University of Zurich[UZH] | ETH zürich
Institute of Informatics – Institute of Neuroinformatics
ROBOTICS & PERCEPTION GROUP

Vision Algorithms for Mobile Robotics

Lecture 01
Introduction

Davide Scaramuzza
http://rpg.ifi.uzh.ch

Today's Class

- About me and my research lab
- What is Computer Vision?
- Why study computer vision?
- Example of Vision Applications
- Live Demos!
- Specifics of this course
- Overview of Visual Odometry

Who am I?

Current positions

- Professor of Robotics, Dep. of Informatics and Neuroinformatics (UZH & ETH)

Education

- PhD from ETH Zurich with Roland Siegwart
- Post-doc at the University of Pennsylvania with Vijay Kumar & Kostas Daniilidis

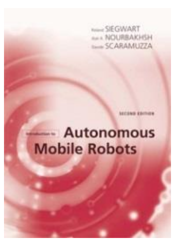

Highlights

- Coordinator of the European project *sFly* on visual navigation of micro drones
  - Which introduced the PX4 autopilot and visual navigation of drones
- Book “Autonomous Mobile Robots,” 2011, MIT Press

Spinoffs & Tech Transfer

- **Zurich-Eye,** enabling machines to see, now **Facebook-Oculus Zurich**
- Former strategic advisor of **Dacuda**, now **Magic Leap Zurich**
- **Fotokite**, aerial filming made simple, incubated in my lab

My Research Background

Computer Vision

- Visual Odometry and SLAM
- Sensor fusion
- Camera calibration

Autonomous Robot Navigation

- Self driving cars
- Micro Flying Robots

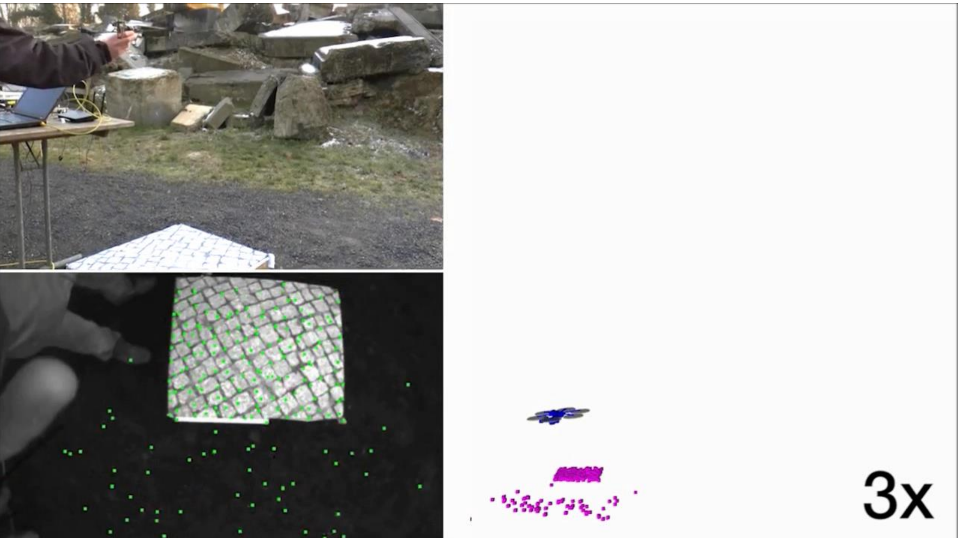

My lab

ROBOTICS & PERCEPTION GROUP

http://rpg.ifi.uzh.ch

Closed to bahnhof Oerlikon, Andreasstrasse 15, 2nd floor

Research Overview

Real-time, Onboard Computer Vision and Control for **Autonomous, Agile** Drone **Flight**

Falanga et al., **The Foldable Drone: A Morphing Quadrotor that can Squeeze and Fly**, RAL'19. PDF. Videos.
Featured in IEEE Spectrum.

Research Overview

Real-time, Onboard Computer Vision and Control for **Autonomous, Agile** Drone **Flight**

Kaufmann, Loquercio, Dosovitskiy, Ranftl, Koltun, Scaramuzza, *Deep Drone Racing: Learning Agile Flight in Dynamic Environments*, Conference on Robot Learning (CORL), Zurich, Oct. 29-31, 2018. PDF, YouTube

Student Projects: http://rpg.ifi.uzh.ch/student_projects.php

University of Zurich[UZH]

ETH zürich

ROBOTICS & PERCEPTION GROUP

Department of Informatics - Institute of Neuroinformatics - Robotics and Perception Group

- News
- People
- Research
- Publications
- Software/Datasets
- Open Positions
- **Student Projects**
- Teaching
- Media Coverage
- Awards
- Gallery
- Contact

Student Projects

How to apply

To apply, please send **your CV, your Ms and Bs transcripts** by email to **all the contacts** indicated below the project description. **Do not apply on SiROP . Since Prof. Davide Scaramuzza is affiliated with ETH, there is no organizational overhead for ETH students.** Custom projects are occasionally available. If you would like to do a project with us but could not find an advertized project that suits you, please contact Prof. Davide Scaramuzza directly to ask for a **tailored project** (sdavide at ifi.uzh.ch).

Upon successful completion of a project in our lab, students may also have the opportunity to get an **internship at one of our numerous industrial and academic partners worldwide** (e.g., NASA/JPL, University of Pennsylvania, UCLA, MIT, Stanford, ...).

Probabilistic System Identification of a Quadrotor Platform - Available

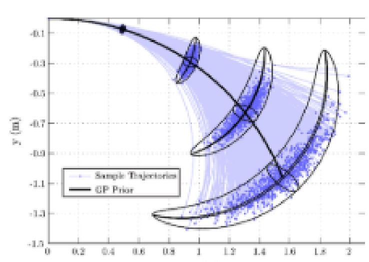

Description: Most planning & control algorithms used on quadrotors make use of a nominal model of the platform dynamics to compute feasible trajectories or generate control commands. Such models are derived using first principles and typically cannot fully capture the true dynamics of the system, leading to sub-optimal performance. One appealing approach to overcome this limitation is to use Gaussian Processes for system modeling. Gaussian Process regression has been widely used in supervised machine learning due to its flexibility and inherent ability to describe uncertainty in the prediction. This work investigates the usage of

Successful Startups

Fotokite (2014) – Power-over-tether drone for aerial filming

- Pilot-free tethered aerial camera system with limitless flight time and data bandwidth
- 1st and only system approved by the FAA for Public Safety teams to use without a pilot license

Zurich-Eye (2015) - now Oculus Zurich

Vision-based Localization and Mapping Solutions for Mobile Robots

Created in Sep. 2015, **became Facebook-Oculus Zurich in Sep. 2016**

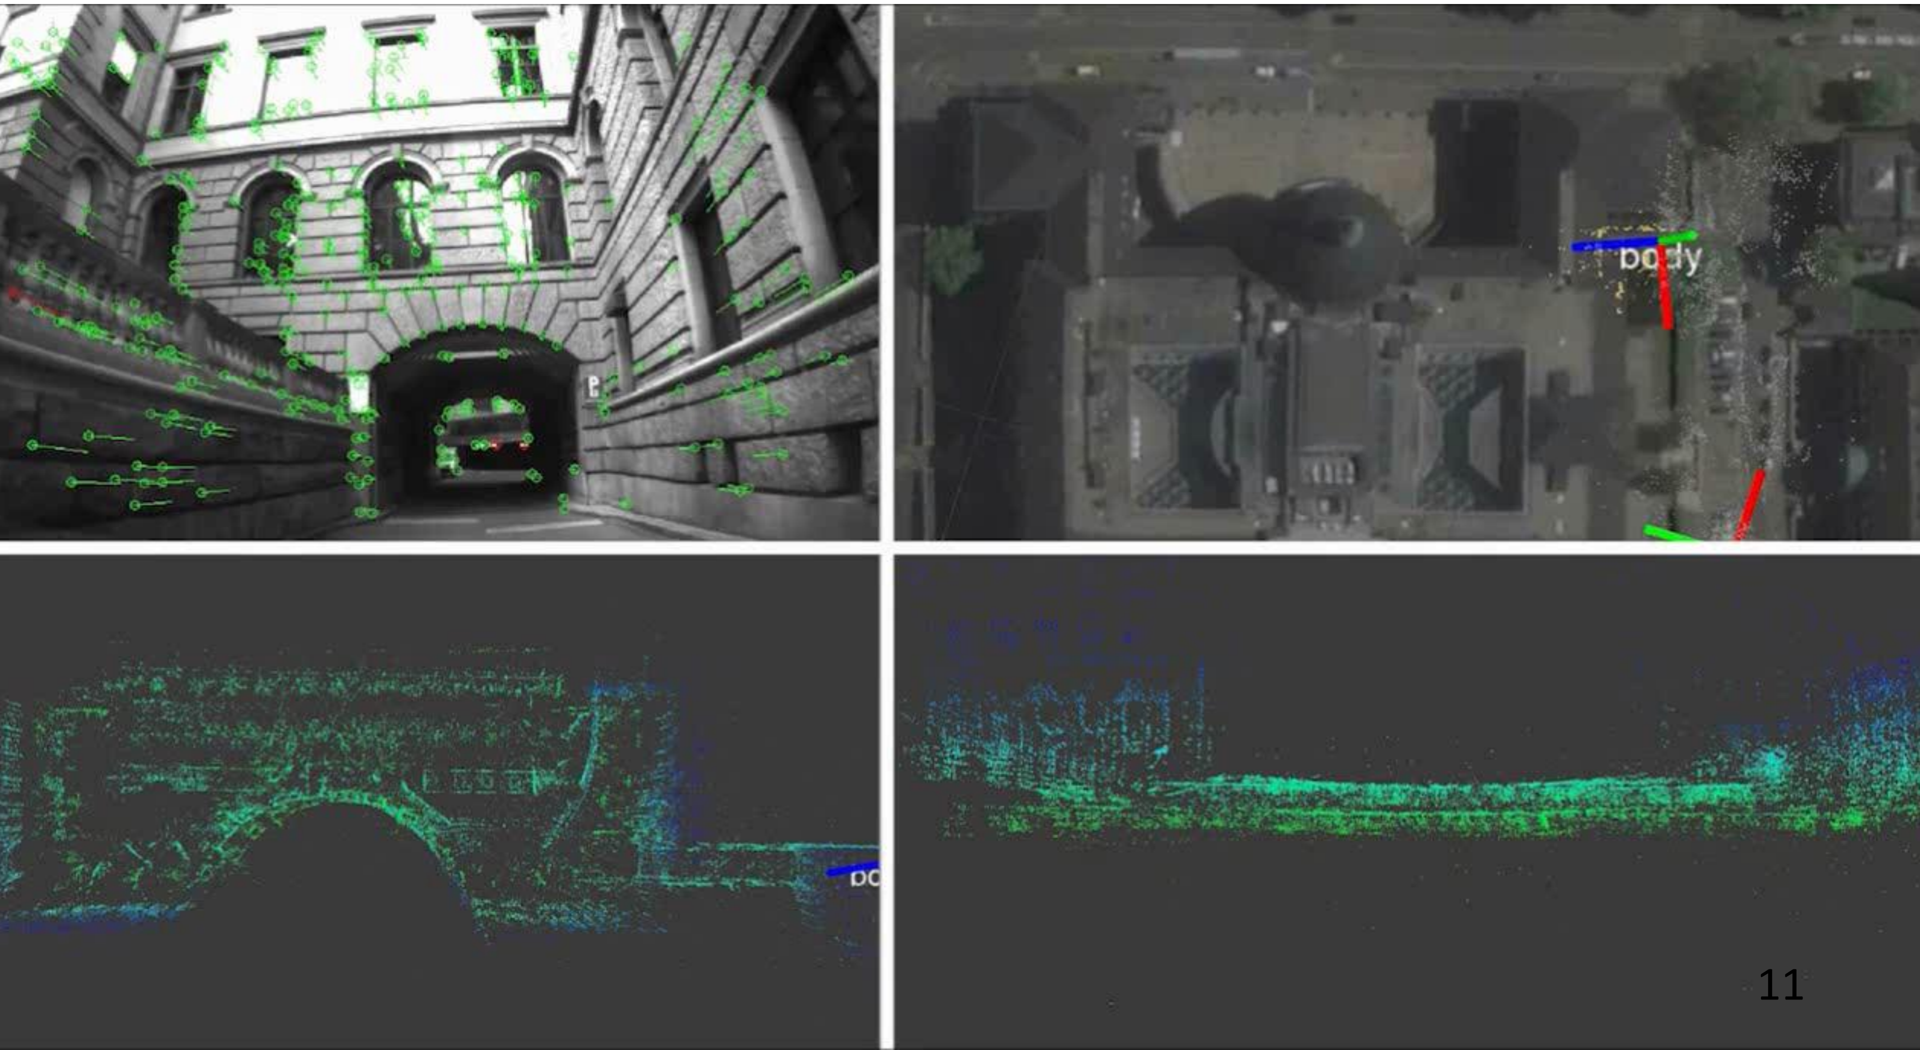

Zurich-Eye (2015) - now Oculus Zurich

Vision-based Localization and Mapping Solutions for Mobile Robots

Created in Sep. 2015, **became Facebook-Oculus Zurich in Sep. 2016**

The Zurich Eye team is behind the new [**Oculus Quest**]()

We will have a lecture by Christian Forster, from Oculus Zurich end of November!

Today's Class

- About me and my research lab
- What is Computer Vision?
- Why study computer vision?
- Example of Vision Applications
- Live Demos!
- Specifics of this course
- Overview of Visual Odometry

What is computer vision?

Automatic extraction of "meaningful" information from images and videos

Semantic information

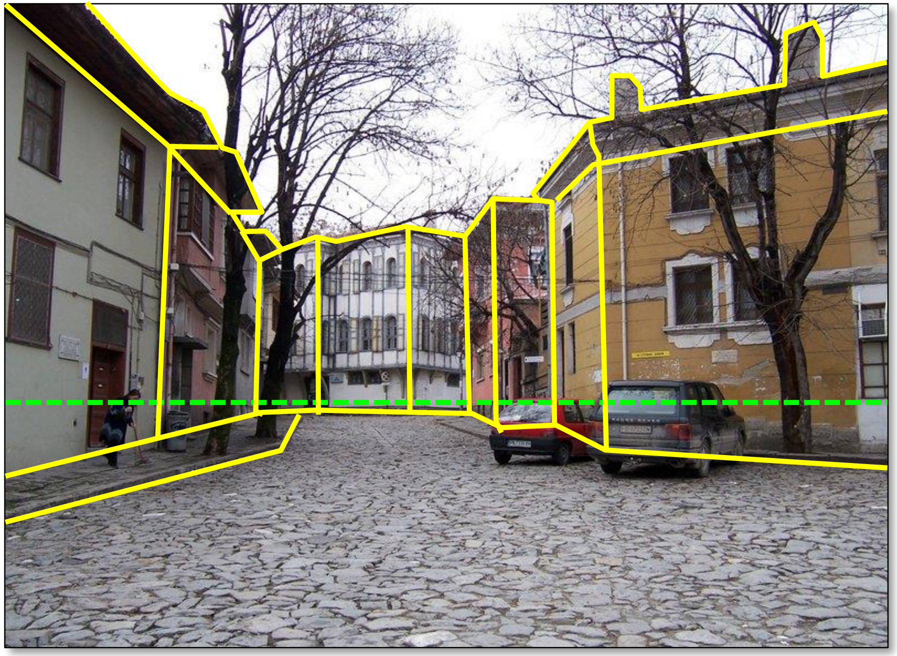

Geometric information (this course)

Vision Demo?

Terminator 2

We are almost there!

Google Lens

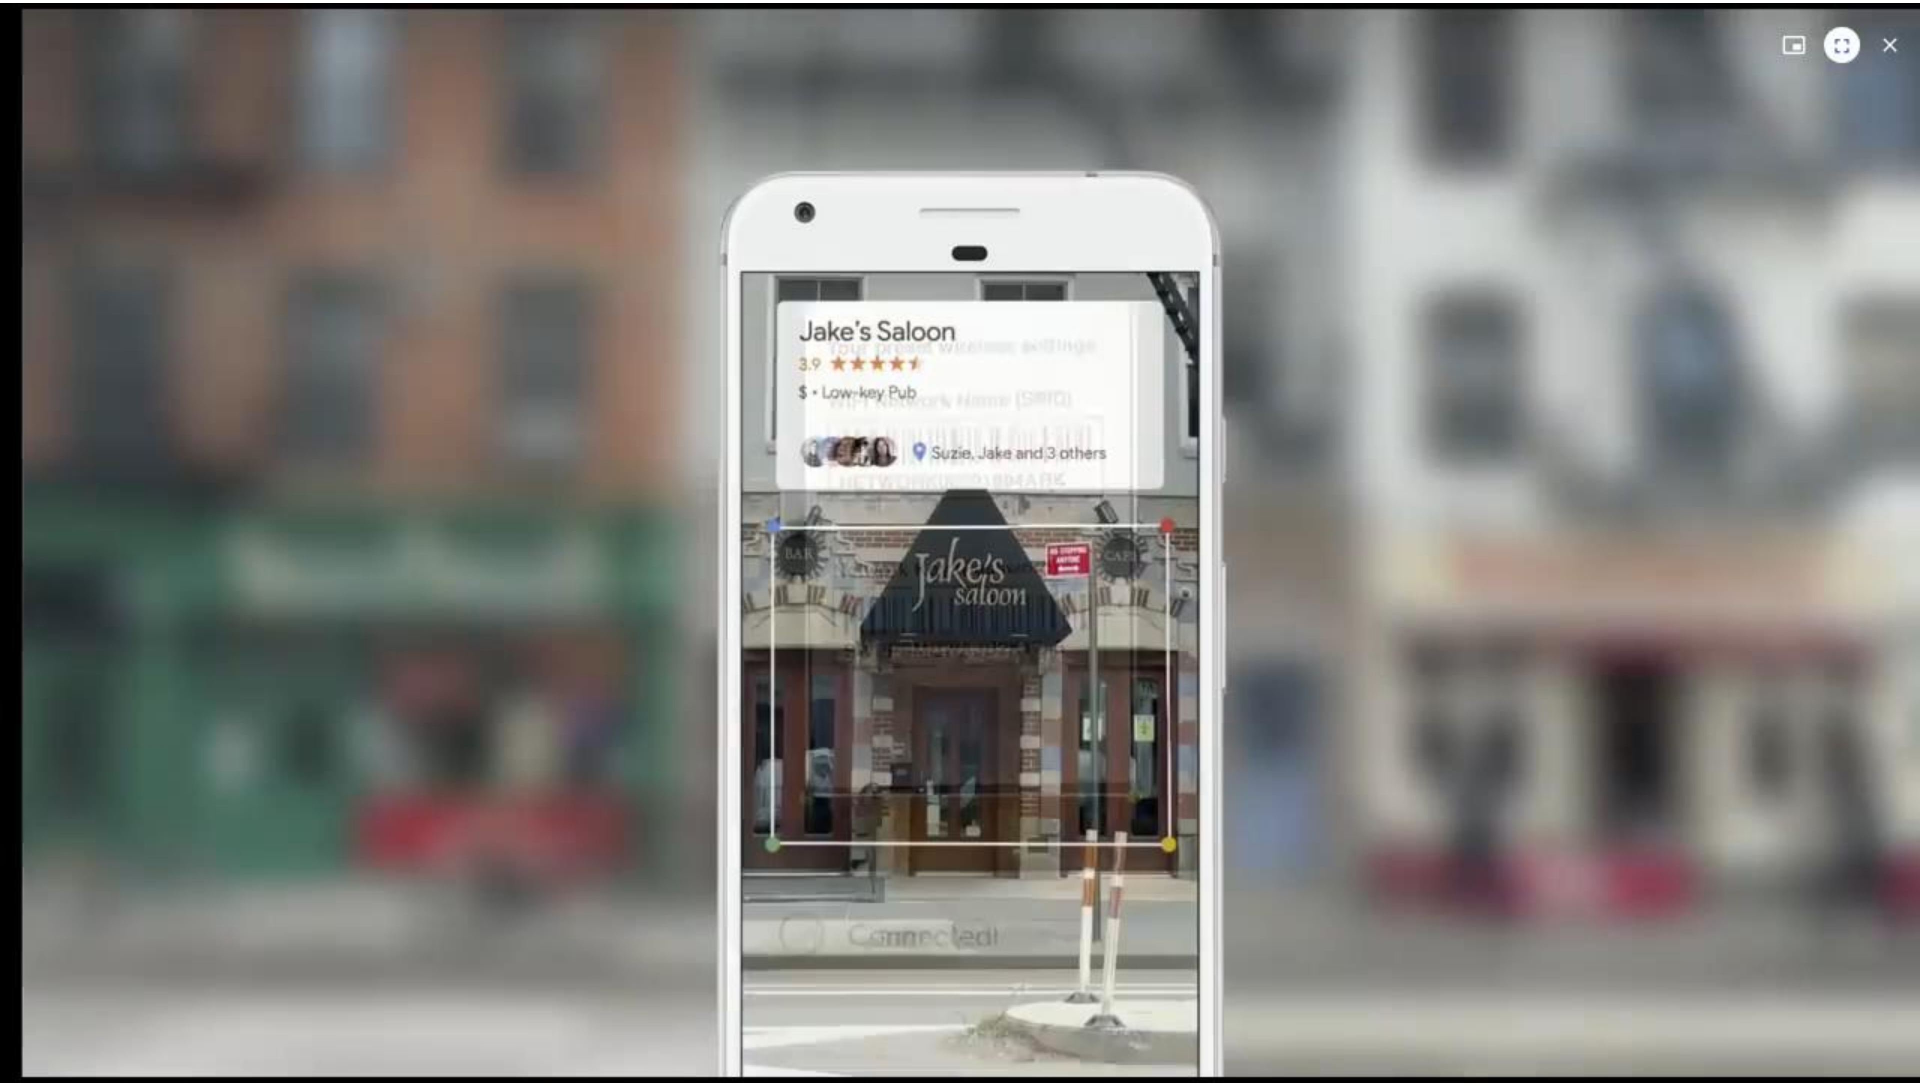

Jake's Saloon
3.9
$ • Low-key Pub
Suzie, Jake and 3 others
Jake's saloon

Today's Class

- About me and my research lab
- What is Computer Vision?
- Why study computer vision?
- Example of Vision Applications
- Live Demos!
- Specifics of this course
- Overview of Visual Odometry

Why study computer vision?

- Relieve humans of boring, easy tasks
- Enhance human abilities: human-computer interaction, visualization, augmented reality (AR)
- Perception for autonomous robots
- Organize and give access to visual content
- Lots of computer-vision companies and jobs in Switzerland (Zurich & Lausanne):
  - Facebook-Oculus (Zurich): AR/VR
  - Magic-Leap (Zurich & Lausanne): AR/VR
  - Microsoft Research (Zurich): Robotics and Hololens
  - Google (Zurich): Brain, ARCore, Street View, YouTube
  - Apple (Zurich): Autonomous Driving, face tracking
  - NVIDIA (Zurich): simulation, autonomous driving
  - Logitech (Zurich, Lausanne)
  - Disney-Research (Zurich)
  - Pix4D (Lausanne)
  - VIZRT (Zurich): sport broadcasting, 3D replay
  - More: https://de.glassdoor.ch/Job/z%C3%BCrich-computer-vision-jobs-SRCH_IL.0,6_IC3297851_KO7,22.htm

Vision in humans

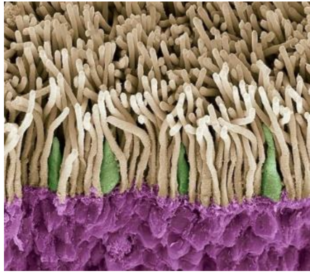

- **Vision** is our most powerful sense. **Half of primate cerebral cortex** is devoted to visual processing
- Retina is ~1,000 mm^2. Contains 130 million **photoreceptors** (120 mil. rods (low light vision) and 10 mil. cones for color sampling)
- Provides **enormous** amount of information: data-rate of ~**3GBytes/s**
- To match the eye resolution we would need a **500 Megapixel** camera. But in practice the acuity of an eye is **8 Megapixels** within a **18-degree field of view** (5.5 mm diameter) region called **fovea**

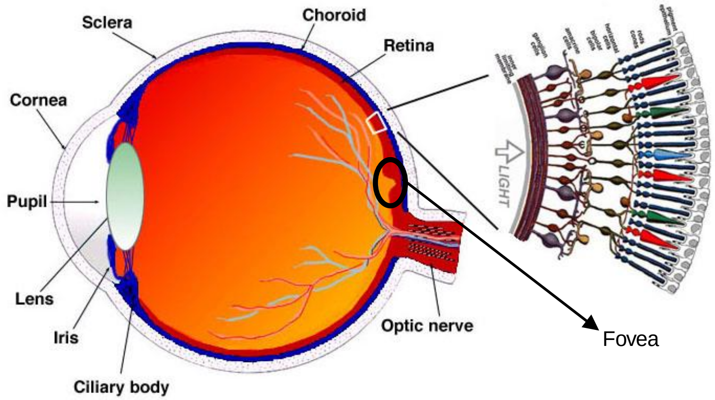

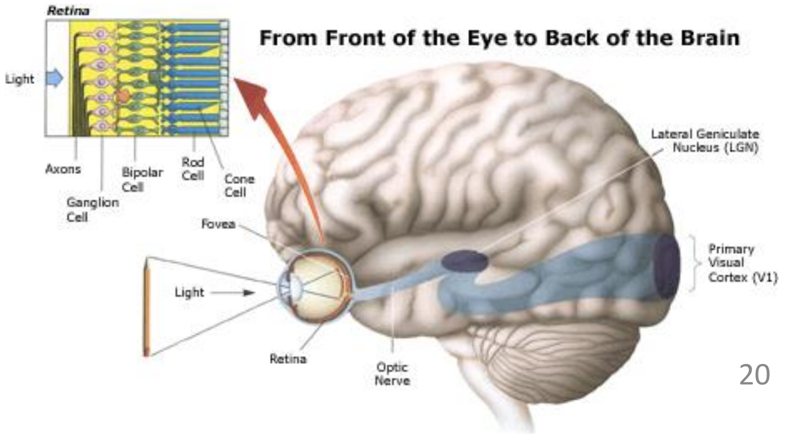

What a newborn sees every month in the first year

"Your baby sees things best from 15 to 30 cm away. This is the perfect distance for gazing up into the eyes of mom or dad. Any farther than that, and the newborn sees mostly blurry shapes because they're nearsighted. At birth, a newborn's eyesight is between 20/200 and 20/400."

http://uk.businessinsider.com/what-a-baby-can-see-every-month-for-the-first-year-of-its-life-2017-1?r=US&IR=T

Why is vision hard?

How do we go from an array of number to recognizing a fruit?

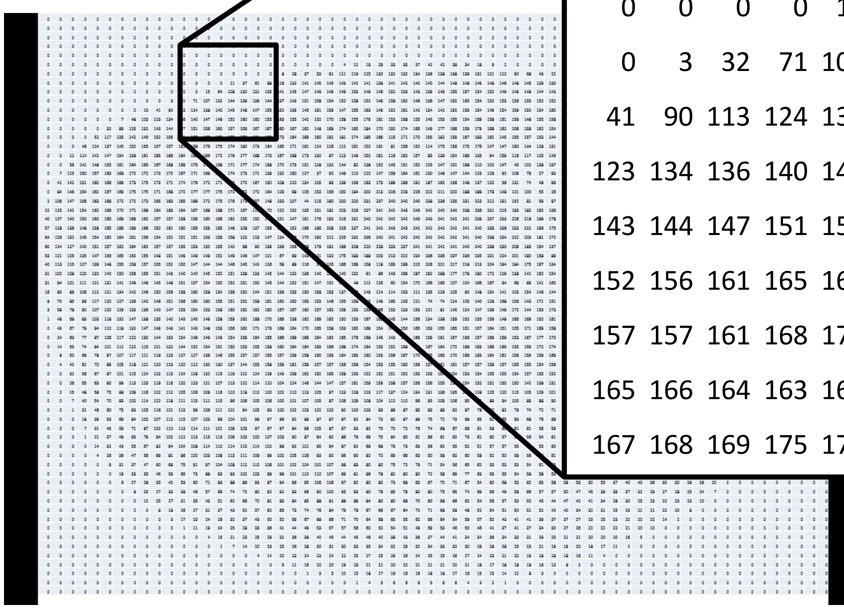

0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	8	28
0	0	0	0	0	0	11	37	61	88	116	132
0	0	0	0	15	64	108	130	131	135	141	145
0	3	32	71	107	132	144	139	139	144	137	143
41	90	113	124	138	140	145	148	147	155	152	139
123	134	136	140	147	149	152	160	160	155	163	155
143	144	147	151	156	160	157	159	167	167	160	167
152	156	161	165	166	169	170	170	164	169	173	164
157	157	161	168	176	175	174	180	173	164	165	171
165	166	164	163	166	172	179	177	168	173	167	168
167	168	169	175	173	168	171	177	174	168	172	173

What a computer sees

What we see

Origins of computer vision

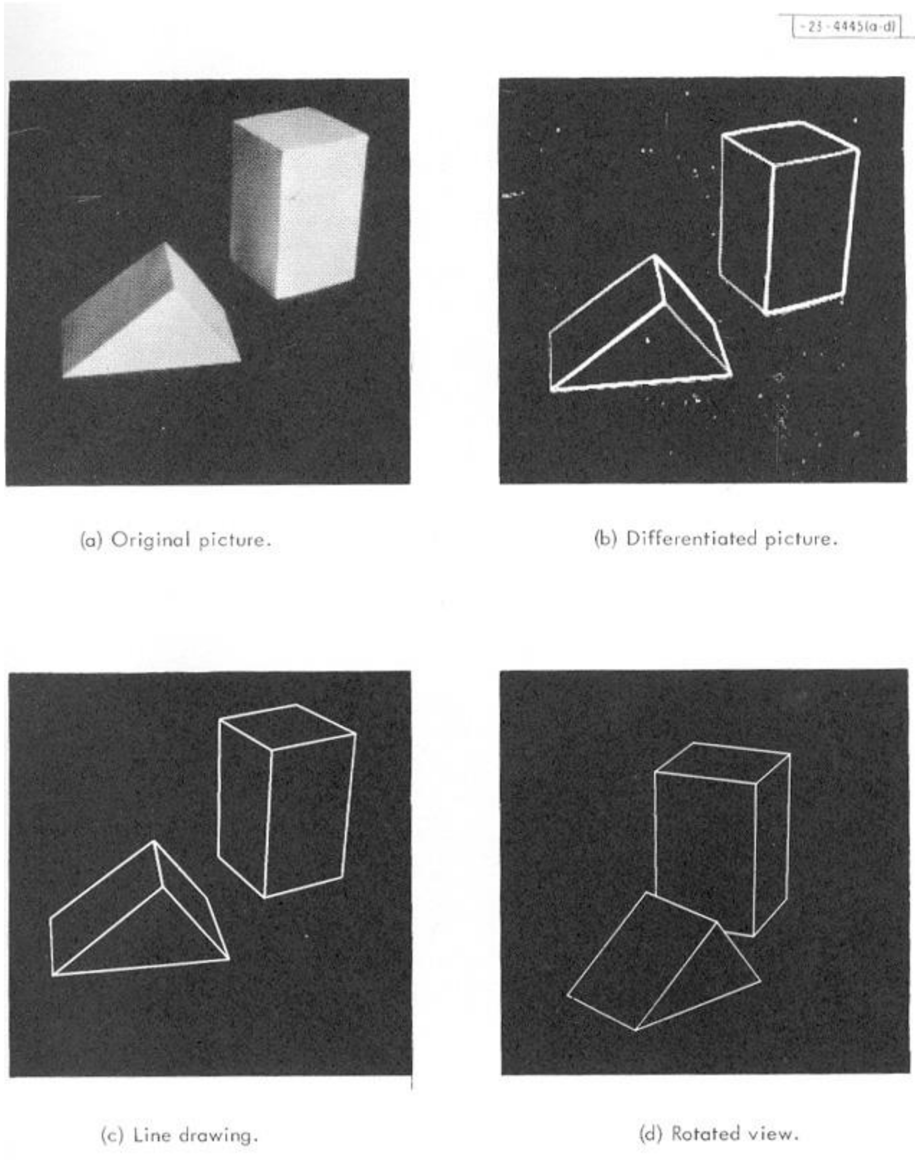

L. G. Roberts, *Machine Perception of Three Dimensional Solids*, Ph.D. thesis, MIT Department of Electrical Engineering, 1963.

He is the **inventor of ARPANET, the current Internet**

Related disciplines

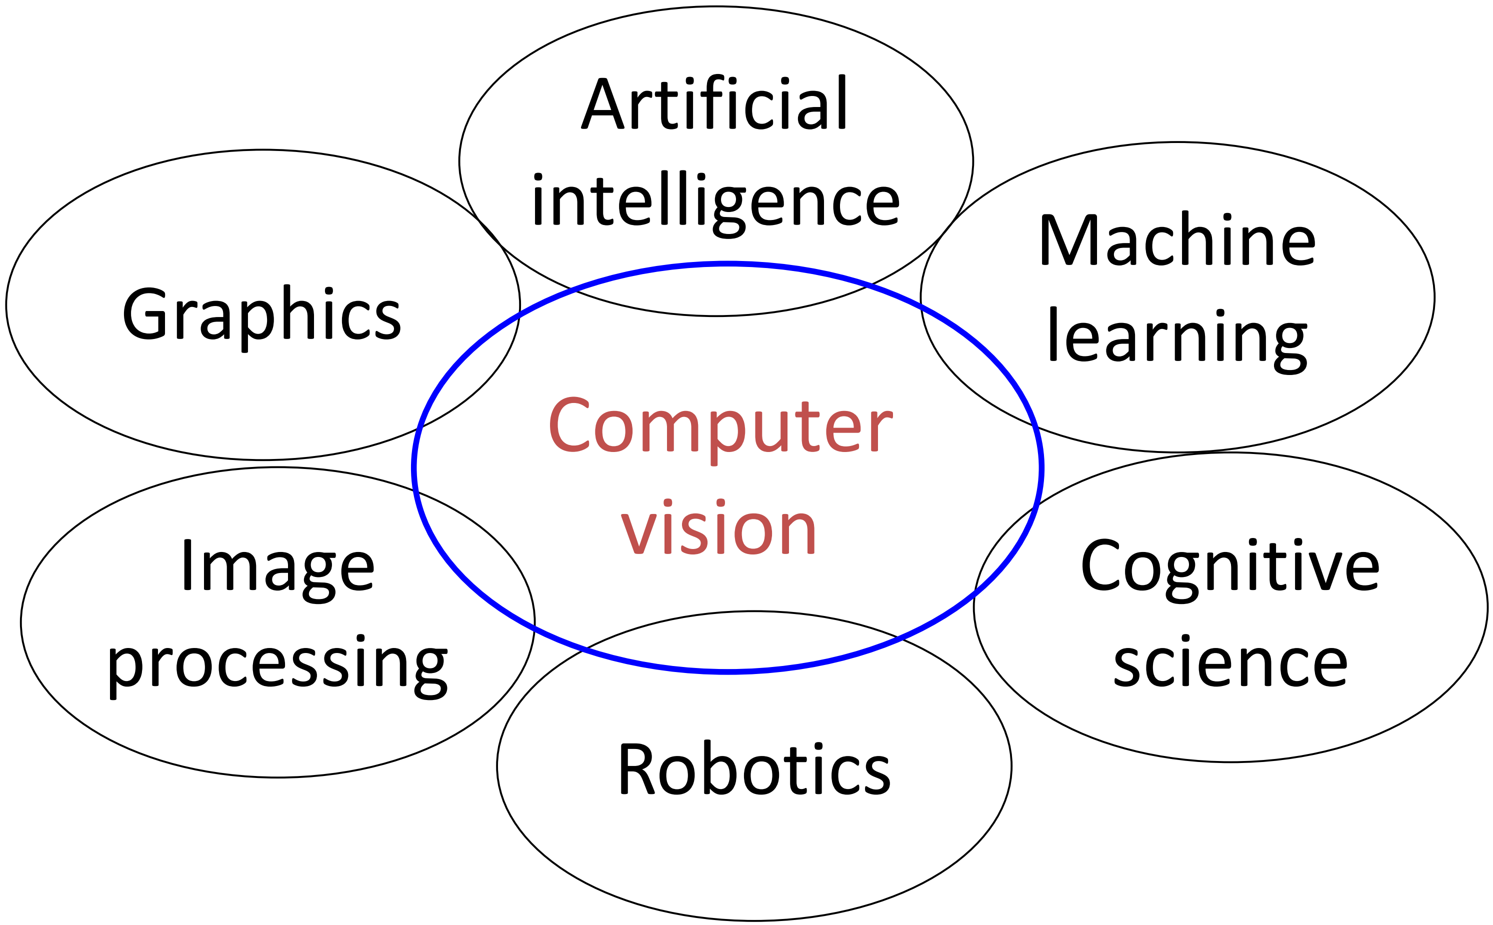

Computer Vision vs Computer Graphics

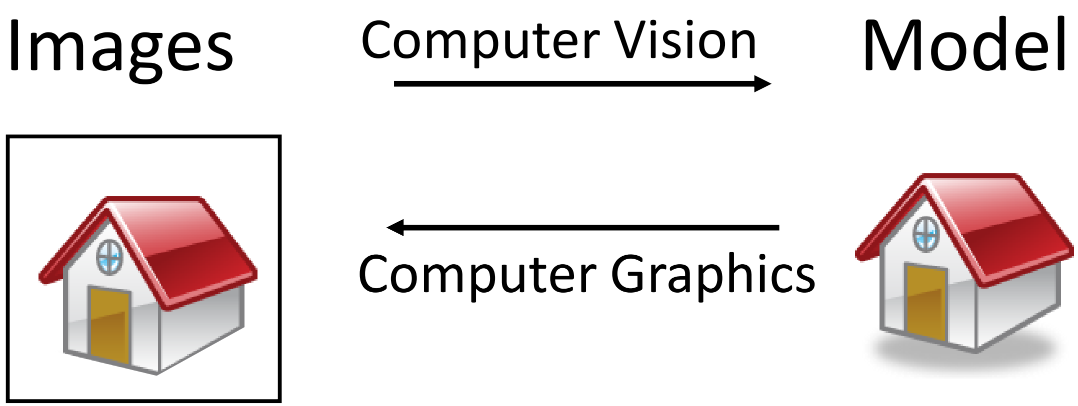

Inverse problems: analysis and synthesis.

Today's Class

- About me and my research lab
- What is Computer Vision?
- Why study computer vision?
- Example of Vision Applications
- Live Demos!
- Specifics of this course
- Overview of Visual Odometry

Optical character recognition (OCR)

Technology to convert scanned docs to text

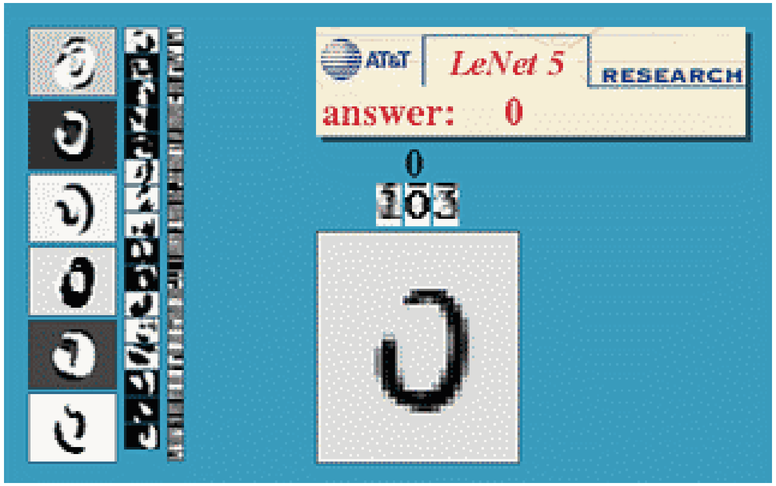

Digit recognition, AT&T labs, using CNN,
by Yann LeCun (1993),
now head of Deep Learning at Facebook
http://yann.lecun.com/

License plate readers
http://en.wikipedia.org/wiki/Automatic_number_plate_recognition

Face detection

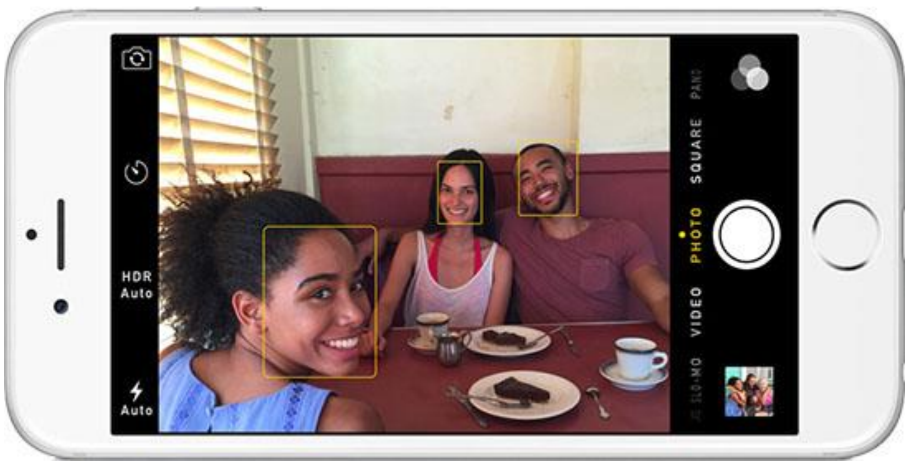

nethods such as Support Vector Machines, Prin-
nponent Analysis and Linear Discriminant Analy-
imited capacity to leverage large volumes of data,
al networks have shown better scaling properties.
ly, there has been a surge of interest in neu-
rks [19, 21]. In particular, deep and large net-
ve exhibited impressive results once: (1) they
applied to large amounts of training data and (2)
omputation resources such as thousands of CPU
] and/or GPU's [19] have become available. Most
Krizhevsky et al. [19] showed that very large and
volutional networks [21] trained by standard back-
on [25] can achieve excellent recognition accuracy
ned on a large dataset.

ecognition state of the art Face recognition er-
have decreased over the last twenty years by three
magnitude [12] when recognizing frontal faces in
es taken in consistently controlled (constrained)
ents. Many vendors deploy sophisticated systems
pplication of border-control and smart biometric
tion. However, these systems have shown to be
o various factors, such as lighting, expression, oc-

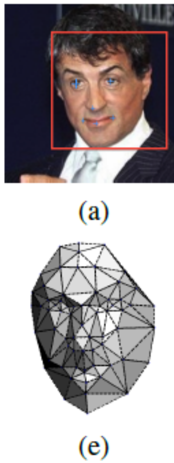

Figure 1. **Alignm**
cial points. (b) Th
the 2D-aligned cro
added triangles on
3D shape transfor
visibility w.r.t. to th
(f) The 67 fiducial
the piece-wise affi
view generated by

Now in all new digital cameras and smartphon

- Taigman, Yang, Ranzato, Wolf, DeepFace: Closing the Gap to Human-Level Performance in Face Verification, CVPR'14, **by Facebook.**
- Schroff, Kalenichenko, Philbin, FaceNet: A Unified Embedding for Face Recognition and Clustering, CVPR'15, **by Google.**
- Before 2012:
  Paul Viola, Michael Jones: Robust Real-time Object Detection, Int. Journal of Computer Vision 2001
  (NB. Paul Viola is now Vice President of Amazon Prime Air)

Pix4D

- EPFL startup – Now a company

Automotive safety

- Mobileye: Vision systems in high-end Tesla, BMW, GM, Volvo models. Bought by **Intel in 2017 for 15 billion USD**!
  - Pedestrian collision warning
  - Forward collision warning
  - Lane departure warning
  - Headway monitoring and warning

Video gaming: Xbox Kinect

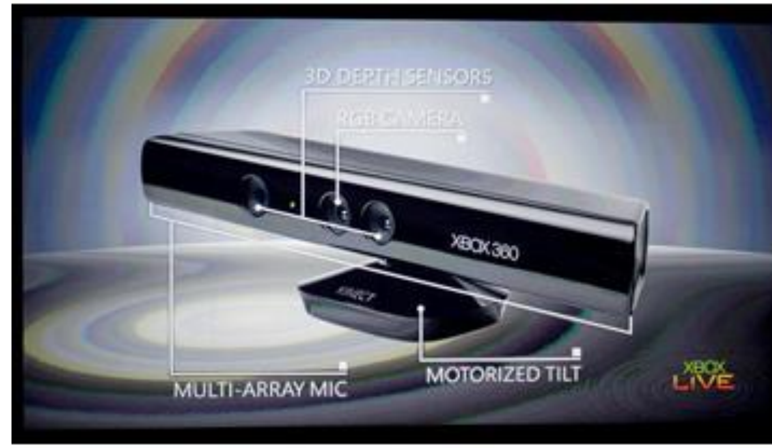

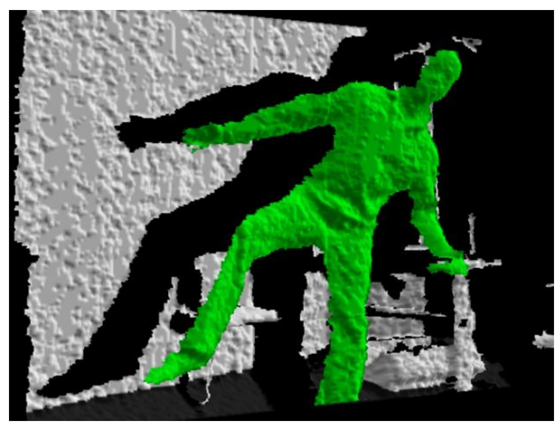

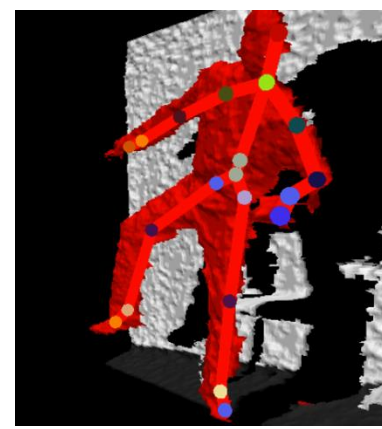

Lot of Computer Vision in Modern Smartphones

iPhone X

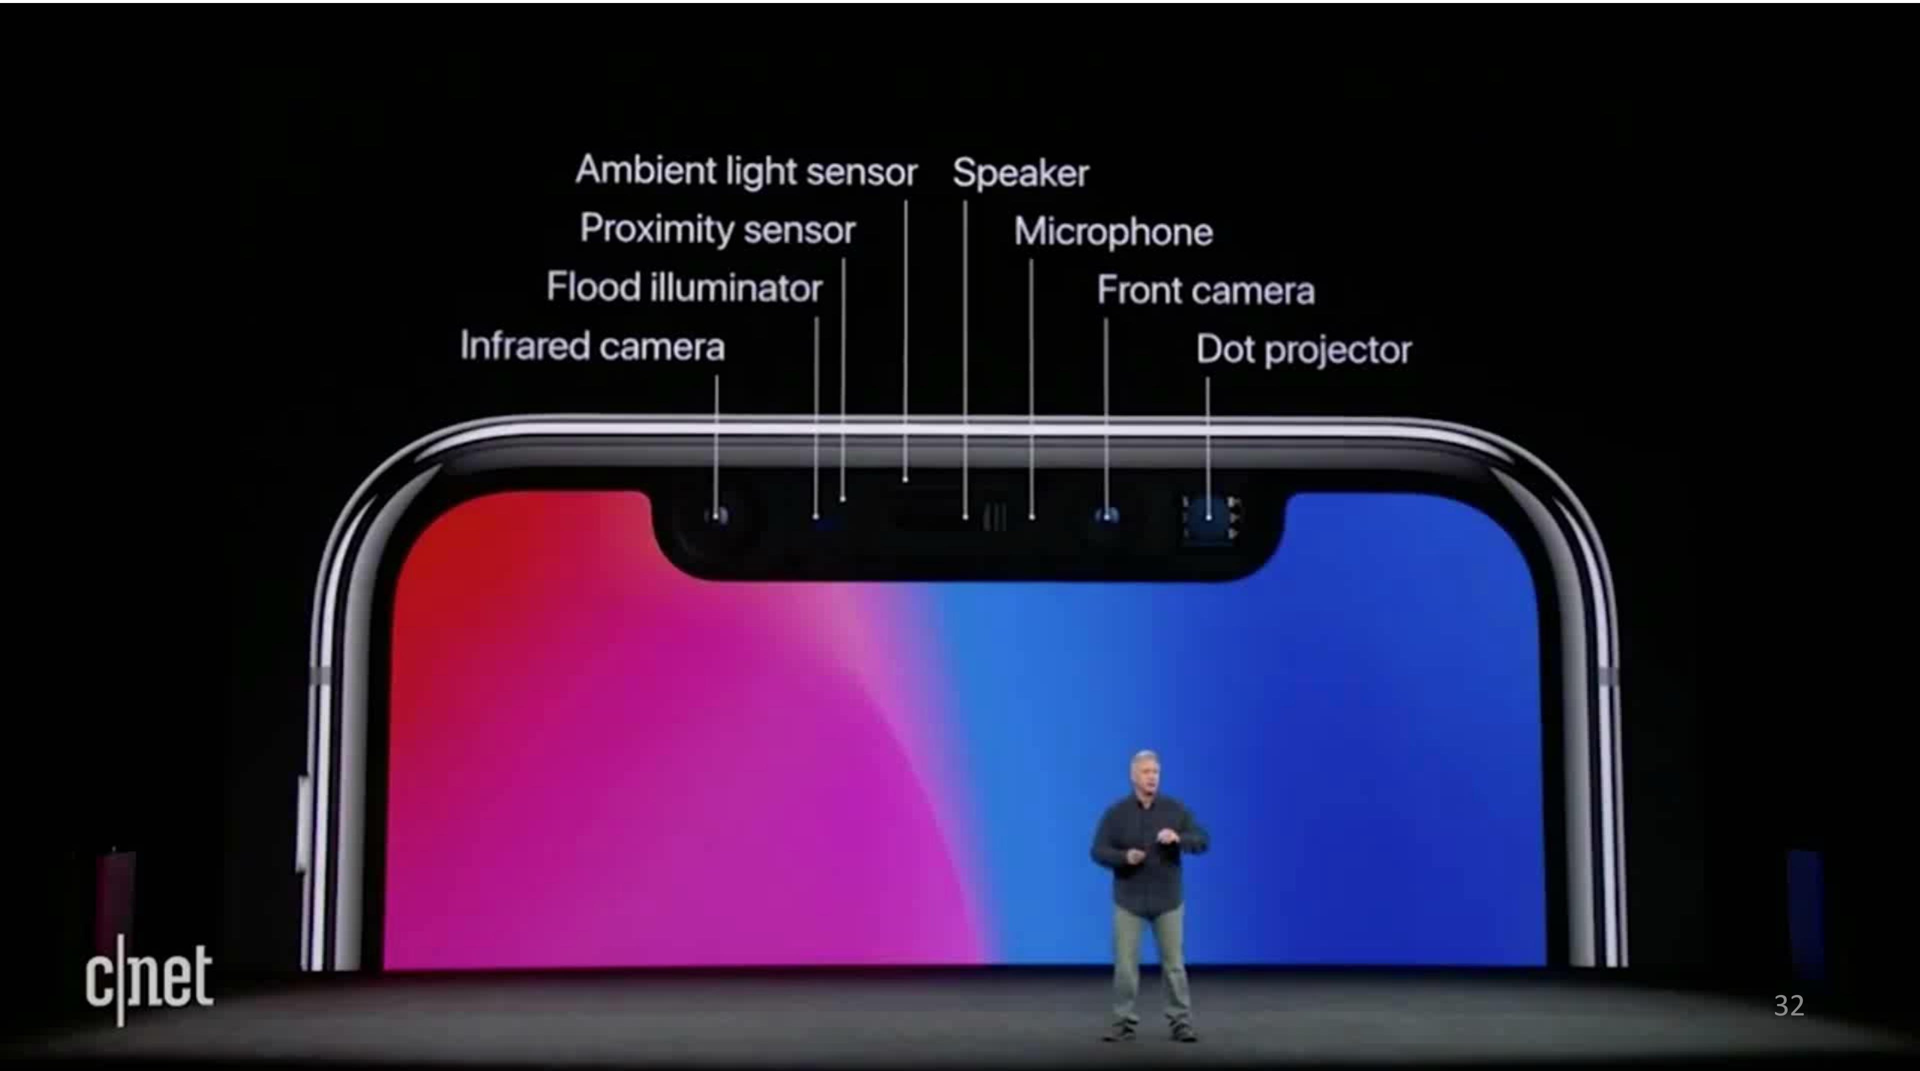

Vision in space

NASA'S Mars Exploration Rover Spirit captured this westward view from atop a low plateau where Spirit spent the closing months of 2007.

Vision systems (made by JPL) used for several tasks

- Panorama stitching
- 3D terrain modeling
- Obstacle detection, position tracking
- For more, read “Computer Vision on Mars” by Matthies et al.

Vision in space

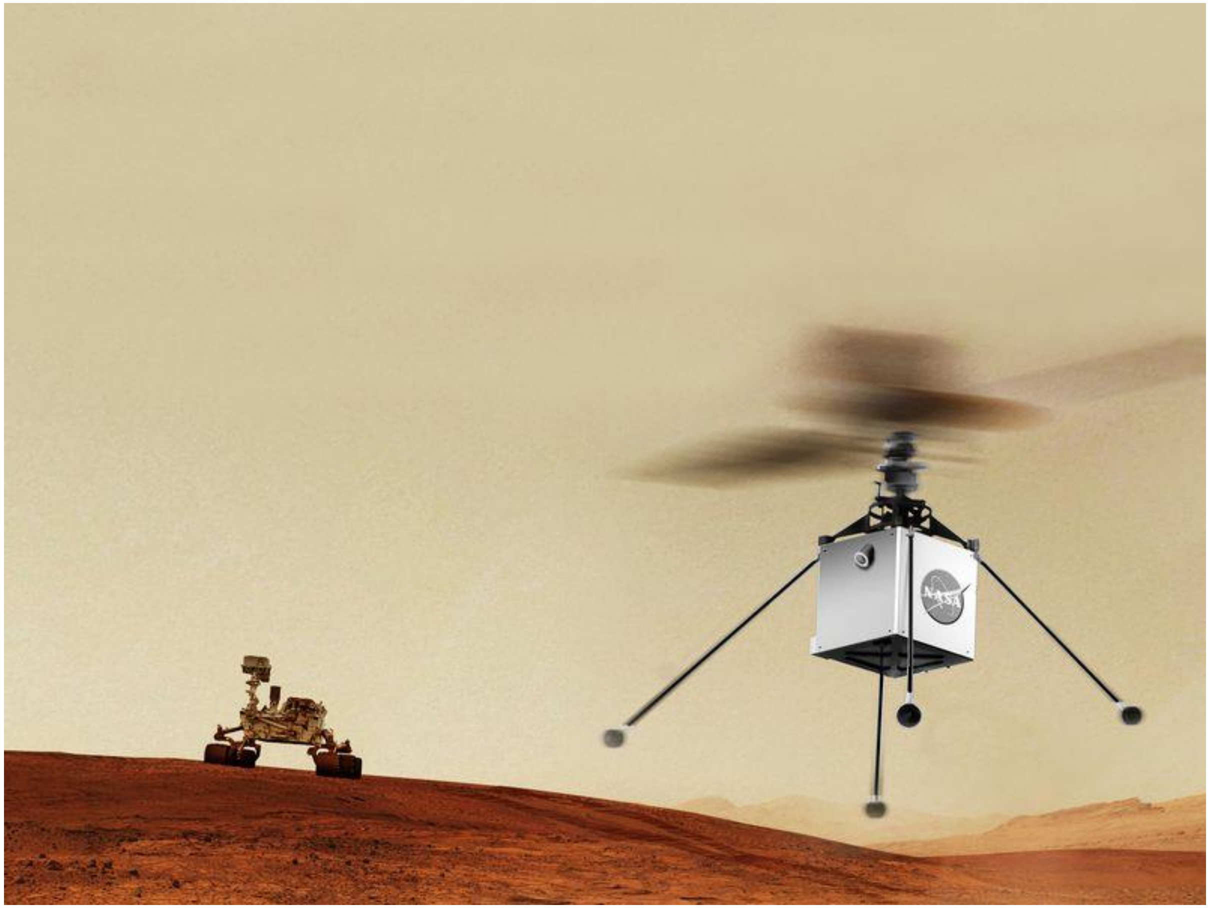

The NASA Mars Helicopter concept. Mission scheduled for 2020.
It will use visual inertial odometry.
https://en.wikipedia.org/wiki/JPL_Mars_Helicopter_Scout

Skydio R1

- 13 cameras for obstacle avoidance, VIO, and "follow me"
- SDK and simulator to be released soon

Skydio R1

- 13 cameras for obstacle avoidance, VIO, and “follow me”
- SDK and simulator to be released soon

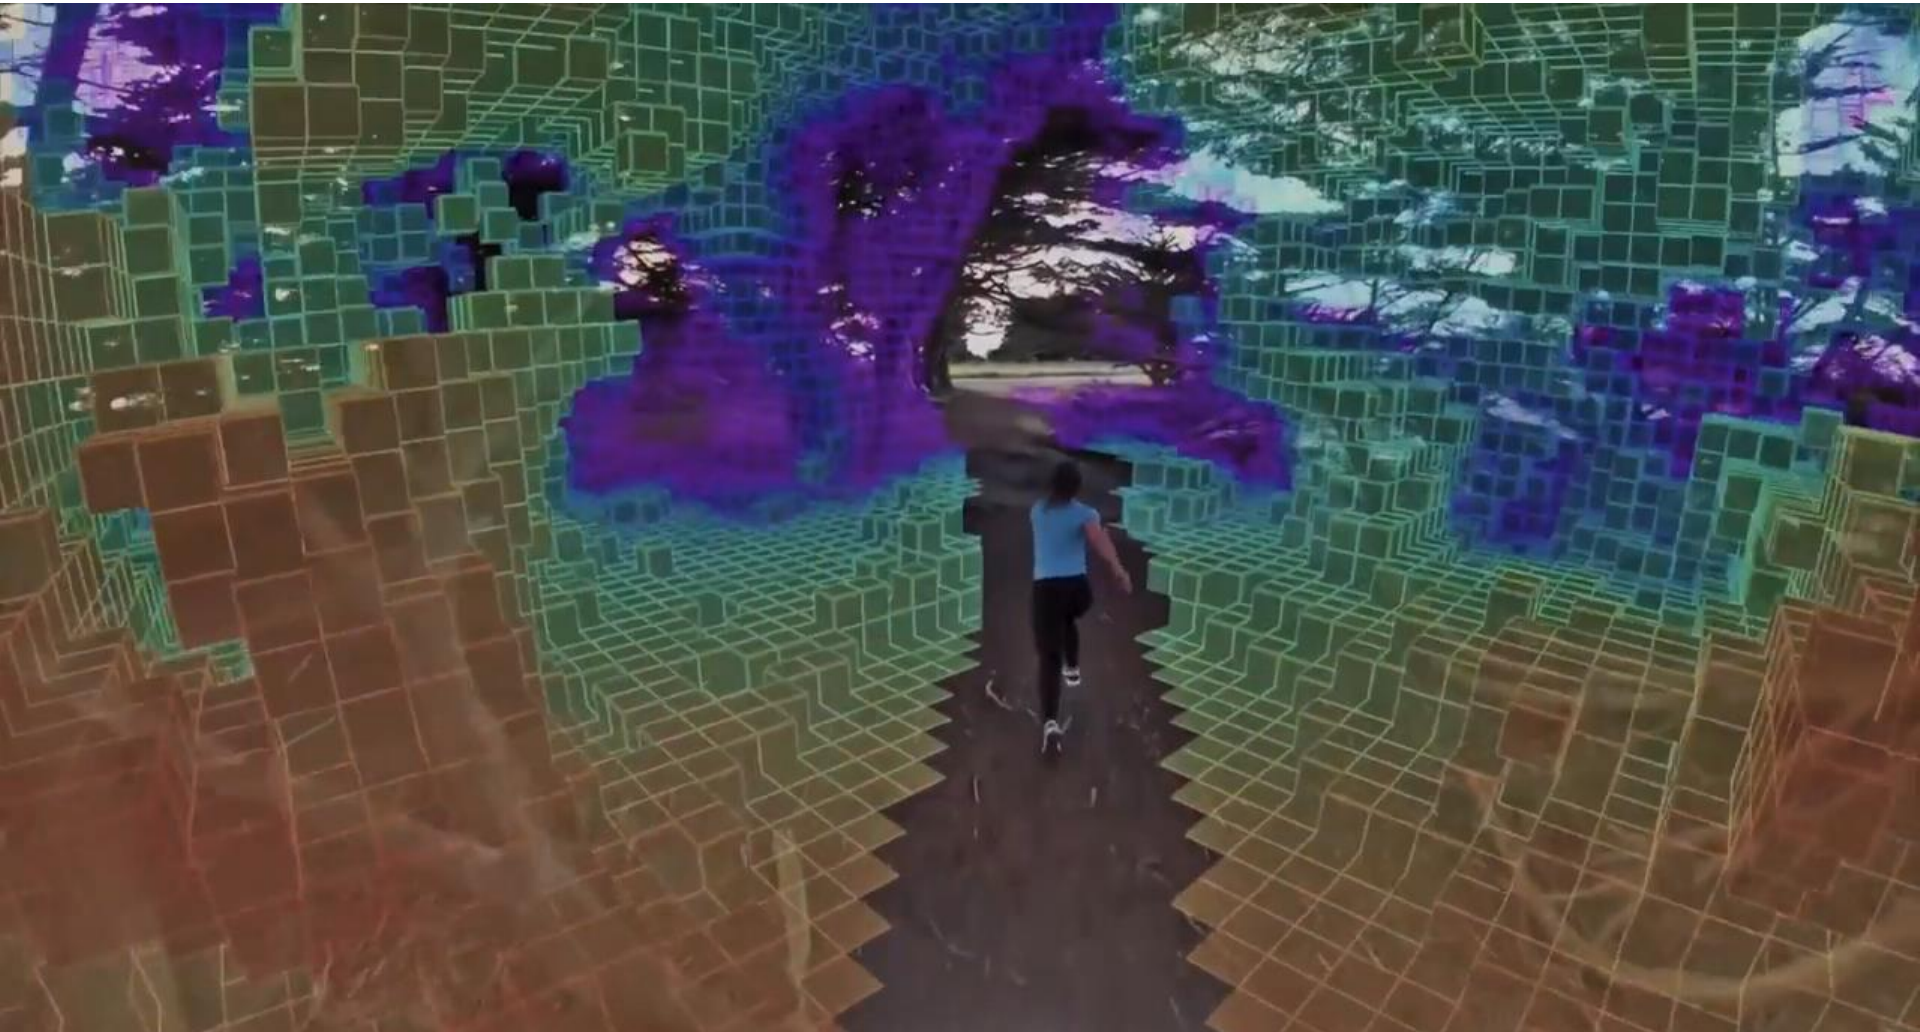

Dacuda's mouse scanner

Dacuda's
3D divison

- World's first mouse scanner,
  Distributed by LG, Logitech, etc.
- Dacuda was bought by Magic Leap in 2017 and is now Magic Leap Zurich (focusing on Augmented Reality)

Microsoft HoloLens

Google Visual Positioning Service (integrated into Maps and Street View)

Instructors

- **Lecturer**
  - **Davide Scaramuzza**: sdavide (at) ifi (dot) uzh (dot) ch
  - When away, I will be replaced by Antonio Loquercio
  - Office hours: **every Thursday from 15:30 to 17:30 (please announce yourself by email)**

- **Teaching Assistants** (exercises)

Mathias Gehrig
mgehrig (at) ifi (dot) uzh (dot) ch

Daniel Gehrig
dgehrig (at) ifi (dot) uzh (dot) ch

http://rpg.ifi.uzh.ch/people.html

Let’s have a 15 minute break
with
Oculus Quest
and
Microsoft Hololens Demos

Today's Class

- About me and my research lab
- What is Computer Vision?
- Why study computer vision?
- Example of Vision Applications
- Live Demos!
- Specifics of this course
- Overview of Visual Odometry

Organization of this Course

- **Lectures**:
  - 10:15 to 12:00 every week
  - Room: ETH LFW C5, Universitätstrasse 2, 8092 Zurich.

- **Exercises:**
  - 13:15 to 15:00: Starting from next week (Lecture 02). Then almost every week.
  - Room: ETH HG E 1.1, Rämistrasse 101, 8092 Zurich.

- **Official course website: http://rpg.ifi.uzh.ch/teaching.html**
  - Check it out for the PDFs of the lecture slides and updates

Learning Objectives

- **High-level goal:** learn to implement current visual odometry pipelines used in mobile robots (drones, cars, Mars rovers), and Virtual-reality (VR) and Augmented reality (AR) products: e.g., Google Visual Positioning Service, Oculus Quest, Microsoft HoloLens, Magic Leap.

- You will also learn **to implement the fundamental computer vision algorithms** used in mobile robotics, in particular: feature extraction, multiple view geometry, dense reconstruction, object tracking, image retrieval, visual-inertial fusion, event-based vision.

Course Schedule

For updates, slides, and additional material: http://rpg.ifi.uzh.ch/teaching.html

19.09.2019	**Lecture 01 - Introduction to Computer Vision and Visual Odometry**	Davide Scaramuzza
26.09.2019	**Lecture 02 - Image Formation 1: perspective projection and camera models** **Exercise 01 - Augmented reality wireframe cube**	Davide Scaramuzza Daniel & Mathias Gehrig
03.10.2019	**Lecture 03 - Image Formation 2: camera calibration algorithms** **Exercise 02 - PnP problem**	Davide Scaramuzza Daniel & Mathias Gehrig
10.10.2019	**Lecture 04 - Filtering & Edge detection**	Davide Scaramuzza
17.10.2019	**Lecture 05 - Point Feature Detectors, Part 1** **Exercise 03 - Harris detector + descriptor + matching**	Davide Scaramuzza Daniel & Mathias Gehrig
24.10.2019	**Lecture 06 - Point Feature Detectors, Part 2** **Exercise 04 - SIFT detector + descriptor + matching**	Davide Scaramuzza Daniel & Mathias Gehrig
31.10.2019	**Lecture 07 - Multiple-view geometry** **Exercise 05 - Stereo vision: rectification, epipolar matching, disparity, triangulation**	Davide Scaramuzza Daniel & Mathias Gehrig
07.11.2019	**Lecture 08 - Multiple-view geometry 2** **Exercise 06 - Eight-Point Algorithm**	Antonio Loquercio Daniel & Mathias Gehrig
14.11.2019	**Lecture 09 - Multiple-view geometry 3 (Part 1)**	Antonio Loquercio
21.11.2019	**Lecture 10 - Multiple-view geometry 3 (Part 2)** **Exercise session: Intermediate VO Integration**	Davide Scaramuzza Daniel & Mathias Gehrig
28.11.2019	**Lecture 11 - Optical Flow and Tracking (Lucas-Kanade)** **Exercise 08 - Lucas-Kanade tracker**	Davide Scaramuzza Daniel & Mathias Gehrig
05.12.2019	**Lecture 12 - Place recognition and 3D Reconstruction** **Exercise session: Deep Learning Tutorial**	Davide Scaramuzza Daniel & Mathias Gehrig
12.12.2019	**Lecture 13 - Visual inertial fusion** **Exercise 09 - Bundle Adjustment**	Davide Scaramuzza Daniel & Mathias Gehrig
19.12.2019	**Lecture 14 - Event based vision** After the lecture, we will Scaramuzza's lab. Departure from lecture room at 12:00 via tram 10. **Exercise session: Final VO Integration**	Davide Scaramuzza Daniel & Mathias Gehrig

Exercises

- Almost every week starting from next week (check out course schedule)
- **Participation in the exercise sessions is mandatory. Questions about the implementation details might be asked at the exam.**
- **Bring your own laptop**
- Each exercise will consist of coding a building block of a visual odometry pipeline. There will be two exercises dedicated to integrating these blocks together.
- Have **Matlab** pre-installed!
  - **ETH:** Download: https://idesnx.ethz.ch/
  - **UZH:** Download: https://www.zi.uzh.ch/de/students/software-elearning/softwareinstructions/Matlab.html
  - An introductory tutorial on Matlab can be found here: http://rpg.ifi.uzh.ch/docs/teaching/2019/MatlabPrimer.pdf
  - **Please install all the toolboxes included in the license.**

Exercises

- **Learning Goal** of the exercises: **Implement a full visual odometry pipeline** (similar to that running on Mars rovers and on current AR/VR devices (but actually much better ☺)).
- **Each week** you will learn how to implement a **building block** of visual odometry. The building blocks are:

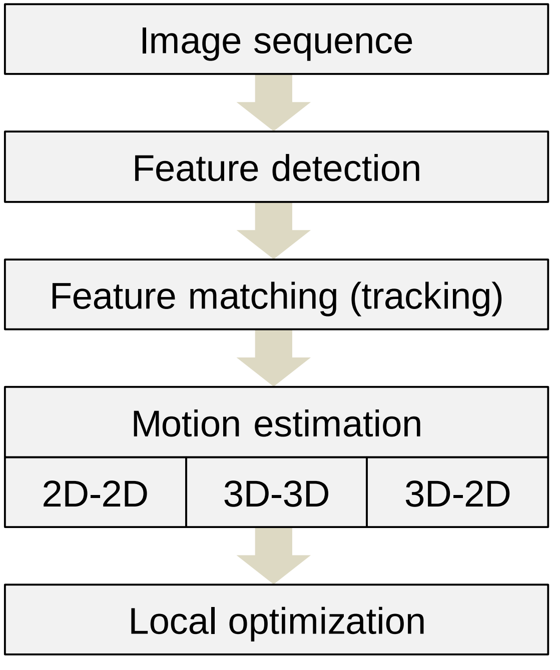

Outcome of last year exercises

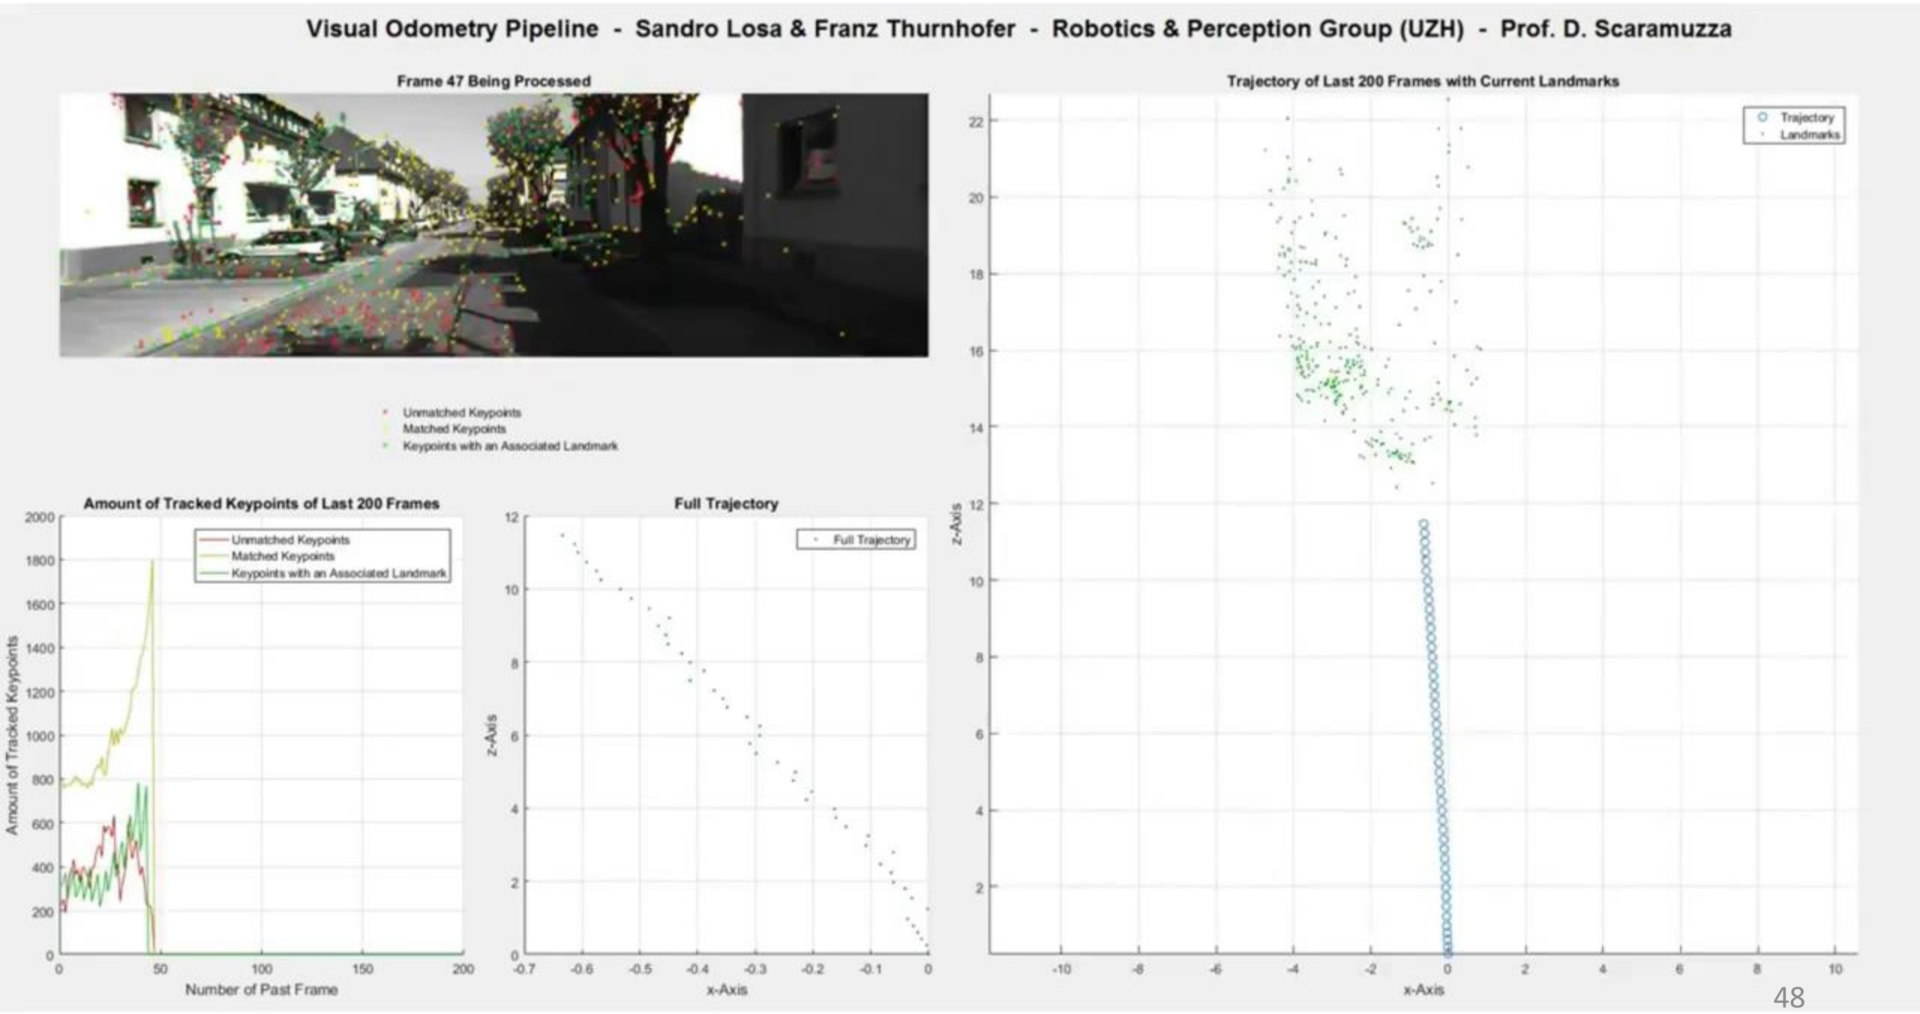

Recommended Textbooks

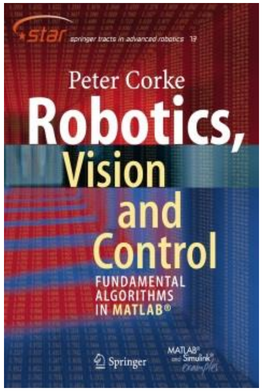

- **Robotics, Vision and Control: Fundamental Algorithms**, by Peter Corke 2011. The PDF of the book can be freely downloaded (only with ETH VPN) from Springer or alternatively from Library Genesys

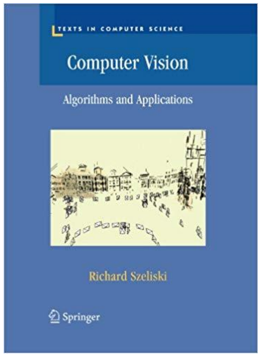

- **Computer Vision: Algorithms and Applications**, by Richard Szeliski, 2009. Can be freely downloaded from the author webpage: http://szeliski.org/Book/

- Other books:
  - *An Invitation to 3D Vision*: Y. Ma, S. Soatto, J. Kosecka, S.S. Sastry
  - *Multiple view Geometry*: R. Hartley and A. Zisserman (Library Genesys)

Prerequisites

- Linear algebra
- Matrix calculus
- No prior knowledge of computer vision and image processing required

Grading and Exam

- The **final grade is based on an oral exam** (30 minutes). Example exam questions here.
  - Exam dates:
    - **UZH**: January 9, 2020
    - **ETH**: from January 20 to February 14, 2020 (the dates are handled by ETH Exam Center and are usually communicated before Christmas)
- **Strong class participation can offset negative performance** at the oral exam.
- **Optional mini project:**
  - you have the **option** (i.e., not mandatory) to do a **mini project**, which consists of implementing a working visual odometry algorithm in **Matlab** (C++ or Python are also acceptable)
  - If the algorithm runs smoothly, producing a reasonable result, you will be rewarded with an **up to 0.5 grade increase on the final grade**. However, notice that the mini project can be quite time consuming!
  - The **deadline** to hand in the mini project is 5.01.2020.
  - **Group work (up to 4) possible.**

Class Participation

- Class participation includes
  - showing up
  - being able to articulate key points from last lecture
  - ask and answer questions

Today's Class

- About me and my research lab
- What is Computer Vision?
- Why study computer vision?
- Example of Vision Applications
- Live Demos!
- Specifics of this course
- Overview of Visual Odometry

What is Visual Odometry (VO) ?

VO is the process of incrementally estimating the pose of the vehicle by examining the changes that motion induces on the images of its onboard cameras

input

Image sequence (or video stream)
from one or more cameras attached to a moving vehicle

output

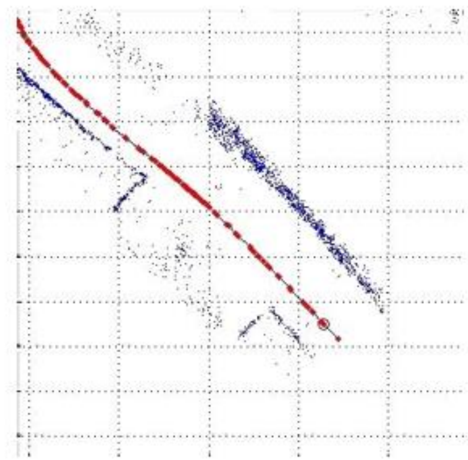

$$R_0, R_1, \dots, R_i$$

$$t_0, t_1, \dots, t_i$$

Camera trajectory (3D structure is a plus)

Why VO?

- Contrary to wheel odometry, VO is **not affected by wheel slippage** on uneven terrain or other adverse conditions.

- More accurate trajectory estimates compared to wheel odometry (**relative position error 0.1% – 2%**)

- VO can be used as a complement to
  - wheel encoders (wheel odometry)
  - GPS
  - inertial measurement units (IMUs)
  - laser odometry

- Crucial for flying, walking, and underwater robots

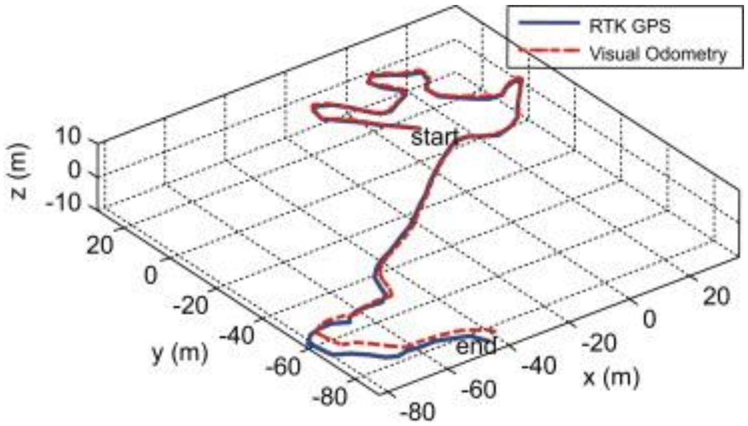

Assumptions

- **Sufficient illumination** in the environment
- **Dominance of static scene** over moving objects
- **Enough texture** to allow apparent motion to be extracted
- Sufficient **scene overlap** between consecutive frames

Is any of these scenes good for VO? Why?

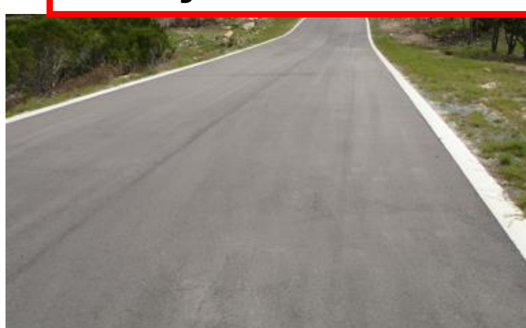

A Brief history of VO

- **1980**: First known VO real-time implementation on a robot by **Hans Moraveck** PhD thesis **(NASA/JPL)** for Mars rovers using one sliding camera (*sliding stereo*).

A Brief history of VO

- **1980**: First known VO real-time implementation on a robot by **Hans Moraveck** PhD thesis **(NASA/JPL)** for Mars rovers using one sliding camera (*sliding stereo*).

- **1980 to 2000**: The VO research was dominated by **NASA/JPL** in preparation of the **2004 mission to Mars**

- **2004**: VO was used on a robot on another planet: Mars rovers Spirit and Opportunity (see seminal paper from NASA/JPL, 2007)

- **2004.** VO was revived in the academic environment by **David Nister**'s «Visual Odometry» paper. The term VO became popular.

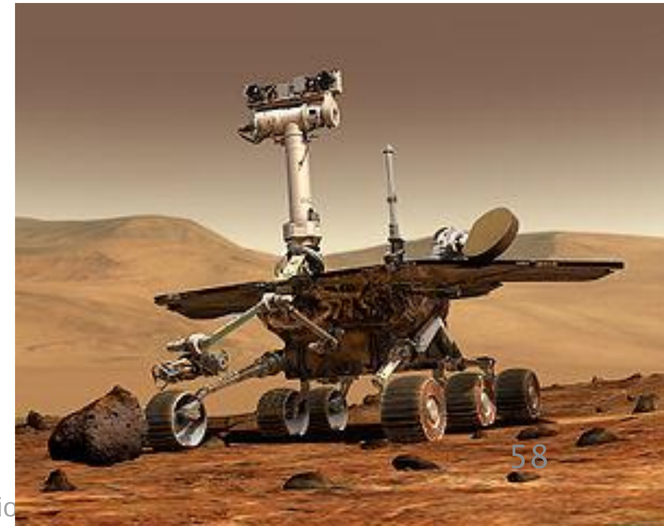

More about history and tutorials

- Scaramuzza, D., Fraundorfer, F., **Visual Odometry: Part I** - The First 30 Years and Fundamentals, *IEEE Robotics and Automation Magazine*, Volume 18, issue 4, 2011. PDF

- Fraundorfer, F., Scaramuzza, D., **Visual Odometry: Part II** - Matching, Robustness, and Applications, *IEEE Robotics and Automation Magazine*, Volume 19, issue 1, 2012. PDF

- C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I.D. Reid, J.J. Leonard, **Past, Present, and Future of Simultaneous Localization and Mapping: Toward the Robust-Perception Age**, *IEEE Transactions on Robotics*, Vol. 32, Issue 6, 2016. PDF

VO vs VSLAM vs SFM

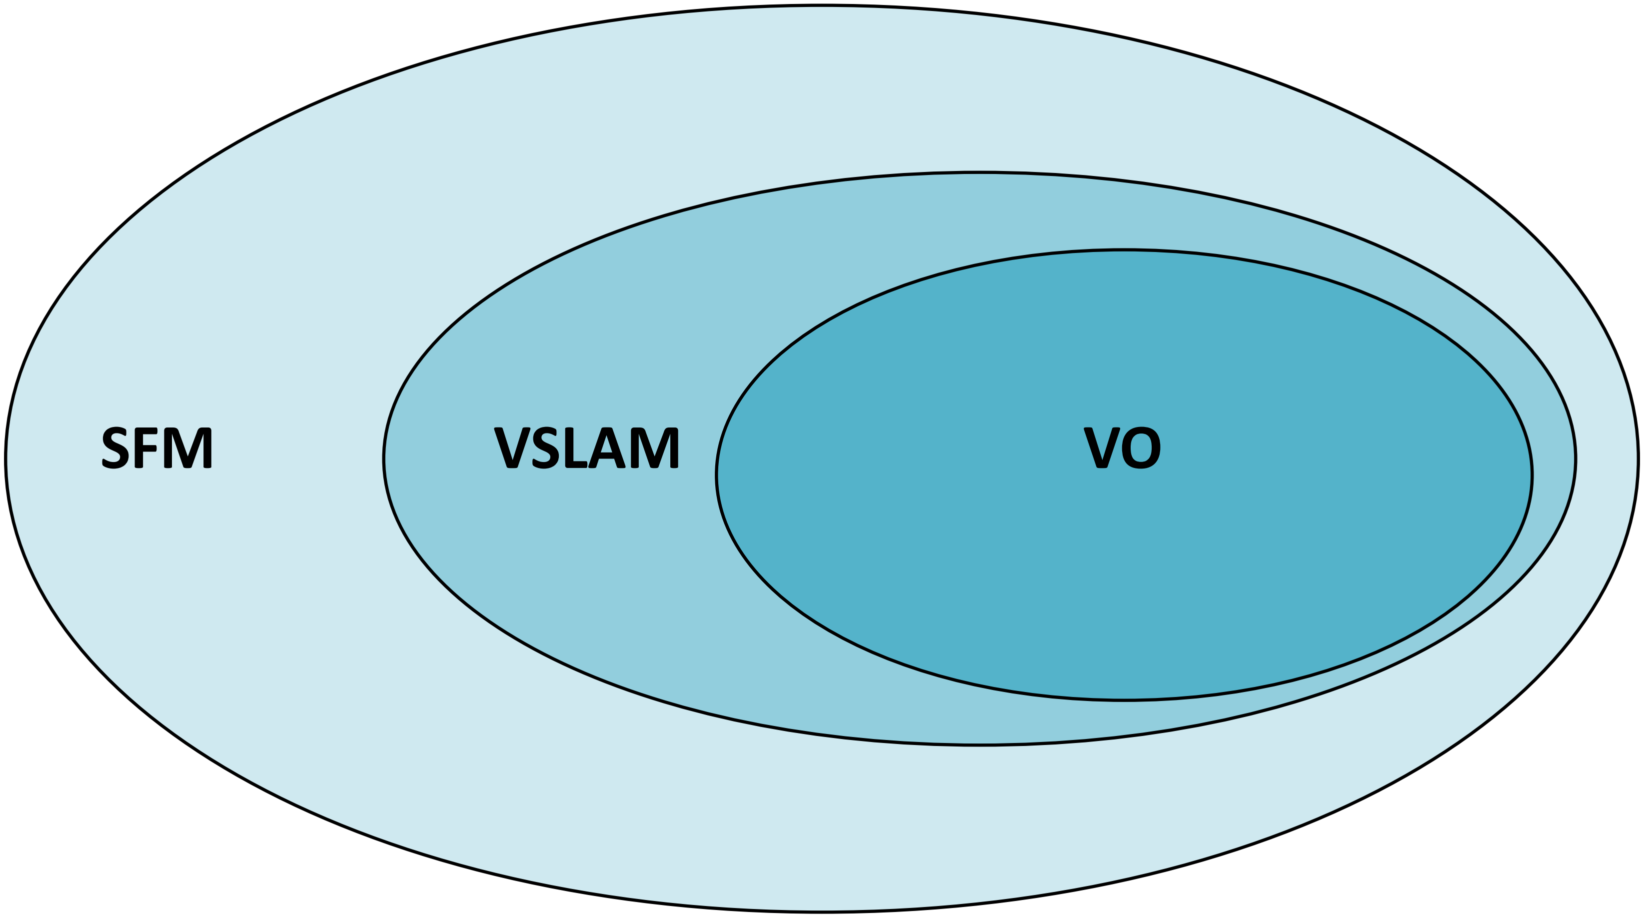

Structure from Motion (SFM)

SFM is more general than VO and tackles the problem of 3D reconstruction and 6DOF pose estimation from **unordered image sets**

Reconstruction from 3 million images from Flickr.com
Cluster of 250 computers, 24 hours of computation!
Paper: “Building Rome in a Day”, ICCV’09

VO vs SFM

- VO is a **particular case** of SFM
- VO focuses on estimating the 6DoF motion of the camera **sequentially** (as a new frame arrives) and in **real time**.
- Terminology: sometimes SFM is used as a synonym of VO

VO vs. Visual SLAM

- **Visual Odometry**
  - Focus on incremental estimation
  - **Guarantees local consistency**

- **Visual SLAM** (Simultaneous Localization And Mapping)
  - **SLAM = visual odometry + loop detection & closure**
  - **Guarantees global consistency**

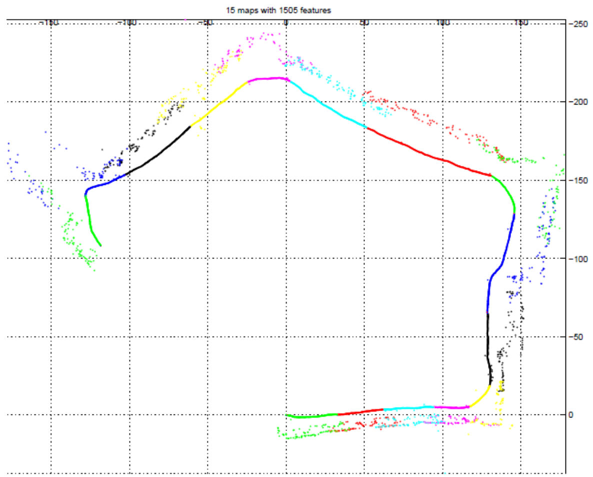

Visual odometry

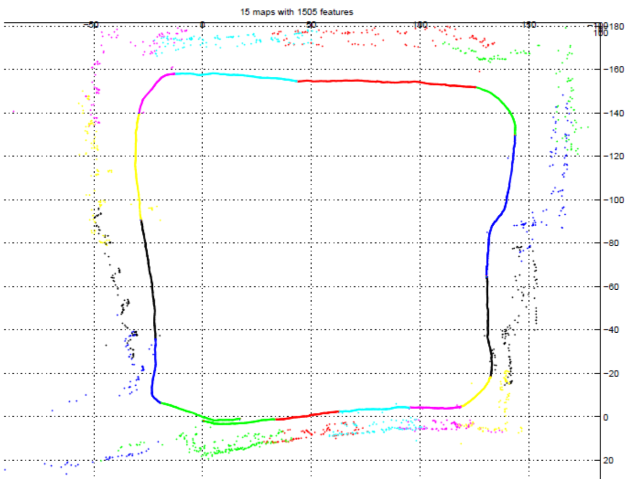

Visual SLAM

Image courtesy of [Clemente et al., RSS'07]

VO Flow Chart

VO computes the camera path incrementally (pose after pose)

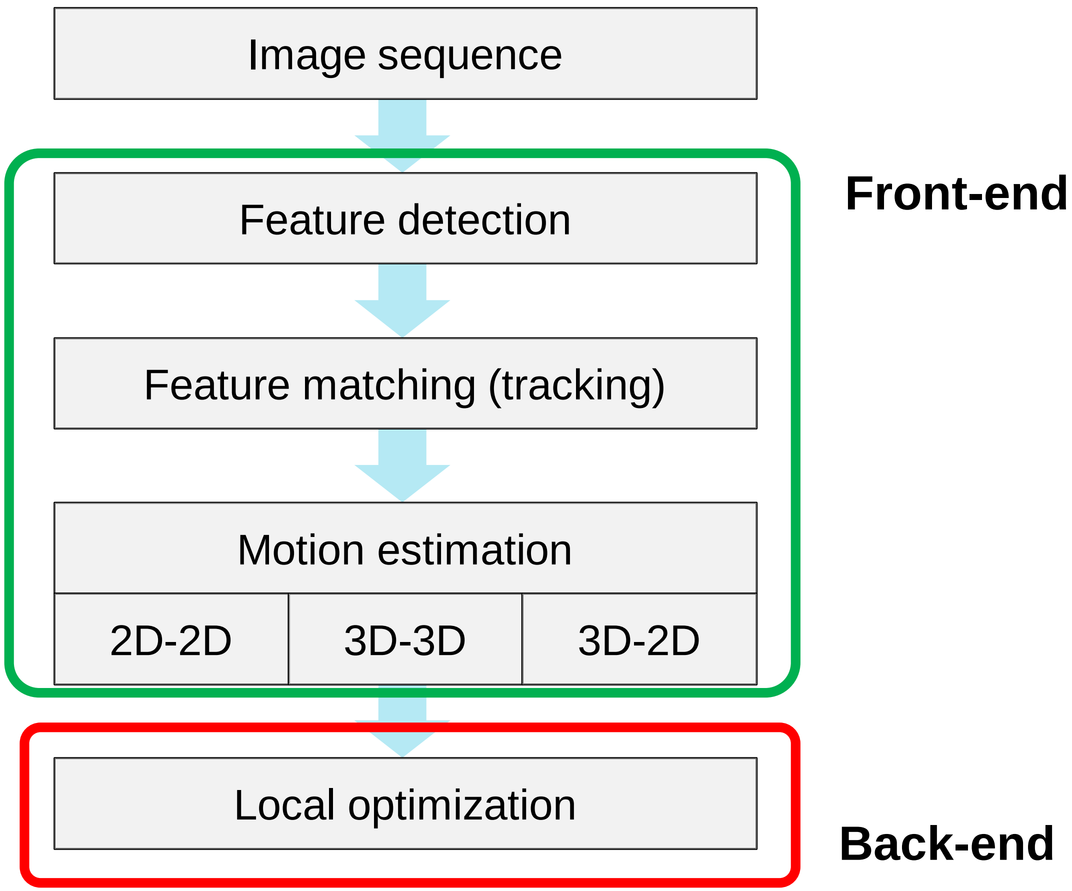

VO Flow Chart

VO computes the camera path incrementally (pose after pose)

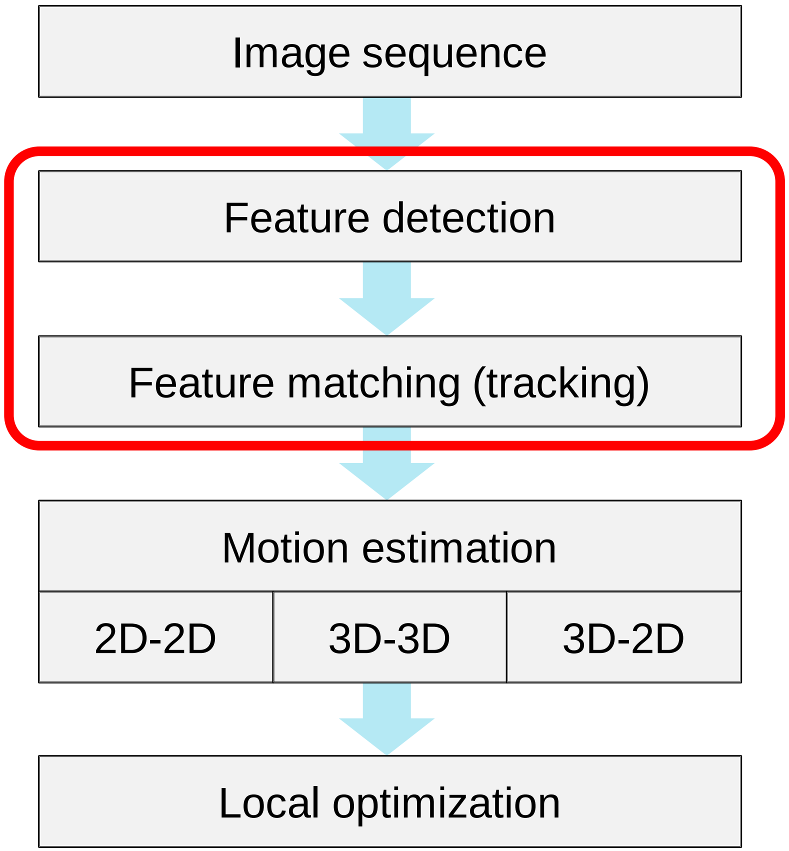

Example features tracks

VO Flow Chart

VO computes the camera path incrementally (pose after pose)

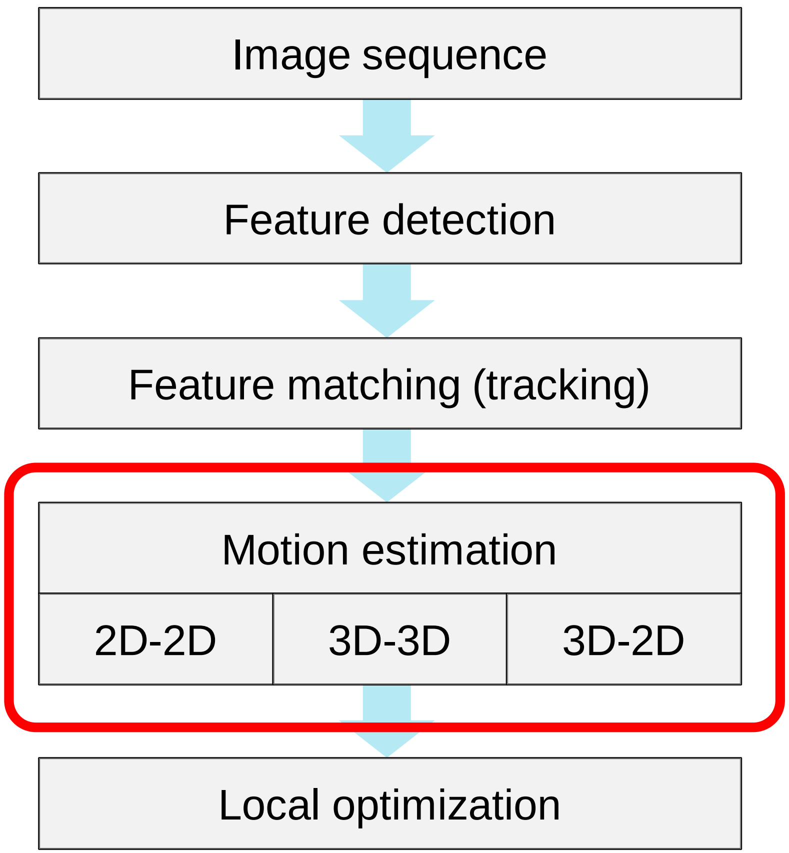

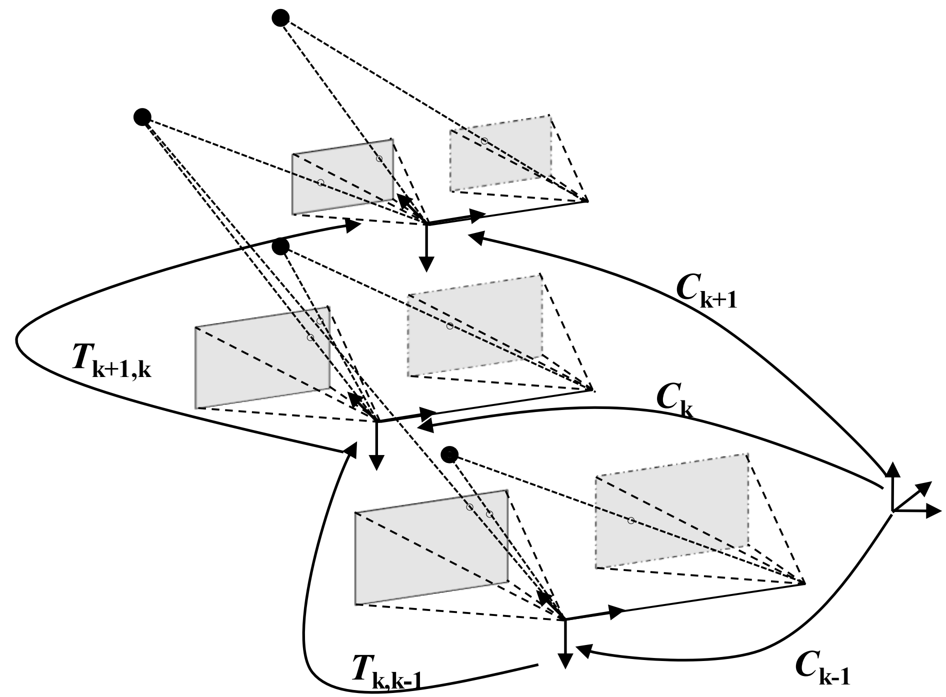

VO Flow Chart

VO computes the camera path incrementally (pose after pose)

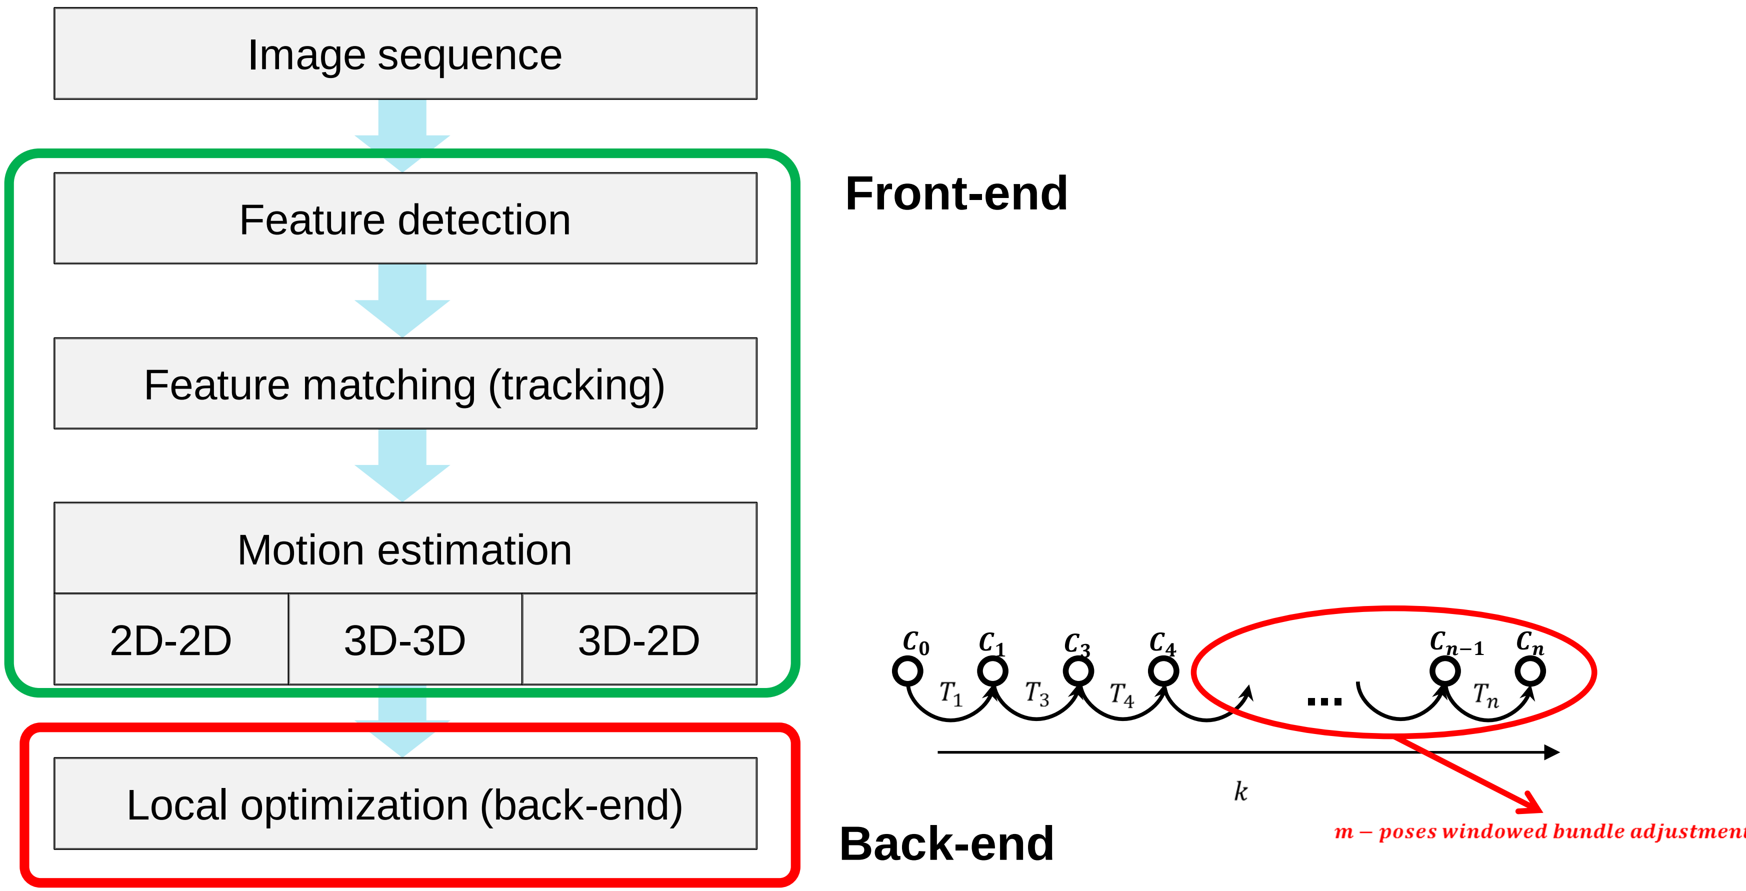

Course Topics

- Principles of image formation
- Image Filtering
- Feature detection and matching
- Multi-view geometry
- Dense reconstruction
- Visual place recognition
- Visual inertial fusion
- Event-based Vision

Course Topics

- Principles of image formation
  - Perspective projection
  - Camera calibration

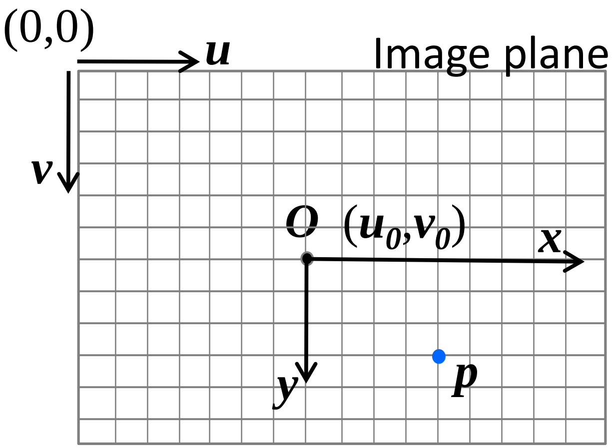

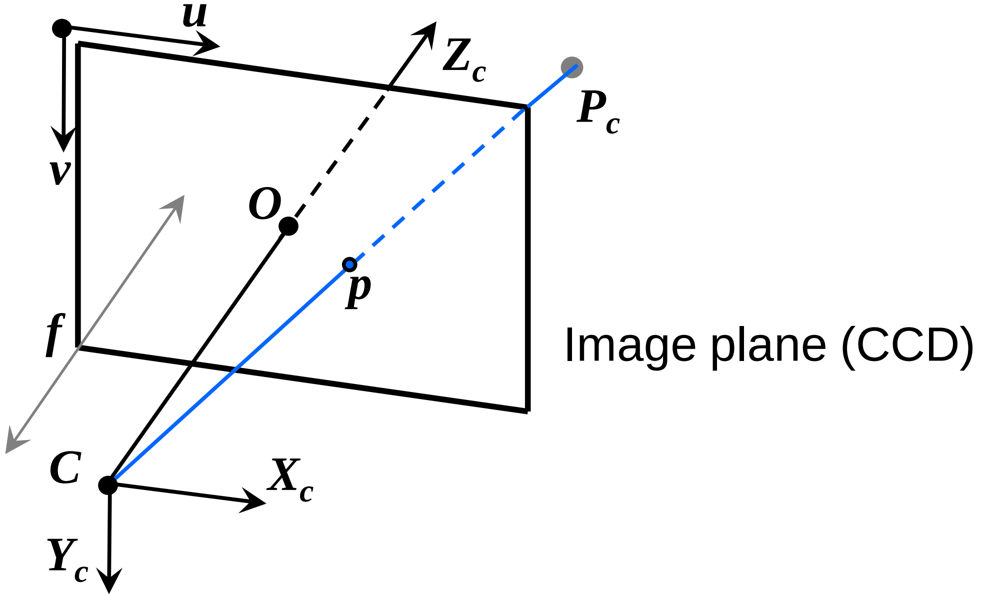

Course Topics

- Feature detection and matching

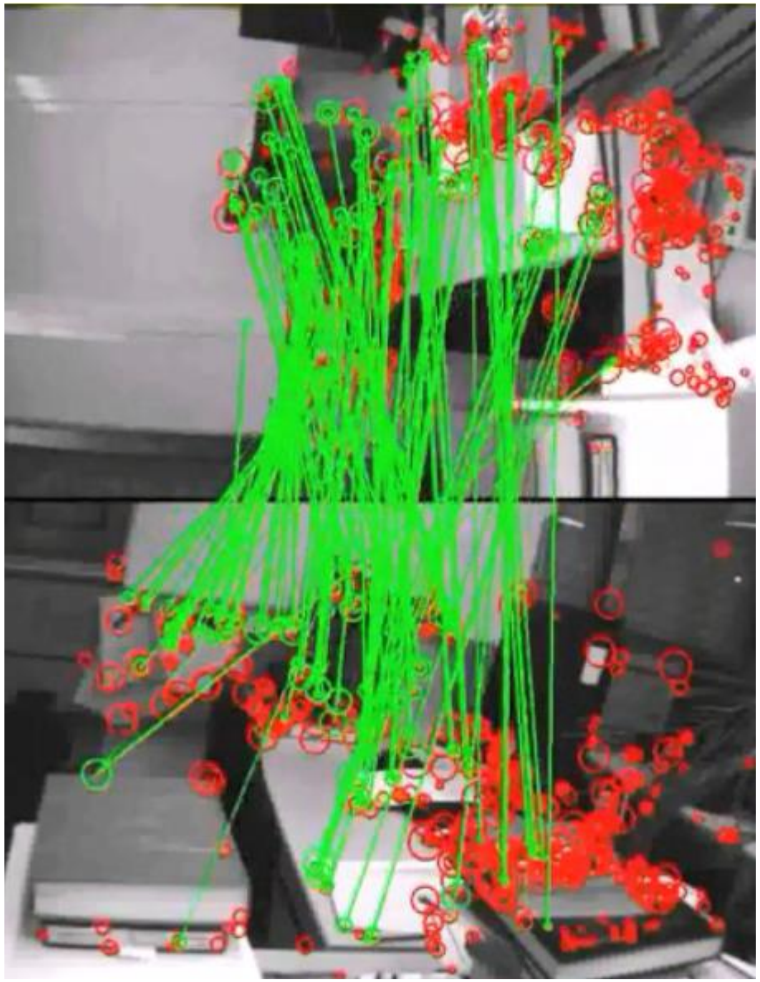

Course Topics

- Multi-view geometry and sparse 3D reconstruction

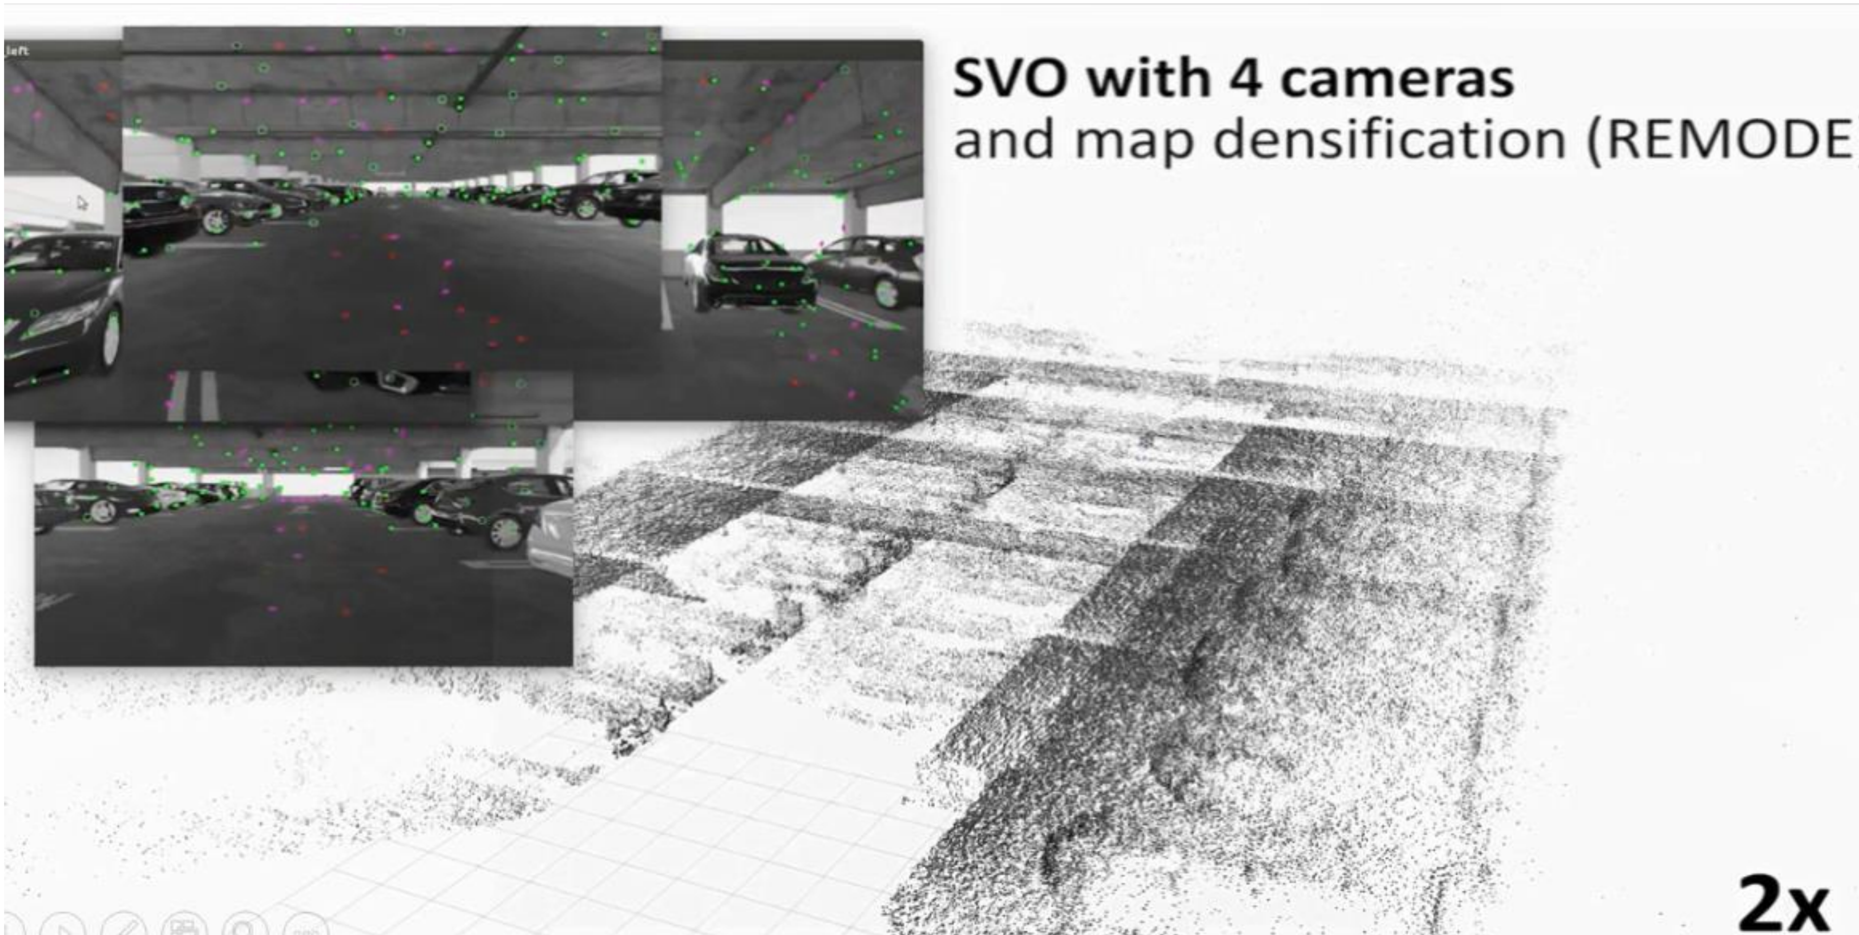

Course Topics

- Dense 3D reconstruction

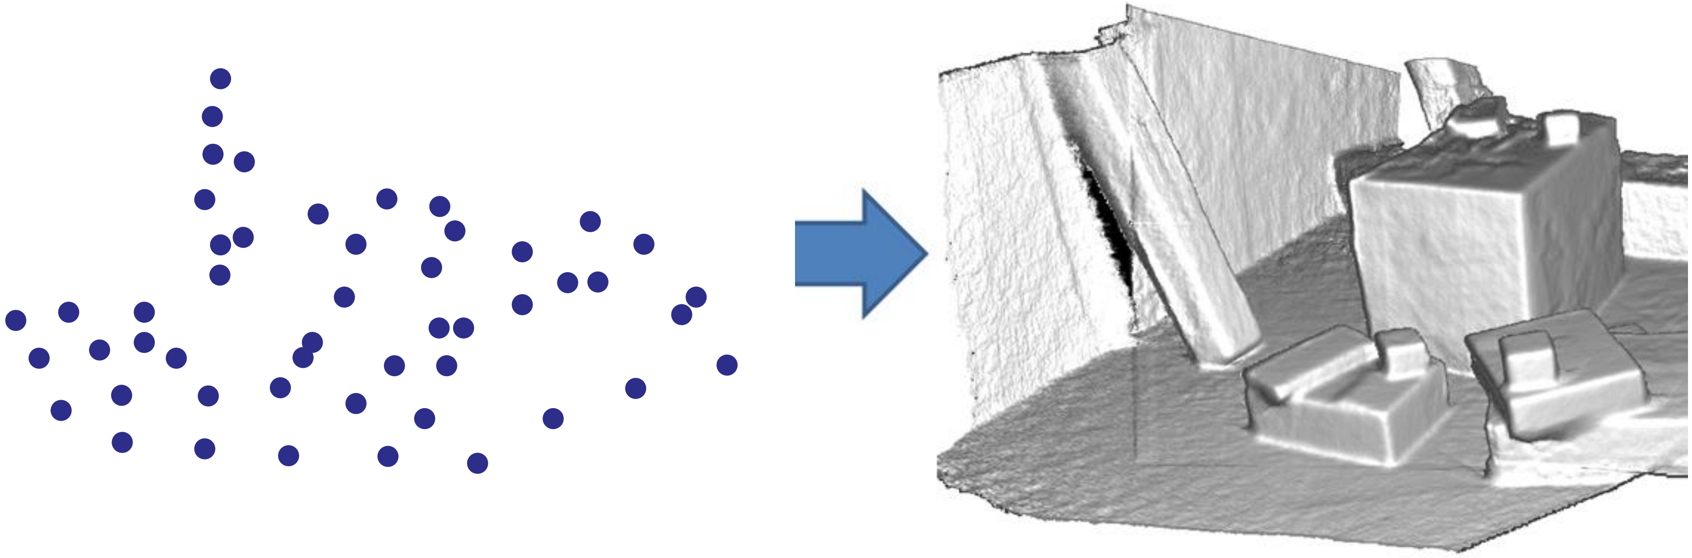

Course Topics

- Dense 3D reconstruction

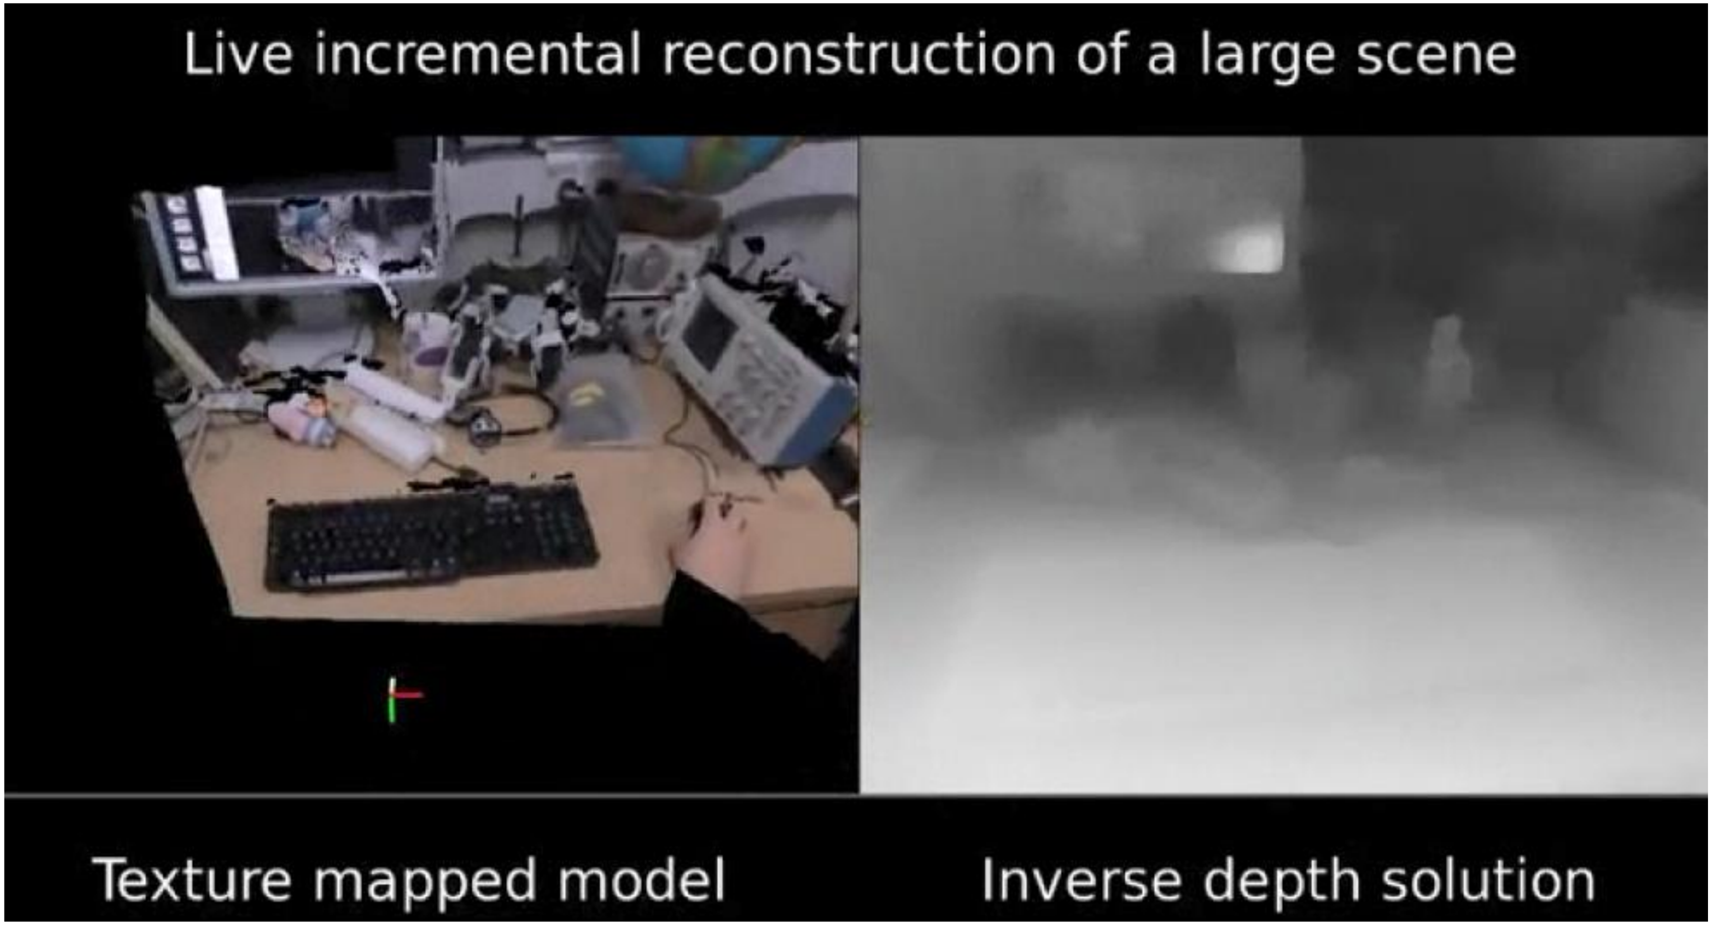

Course Topics

- Place recognition

Query image

Most similar places from a database of millions of images

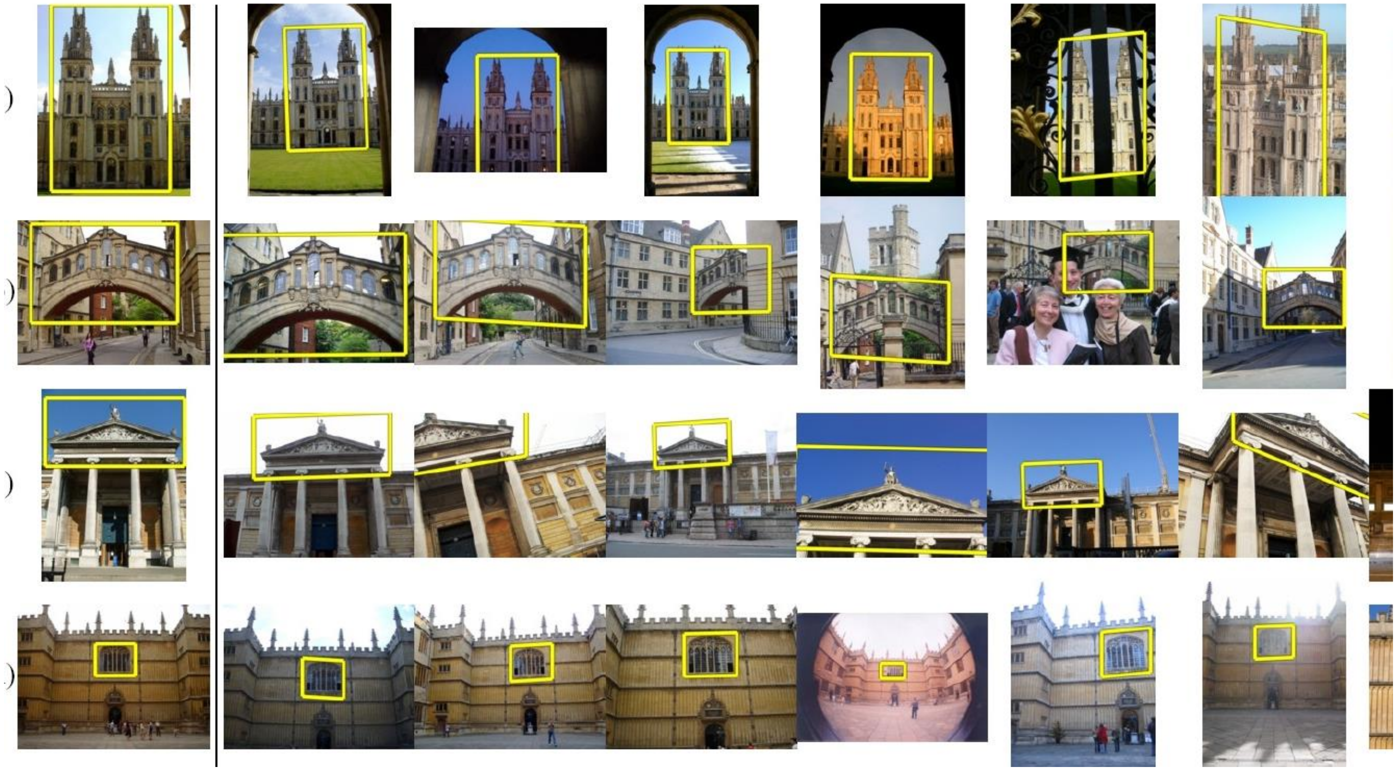

Course Topics

- Visual-inertial fusion

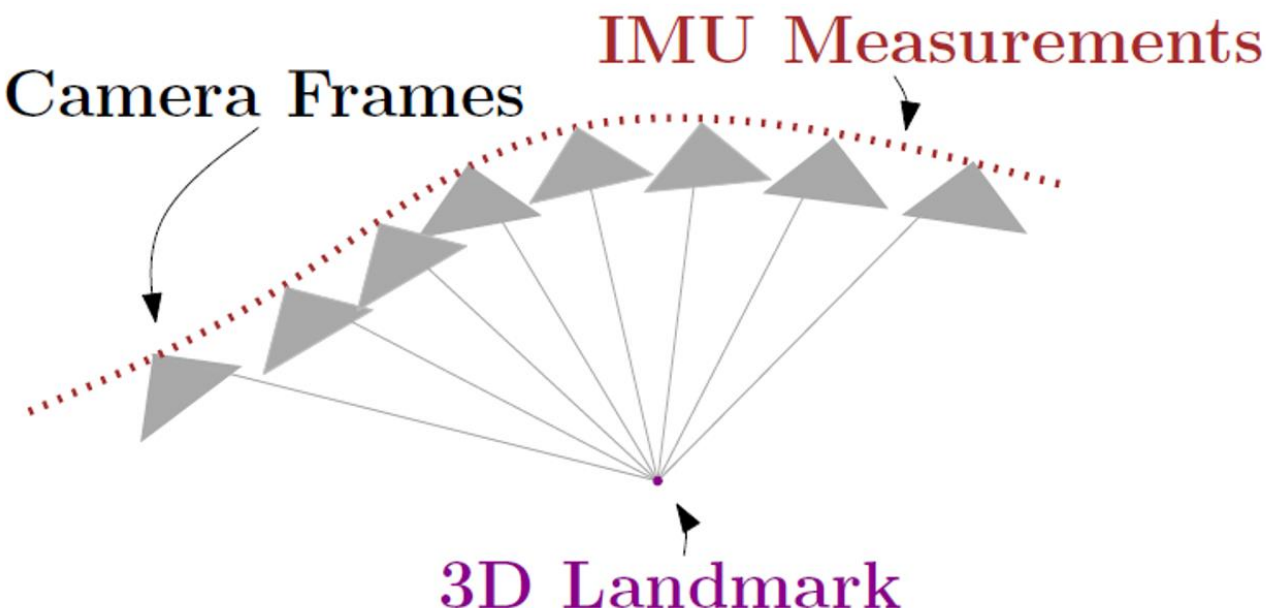

Course Topics

- Event-based vision

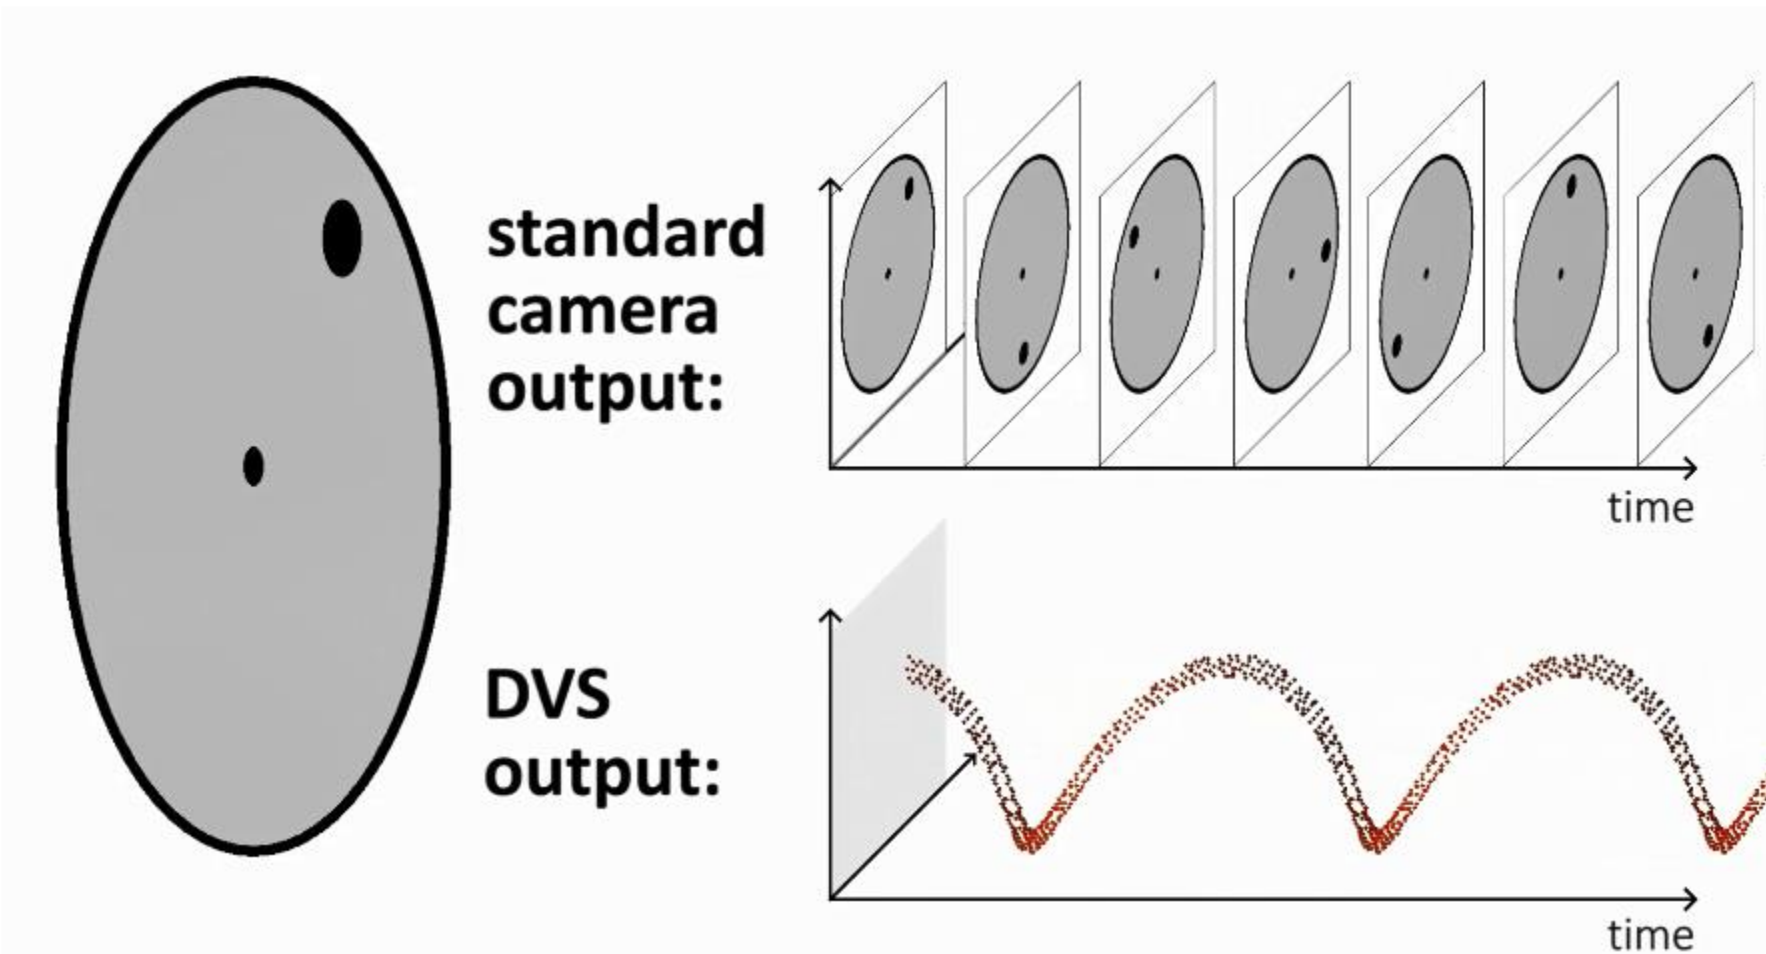

Application: High speed VO

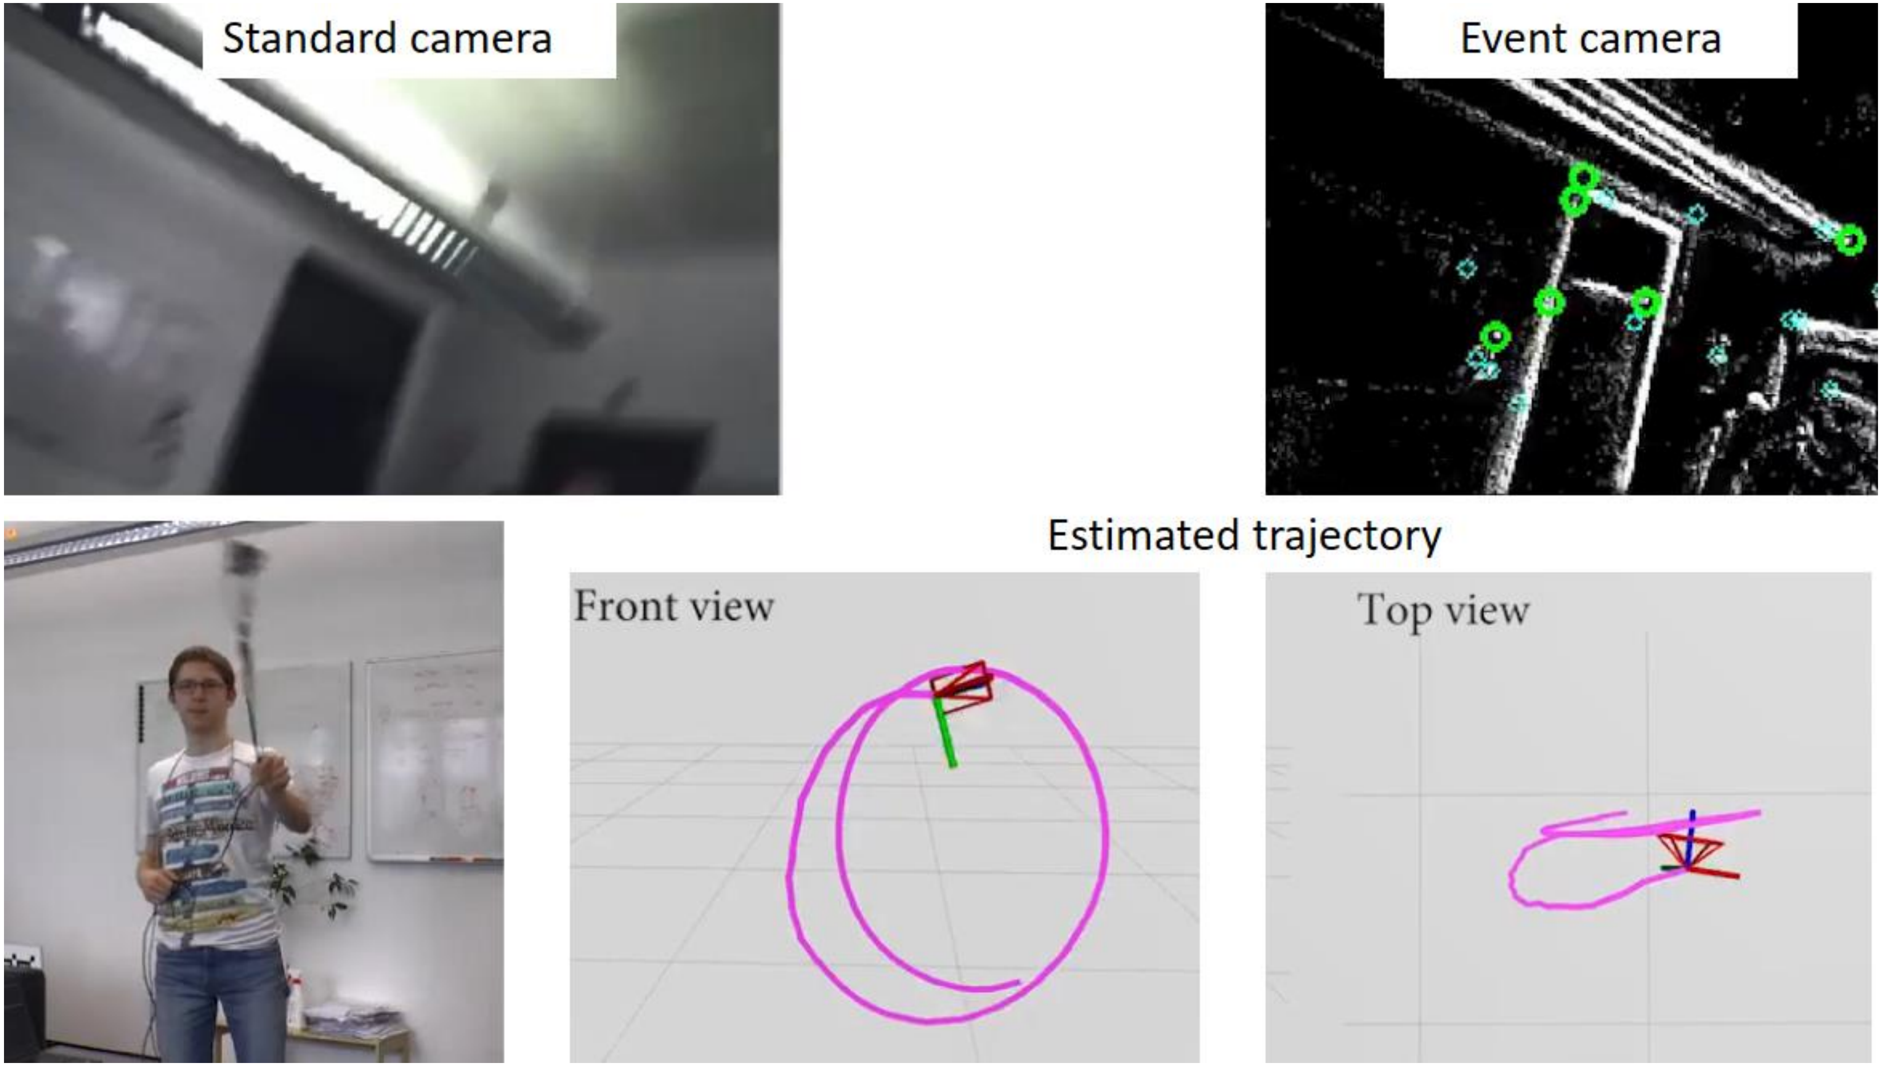

Rosinol et al., *Ultimate SLAM?* **IEEE RAL'18 best Paper Award Honorable Mention** PDF. Video. IEEE Spectrum.

Understanding Check

Are you able to:

- Provide a definition of Visual Odometry?
- Explain the most important differences between VO, VSLAM and SFM?
- Describe the needed assumptions for VO?
- Illustrate its building blocks?