{L}}{\partial q}\frac{\partial q}{\partial \mathbf{z}}\frac{\partial \mathbf{z}}{\partial \mathrm{y}}\frac{\partial \mathbf{y}}{\partial \mathtt{W}}$

- We will understand why NN does not work
- We will study the cat experiment [Hubel&Wiesel 1960]
- Justify using the convolutional prior/layer in neural nets

We will dive deep into learning

You will exploit what you have learnt in the optimization course

$$\mathbf{w}_{t+1} \approx \mathbf{w}_t - \alpha \left[\mathrm{diag}\left(\nabla \mathbf{w}_t \nabla \mathbf{w}_t^\top\right)^{1/2}\right]^{-1} \left.\frac{\partial f(\mathbf{w})}{\partial \mathbf{w}}\right|_{\mathbf{w}=\mathbf{w}_t}$$

$$\mathbf{w}_{t+1} \approx \mathbf{w}_t - \frac{\alpha}{\sqrt{\nabla \mathbf{w}_t^2 + \epsilon}} \odot \nabla \mathbf{w}_t$$

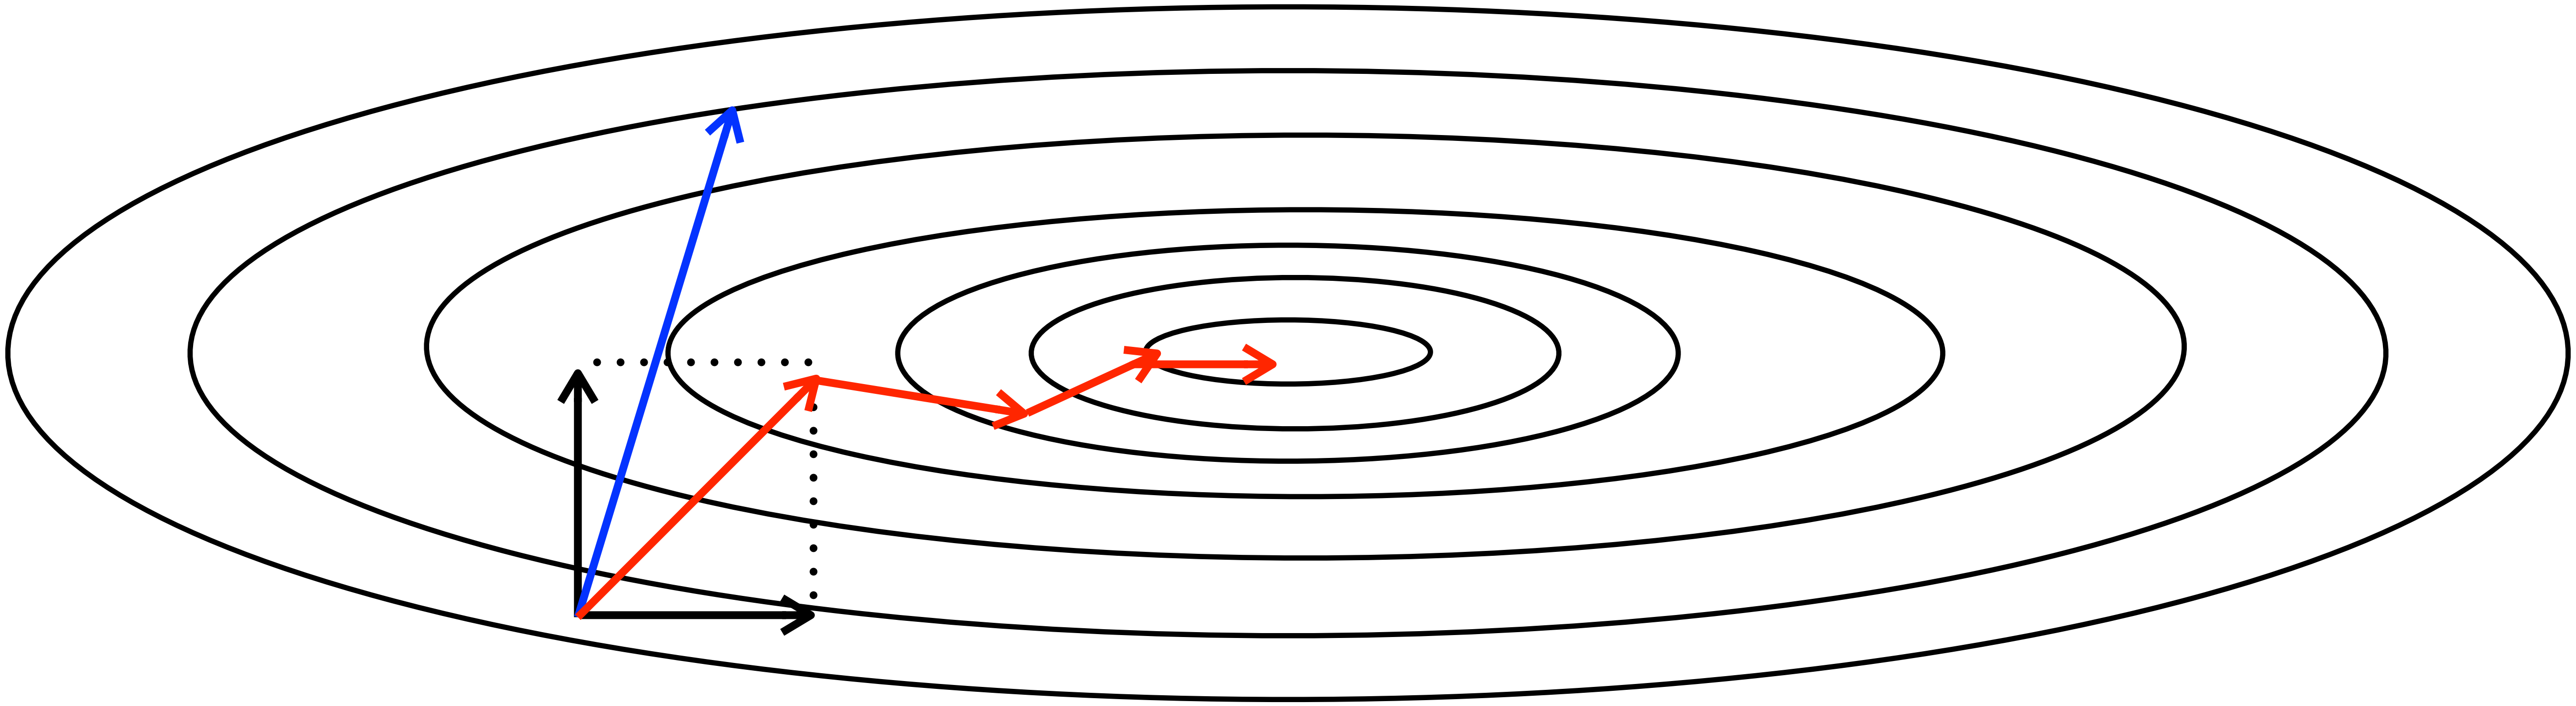

http://www.jmlr.org/papers/volume12/duchi11a/duchi11a.pdf

We will study influence of different layers on the learning

Sigmoid

$\sigma(x) = \frac{1}{1+e^{-x}}$

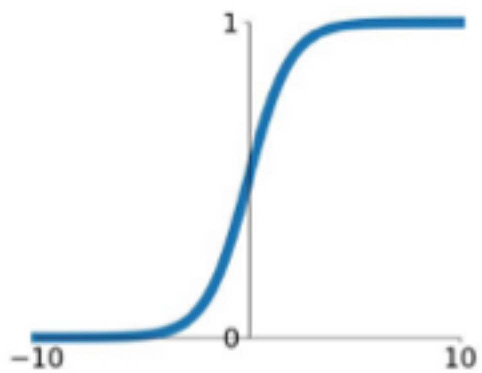

- zero gradient when saturated
- undesired zig-zag behaviour

$$\frac{\partial p}{\partial w_1} = x_1 \cdot 1 \cdot \frac{\partial p}{\partial v} \quad \begin{matrix} >0 \\ <0 \end{matrix}$$

$$\frac{\partial p}{\partial w_2} = x_2 \cdot 1 \cdot \frac{\partial p}{\partial v} \quad \begin{matrix} >0 \\ <0 \end{matrix}$$

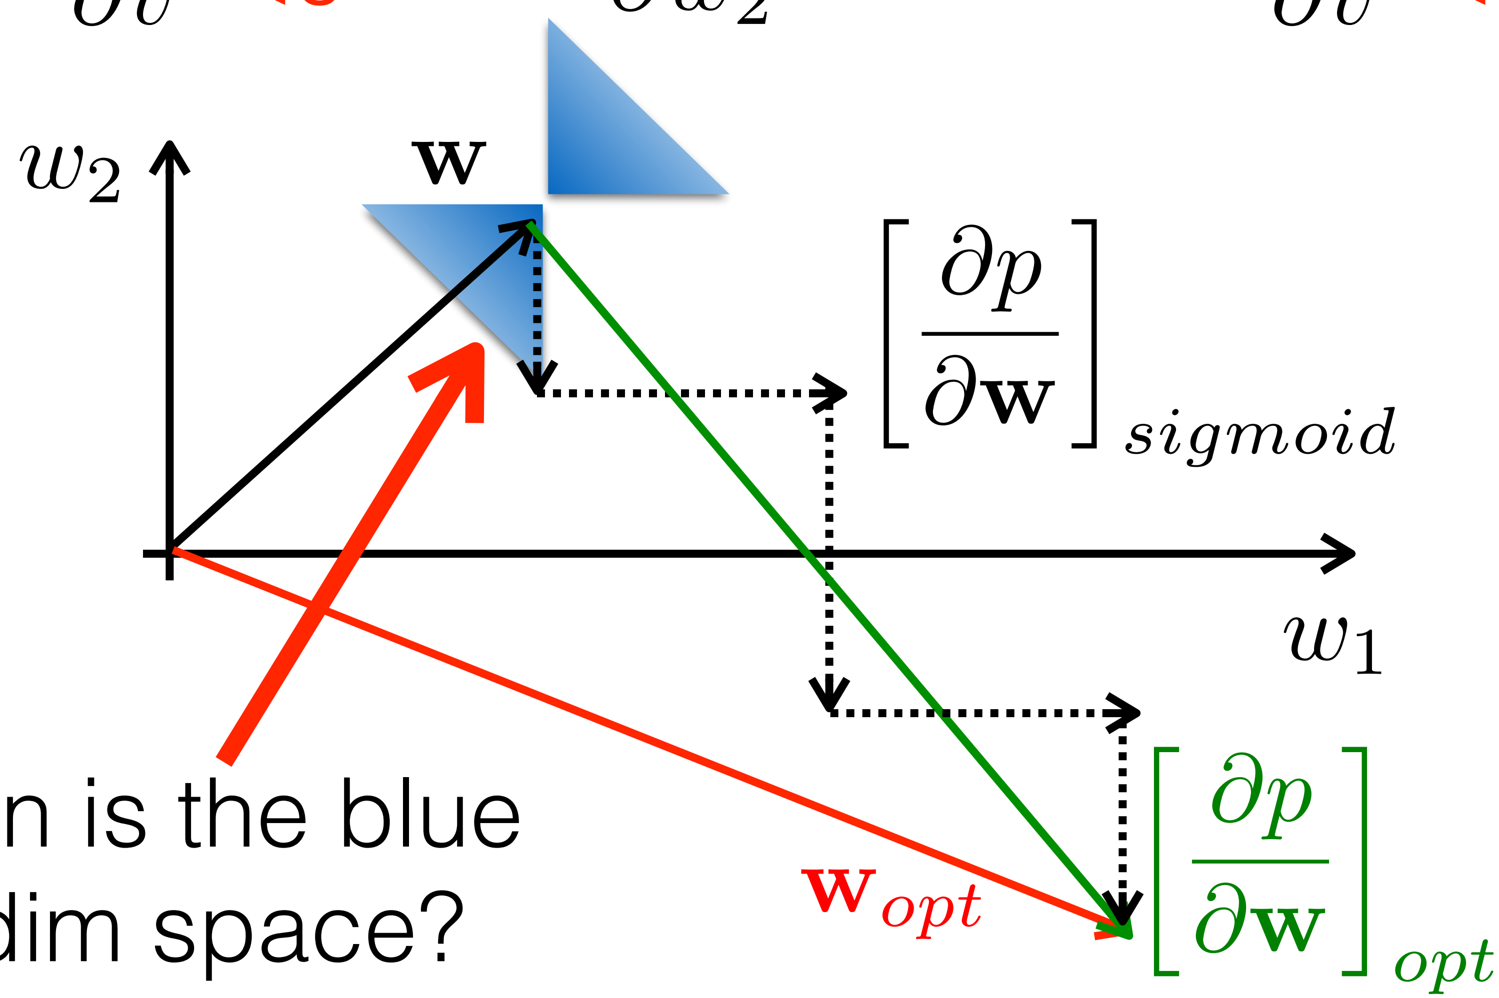

how big fraction is the blue region in 10-dim space?

Outline

date	wee	due	Topic
01.10.20	1	T0	Outline, organization of the course.
01.10.20	1		Linear regression and classification + priors
08.10.20	2	T1	Neural nets + backprop
08.10.20	2		Convolutional network + backprop
15.10.20	3	T2	Learning (gradients, adaptive methods) + layers (activation functions,
15.10.20	3		Frameworks for learning
22.10.20	4	HW1	Recognition, detection, segmentation
22.10.20	4		Camera
29.10.20	5	T3	Depth sensors (stereo, kinect, time-of-flight)
29.10.20	5		Mapping I (basic SLAM from camera and depth sensor)
05.11.20	6	HW2	Mapping II (recurrent nets, learnable SLAM, stereo and monodepth)
05.11.20	6		Motion control from images
12.11.20	7		Domain transfer (GANs, time-constrastive nets, style transfer nets,
12.11.20	7		Semestral work assignment
19.11.20	8		— supervisor consultations —
19.11.20	8		— supervisor consultations —
26.11.20	9		— supervisor consultations —
26.11.20	9		— supervisor consultations —
03.12.20	10		— supervisor consultations —
03.12.20	10		— supervisor consultations —
10.12.20	11		— supervisor consultations —
10.12.20	11		— supervisor consultations —
17.12.20	12		— supervisor consultations —
17.12.20	12		— supervisor consultations —
07.01.20	15	SW	Semestral work presentations
07.01.20	15	T4	Exam test

Frameworks

- Interactive lecture given by Teymur
- Personal / school computers
- You will use Python and PyTorch (install it in advance)
- Semestral work and the urgent need for GPU?:
  - Google Colab (12hours of GPU/day free for students)
  - GPUs in our grid (metacentrum??)

```
import numpy as np
```

http://www.numpy.org

```
import torch
```

https://pytorch.org/

PyCharm https://www.jetbrains.com/pycharm/

https://colab.research.google.com/

Outline

date	wee	due	Topic
01.10.20	1	T0	Outline, organization of the course.
01.10.20	1		Linear regression and classification + priors
08.10.20	2	T1	Neural nets + backprop
08.10.20	2		Convolutional network + backprop
15.10.20	3	T2	Learning (gradients, adaptive methods) + layers (activation functions,
15.10.20	3		Frameworks for learning
22.10.20	4	HW1	Recognition, detection, segmentation
22.10.20	4		Camera
29.10.20	5	T3	Depth sensors (stereo, kinect, time-of-flight)
29.10.20	5		Mapping I (basic SLAM from camera and depth sensor)
05.11.20	6	HW2	Mapping II (recurrent nets, learnable SLAM, stereo and monodepth)
05.11.20	6		Motion control from images
12.11.20	7		Domain transfer (GANs, time-constrastive nets, style transfer nets,
12.11.20	7		Semestral work assignment
19.11.20	8		— supervisor consultations —
19.11.20	8		— supervisor consultations —
26.11.20	9		— supervisor consultations —
26.11.20	9		— supervisor consultations —
03.12.20	10		— supervisor consultations —
03.12.20	10		— supervisor consultations —
10.12.20	11		— supervisor consultations —
10.12.20	11		— supervisor consultations —
17.12.20	12		— supervisor consultations —
17.12.20	12		— supervisor consultations —
07.01.20	15	SW	Semestral work presentations
07.01.20	15	T4	Exam test

We will study winning architectures in recognition, object detection and semantic segmentation

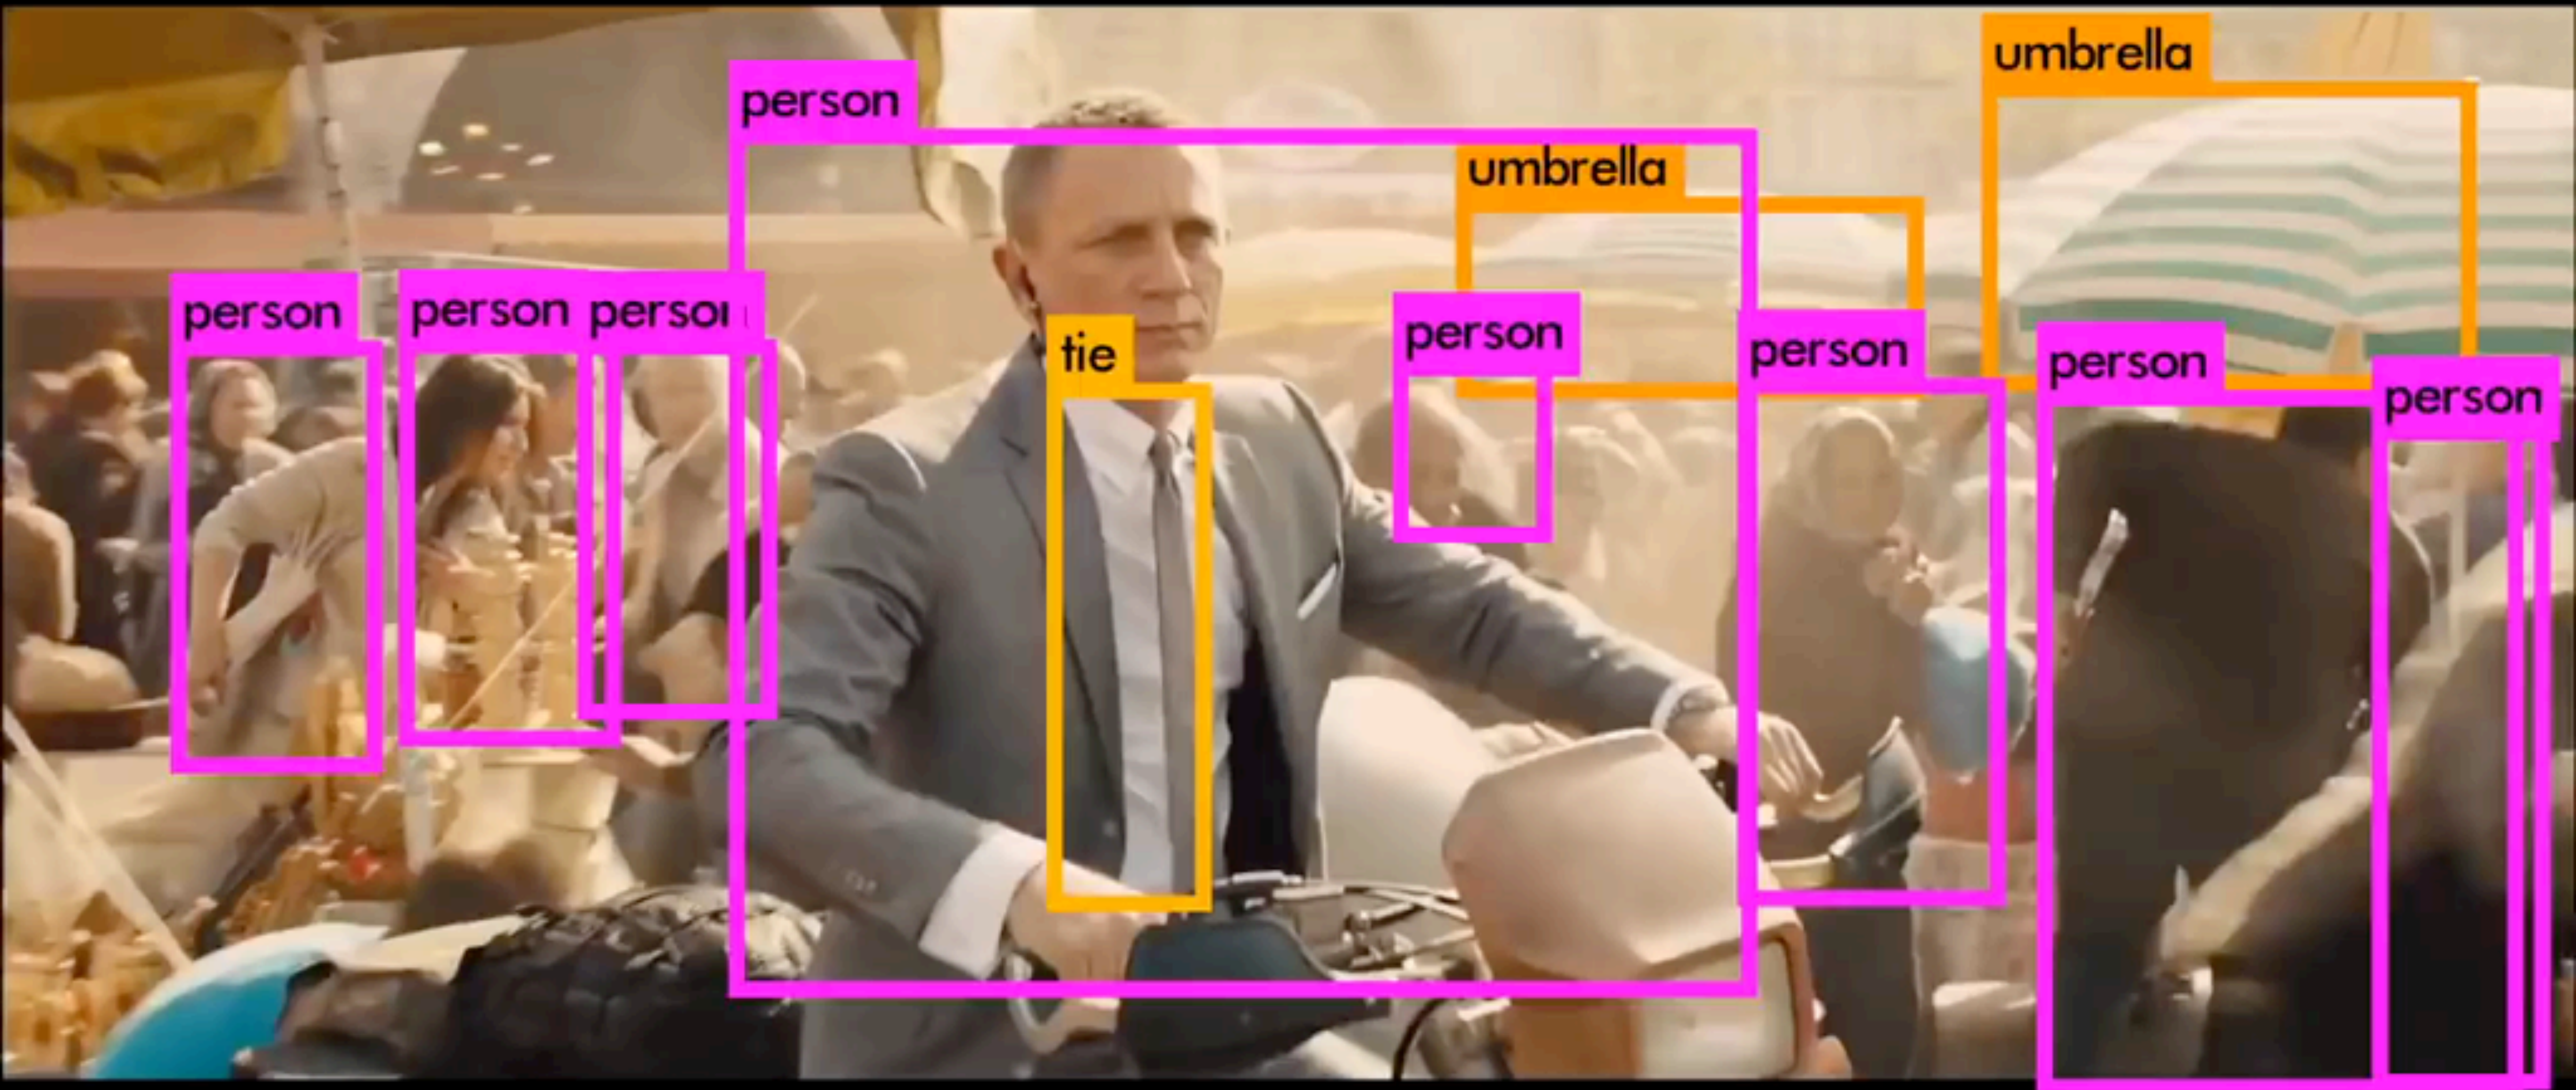

Outline

date	wee	due	Topic
01.10.20	1	T0	Outline, organization of the course.
01.10.20	1		Linear regression and classification + priors
08.10.20	2	T1	Neural nets + backprop
08.10.20	2		Convolutional network + backprop
15.10.20	3	T2	Learning (gradients, adaptive methods) + layers (activation functions,
15.10.20	3		Frameworks for learning
22.10.20	4	HW1	Recognition, detection, segmentation
22.10.20	4		Camera
29.10.20	5	T3	Depth sensors (stereo, kinect, time-of-flight)
29.10.20	5		Mapping I (basic SLAM from camera and depth sensor)
05.11.20	6	HW2	Mapping II (recurrent nets, learnable SLAM, stereo and monodepth)
05.11.20	6		Motion control from images
12.11.20	7		Domain transfer (GANs, time-constrastive nets, style transfer nets,
12.11.20	7		Semestral work assignment
19.11.20	8		— supervisor consultations —
19.11.20	8		— supervisor consultations —
26.11.20	9		— supervisor consultations —
26.11.20	9		— supervisor consultations —
03.12.20	10		— supervisor consultations —
03.12.20	10		— supervisor consultations —
10.12.20	11		— supervisor consultations —
10.12.20	11		— supervisor consultations —
17.12.20	12		— supervisor consultations —
17.12.20	12		— supervisor consultations —
07.01.20	15	SW	Semestral work presentations
07.01.20	15	T4	Exam test

Meaning of a map

- Many robotic problems need to summarize captured knowledge about the environment in some kind of “map”
  - Explicit map (e.g. point-cloud with detected objects, or traversability estimated from optimal trajectories)

Semantic segmentation [Tateno CVPR 2017]

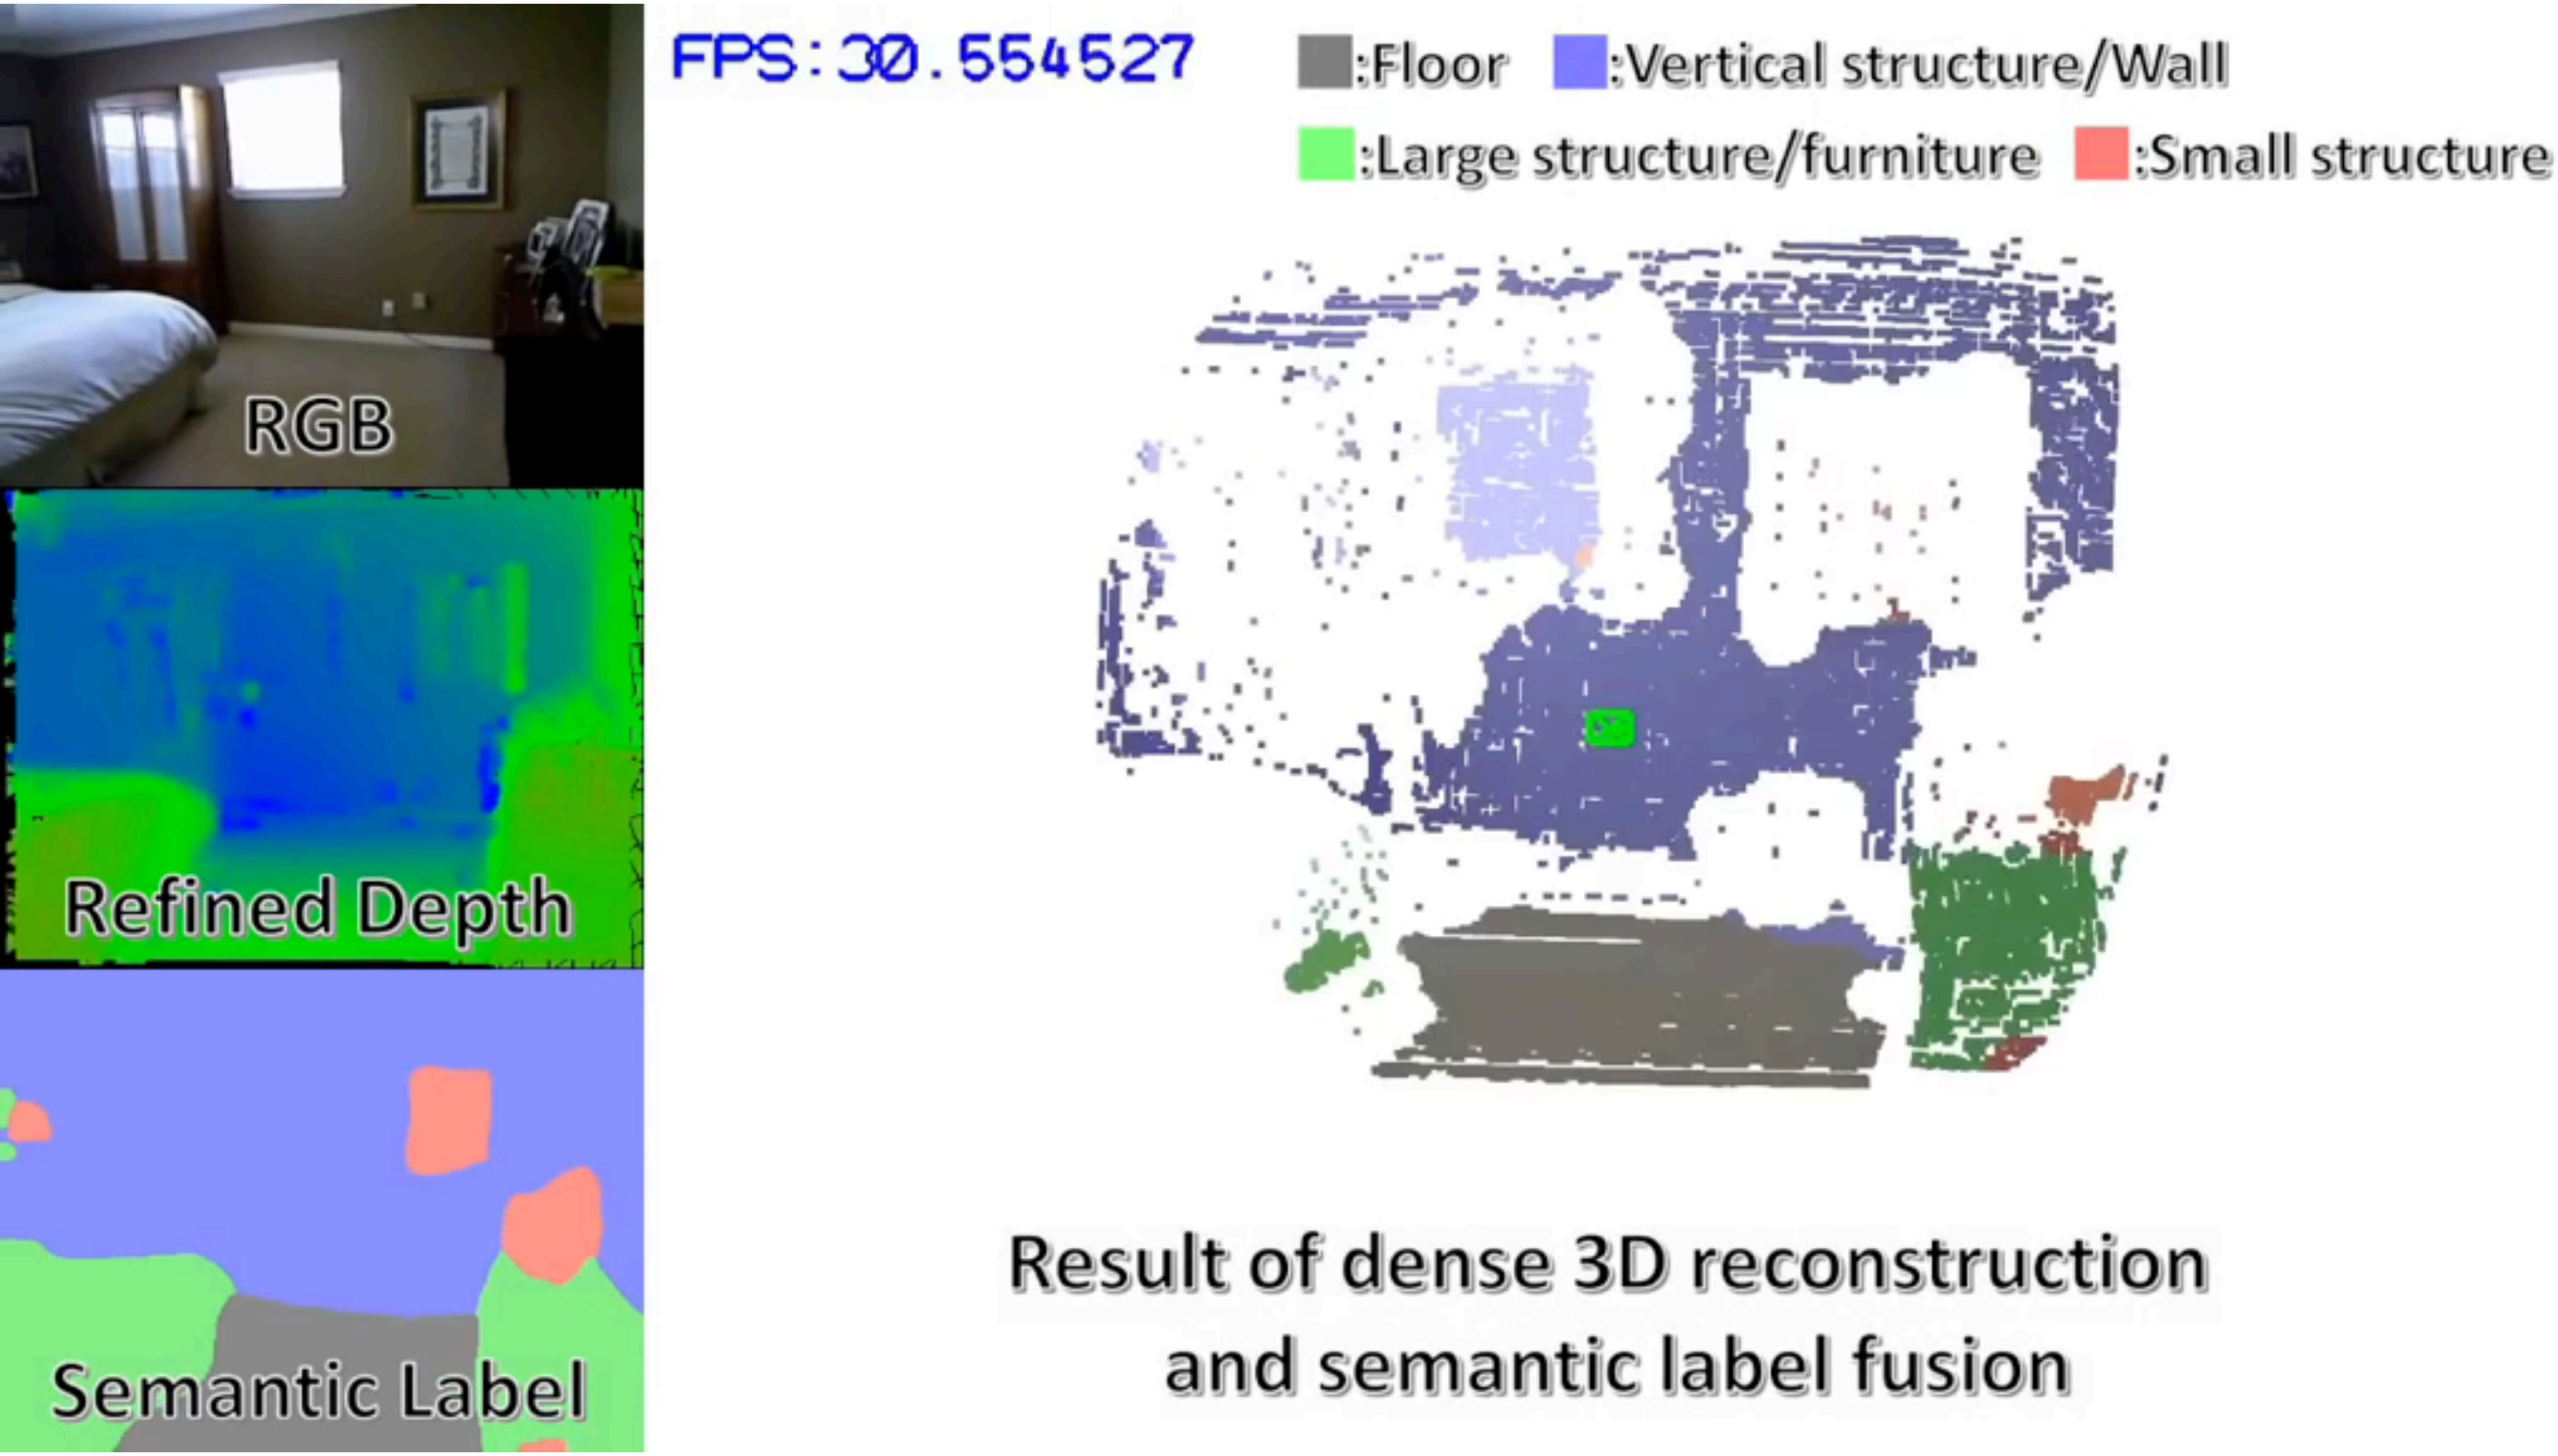

https://pdfs.semanticscholar.org/665f/708c1e6c1caec4b1aa8a8a9d01052d7a5dcd.pdf

Meaning of a map

- Many robotic problems need to summarize captured knowledge about the environment in some kind of “map”
  - Explicit map (e.g. point-cloud with detected objects, or traversability estimated from optimal trajectories)
  - Implicit map (e.g. memory of recurrent neural net for motion control)

Meaning of a map

- Many robotic problems need to summarize captured knowledge about the environment in some kind of “map”
  - Explicit map (e.g. point-cloud with detected objects, or traversability estimated from optimal trajectories)
  - Implicit map (e.g. memory of recurrent neural net for motion control)
- Maps are usually good for:
  - visualization to the operator
  - motion planning (e.g. cheapest path, travelling salesman, exploration)
  - reactive motion control (e.g. locomotion, predictive active damping)

Meaning of a map

- Many robotic problems need to summarize captured knowledge about the environment in some kind of "map"
  - Explicit map (e.g. point-cloud with detected objects, or traversability estimated from optimal trajectories)
  - Implicit map (e.g. memory of recurrent neural net for motion control)
- Maps are usually good for:
  - visualization to the operator
  - motion planning (e.g. cheapest path, travelling salesman, exploration)
  - reactive motion control (e.g. locomotion, predictive active damping)
  - In order to build maps we need to understand geometry of exteroceptive sensors.

- Derive camera motion and projective transformation models
- Prerequisite: constrained least square problem

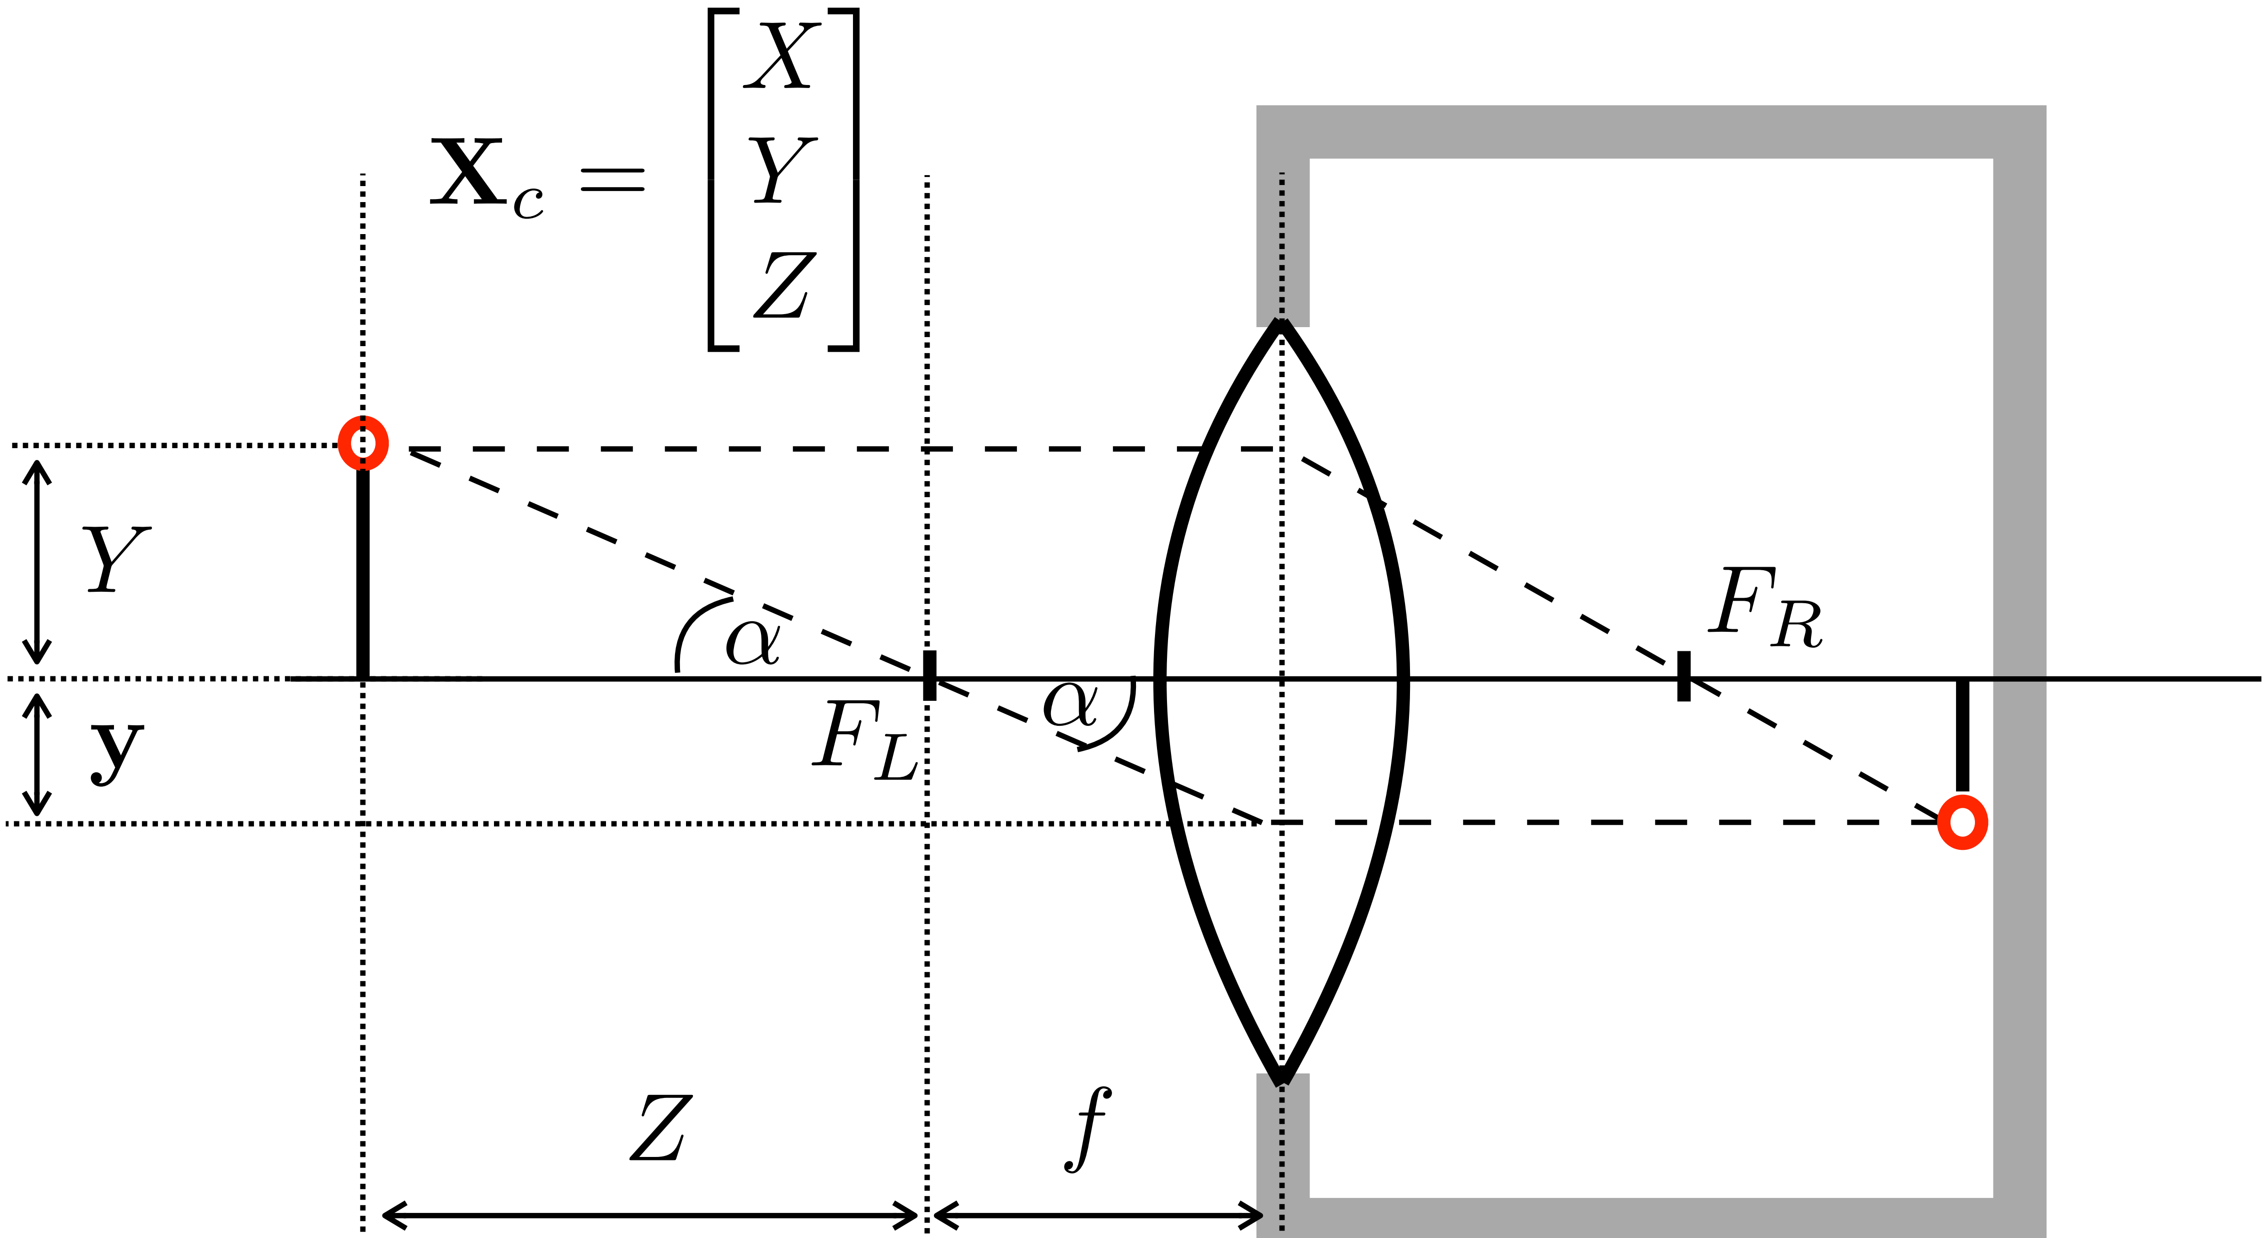

$$\mathrm{tg}(\alpha) = \frac{Y}{Z} = \frac{\mathbf{y}}{f}$$

$$\begin{aligned} \lambda \mathbf{x} &= f X \\ \lambda \mathbf{y} &= f Y \\ \lambda &= Z \end{aligned} \Rightarrow \lambda \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \\ 1 \end{bmatrix} = \begin{bmatrix} f & 0 & 0 \\ 0 & f & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}$$

Pinhole camera model experiment

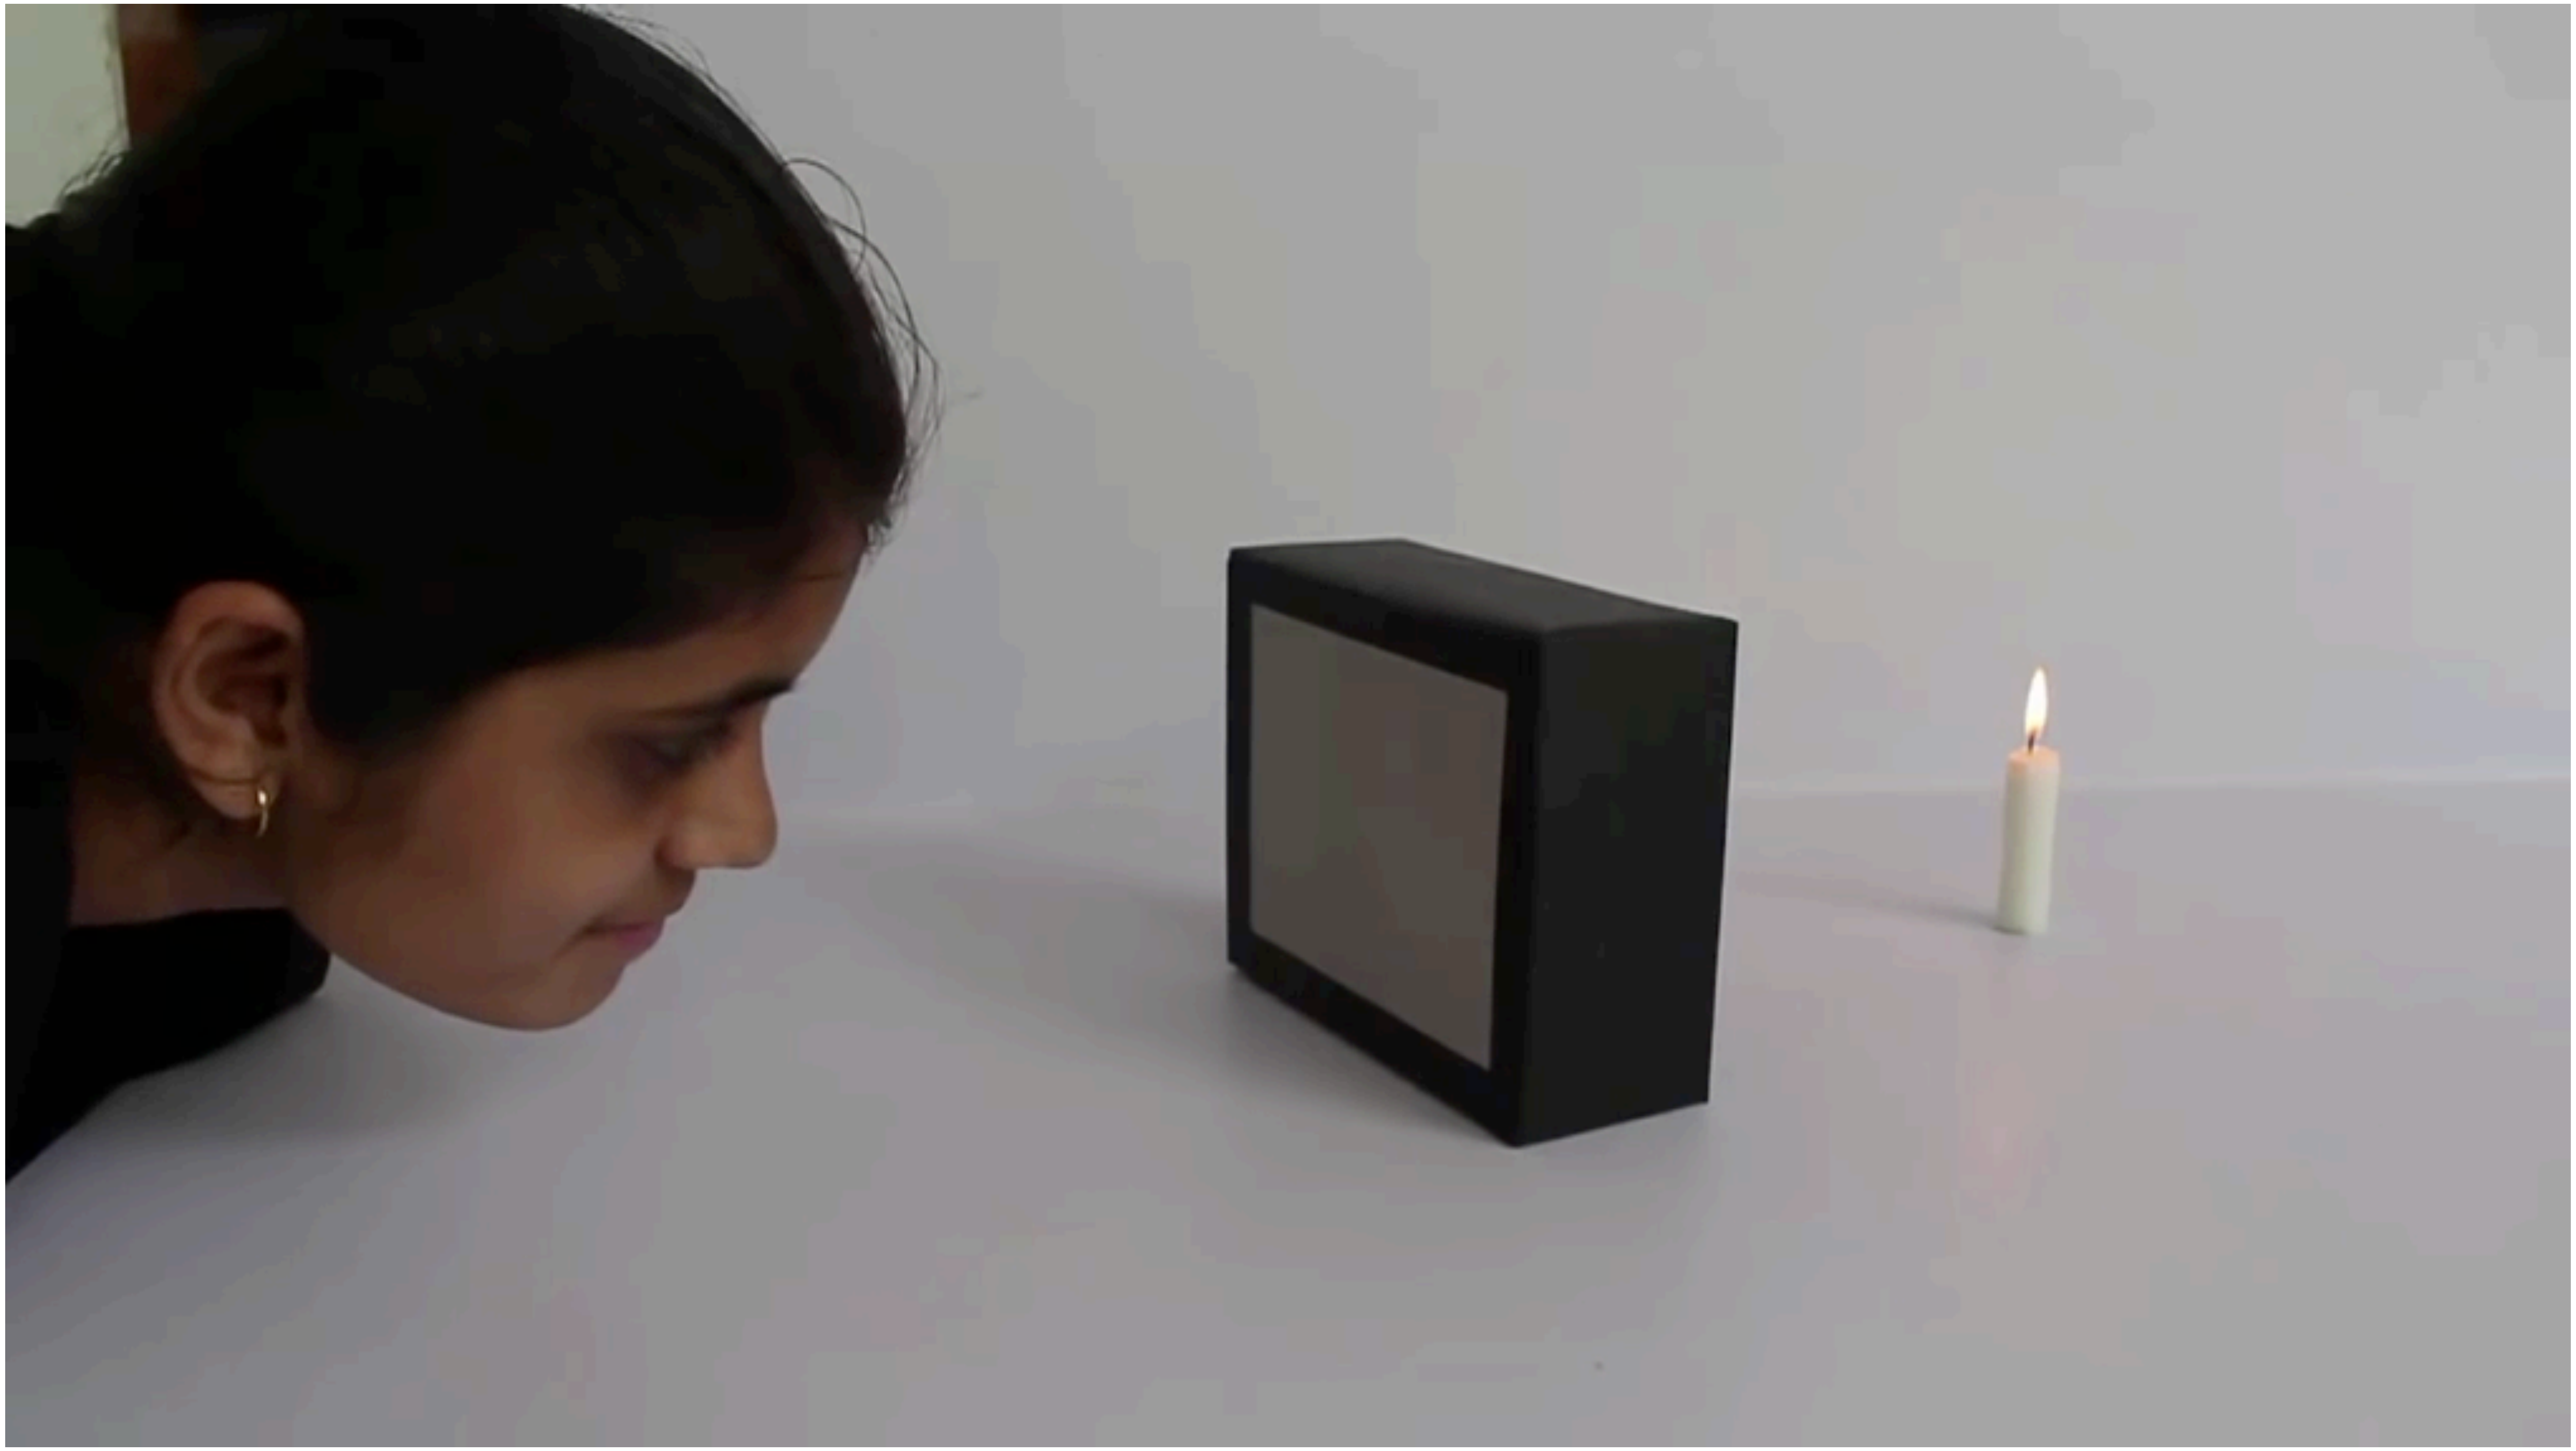

Outline

date	wee	due	Topic
01.10.20	1	T0	Outline, organization of the course.
01.10.20	1		Linear regression and classification + priors
08.10.20	2	T1	Neural nets + backprop
08.10.20	2		Convolutional network + backprop
15.10.20	3	T2	Learning (gradients, adaptive methods) + layers (activation functions,
15.10.20	3		Frameworks for learning
22.10.20	4	HW1	Recognition, detection, segmentation
22.10.20	4		Camera
29.10.20	5	T3	Depth sensors (stereo, kinect, time-of-flight)
29.10.20	5		Mapping I (basic SLAM from camera and depth sensor)
05.11.20	6	HW2	Mapping II (recurrent nets, learnable SLAM, stereo and monodepth)
05.11.20	6		Motion control from images
12.11.20	7		Domain transfer (GANs, time-constrastive nets, style transfer nets,
12.11.20	7		Semestral work assignment
19.11.20	8		— supervisor consultations —
19.11.20	8		— supervisor consultations —
26.11.20	9		— supervisor consultations —
26.11.20	9		— supervisor consultations —
03.12.20	10		— supervisor consultations —
03.12.20	10		— supervisor consultations —
10.12.20	11		— supervisor consultations —
10.12.20	11		— supervisor consultations —
17.12.20	12		— supervisor consultations —
17.12.20	12		— supervisor consultations —
07.01.20	15	SW	Semestral work presentations
07.01.20	15	T4	Exam test

Epipolar geometry: Depth from stereo, kinect and time-of-flight

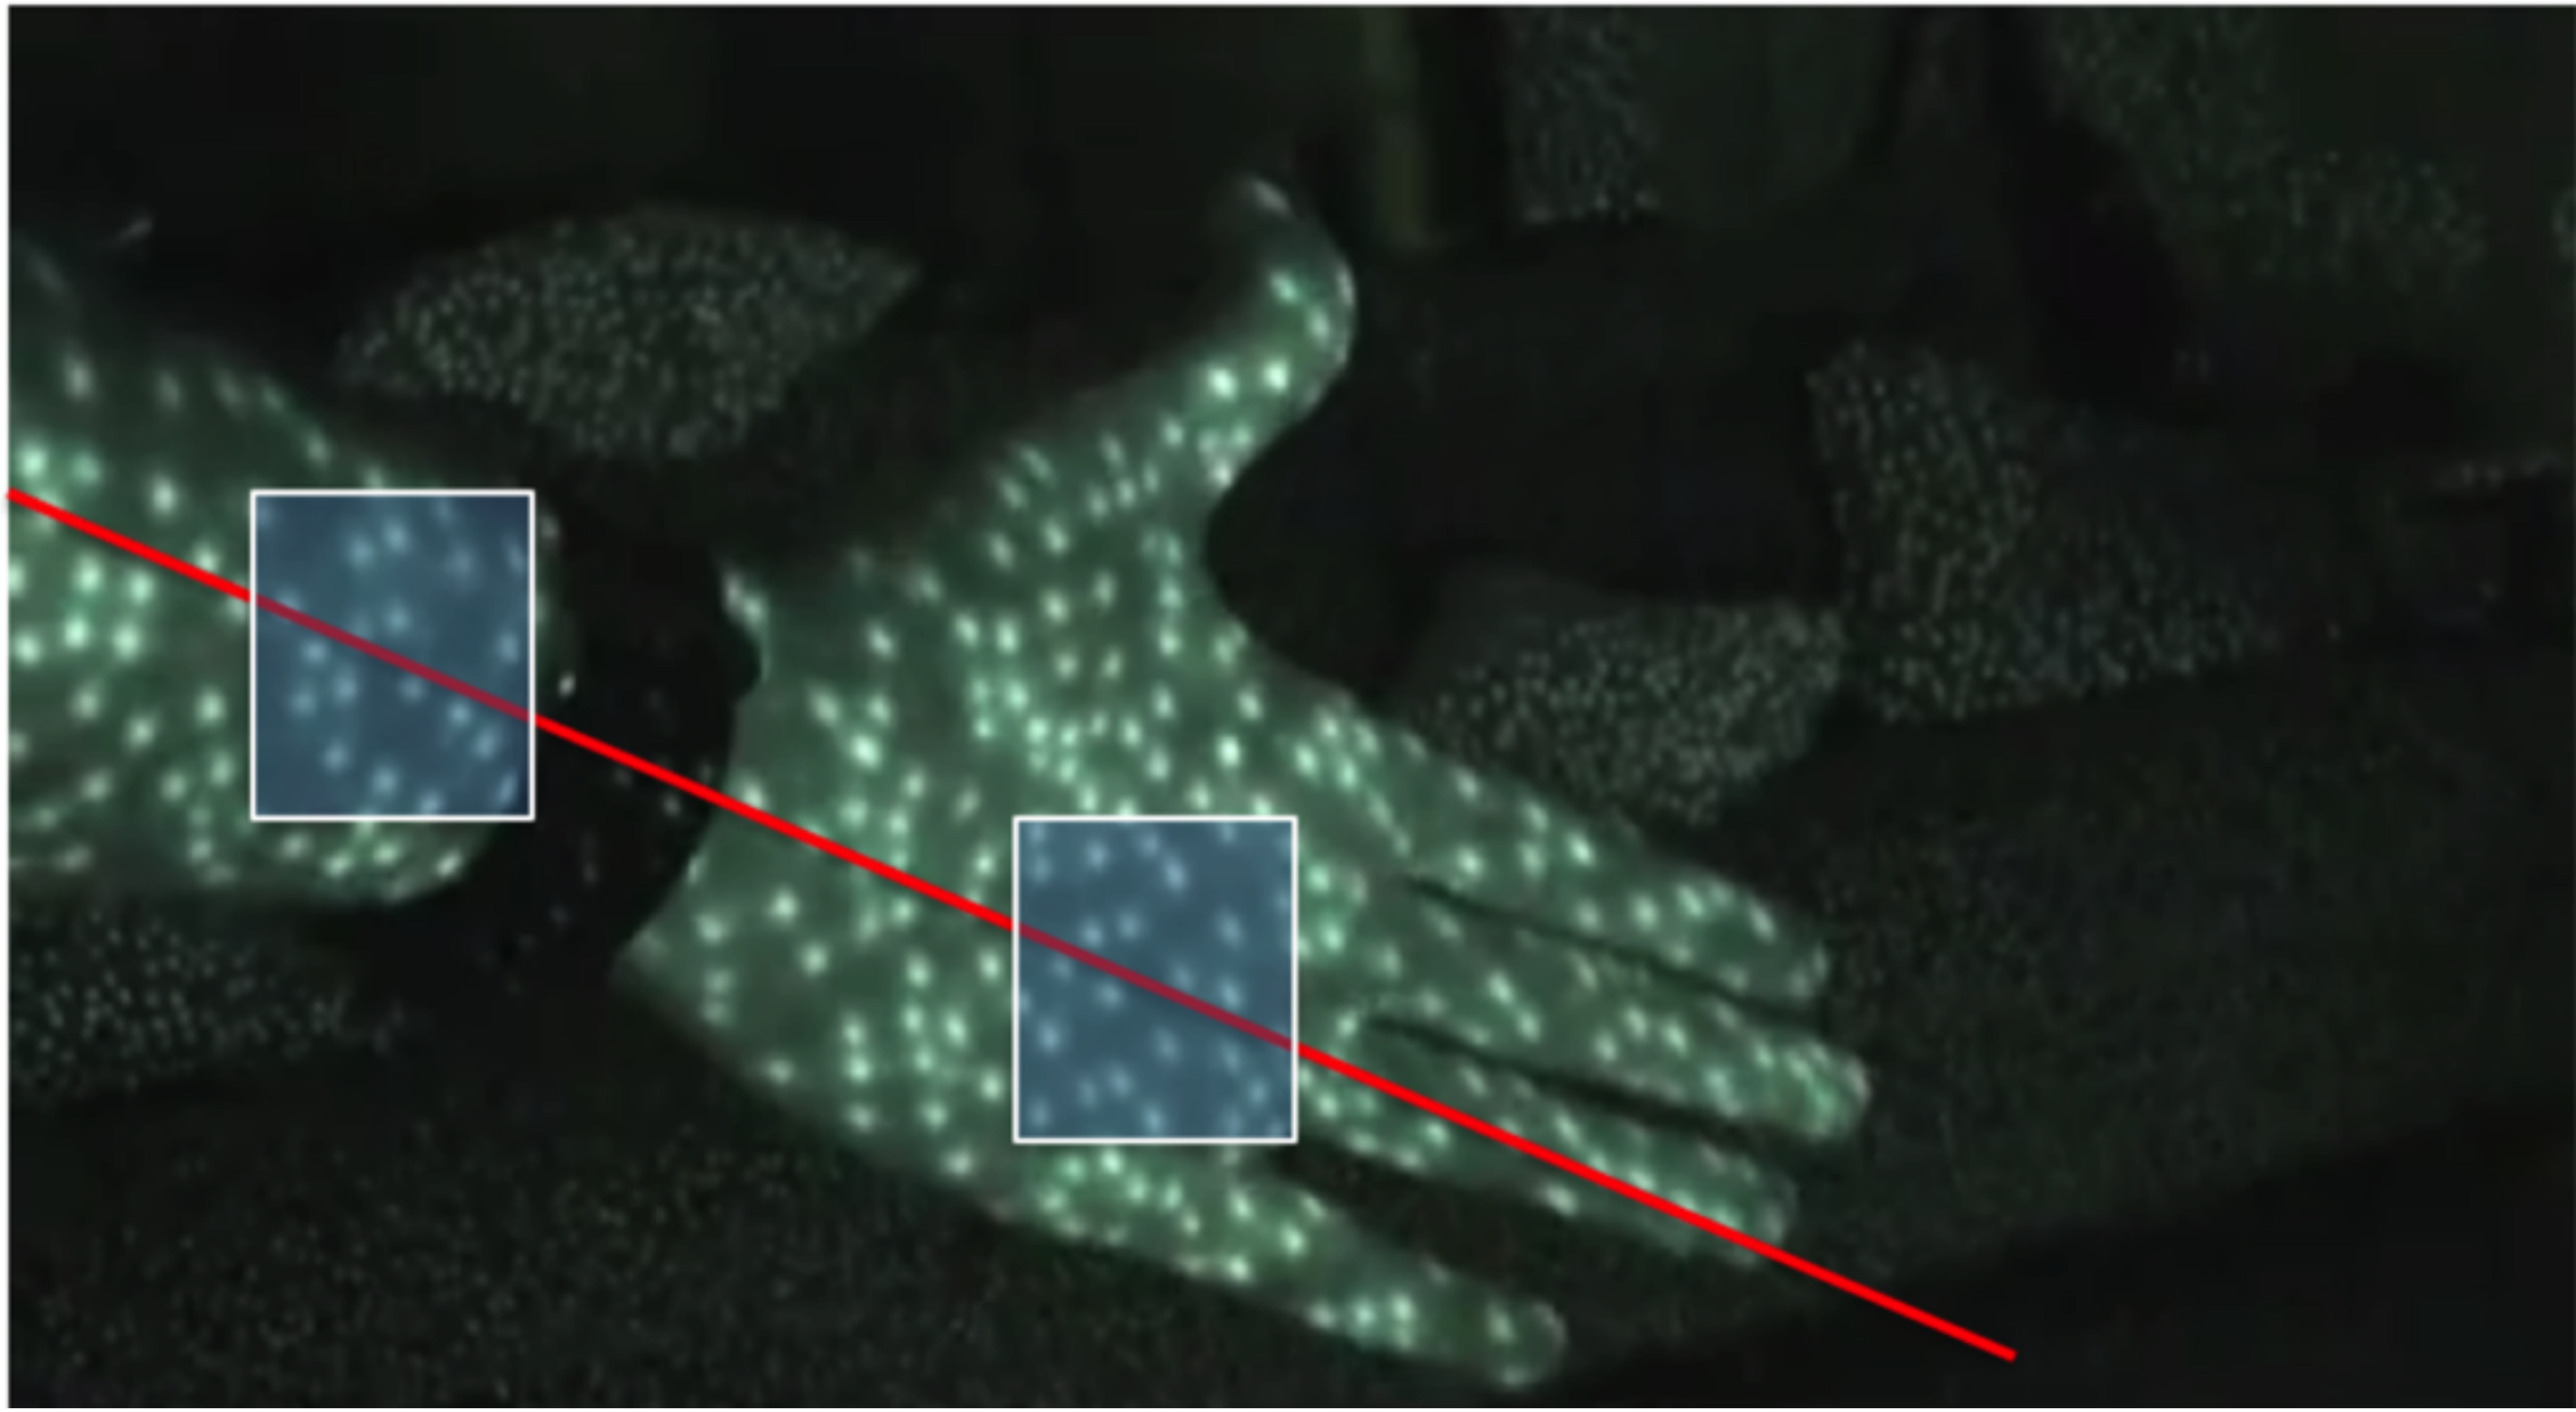

Czech Technical University in Prague
Faculty of Electrical Engineering, Department of Cybernetics

Outline

date	wee	due	Topic
01.10.20	1	T0	Outline, organization of the course.
01.10.20	1		Linear regression and classification + priors
08.10.20	2	T1	Neural nets + backprop
08.10.20	2		Convolutional network + backprop
15.10.20	3	T2	Learning (gradients, adaptive methods) + layers (activation functions,
15.10.20	3		Frameworks for learning
22.10.20	4	HW1	Recognition, detection, segmentation
22.10.20	4		Camera
29.10.20	5	T3	Depth sensors (stereo, kinect, time-of-flight)
29.10.20	5		Mapping I (basic SLAM from camera and depth sensor)
05.11.20	6	HW2	Mapping II (recurrent nets, learnable SLAM, stereo and monodepth)
05.11.20	6		Motion control from images
12.11.20	7		Domain transfer (GANs, time-constrastive nets, style transfer nets,
12.11.20	7		Semestral work assignment
19.11.20	8		— supervisor consultations —
19.11.20	8		— supervisor consultations —
26.11.20	9		— supervisor consultations —
26.11.20	9		— supervisor consultations —
03.12.20	10		— supervisor consultations —
03.12.20	10		— supervisor consultations —
10.12.20	11		— supervisor consultations —
10.12.20	11		— supervisor consultations —
17.12.20	12		— supervisor consultations —
17.12.20	12		— supervisor consultations —
07.01.20	15	SW	Semestral work presentations
07.01.20	15	T4	Exam test

How to build a map without learning

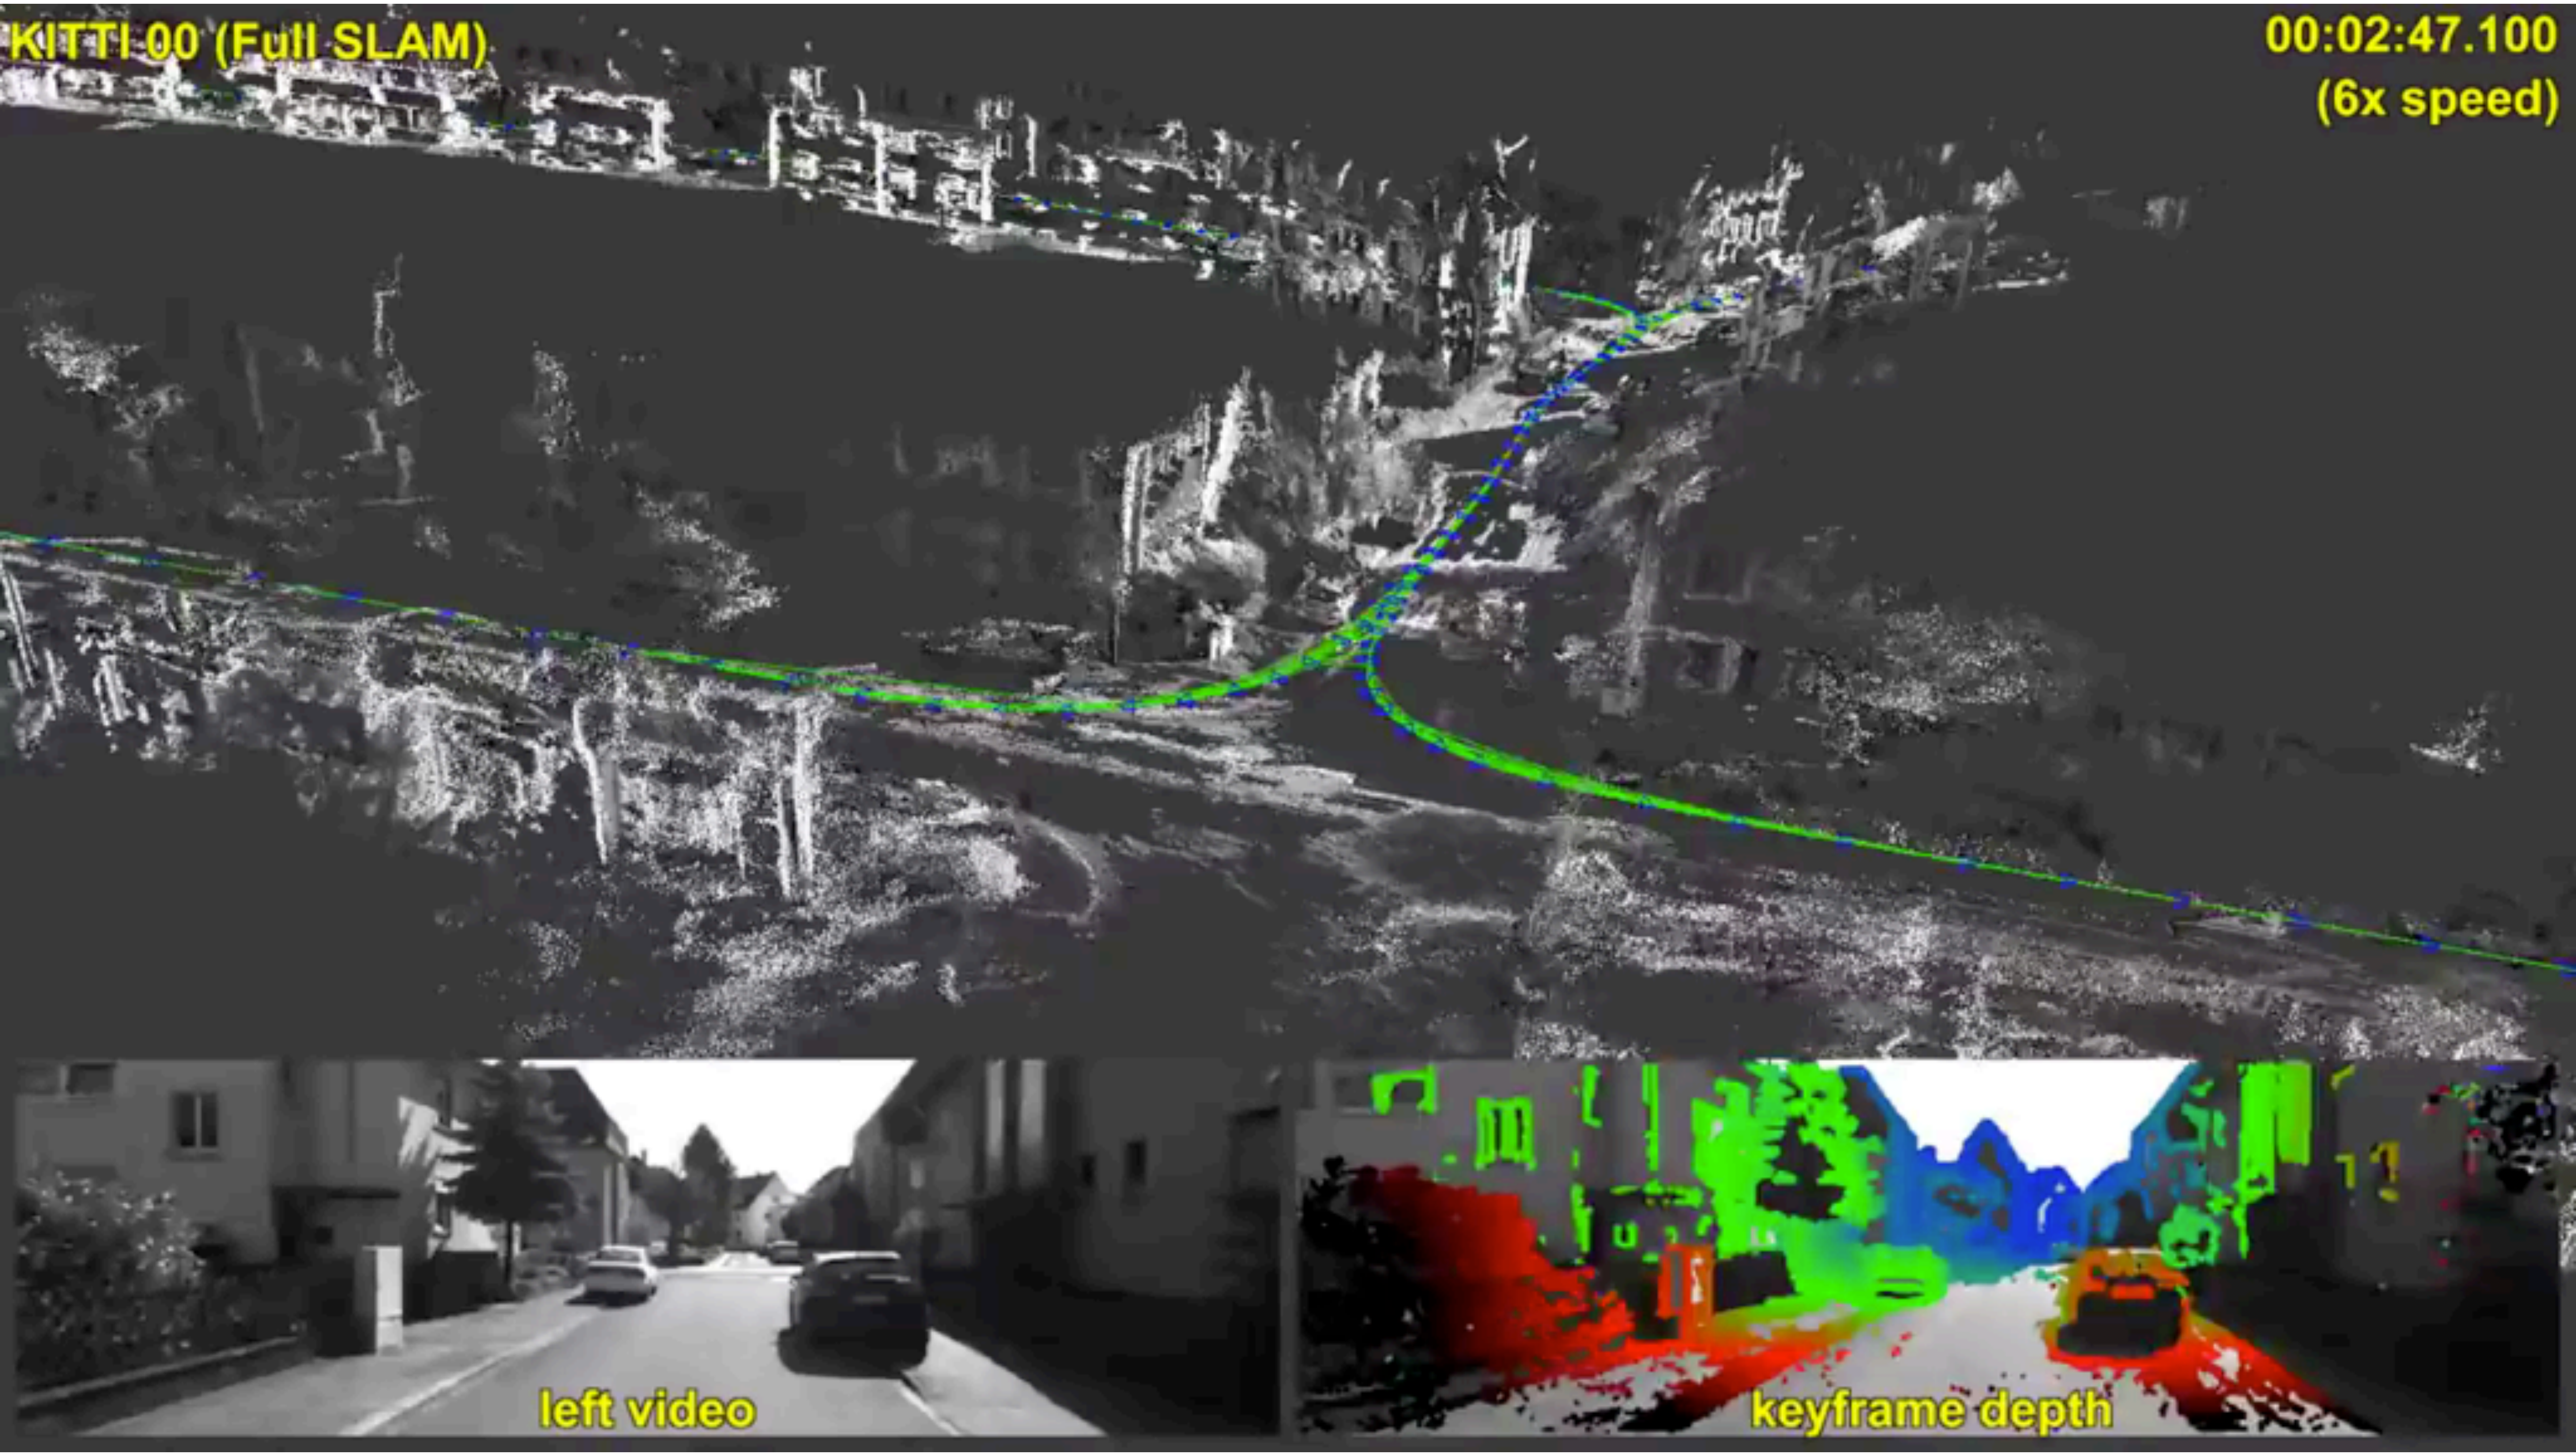

https://vision.in.tum.de/research/vslam/lsdslam

How to build a map with deep recurrent learning

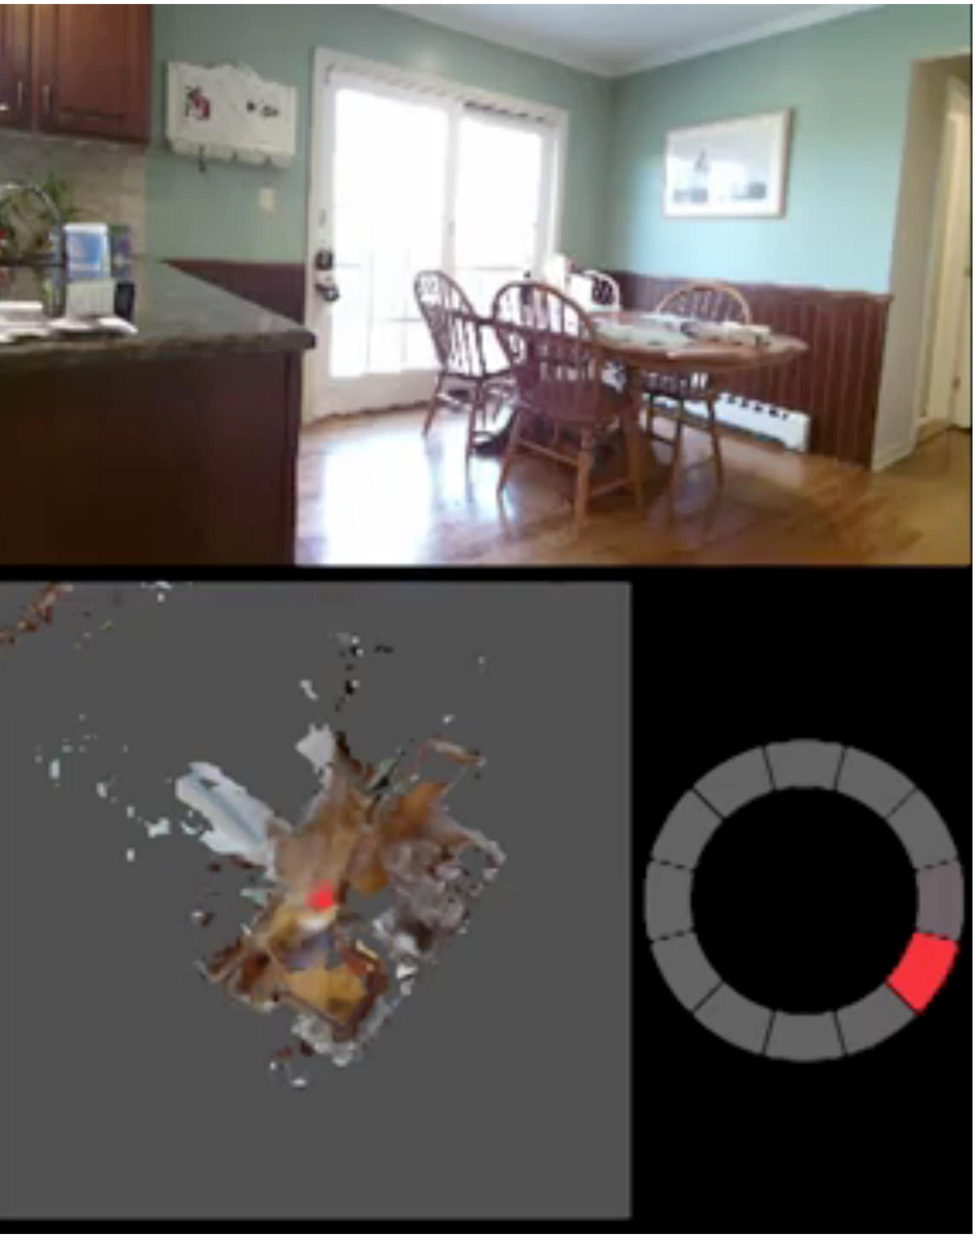

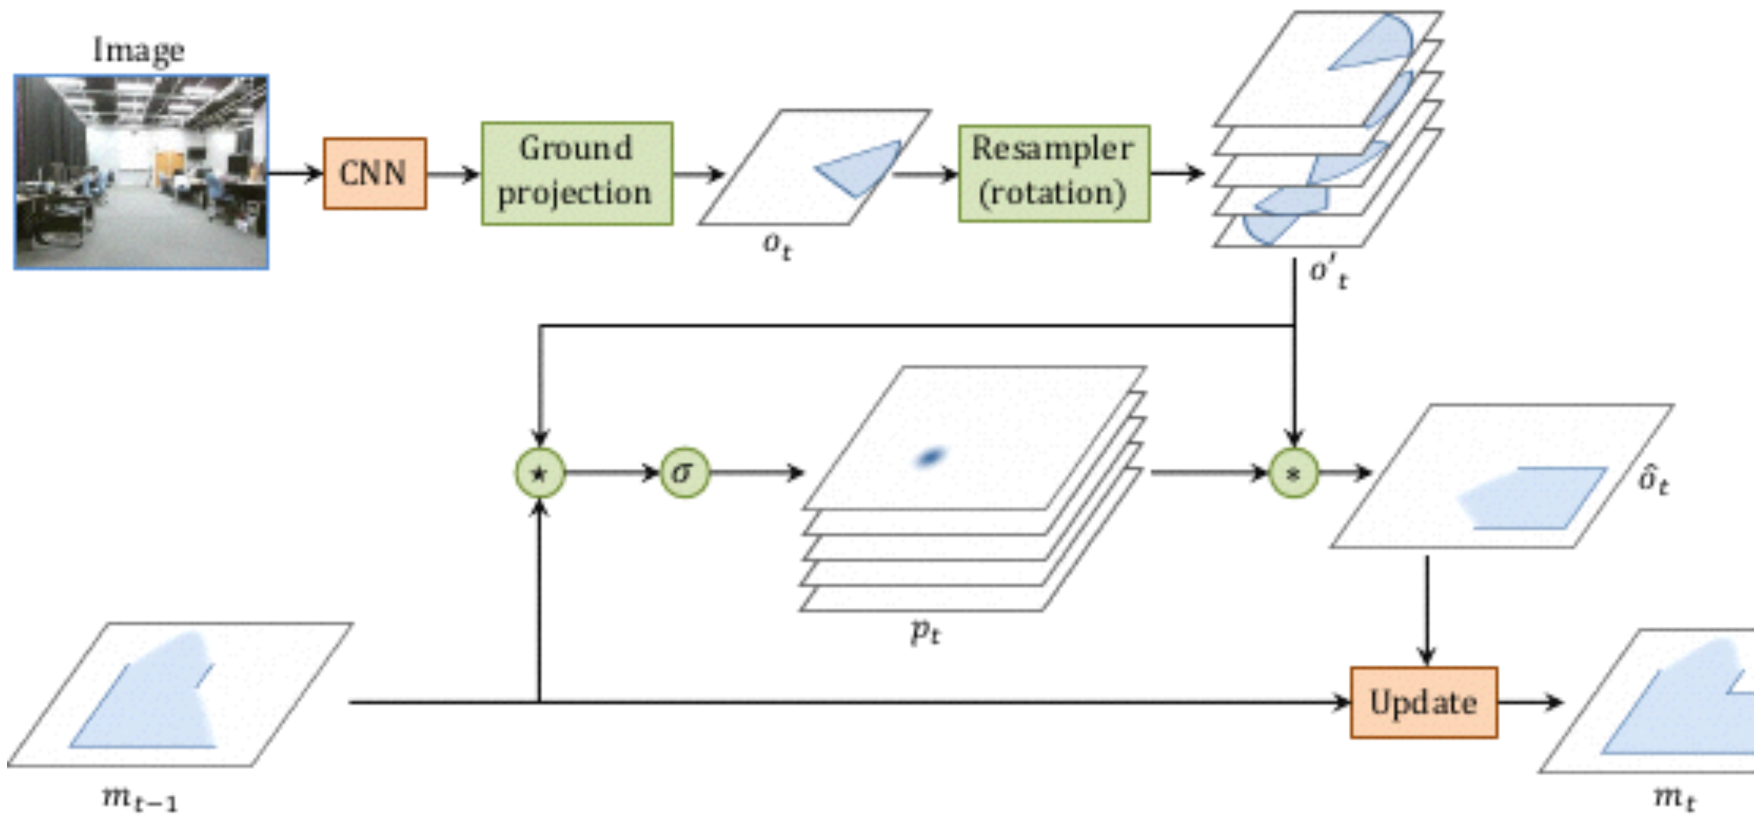

http://www.robots.ox.ac.uk/~joao/mapnet/

Outline

date	wee	due	Topic
01.10.20	1	T0	Outline, organization of the course.
01.10.20	1		Linear regression and classification + priors
08.10.20	2	T1	Neural nets + backprop
08.10.20	2		Convolutional network + backprop
15.10.20	3	T2	Learning (gradients, adaptive methods) + layers (activation functions,
15.10.20	3		Frameworks for learning
22.10.20	4	HW1	Recognition, detection, segmentation
22.10.20	4		Camera
29.10.20	5	T3	Depth sensors (stereo, kinect, time-of-flight)
29.10.20	5		Mapping I (basic SLAM from camera and depth sensor)
05.11.20	6	HW2	Mapping II (recurrent nets, learnable SLAM, stereo and monodepth)
05.11.20	6		Motion control from images
12.11.20	7		Domain transfer (GANs, time-constrastive nets, style transfer nets,
12.11.20	7		Semestral work assignment
19.11.20	8		— supervisor consultations —
19.11.20	8		— supervisor consultations —
26.11.20	9		— supervisor consultations —
26.11.20	9		— supervisor consultations —
03.12.20	10		— supervisor consultations —
03.12.20	10		— supervisor consultations —
10.12.20	11		— supervisor consultations —
10.12.20	11		— supervisor consultations —
17.12.20	12		— supervisor consultations —
17.12.20	12		— supervisor consultations —
07.01.20	15	SW	Semestral work presentations
07.01.20	15	T4	Exam test

Czech Technical University in Prague
Faculty of Electrical Engineering, Department of Cybernetics

[Heess 2017] https://arxiv.org/abs/1707.02286

[Levine IJRR 2017] https://arxiv.org/abs/1603.02199

Outline

date	wee	due	Topic
01.10.20	1	T0	Outline, organization of the course.
01.10.20	1		Linear regression and classification + priors
08.10.20	2	T1	Neural nets + backprop
08.10.20	2		Convolutional network + backprop
15.10.20	3	T2	Learning (gradients, adaptive methods) + layers (activation functions,
15.10.20	3		Frameworks for learning
22.10.20	4	HW1	Recognition, detection, segmentation
22.10.20	4		Camera
29.10.20	5	T3	Depth sensors (stereo, kinect, time-of-flight)
29.10.20	5		Mapping I (basic SLAM from camera and depth sensor)
05.11.20	6	HW2	Mapping II (recurrent nets, learnable SLAM, stereo and monodepth)
05.11.20	6		Motion control from images
12.11.20	7		Domain transfer (GANs, time-constrastive nets, style transfer nets,
12.11.20	7		Semestral work assignment
19.11.20	8		— supervisor consultations —
19.11.20	8		— supervisor consultations —
26.11.20	9		— supervisor consultations —
26.11.20	9		— supervisor consultations —
03.12.20	10		— supervisor consultations —
03.12.20	10		— supervisor consultations —
10.12.20	11		— supervisor consultations —
10.12.20	11		— supervisor consultations —
17.12.20	12		— supervisor consultations —
17.12.20	12		— supervisor consultations —
07.01.20	15	SW	Semestral work presentations
07.01.20	15	T4	Exam test

Annotations costs 10$ per image - can we get it for free?

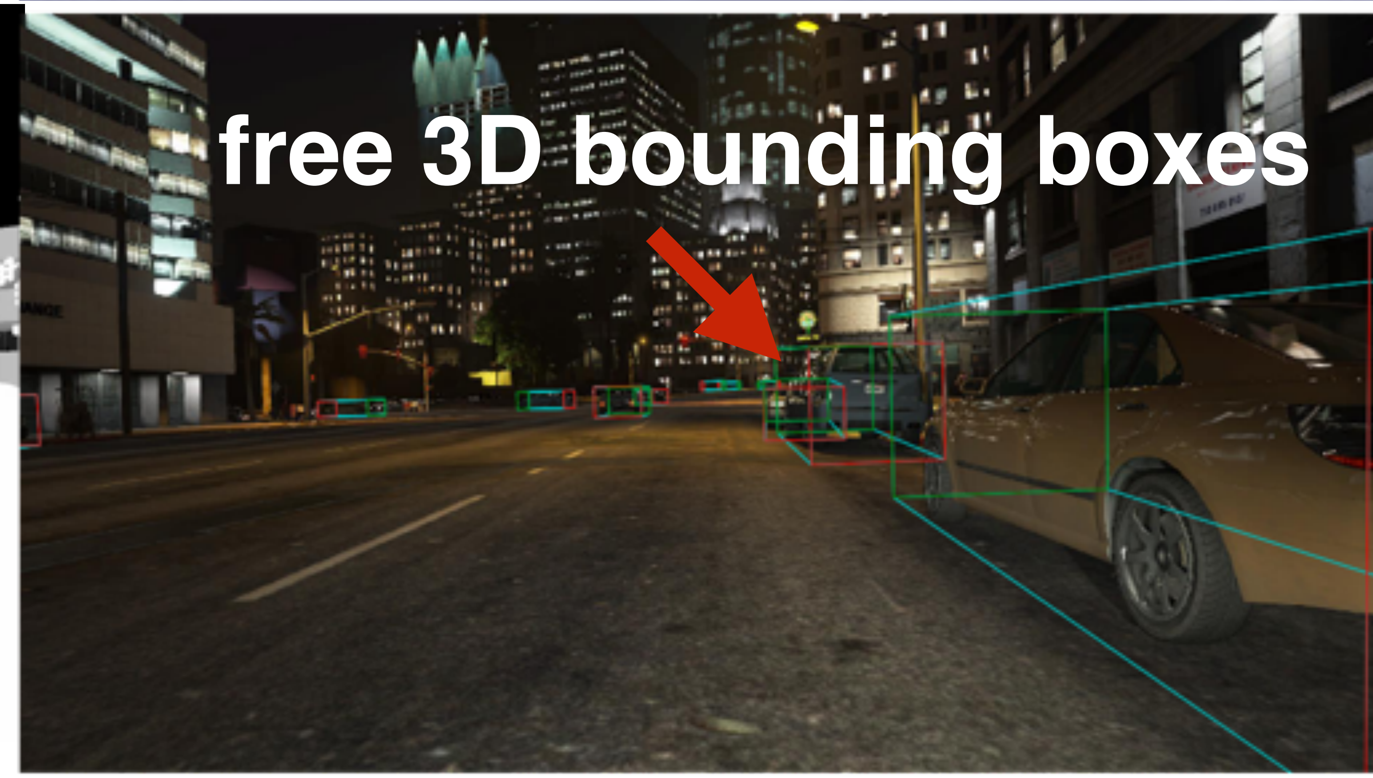

Czech Technical University in Prague
Faculty of Electrical Engineering, Department of Cybernetics

Domain transfer - learning to simulate

https://github.com/jhoffman/cycada_release

https://people.eecs.berkeley.edu/~jhoffman/papers/2018_cycada.pdf

Semantic segmentation on cityscapes
https://github.com/Eromera/erfnet_pytorch

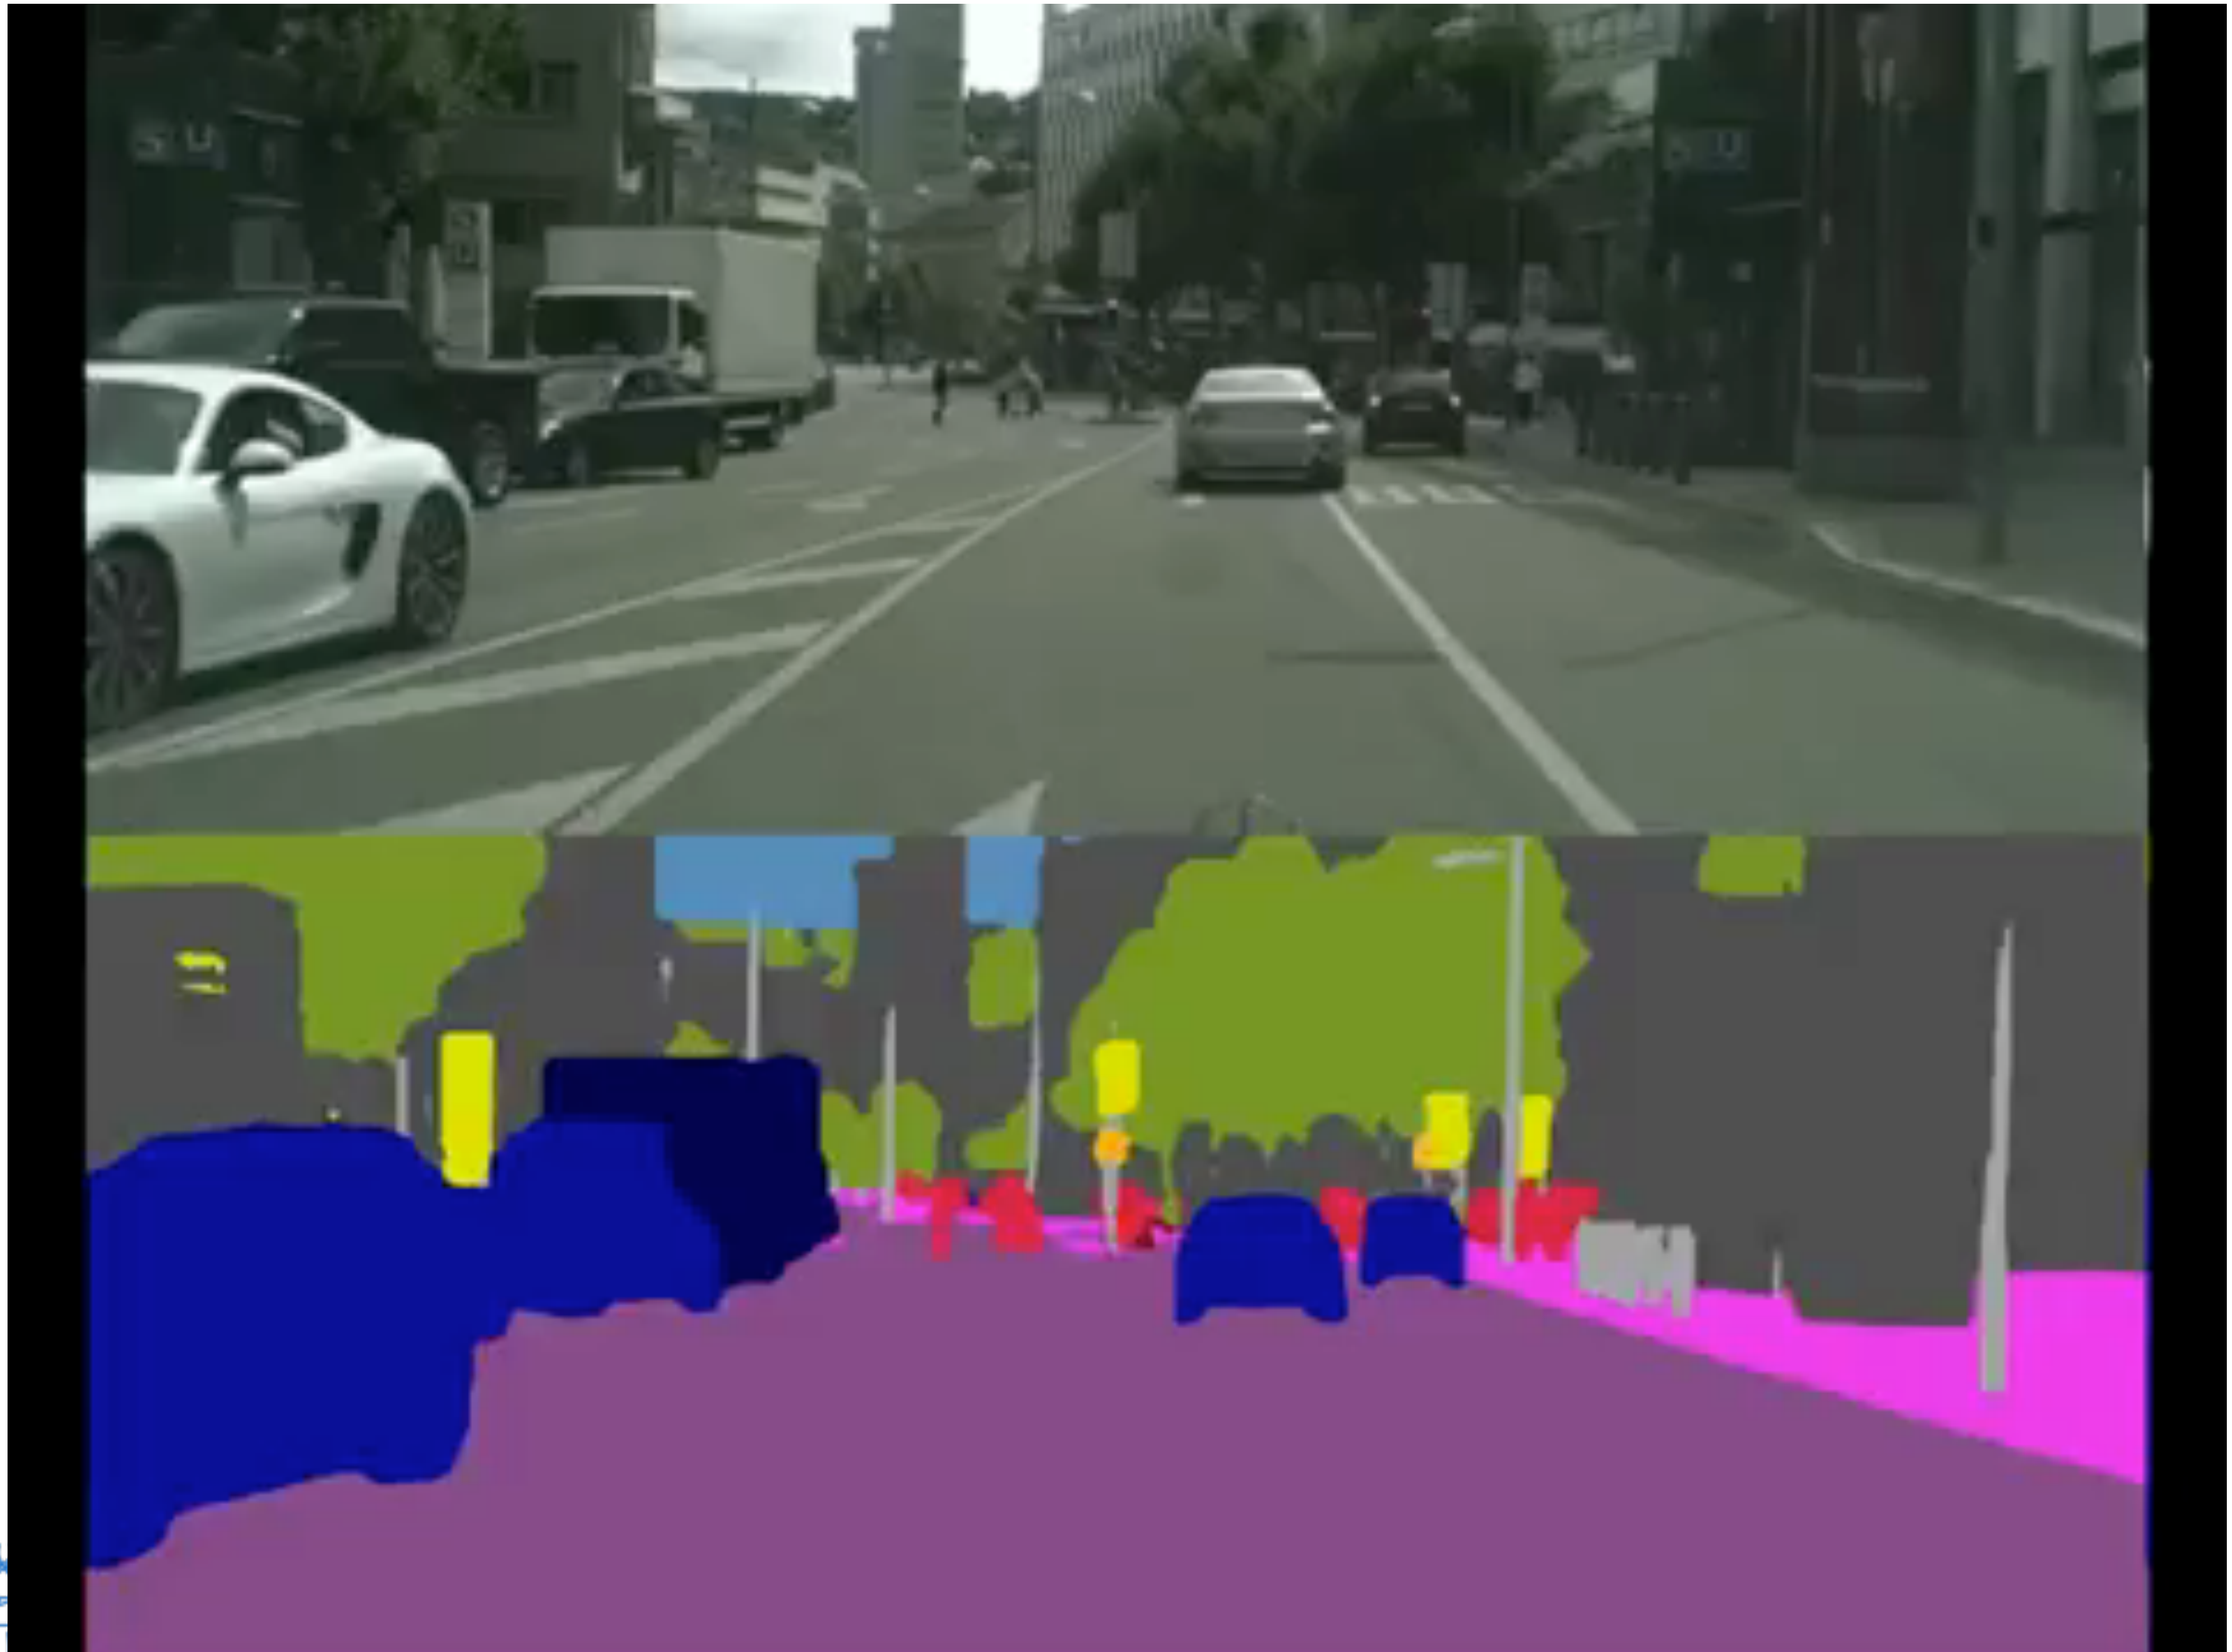

Deep dreaming

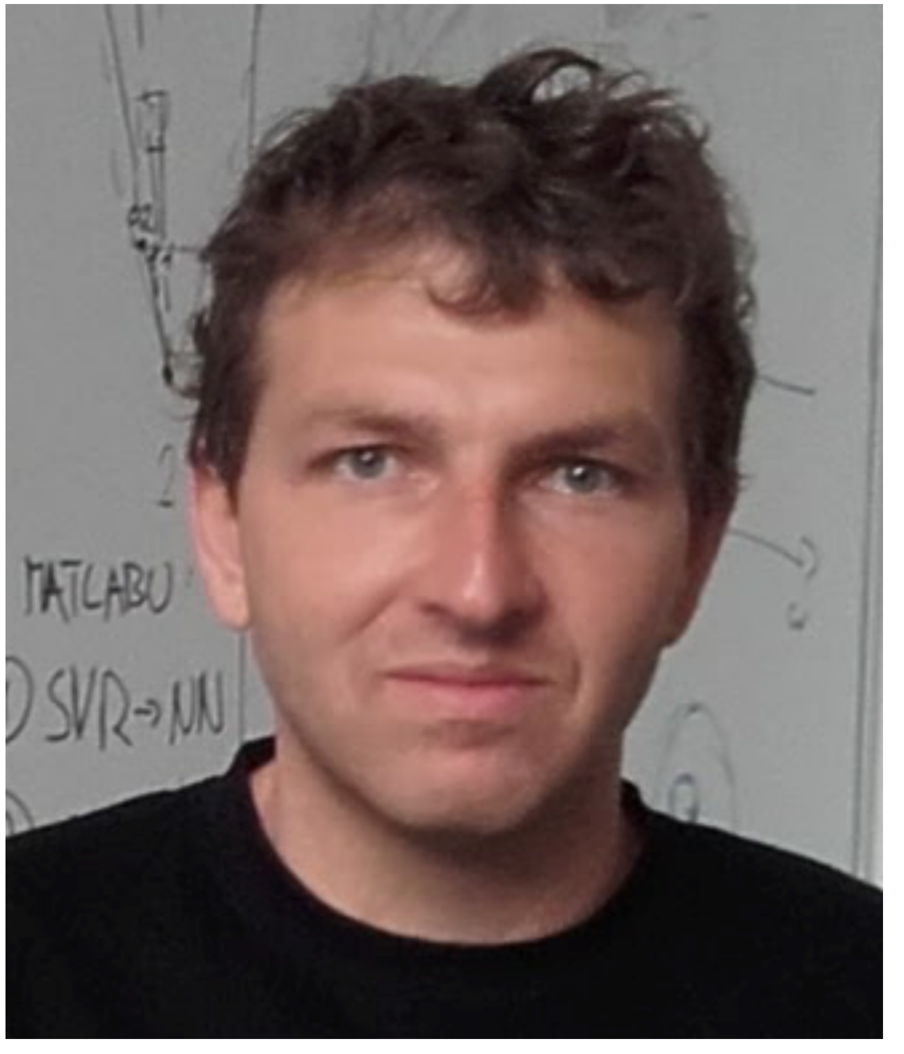

https://deepdreamgenerator.com

Deep dreaming

https://deepdreamgenerator.com

Deep dreaming

https://deepdreamgenerator.com

Deep dreaming

https://deepdreamgenerator.com

Outline

date	wee	due	Topic
01.10.20	1	T0	Outline, organization of the course.
01.10.20	1		Linear regression and classification + priors
08.10.20	2	T1	Neural nets + backprop
08.10.20	2		Convolutional network + backprop
15.10.20	3	T2	Learning (gradients, adaptive methods) + layers (activation functions,
15.10.20	3		Frameworks for learning
22.10.20	4	HW1	Recognition, detection, segmentation
22.10.20	4		Camera
29.10.20	5	T3	Depth sensors (stereo, kinect, time-of-flight)
29.10.20	5		Mapping I (basic SLAM from camera and depth sensor)
05.11.20	6	HW2	Mapping II (recurrent nets, learnable SLAM, stereo and monodepth)
05.11.20	6		Motion control from images
12.11.20	7		Domain transfer (GANs, time-constrastive nets, style transfer nets,
12.11.20	7		Semestral work assignment
19.11.20	8		— supervisor consultations —
19.11.20	8		— supervisor consultations —
26.11.20	9		— supervisor consultations —
26.11.20	9		— supervisor consultations —
03.12.20	10		— supervisor consultations —
03.12.20	10		— supervisor consultations —
10.12.20	11		— supervisor consultations —
10.12.20	11		— supervisor consultations —
17.12.20	12		— supervisor consultations —
17.12.20	12		— supervisor consultations —
07.01.20	15	SW	Semestral work presentations
07.01.20	15	T4	Exam test

Semestral work

- solved by the teams of 3 students
- topics/supervisors will be made available in the 5th week
- limited capacity => tests and homework results considered

Tomas Petricek
CTU
deep slam

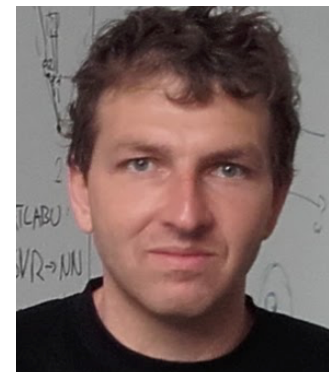

Karel Zimmermann
CTU
domain adaptation

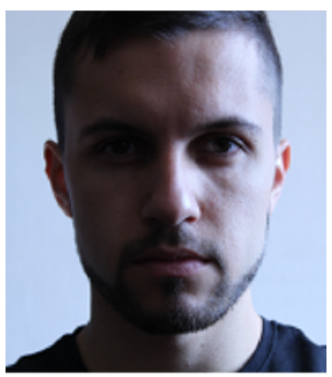

Teymur Azayev
CTU
deep motion control

David Coufal
Ustav informace AV
GANs

Michal Reinstein
SpaceKnow
detection in satellite images

Outline

date	wee	due	Topic
01.10.20	1	T0	Outline, organization of the course.
01.10.20	1		Linear regression and classification + priors
08.10.20	2	T1	Neural nets + backprop
08.10.20	2		Convolutional network + backprop
15.10.20	3	T2	Learning (gradients, adaptive methods) + layers (activation functions,
15.10.20	3		Frameworks for learning
22.10.20	4	HW1	Recognition, detection, segmentation
22.10.20	4		Camera
29.10.20	5	T3	Depth sensors (stereo, kinect, time-of-flight)
29.10.20	5		Mapping I (basic SLAM from camera and depth sensor)
05.11.20	6	HW2	Mapping II (recurrent nets, learnable SLAM, stereo and monodepth)
05.11.20	6		Motion control from images
12.11.20	7		Domain transfer (GANs, time-constrastive nets, style transfer nets,
12.11.20	7		Semestral work assignment
19.11.20	8		— supervisor consultations —
19.11.20	8		— supervisor consultations —
26.11.20	9		— supervisor consultations —
26.11.20	9		— supervisor consultations —
03.12.20	10		— supervisor consultations —
03.12.20	10		— supervisor consultations —
10.12.20	11		— supervisor consultations —
10.12.20	11		— supervisor consultations —
17.12.20	12		— supervisor consultations —
17.12.20	12		— supervisor consultations —
07.01.20	15	SW	Semestral work presentations
07.01.20	15	T4	Exam test

Outline

date	wee	due	Topic
01.10.20	1	T0	Outline, organization of the course.
01.10.20	1		Linear regression and classification + priors
08.10.20	2	T1	Neural nets + backprop
08.10.20	2		Convolutional network + backprop
15.10.20	3	T2	Learning (gradients, adaptive methods) + layers (activation functions,
15.10.20	3		Frameworks for learning
22.10.20	4	HW1	Recognition, detection, segmentation
22.10.20	4		Camera
29.10.20	5	T3	Depth sensors (stereo, kinect, time-of-flight)
29.10.20	5		Mapping I (basic SLAM from camera and depth sensor)
05.11.20	6	HW2	Mapping II (recurrent nets, learnable SLAM, stereo and monodepth)
05.11.20	6		Motion control from images
12.11.20	7		Domain transfer (GANs, time-constrastive nets, style transfer nets,
12.11.20	7		Semestral work assignment
19.11.20	8		— supervisor consultations —
19.11.20	8		— supervisor consultations —
26.11.20	9		— supervisor consultations —
26.11.20	9		— supervisor consultations —
03.12.20	10		— supervisor consultations —
03.12.20	10		— supervisor consultations —
10.12.20	11		— supervisor consultations —
10.12.20	11		— supervisor consultations —
17.12.20	12		— supervisor consultations —
17.12.20	12		— supervisor consultations —
07.01.20	15	SW	Semestral work presentations
07.01.20	15	T4	Exam test

- Introduction of the ViR-team
- Outline of the course
- Organization (homework, tests, semestral work)
- Test0

Semestral work presentations + Exam test

- 50 points from the semestral work
  - evaluation based on students and lecturers voting
  - it is assumed that work will correspond to at least 3*7*6=126 hours of work
  - presentation contains explicit specification of what has been done by each team member
- 50 points from test and homework
  - including the exam test
  - evaluated by Teymur and Tomas ;-)
- minimum credit requirements is 50 points (out of 100)
- final grade determined by the total number of points

- Introduction of the ViR-team
- Outline of the course
- Organization (homework, tests, semestral work)
- Test0

Summary

- We will be happy if you help in any possible way
  - suggesting better course logo
  - suggesting/preparing new homework
  - implementing nice demos (software or hardware)
  - giving any reasonable feedback
  - start your own research with us

What you can do?

- We're looking for students
  - competent in theory and practice (code development,work with real robotic platforms)
  - motivated to write top research papers with us
  - willing to work hard under our guidance
- We're offering:
  - diploma/bachelor theses, semester work or project
  - paid internships / summer jobs
  - international collaboration opportunities

Object detection and tracking

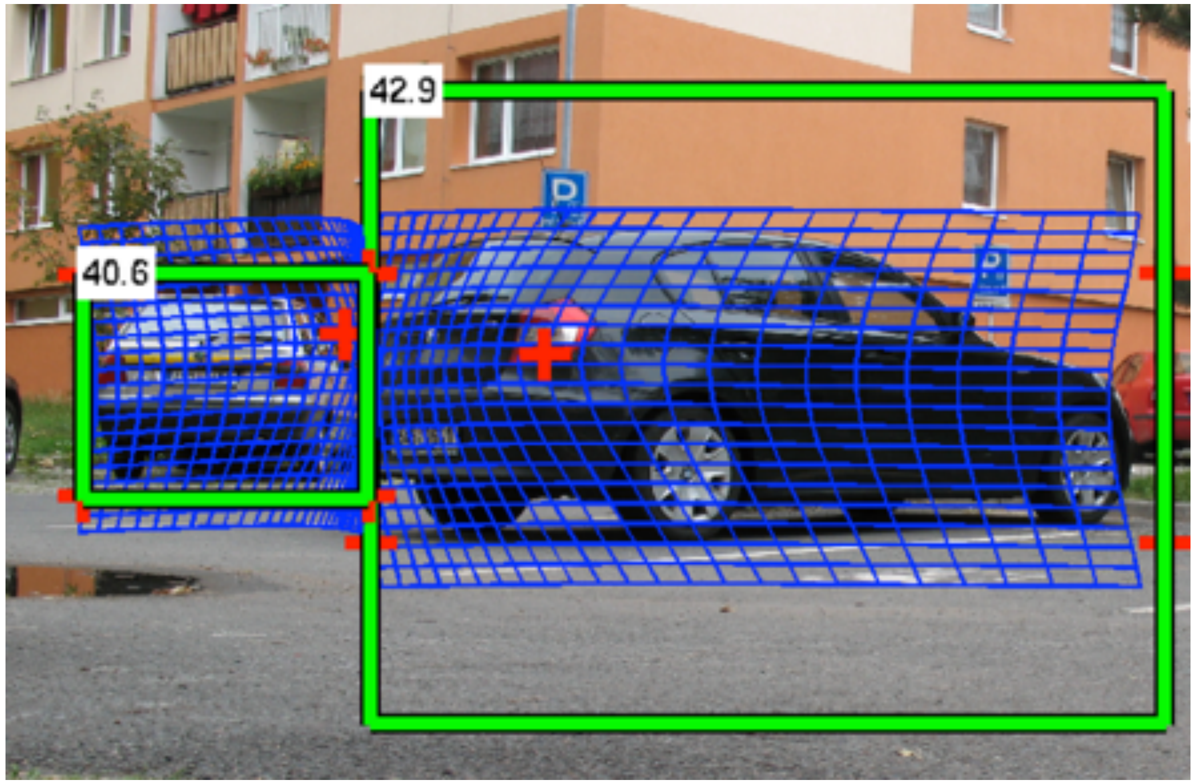

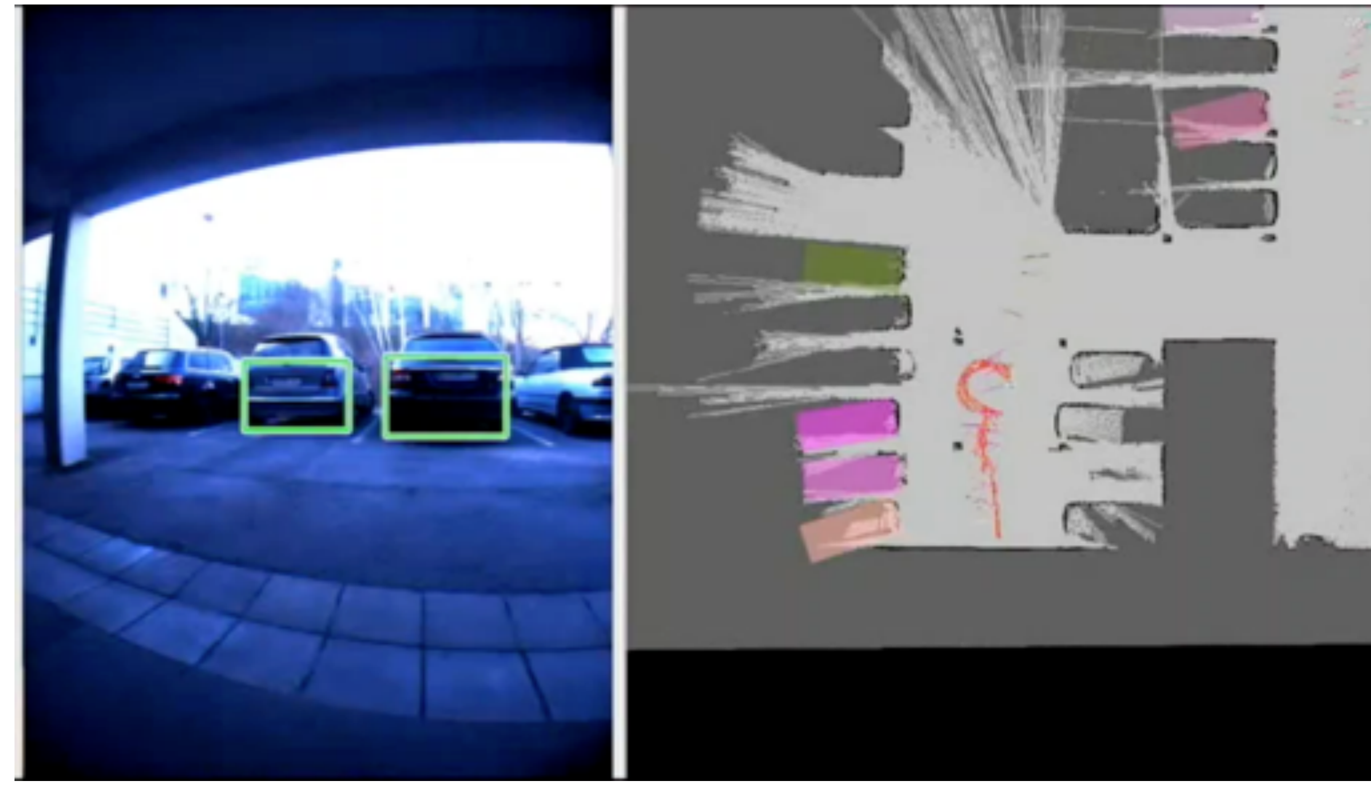

[1] <u>K.Zimmermann</u>, D.Hurych, T.Svoboda, Non-Rigid *Object Detection with Local Interleaved Sequential Alignment (LISA),* **TPAMI (IF=5)**, 2014

[2] <u>K.Zimmermann</u>, J.Matas, T.Svoboda, *Tracking by an Optimal Sequence of Linear Predictors*,**TPAMI (IF=5 selected for II.pillar evaluation)**, 2009.

Motion and compliance control of flippers

[3] Pecka, Zimmermann, Svoboda, Hlavac, et al.
IROS/RAL/TIE(IF=6), 2015-2018

Traffic sign detection and 3D localization

1.5 year PostDoc in Luc van Gool's lab at Katholieke Universiteit Leuven

[4] R.Timofte, K.Zimmermann, Luc van Gool, Multi-view traffic sign detection, recognition, and 3D localisation, **MVA (IF=1.5, over 200 citations), 2011**

Experiment:Active 3D mapping

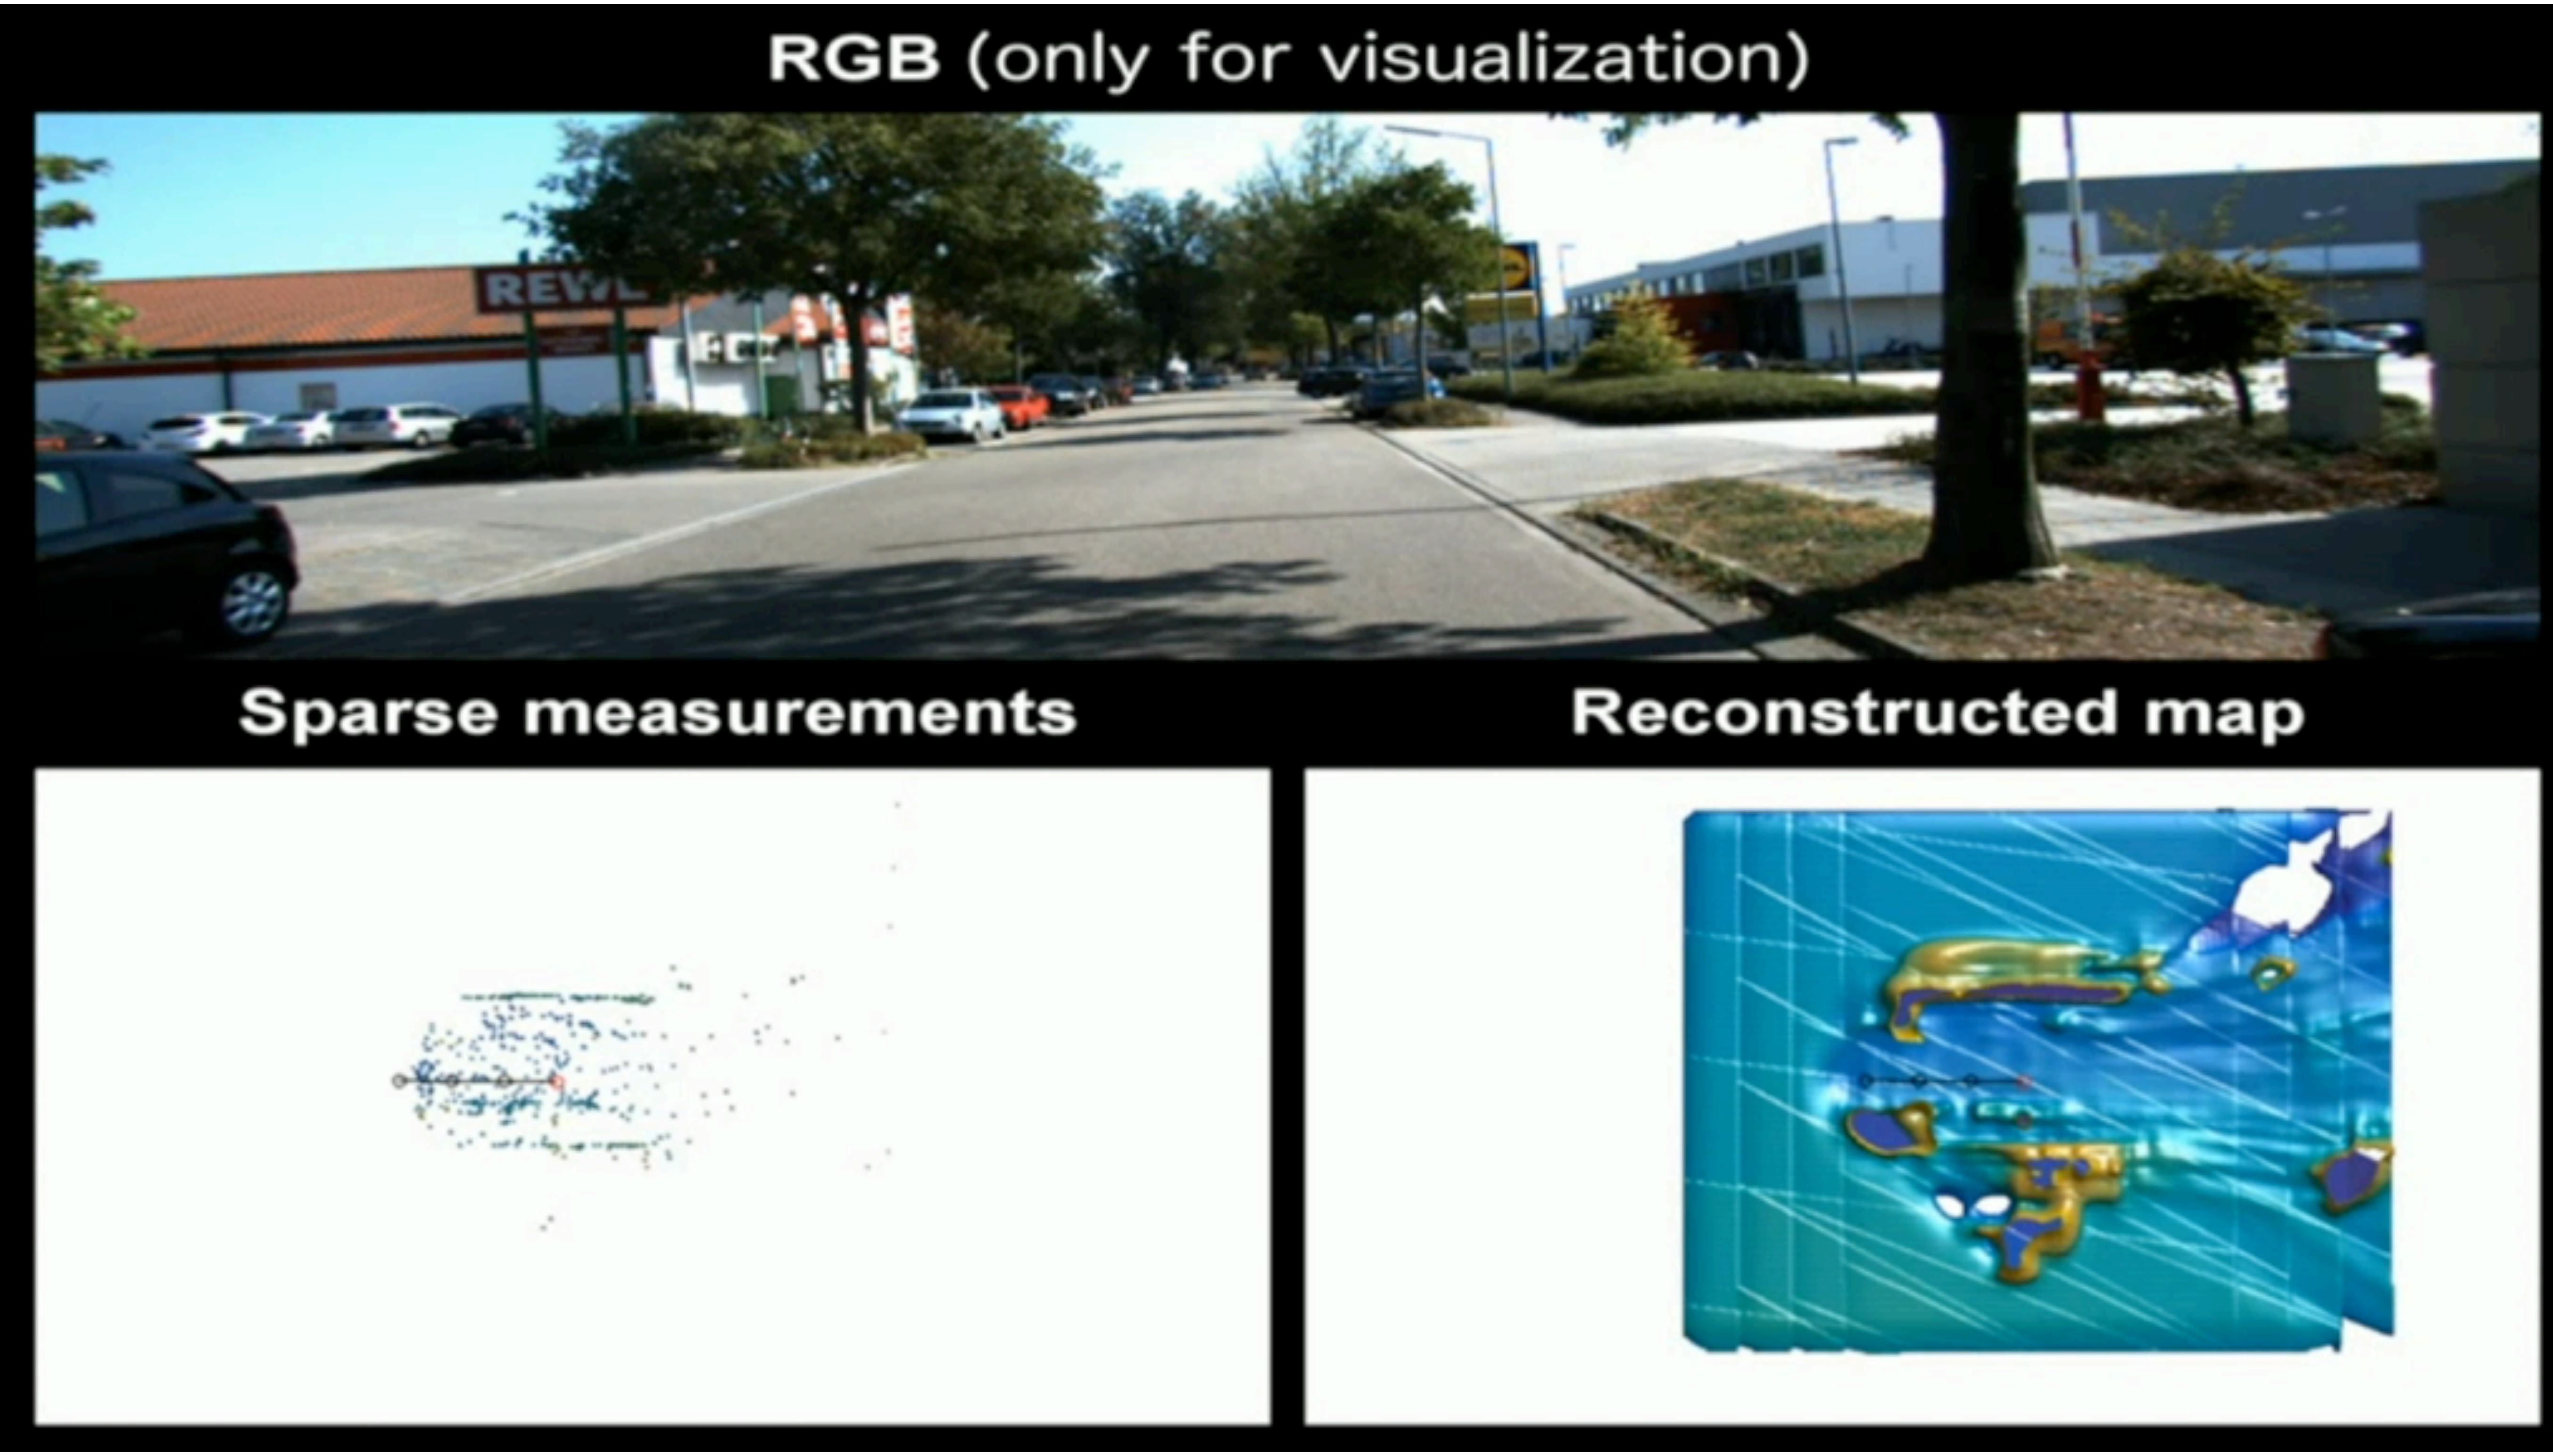

[5] Zimmermann, Petricek, Salansky, Svoboda, Learning for Active 3D Mapping, **ICCV oral (rank A*, AC=2%)**, 2017

What we want to do?

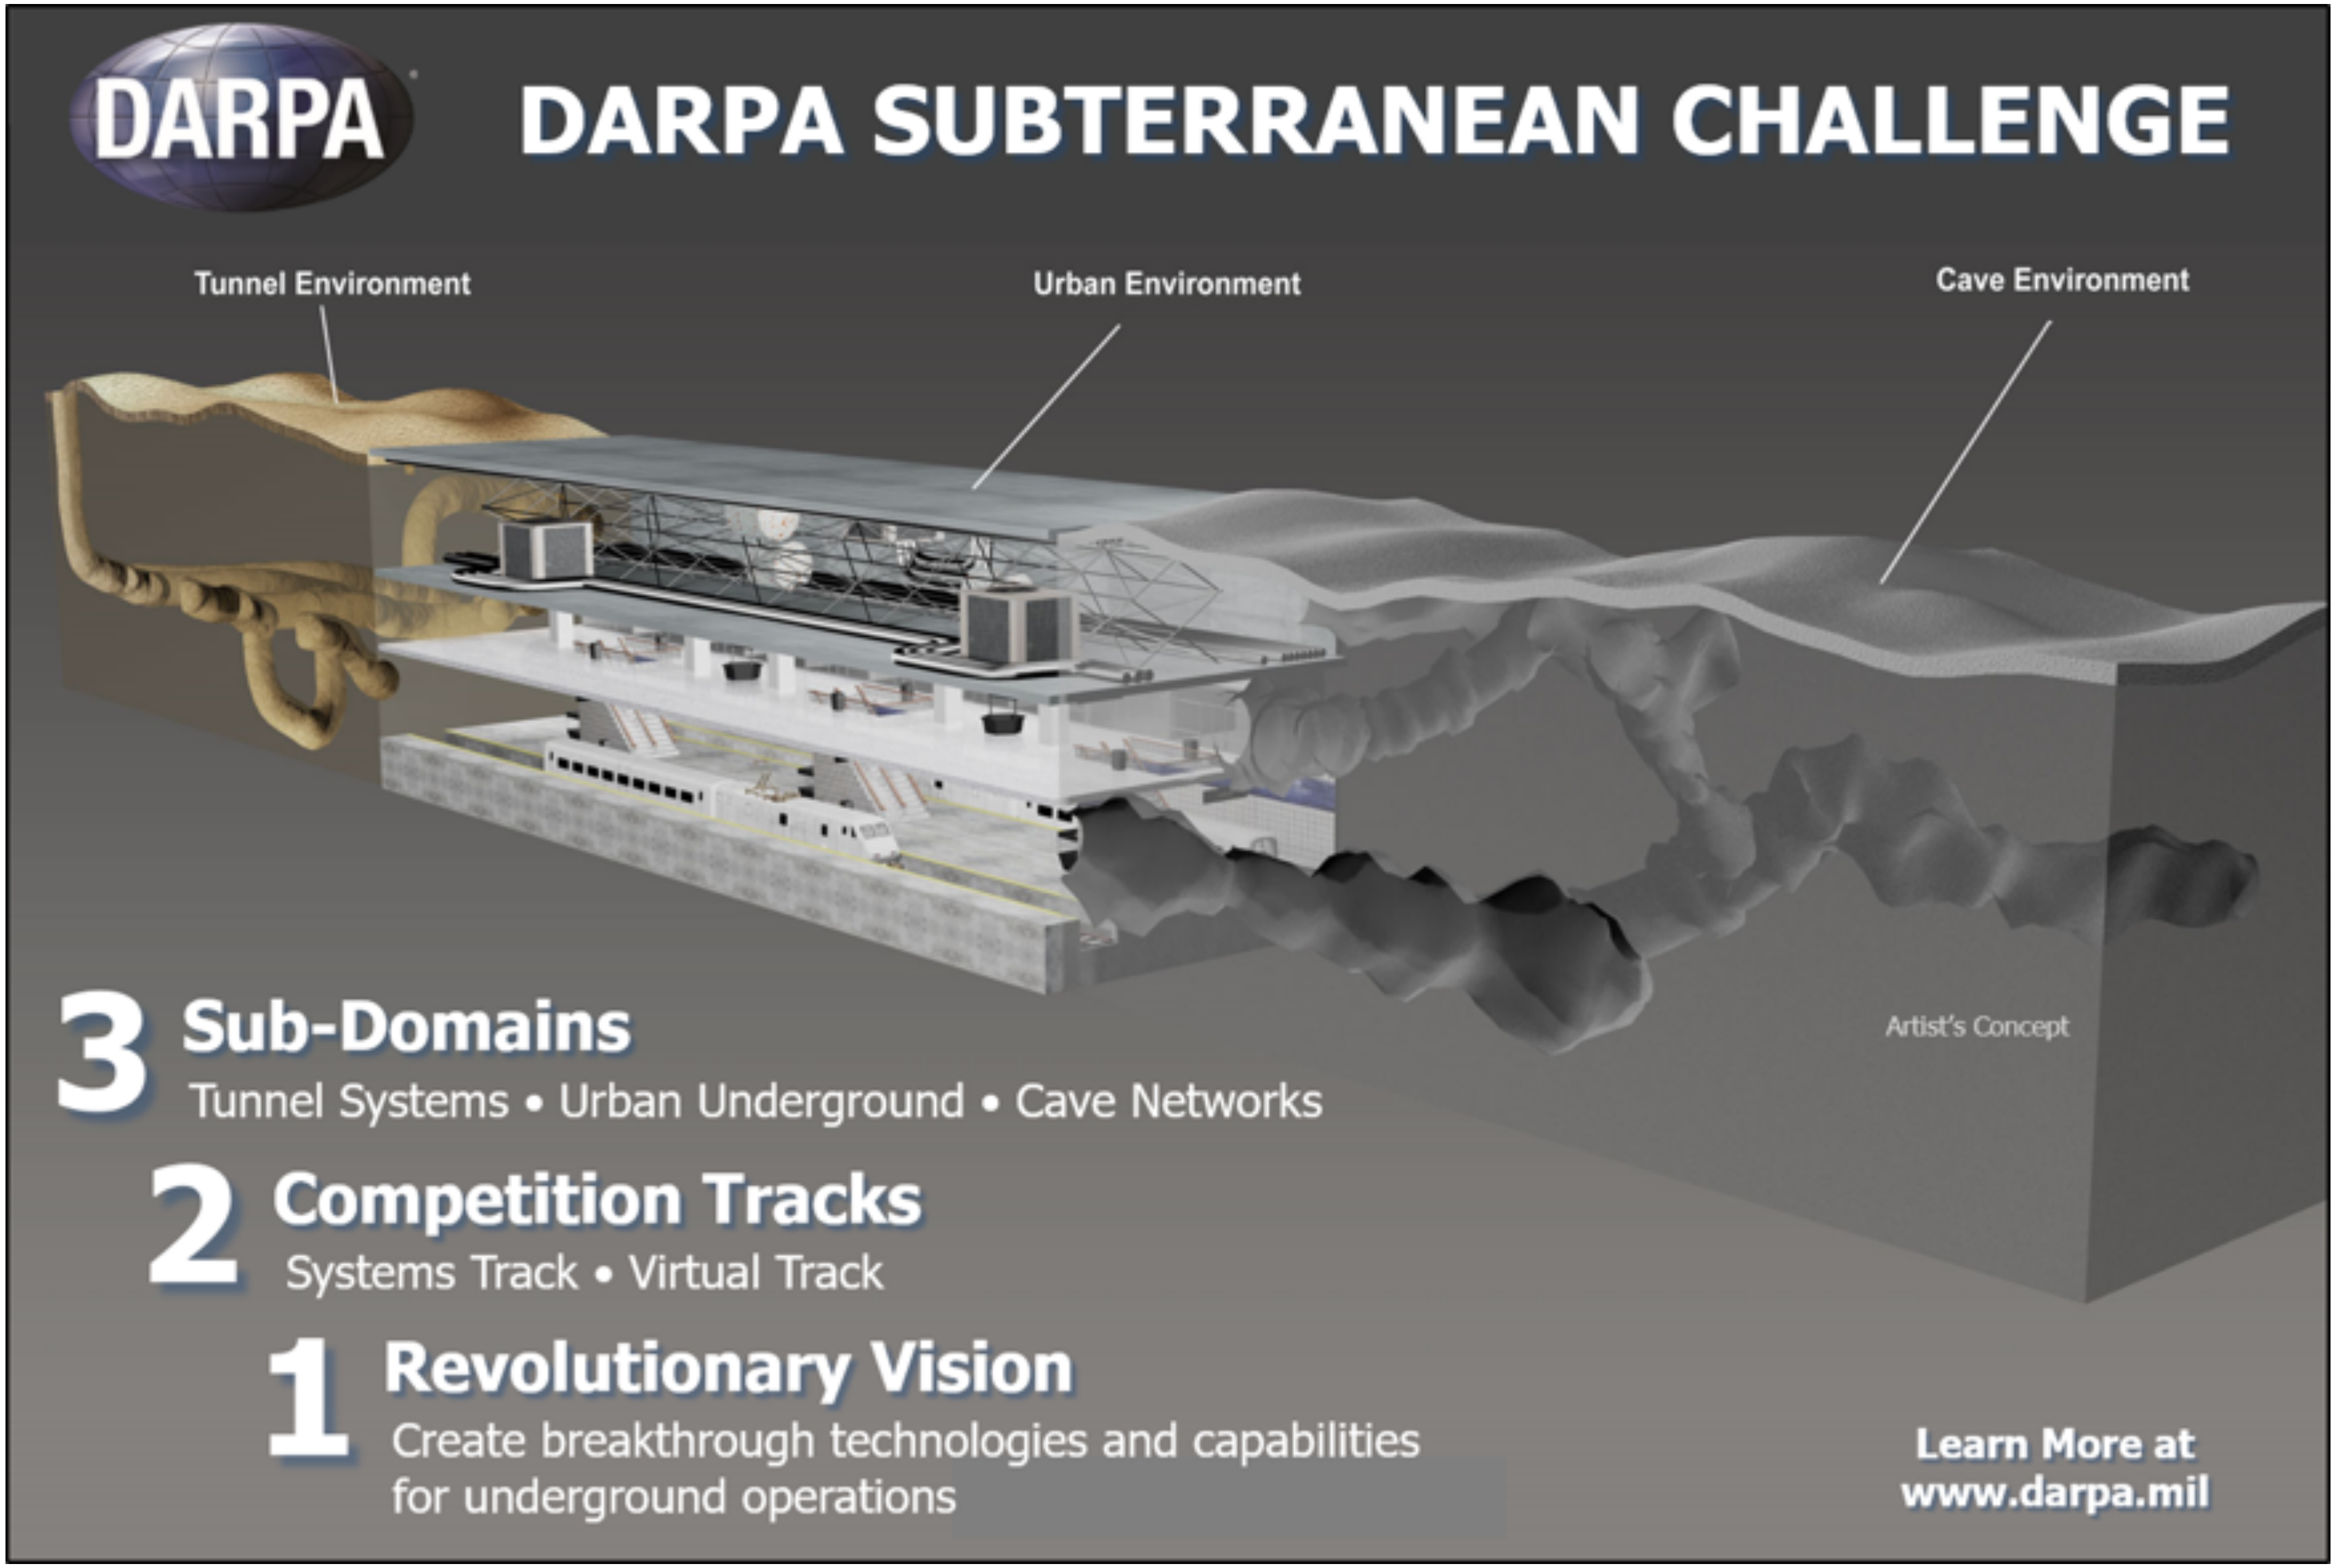

https://www.darpa.mil/news-events/2017-12-21