

University of Zurich[UZH]

ROBOTICS & PERCEPTION GROUP

Vision Algorithms for Mobile Robotics

Lecture 01
Introduction

Davide Scaramuzza

https://rpg.ifi.uzh.ch

Universität Zürich[UZH]

Rämistrasse 69 (SOC)

In diesem Hause lehrten und forschten

1909 – 1911	**Albert Einstein** Nobelpreis für Physik 1921
1911 – 1912	**Peter Debye** Nobelpreis für Chemie 1936
1912 – 1914	**Max von Laue** Nobelpreis für Physik 1914
1913 – 1951	**Walter Rudolf Hess** Nobelpreis für Medizin und Physiologie 1949
1921 – 1927	**Erwin Schrödinger** Nobelpreis für Physik 1933

Today's Outline

- About me and my research lab
- What is Computer Vision?
- Why study Computer Vision?
- Examples of vision applications
- Organization of the course
- Start: Visual Odometry overview

Who am I?

Current position

- Professor of Robotics & Computer Vision since 2012
- Dep. of Informatics (UZH) and Neuroinformatics (UZH & ETH)
- Founder & Director of the Master's program in Artificial Intelligence at the Dep. of Informatics (UZH)
- Adjunct Professor of the ETH Master in Robotics, Systems and Control and Associate faculty of the ETH AI Center

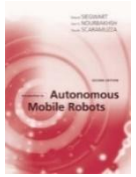
Previously

- Distinguished Visiting Scientist at NASA Jet Propulsion Laboratory, 2025
- Visiting professor at Stanford University, 2019

Education

- Master in Electronics Engineering at the University of Perugia, Italy, 2004
- PhD in Robotics and Computer Vision at ETH Zurich, Switzerland, 2008
- Post-doc at the University of Pennsylvania, Philadelphia, USA

Book



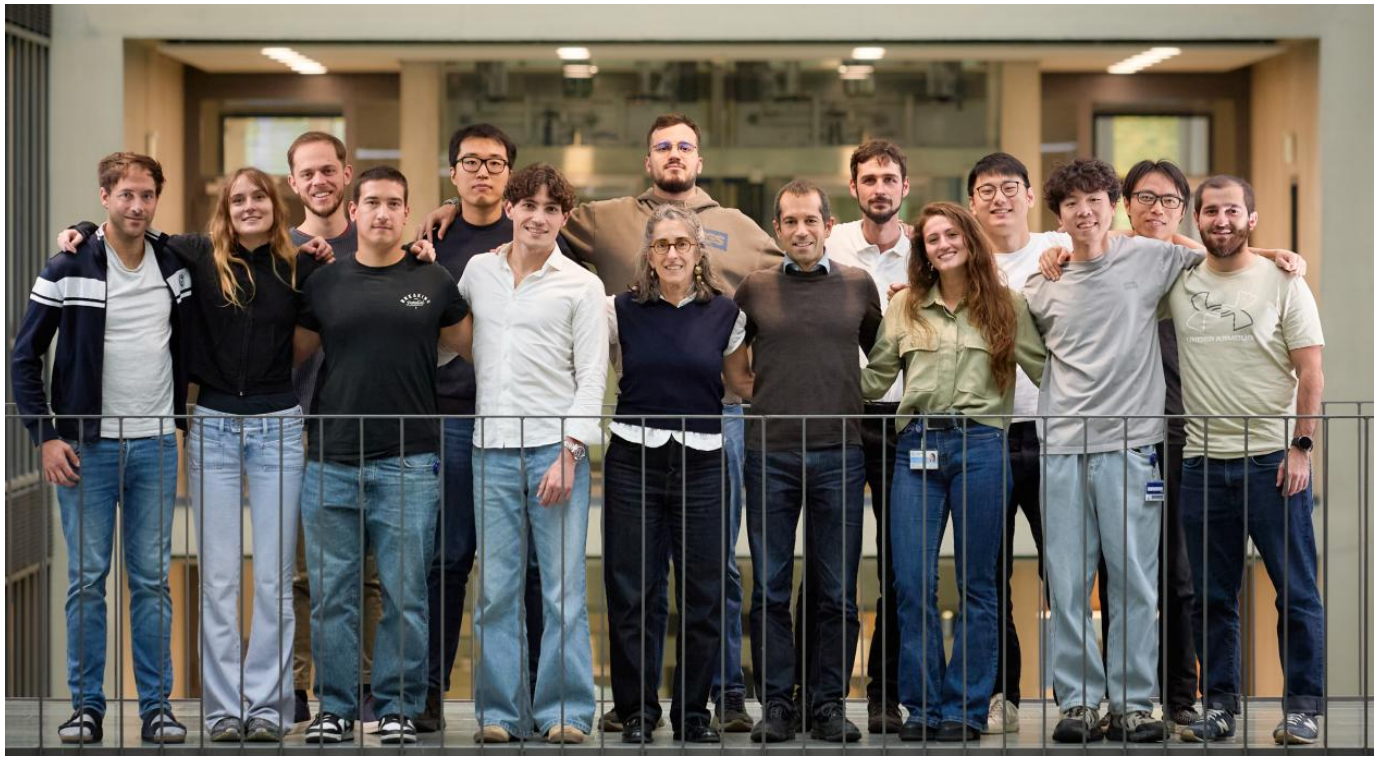
- "Autonomous Mobile Robots, 2nd Edition" MIT Press, 2011

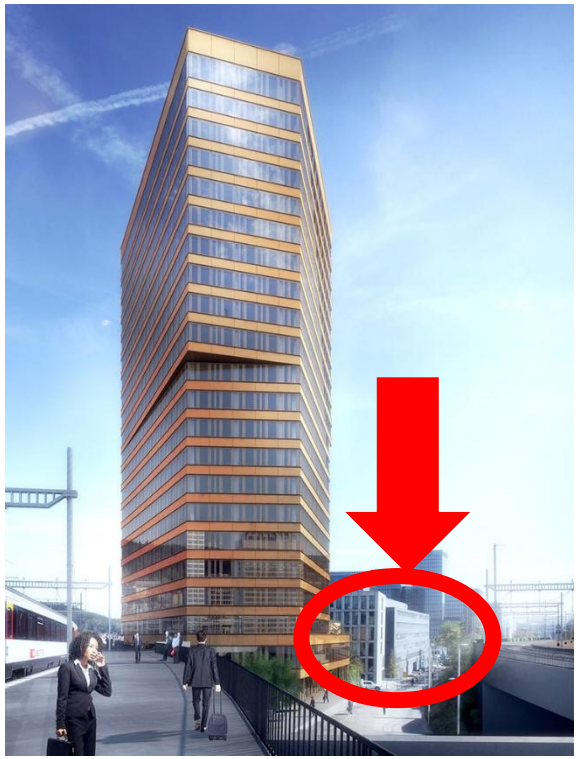
Hobbies

- Running, piano, magic tricks

My lab: the Robotics and Perception Group

- **Address**: Andreasstrasse 15, 2nd floor, next to **Zurich Oerlikon** train station
- **Webpage**: https://rpg.ifi.uzh.ch/
- **Lab Open Door Day**: December 10, at 5pm



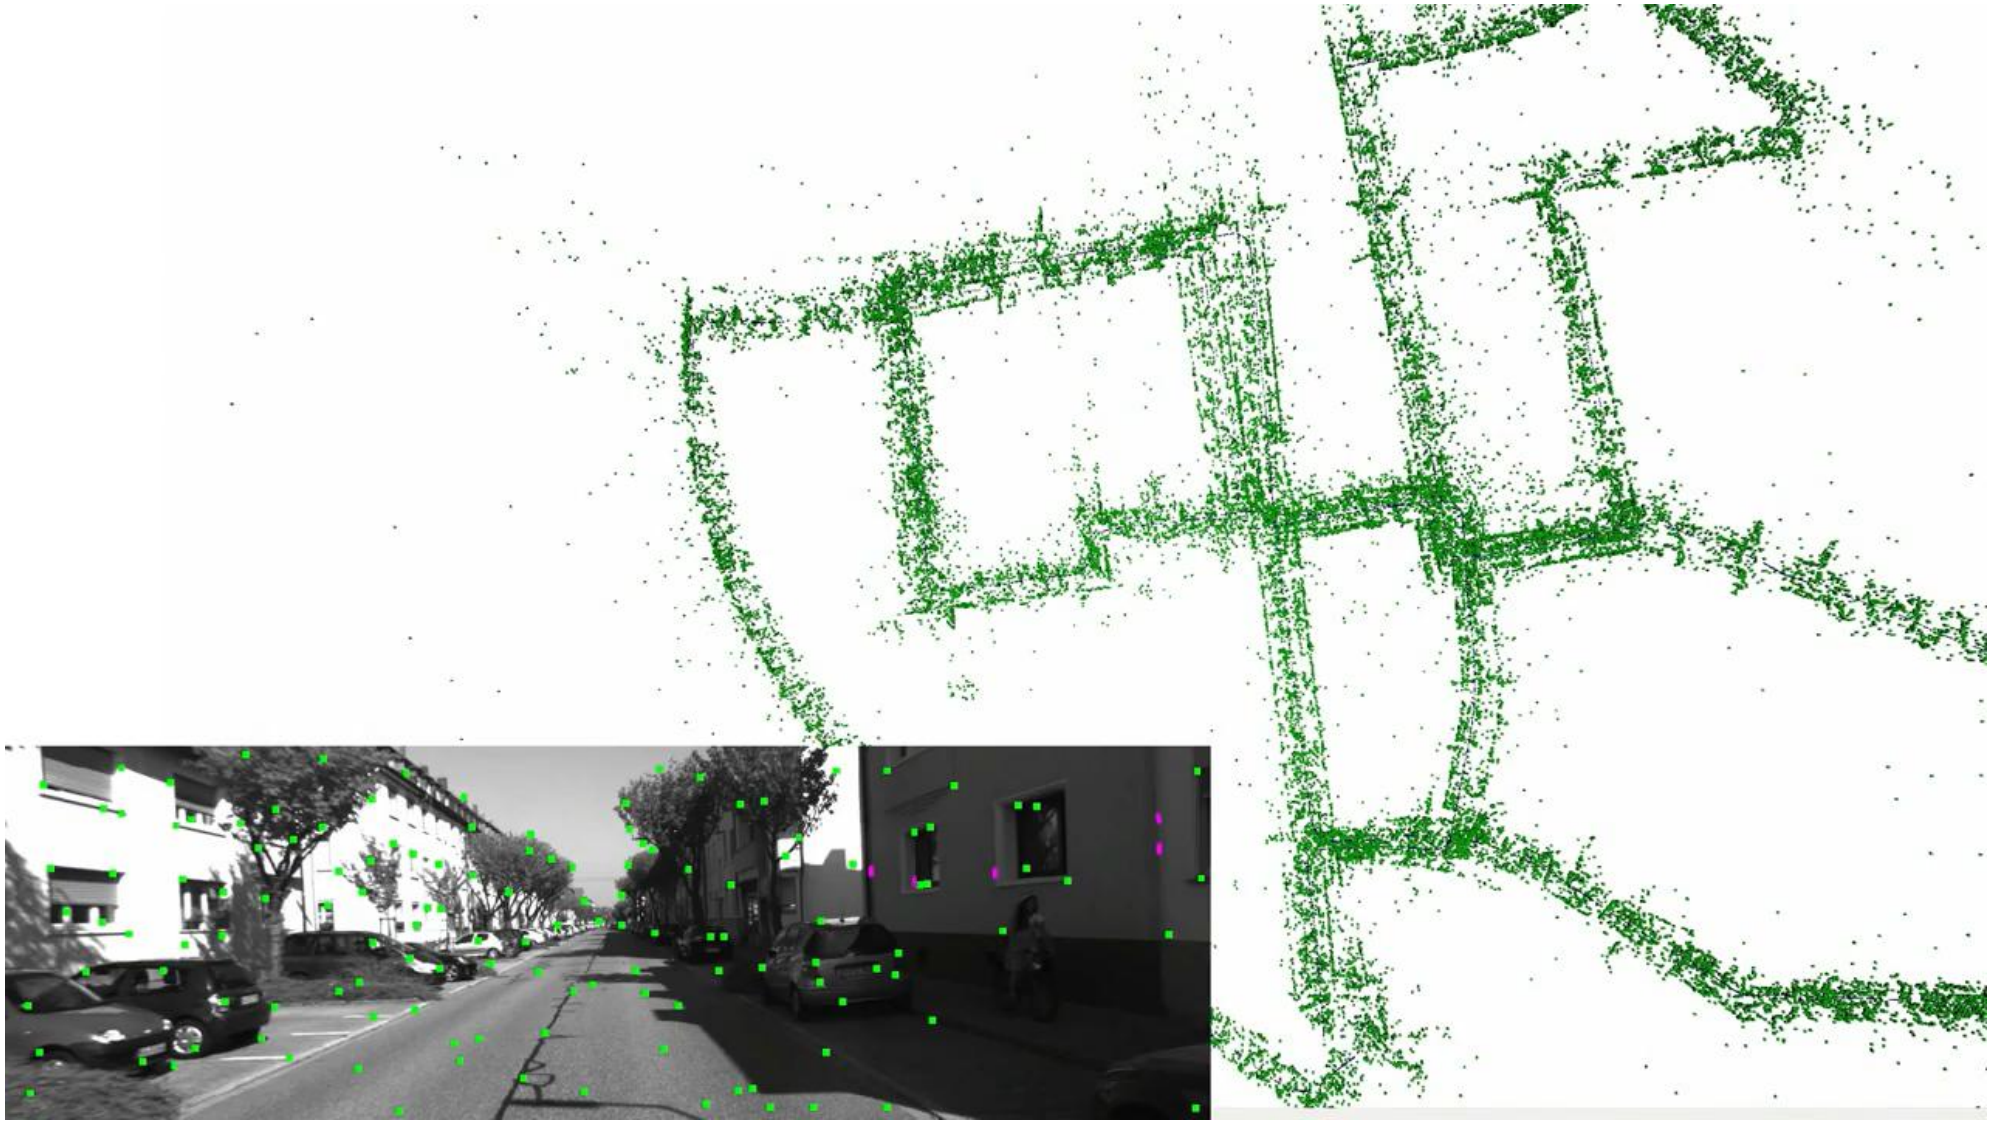


Our Research: Autonomous, Vision-based Navigation

Real-time, onboard, perception, learning, and control for autonomous agile robotics using both standard cameras or neuromorphic event-based cameras



Spatial AI: Visual Simultaneous Localization and Mapping



Agile Robotics





Kaufmann, Bauersfeld, Loquercio, Mueller, Koltun, Scaramuzza,
Champion-Level Drone Racing using Deep Reinforcement Learning, **Nature, 2023**

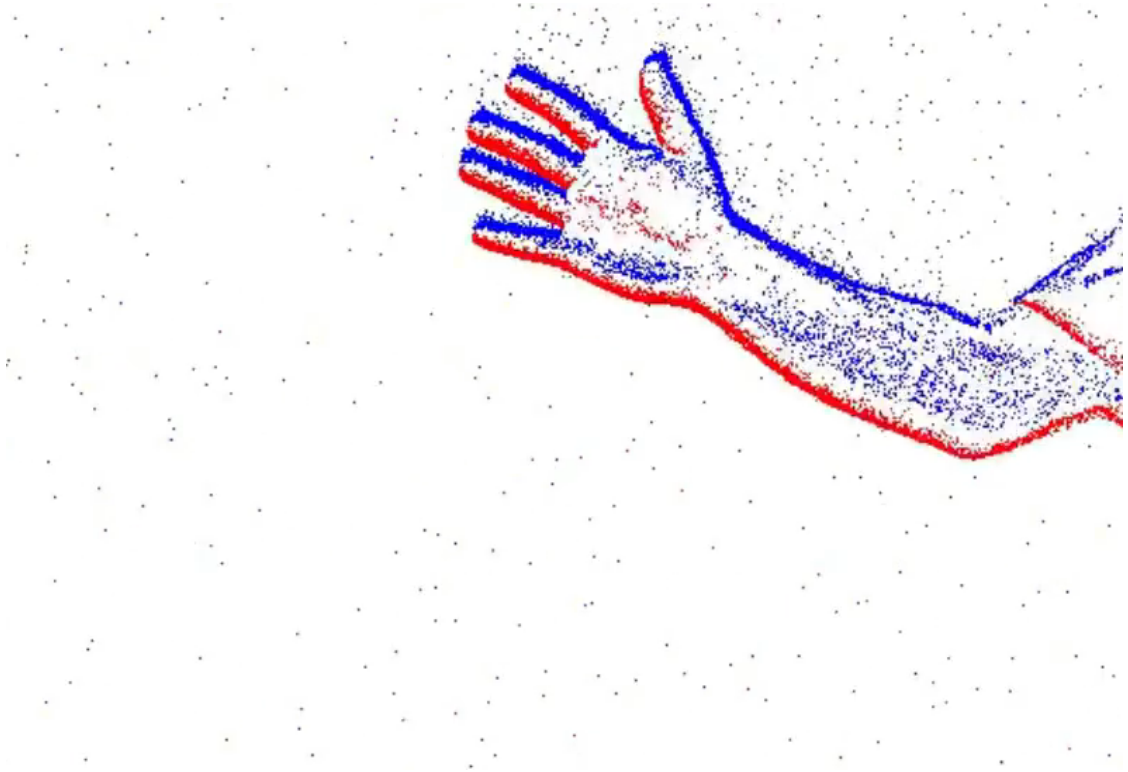
Reinforcement Learning, World Models, Differentiable Simulation

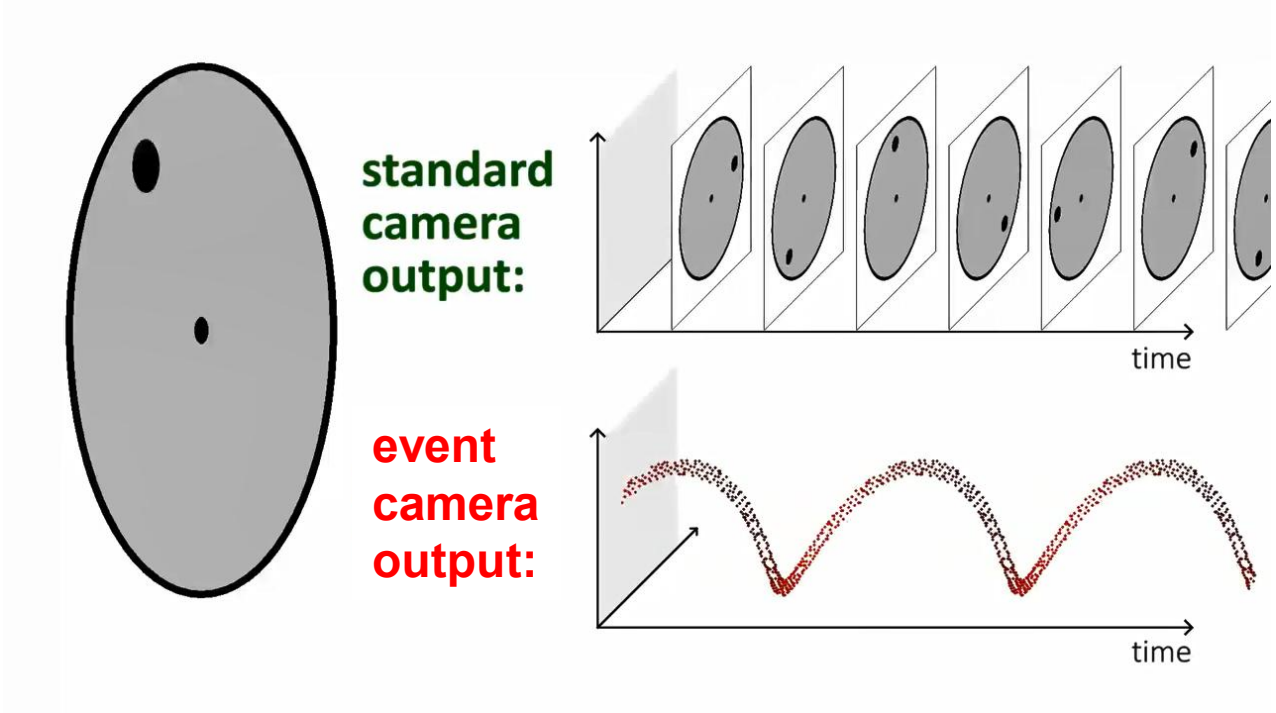




- Song, Romero, Mueller, Koltun, Scaramuzza, *Reaching the Limit in Autonomous Racing: Optimal Control versus Reinforcement Learning*, **Science Robotics**, 2023
- Romero, Shenai, Geles, Aljalbout, Scaramuzza, *Dream to Fly: Model-Based Reinforcement Learning for Vision-Based Drone Flight*, **ICRA**, 2026
- Pan, Xing, Reiter, Zhai, Aljalbout, Scaramuzza, *Learning on the Fly: Rapid Policy Adaptation via Differentiable Simulation,* **RAL**, 2026.

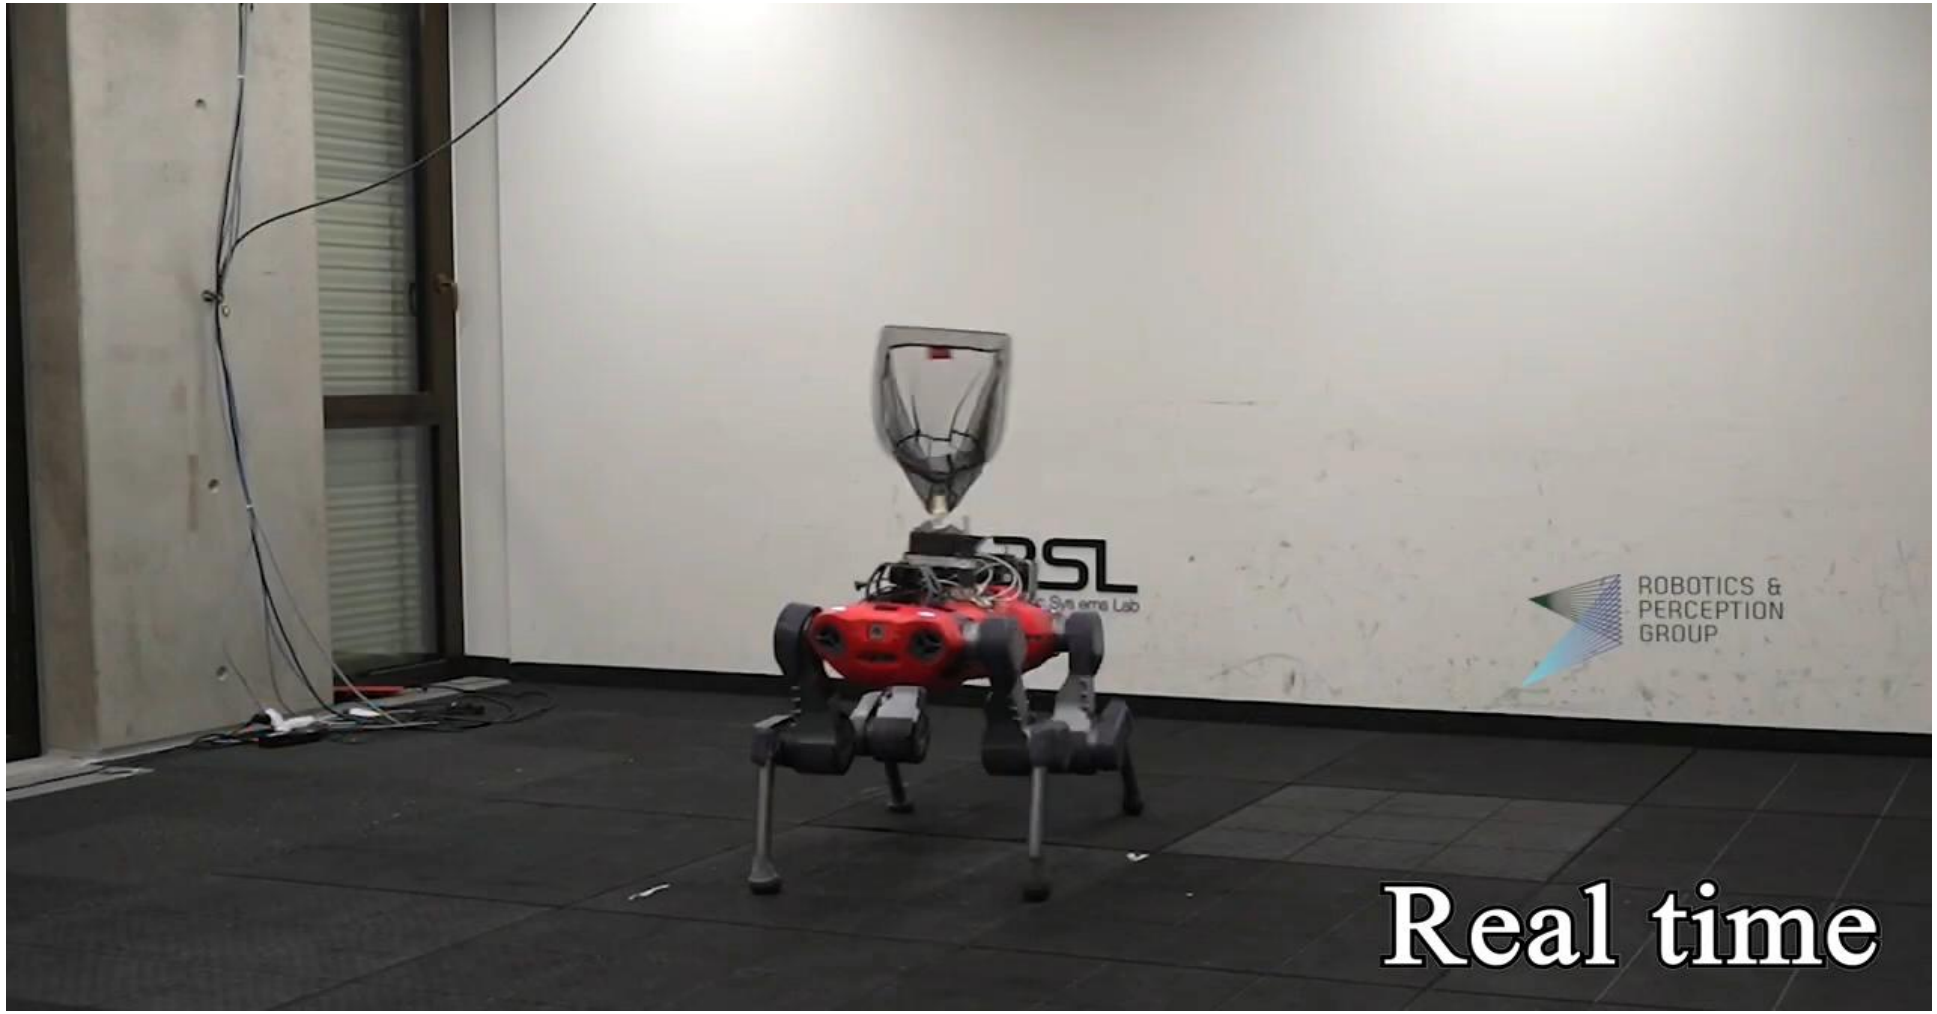
Event Cameras





- Lichtsteiner, Posch, Delbruck, *A 128x128 120 dB 15µs Latency Asynchronous Temporal Contrast Vision Sensor*, **IEEE Journal of Solid-State Circuits**, 2008
- Gallego et al., *Event-based Vision: A Survey*, **T-PAMI, 2020**

Low-Latency Perception and Control



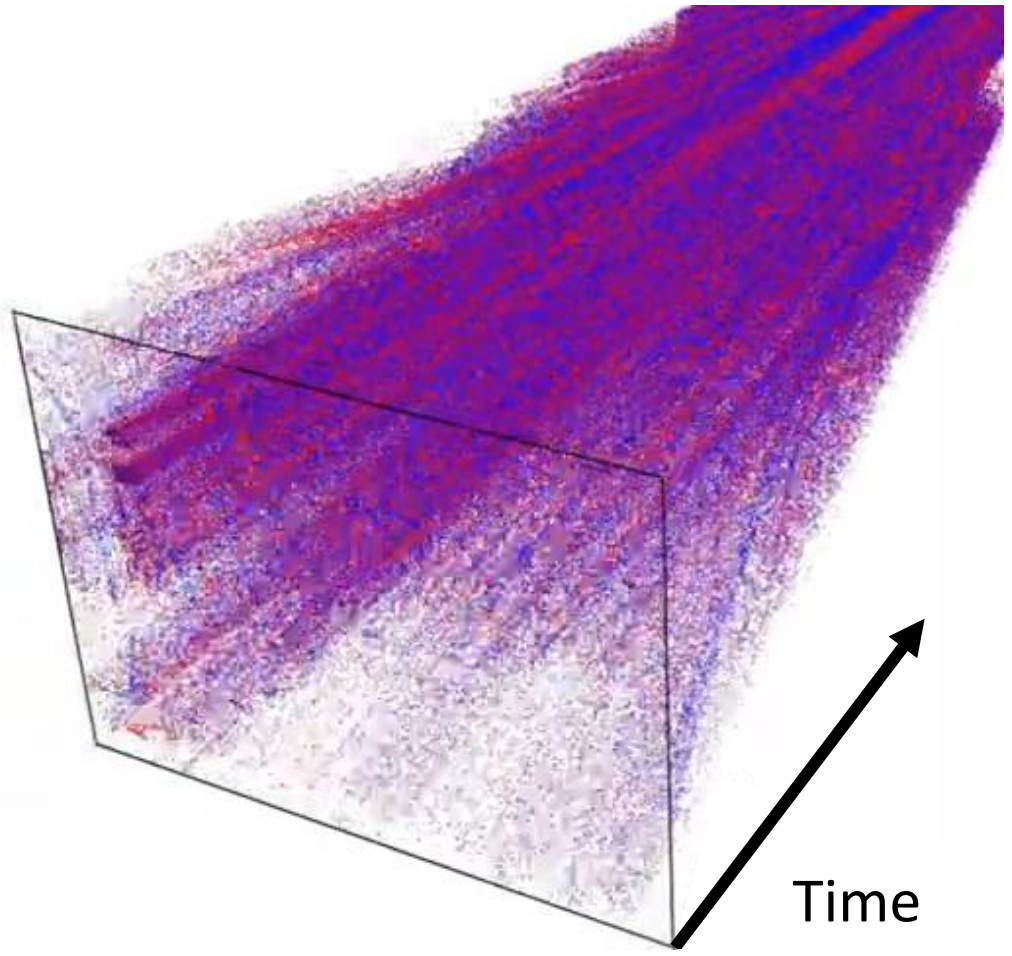
Forrai, Miki, Gehrig, Hutter, Scaramuzza, *Event-based Agile Object Catching with a Quadrupedal Robot*, ICRA'23

Low-Latency Perception for Autonomous Driving

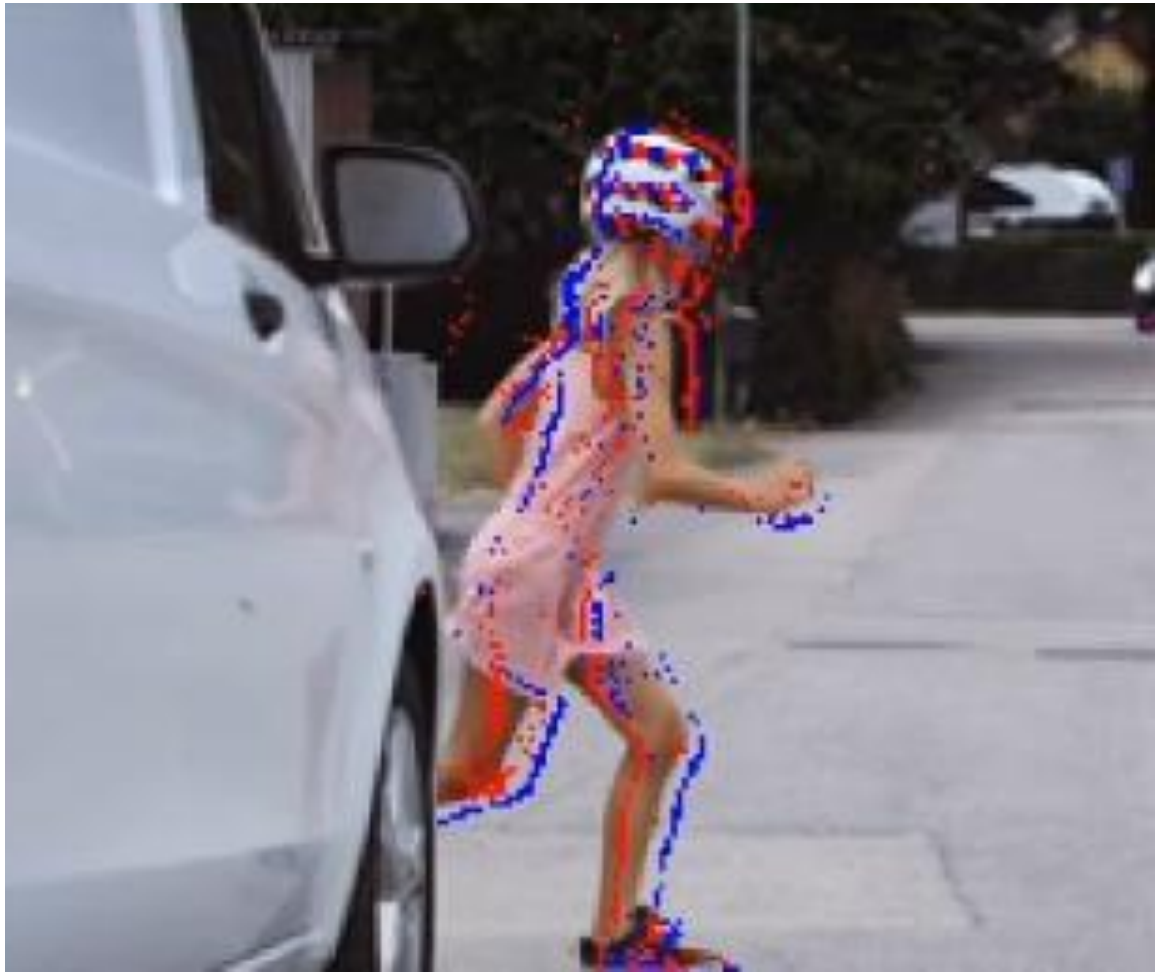
Standard camera



Event camera



Low-Latency Perception and World Models for Autonomous Driving

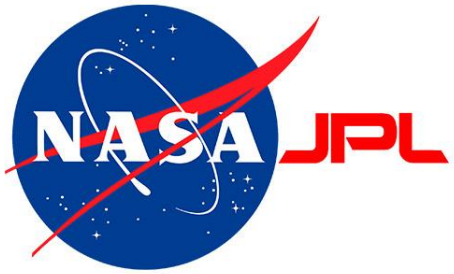




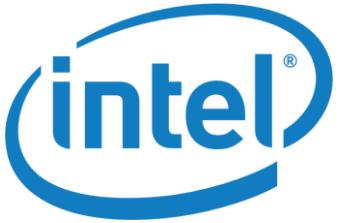
- Gehrig, Scaramuzza, *Low Latency Automotive Vision with Event Cameras*, **Nature**, 2024.
- Cao, Xing, Messikommer, Scaramuzza, *Generative Event Pretraining with Foundation Model Alignment*, **CVPR**, 2026

Tech Transfer and Spin-offs













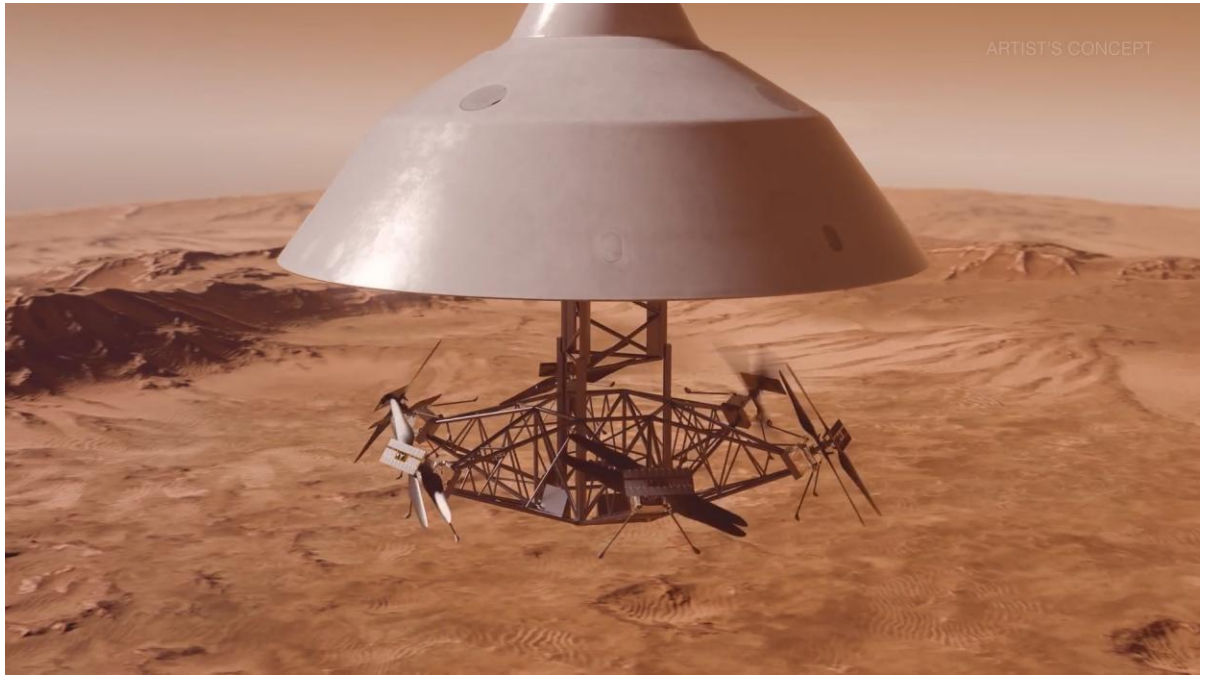








Collaboration with NASA/JPL for Skyfall mission



https://science.nasa.gov/mission/skyfall/

1. Alberico, Delaune, Cioffi, Scaramuzza, Structure-Invariant Range-Visual-Inertial Odometry, IROS'24, PDF, Video
2. Polizzi, Hewitt, Hidalgo-Carrió, Delaune, Scaramuzza, Data-Efficient Collaborative Decentralized Thermal-Inertial Odometry, RAL'22, PDF, Video
3. Mahlknecht, Gehrig, Nash, Rockenbauer, Morrell, Delaune, Scaramuzza, Exploring Event Camera-based Odometry for Planetary Robots, RAL'22, PDF, Video
4. More in this Swissinfo article

Autonomous Crop Spraying and Forest monitoring



Fotokite A tethered drone for first response

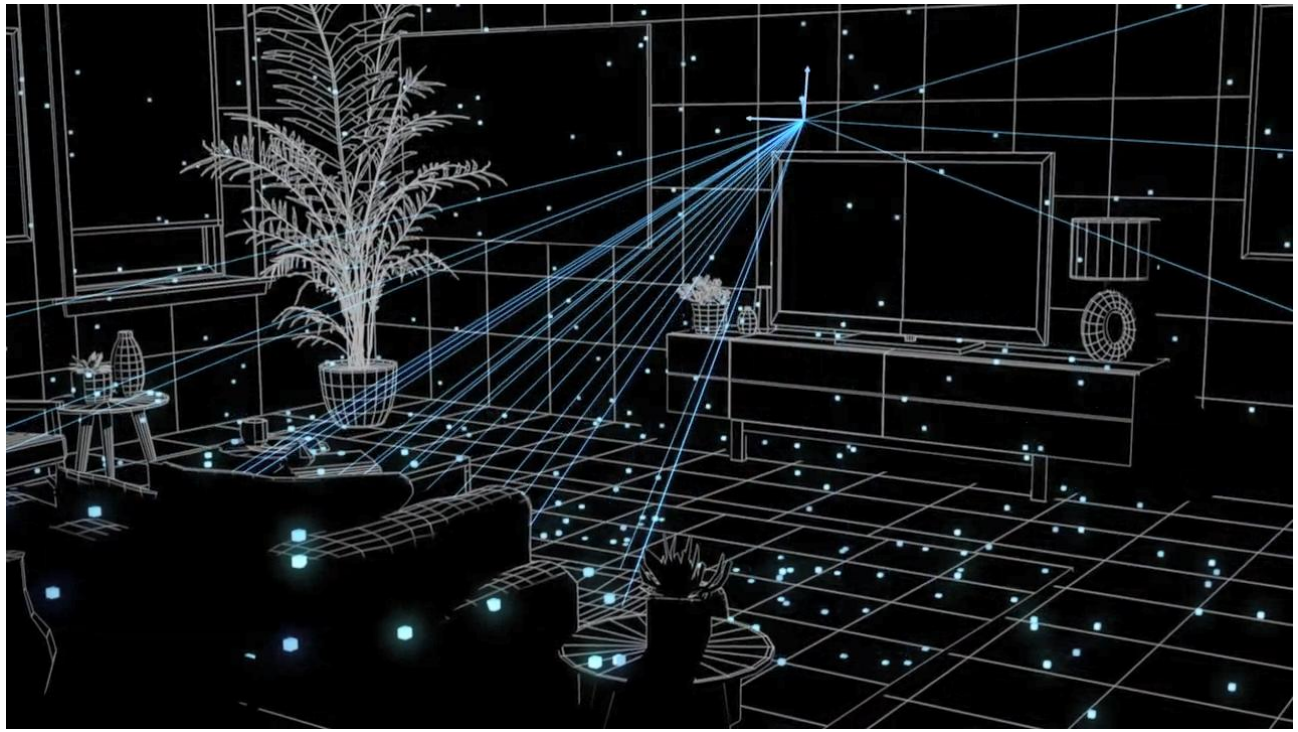
The drone receives electrical power over-tether from the ground so that it can fly “forever”

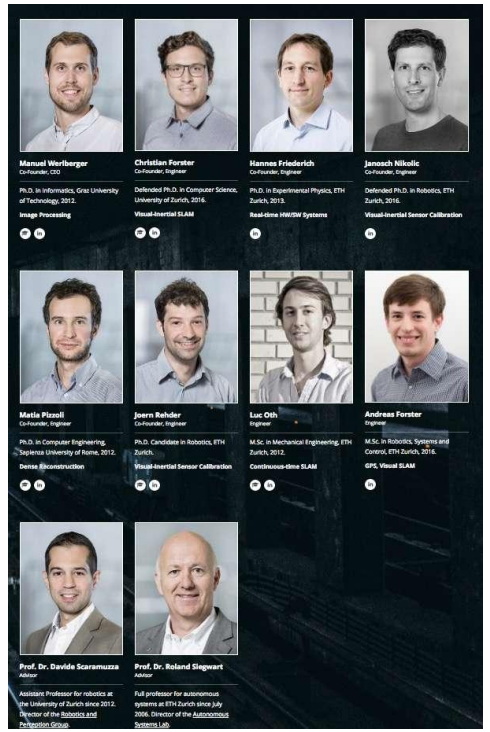




“Zurich-Eye” – Today: Meta Zurich

- **Vision-based Localization and Mapping** systems for mobile robots
- Born in Sep. 2015, became **Facebook-Oculus Zurich** in Sep. 2016. Today, over **400 employees**.
- In 2018, Zurich-Eye launched **Oculus Quest** (25 million units sold so far)

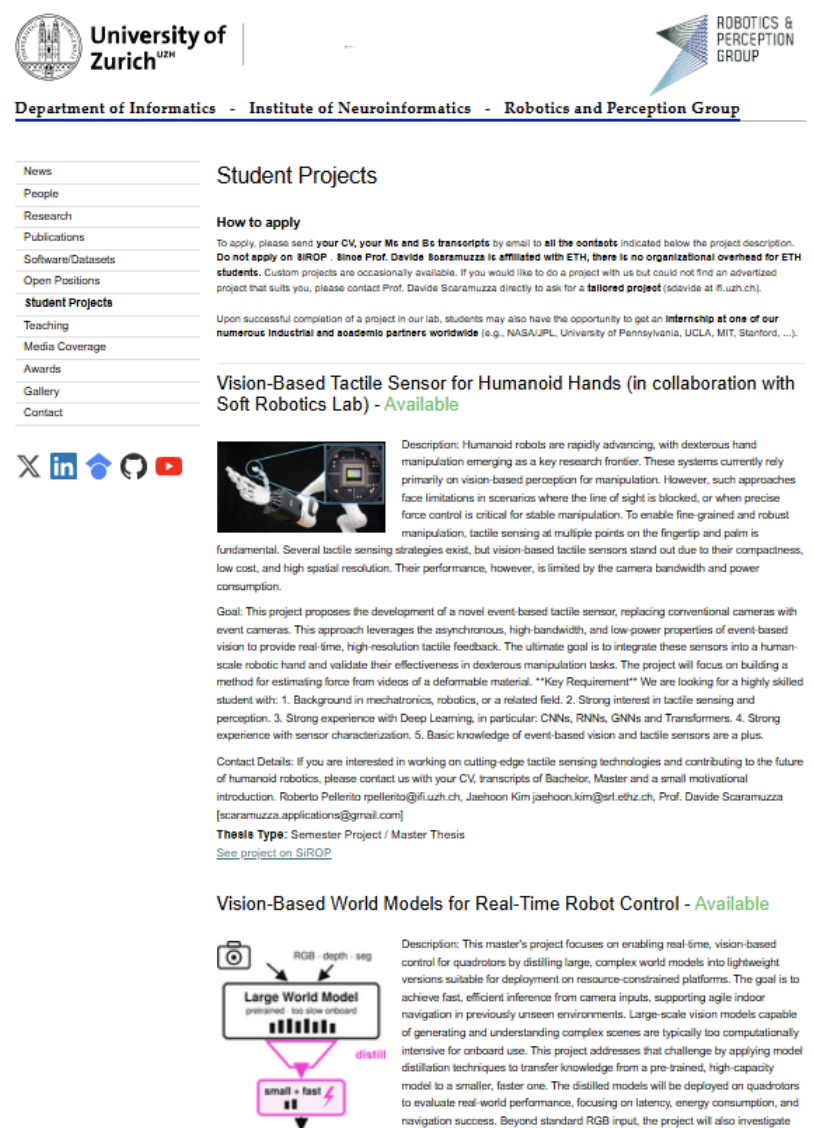




“**From the lab to the living room**”: The story behind Facebook’s Oculus Insight technology from Zurich-Eye to Oculus Quest: https://tech.fb.com/the-story-behind-oculus-insight-technology/

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

- **Topics**: machine learning, computer vision, event cameras, foundation models, world models, reinforcement learning, control, planning, system integration
- **Highlights**: many of our Master students have published their thesis in international conferences, won prestigious awards (e.g., ETH Medal, Fritz-Kutter Award, conference paper awards), were admitted for a PhD at top robotics labs (UZH, ETH, UPenn, CMU, MIT, etc.), or landed jobs at top-tier companies (NASA/JPL, CERN, SONY, Meta, NVIDIA, Microsoft, Skydio, etc.)



University of Zurich[UZH]

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, ...).

Vision-Based Tactile Sensor for Humanoid Hands (in collaboration with Soft Robotics Lab) - Available

Description: Humanoid robots are rapidly advancing, with dexterous hand manipulation emerging as a key research frontier. These systems currently rely primarily on vision-based perception for manipulation. However, such approaches face limitations in scenarios where the line of sight is blocked, or when precise force control is critical for stable manipulation. To enable fine-grained and robust manipulation, tactile sensing at multiple points on the fingertip and palm is fundamental. Several tactile sensing strategies exist, but vision-based tactile sensors stand out due to their compactness, low cost, and high spatial resolution. Their performance, however, is limited by the camera bandwidth and power consumption.

Goal: This project proposes the development of a novel event-based tactile sensor, replacing conventional cameras with event cameras. This approach leverages the asynchronous, high-bandwidth, and low-power properties of event-based vision to provide real-time, high-resolution tactile feedback. The ultimate goal is to integrate these sensors into a human-scale robotic hand and validate their effectiveness in dexterous manipulation tasks. The project will focus on building a method for estimating force from videos of a deformable material. **Key Requirement** We are looking for a highly skilled student with: 1. Background in mechatronics, robotics, or a related field. 2. Strong interest in tactile sensing and perception. 3. Strong experience with Deep Learning, in particular: CNNs, RNNs, GNNs and Transformers. 4. Strong experience with sensor characterization. 5. Basic knowledge of event-based vision and tactile sensors are a plus.

Contact Details: If you are interested in working on cutting-edge tactile sensing technologies and contributing to the future of humanoid robotics, please contact us with your CV, transcripts of Bachelor, Master and a small motivational introduction. Roberto Pellerito rpellerito@ifi.uzh.ch, Jaehoon Kim jaehoon.kim@srl.ethz.ch, Prof. Davide Scaramuzza [scaramuzza.applications@gmail.com]

Thesis Type: Semester Project / Master Thesis

See project on SiROP

Vision-Based World Models for Real-Time Robot Control - Available

Description: This master's project focuses on enabling real-time, vision-based control for quadrotors by distilling large, complex world models into lightweight versions suitable for deployment on resource-constrained platforms. The goal is to achieve fast, efficient inference from camera inputs, supporting agile indoor navigation in previously unseen environments. Large-scale vision models capable of generating and understanding complex scenes are typically too computationally intensive for onboard use. This project addresses that challenge by applying model distillation techniques to transfer knowledge from a pre-trained, high-capacity model to a smaller, faster one. The distilled models will be deployed on quadrotors to evaluate real-world performance, focusing on latency, energy consumption, and navigation success. Beyond standard RGB input, the project will also investigate

Today's Outline

- About me and my research lab
- What is Computer Vision?
- Why study computer vision?
- Examples of Vision Applications
- Organization of the course
- Start: Visual Odometry overview

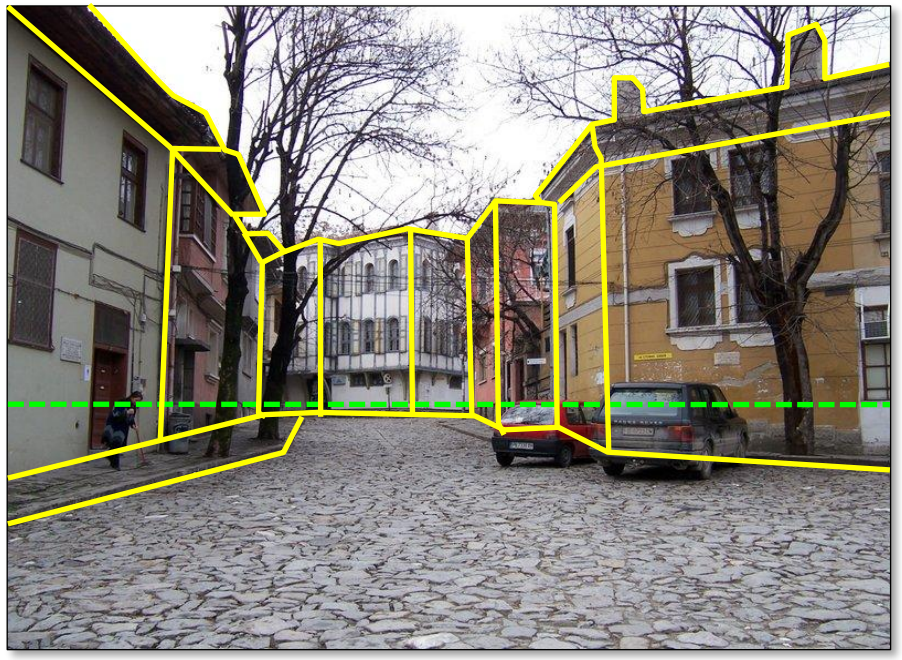
What is computer vision?

Automatic extraction of "**meaningful**" information from images and videos



Semantic information

("Image Analysis and Computer Vision" course)



Geometric information

(this course)

Vision Demo?

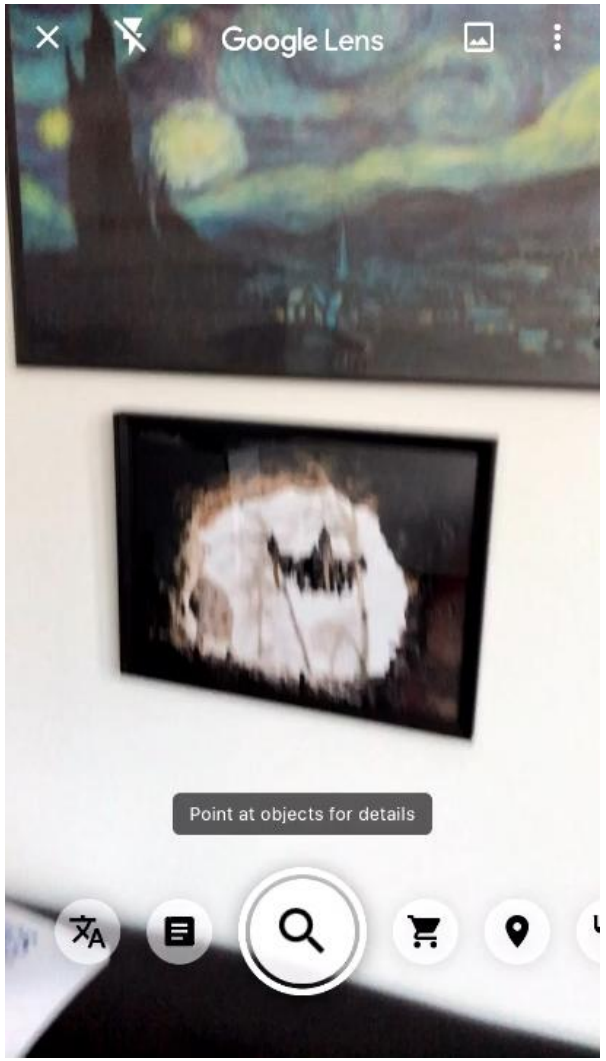


Terminator 2



Are we there? Almost!

Google App



Today's Outline

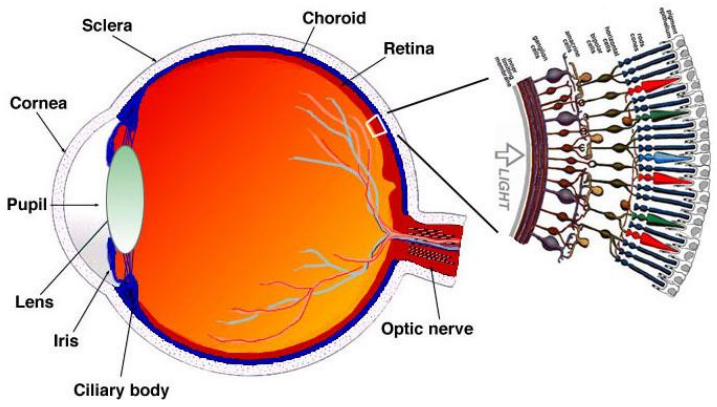
- About me and my research lab
- What is Computer Vision?
- Why study computer vision?
- Examples of Vision Applications
- Organization of the course
- Start: Visual Odometry overview

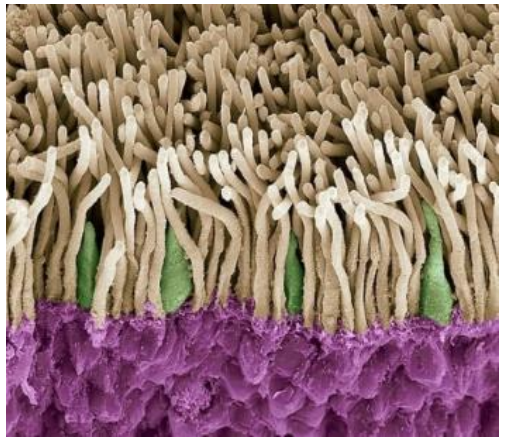
Why study computer vision?

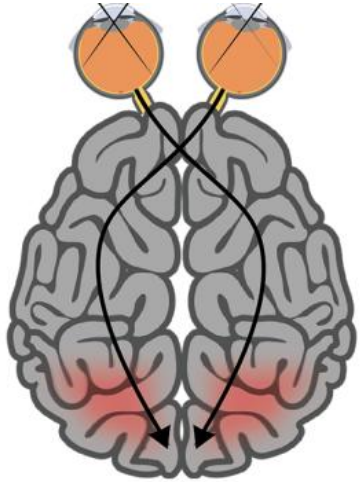
- **Relieve** humans of boring, easy tasks
- **Enhance** human abilities: human-computer interaction, visualization, augmented reality (AR)
- Perception for autonomous **robots**
- **Organize** and give access to visual **content**
- Lots of computer-vision **companies and jobs in Switzerland** (Zurich & Lausanne):
  - Meta (Zurich): AR/VR, Instagram
  - Huawei (Zurich): automotive, autonomous cars, event cameras, computational photography
  - Verity (Zurich): SLAM engineer
  - Perspective Robotics (Zurich): Computer vision engineer
  - Fixposition (Zurich): Sensor fusion engineer
  - Magic-Leap (Zurich): AR/VR
  - Microsoft Research (Zurich): Robotics and Hololens AR
  - Google (Zurich): Brain, Positioning Services, Street View, YouTube
  - Apple (Zurich): Autonomous Driving, face tracking
  - NVIDIA (Zurich): simulation, autonomous driving
  - Logitech (Zurich, Lausanne)
  - Disney-Research (Zurich)
  - VIZRT (Zurich): sport broadcasting, 3D replay
  - Pix4D (Lausanne): 3D reconstruction from drones
  - More on glassdoor.ch

Vision in humans

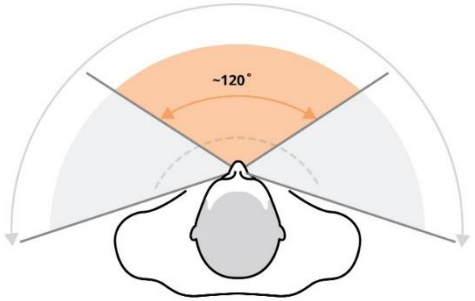
- **Vision** is our most powerful sense. **Half of the primate cerebral cortex** is devoted to visual processing
- The retina is ~1,000 mm^2. Contains **130 million photoreceptors** (120 mil. **rods for low light vision** and 10 mil. **cones for color sampling**) covering a **visual field** of $\mathbf{220 \times 135}$ **degrees**
- 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) around a small depression called **fovea**







visual cortex

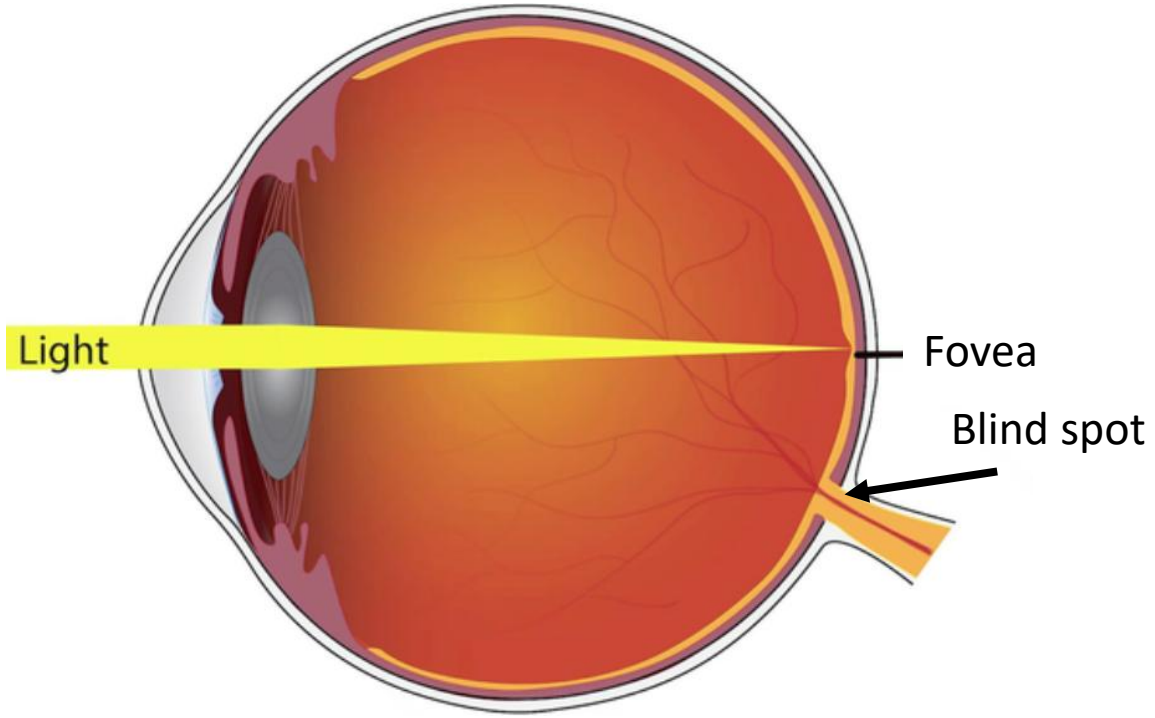


Vision in humans: how we see

- The **area we see in focus** and in **full color** represents the part of the visual field that is covered by the **fovea**
- The **fovea** is 0.35 mm in diameter, covers a visual field of **1-2 degrees**, has **high density of cone cells**
- Within the rest of the **peripheral visual field**, the image we perceive becomes more **blurry** (**rod cells**)



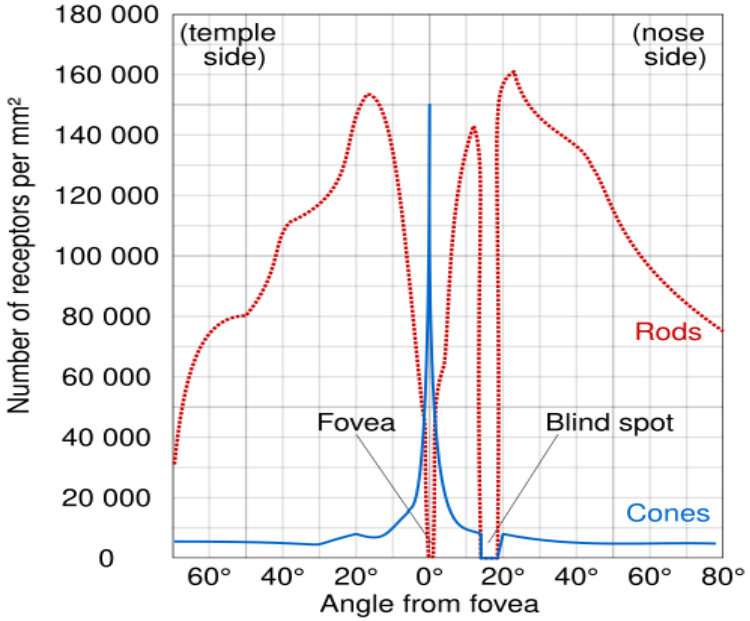
How we actually see. This principle is used in **foveated rendering**



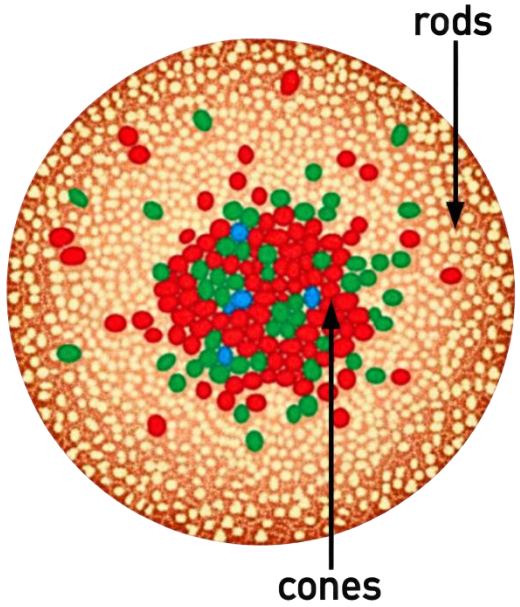
If you are interested in studying human perception, check out the course "Computational Vision" (next semester)

Vision in humans: how we see

- The **area we see in focus** and in **full color** represents the part of the visual field that is covered by the **fovea**
- The **fovea** is 0.35 mm in diameter, covers a visual field of **1-2 degrees**, has **high density of cone cells**
- Within the rest of the **peripheral visual field**, the image we perceive becomes more **blurry** (**rod cells**)



Distribution of rods and cones along a line passing through the fovea and the blind spot of a human eye



Rods are located mainly in the peripheral retina and are absent in the middle of the fovea. Cones are located throughout the retina but are concentrated in the very center of the retina.

Can we really see the reality?

- **Blind spot**: why don’t we see black?
- **Monkey business illusion**: https://www.youtube.com/@DanielSimons
- **Gradual Change blindness illusion**: https://www.youtube.com/watch?v=EARtANyz98Q

Humans do not perceive reality directly; our brain builds a predictive model of it and fills in gaps!

What a newborn sees every month in the first year

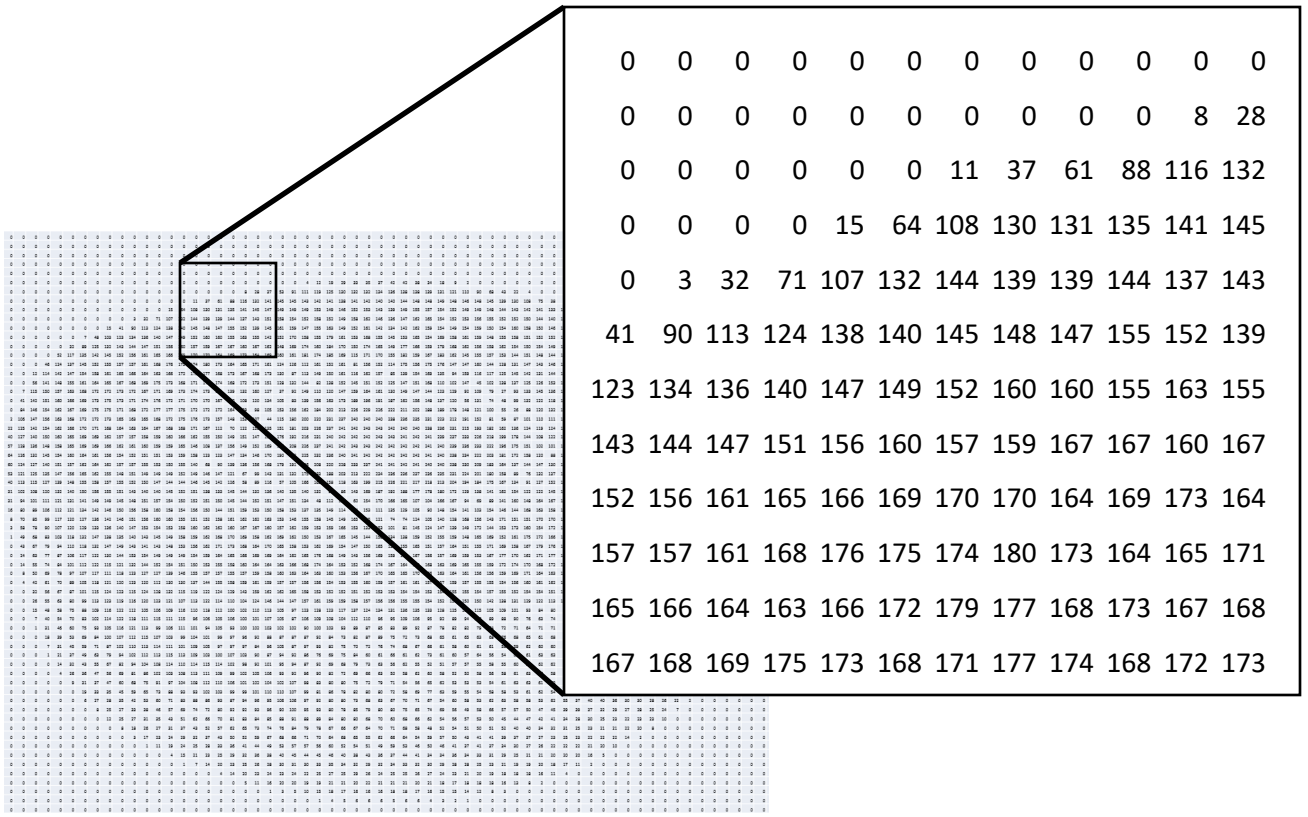
“Your baby sees things best from 15 to 30 cm away. This is the perfect distance for gazing up into the eyes of mom or dad. Any farther than that, and the newborn sees mostly blurry shapes because they're nearsighted. At birth, a newborn's eyesight is between 20/200 and 20/400.”



https://www.dailymail.co.uk/health/article-4133212/What-babies-see.html

Why is vision hard?

How do we go from an array of numbers 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

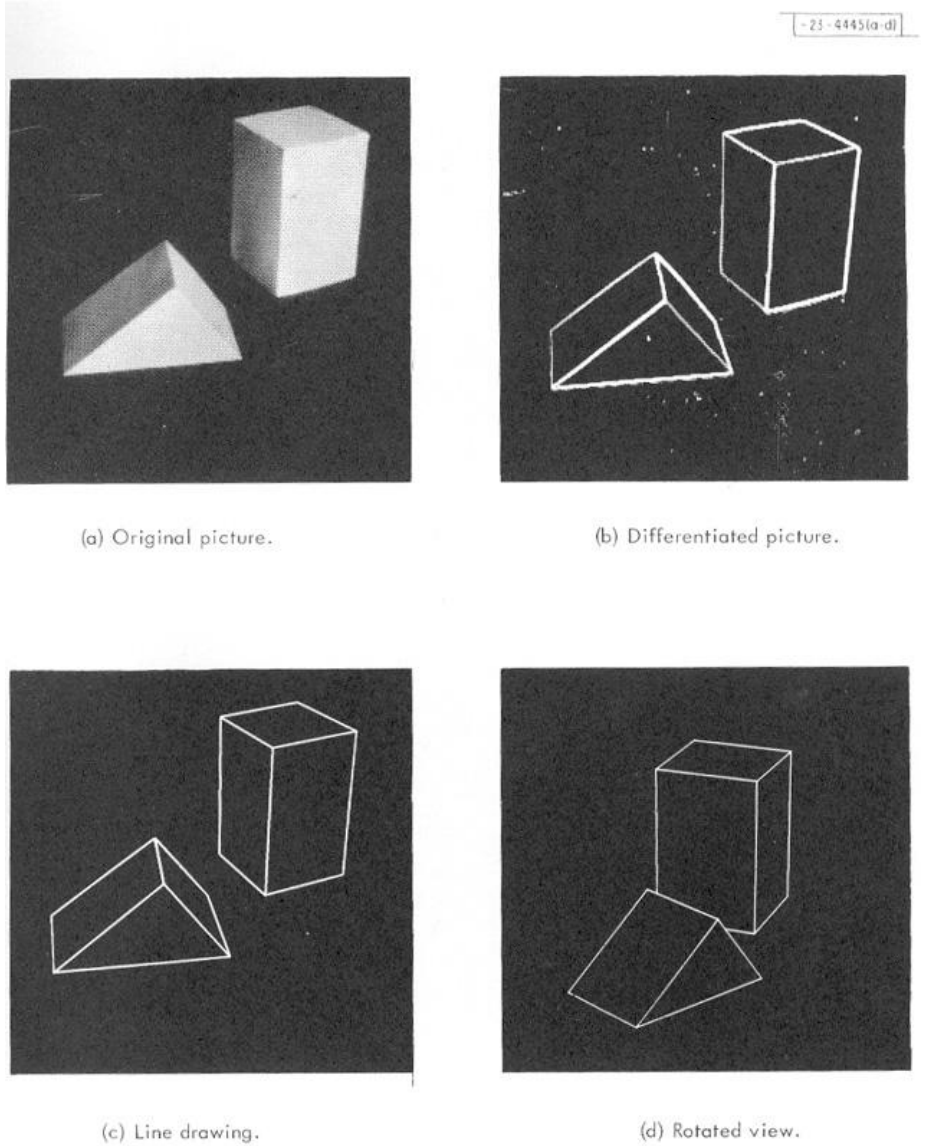
What a computer sees



What we see

Origins of computer vision

- **1963** - L. G. Roberts publishes his PhD thesis on *Machine Perception of Three Dimensional Solids*, thesis, MIT Department of Electrical Engineering
- He is the **inventor of ARPANET, the current Internet**

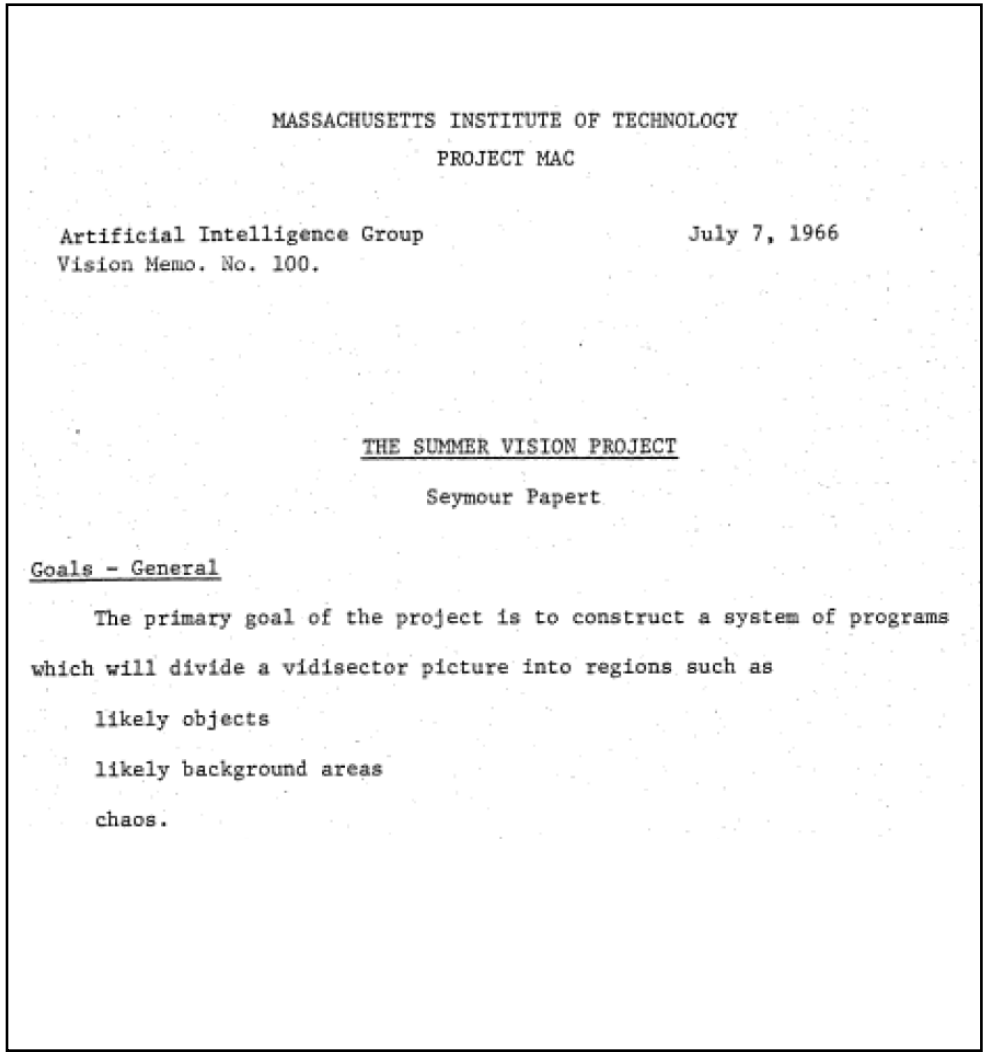


(a) Original picture. (b) Differentiated picture.

(c) Line drawing. (d) Rotated view.

Origins of computer vision

- **1963** - L. G. Roberts publishes his PhD thesis on *Machine Perception of Three Dimensional Solids*, thesis, MIT Department of Electrical Engineering
- He is the **inventor of ARPANET, the current Internet**
- **1966** – Seymour Papert, MIT, publishes the Summer Vision Project asking students to design an algorithm to segment an image into objects and background… within summer!
- **1969** – David Marr starts developing a framework for processing visual information



MASSACHUSETTS INSTITUTE OF TECHNOLOGY
PROJECT MAC

Artificial Intelligence Group
Vision Memo. No. 100.

July 7, 1966

THE SUMMER VISION PROJECT

Seymour Papert

Goals - General

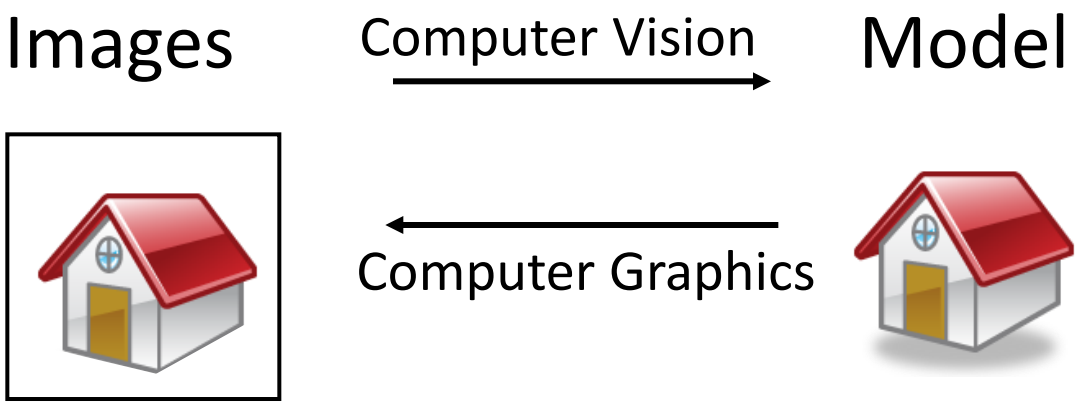
The primary goal of the project is to construct a system of programs which will divide a vidisector picture into regions such as

likely objects

likely background areas

chaos.

Computer Vision vs Computer Graphics

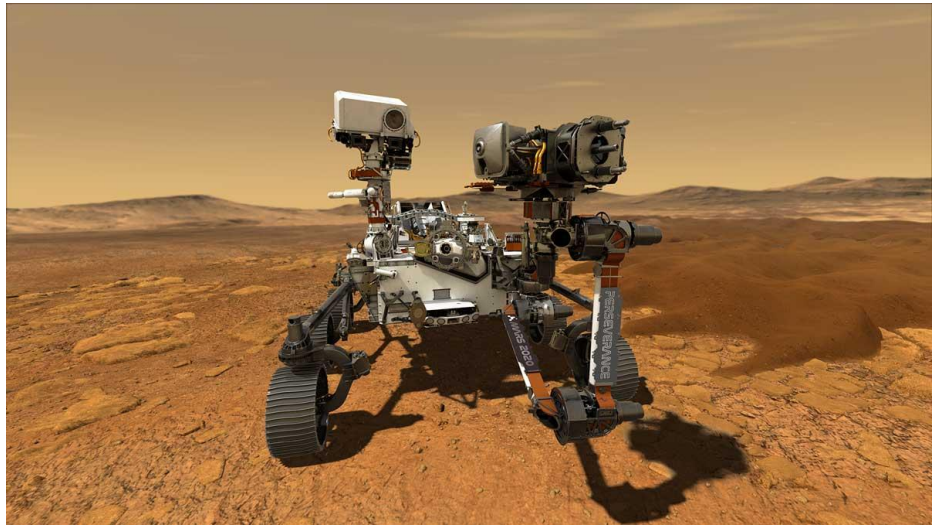


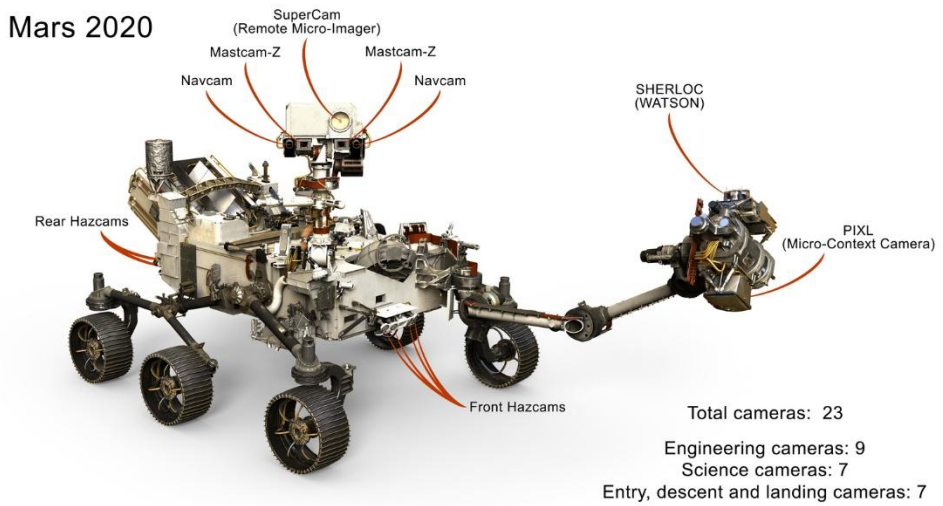
Inverse problems: analysis and synthesis.

Today's Outline

- About me and my research lab
- What is Computer Vision?
- Why study computer vision?
- Examples of Vision Applications
- Organization of the course
- Start: Visual Odometry overview

NASA Mars Rovers - Since 2004





Spirit and Opportunity (landed in 2004) were the first to use advanced computer vision (such as visual odometry) onboard. NASA'S Perseverance Rover landed in 2021 features 23 cameras used for:

- Autonomous landing on Mars (Terrain Relative Navigation, next slide)
- Autonomous navigation and positioning of the robotic arm
- Science: 3D photos and videos, analysis of chemical components of the soil

For more info, watch the RSS'21 keynote by Larry Matthies, head of the Computer Vision Group at NASA/JPL: https://youtu.be/NcI6fJOzBsU

Advanced Driving Assistance Systems (ADAS) - Since 2007



Mobileye started deploying computer vision systems on cars in 2007; now used by **BMW, GM, Volvo** models. Bought by **Intel in 2017 for 15 billion USD,** it is used in more than **200 million cars** worldwide

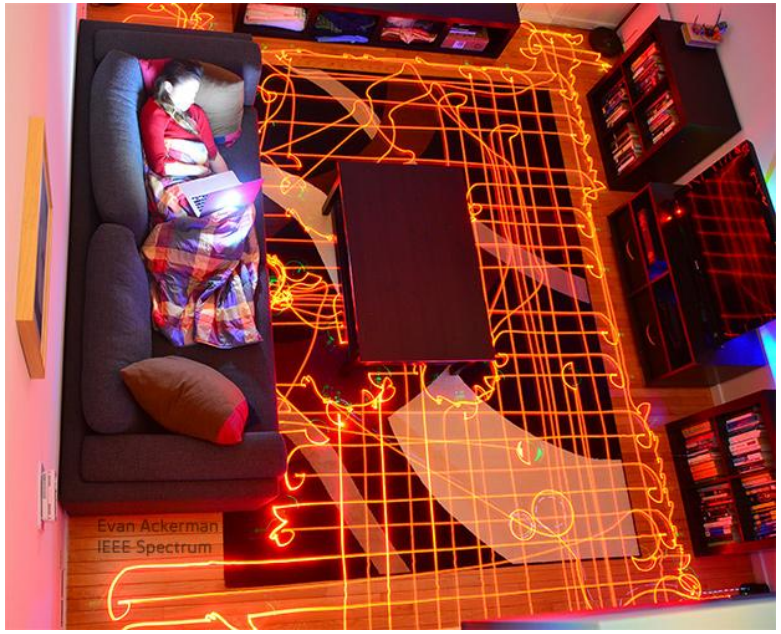
- **Pedestrian & car** collision warning
- **Lane departure** warning
- **Safety distance** monitoring and warning

Roomba Vacuum Cleaner - Since 2015

- Introduced in 2002 by iRobot; more than 50 million Roombas sold so far
- Fully autonomous but blind since 2002; cameras were only introduced in 2015 to recognize places

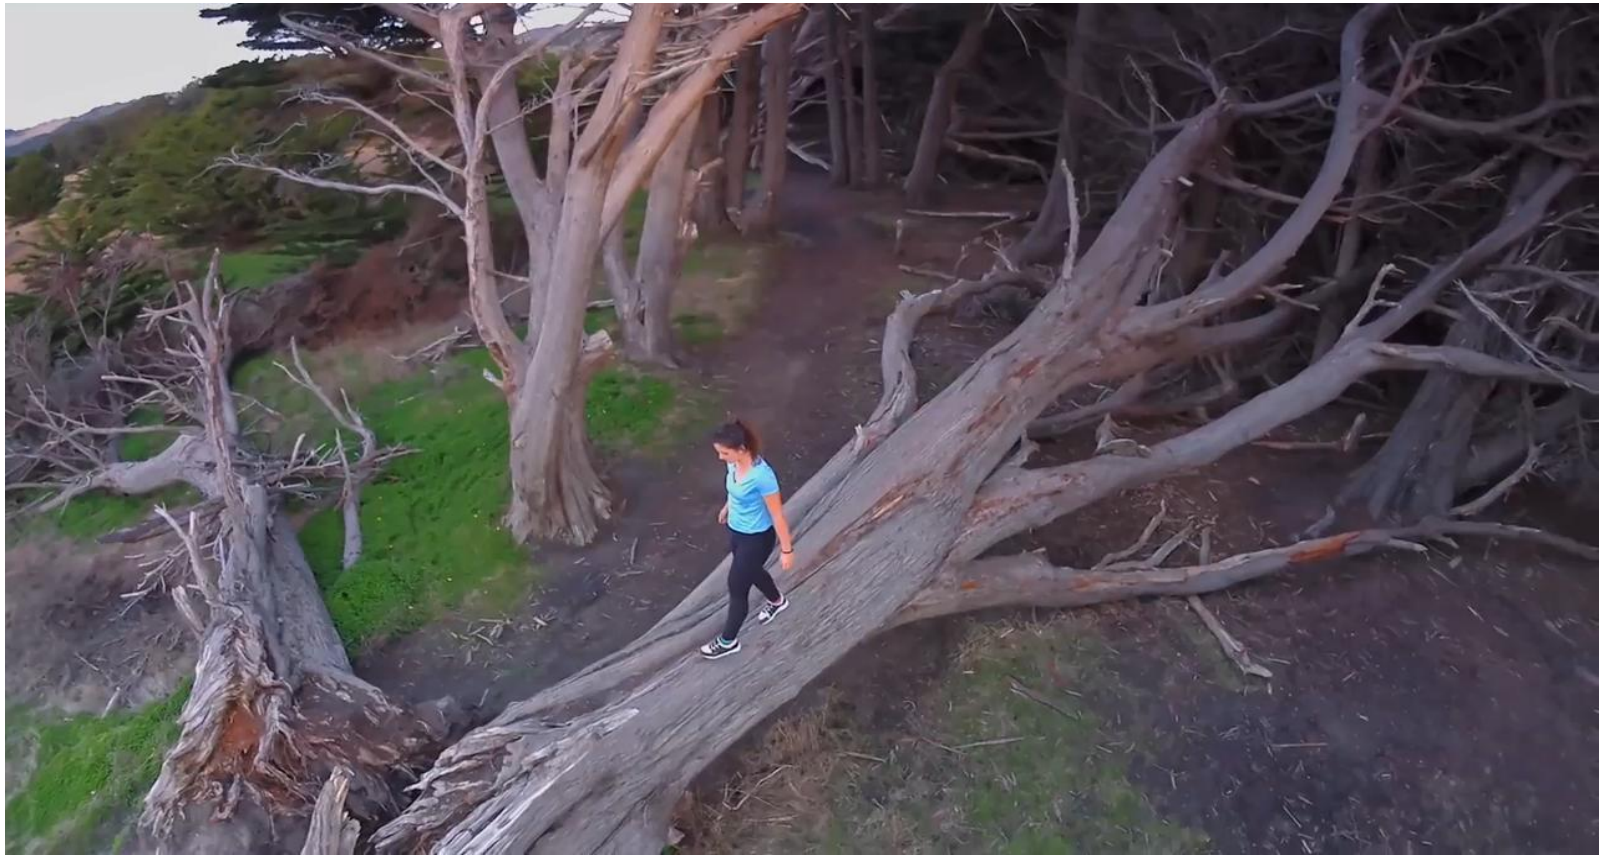


Until 2014, the Roomba robots navigated by zigzagging randomly across the room. Average cleaning time: 20 minutes



In 2015, the Roomba 980 was introduced, which navigated by following a pre-defined path optimized thanks to visual SLAM - Average cleaning time: 5 minutes

Skydio and DJI Drones - Since 2018

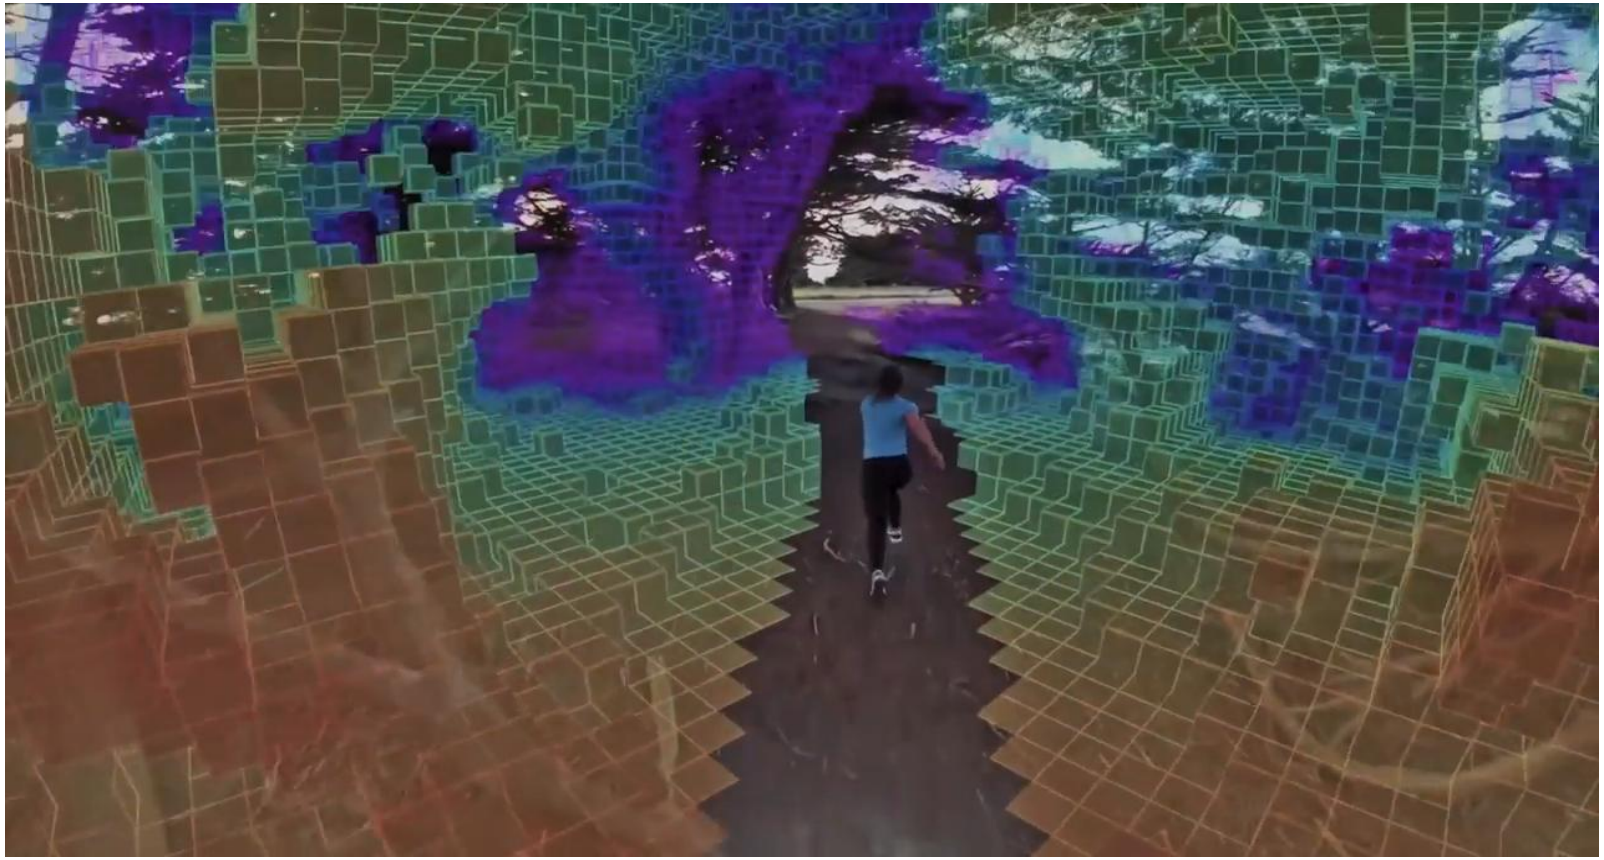


The Skydio drone features 7 cameras for obstacle avoidance, visual odometry, and person following

Watch Skydio keynote at the MIT Robotics Today Seminar Series: https://youtu.be/ncZmnfIRIWE

On November 12, we will have a lecture by Davide Falanga, head of Skydio Zurich

Skydio and DJI Drones - Since 2018

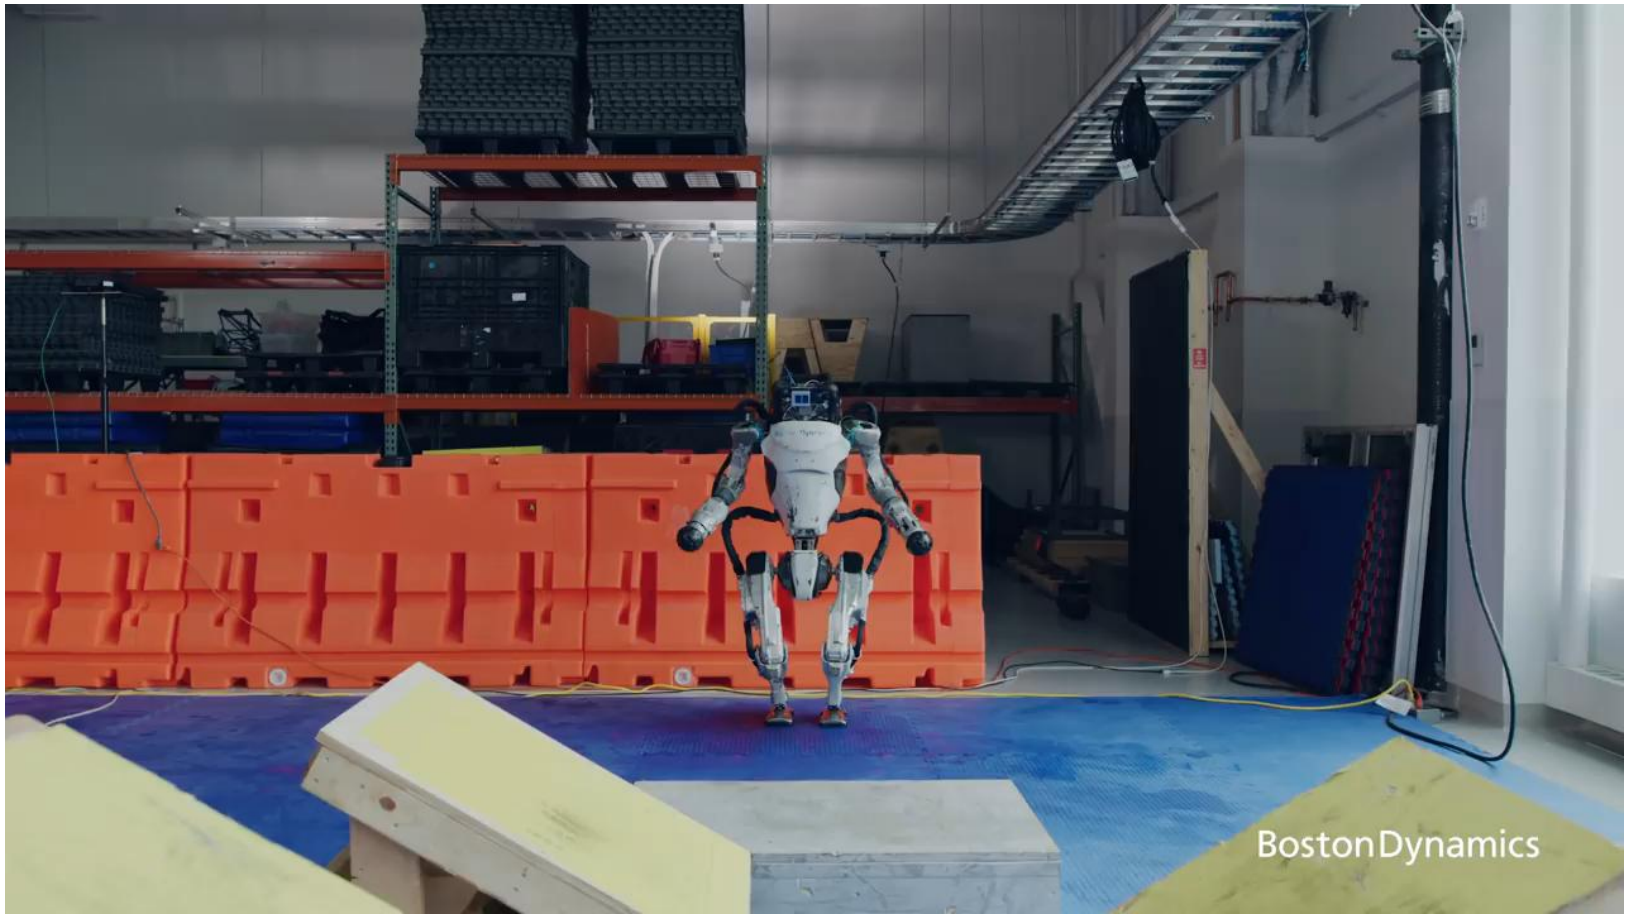


The Skydio drone features 7 cameras for obstacle avoidance, visual odometry, and person following

Watch Skydio keynote at the MIT Robotics Today Seminar Series: https://youtu.be/ncZmnfIRIWE

On November 12, we will have a lecture by Davide Falanga, head of Skydio Zurich

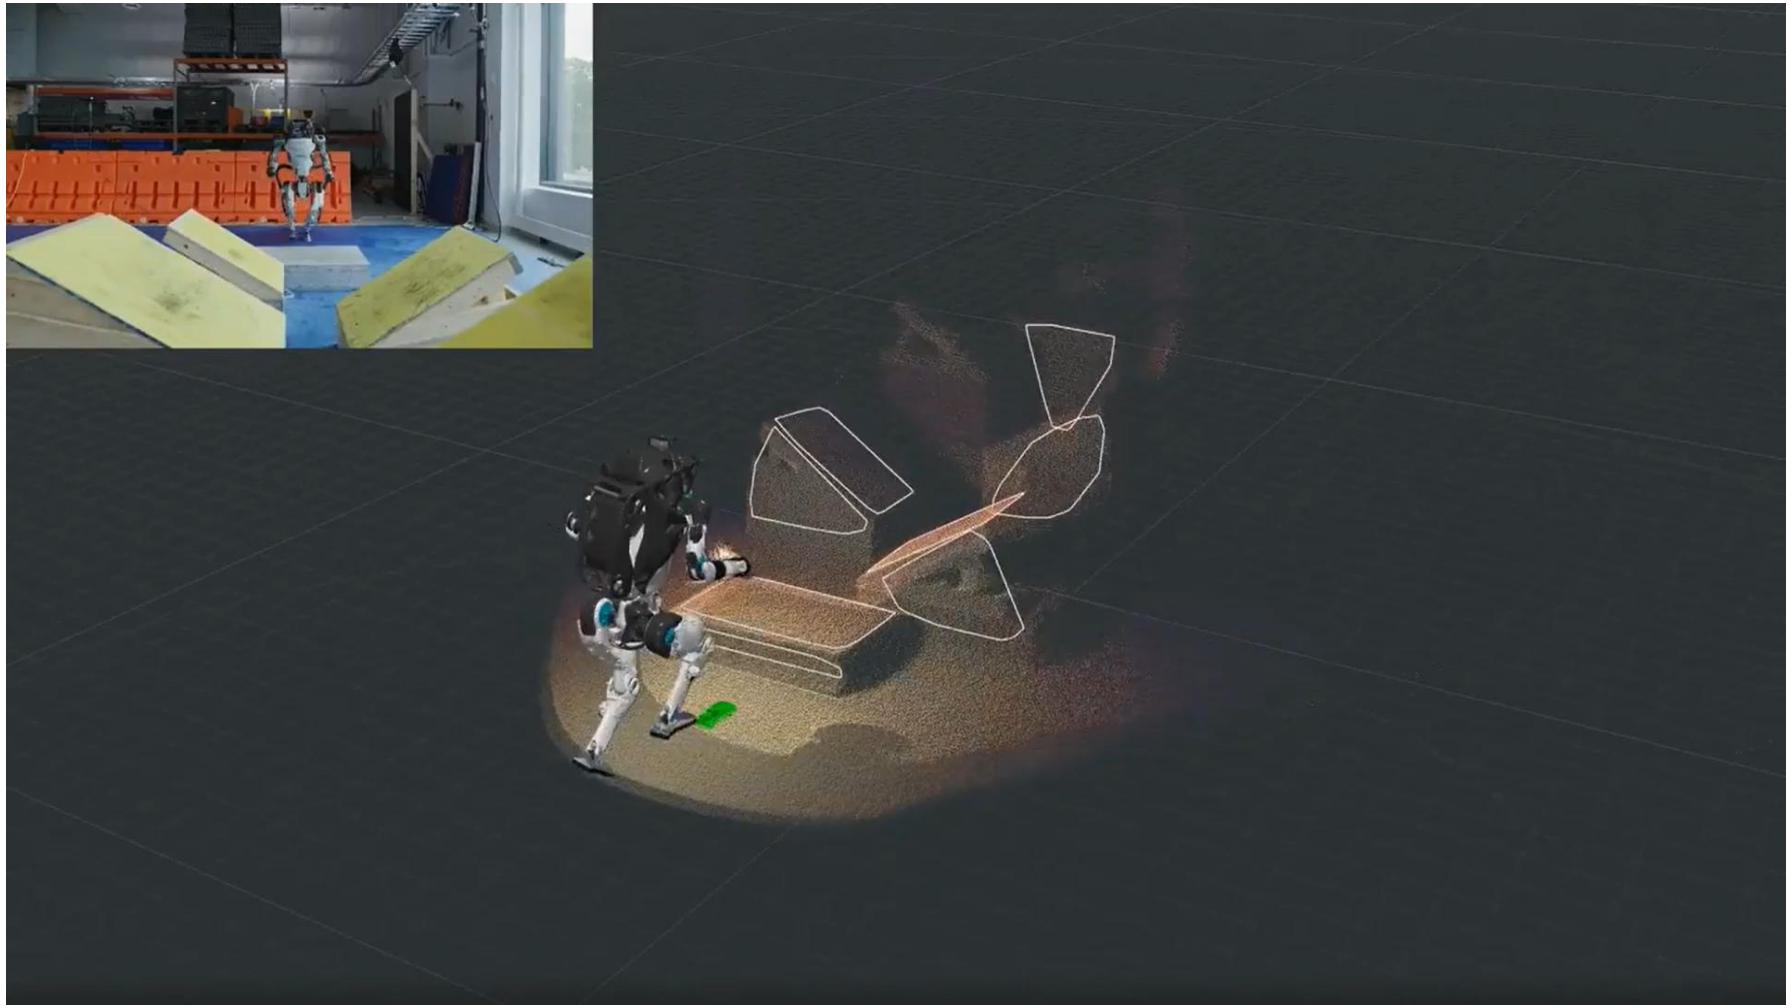
Boston Dynamics ATLAS Robot - Since 2018



https://youtu.be/tF4DML7FIWk

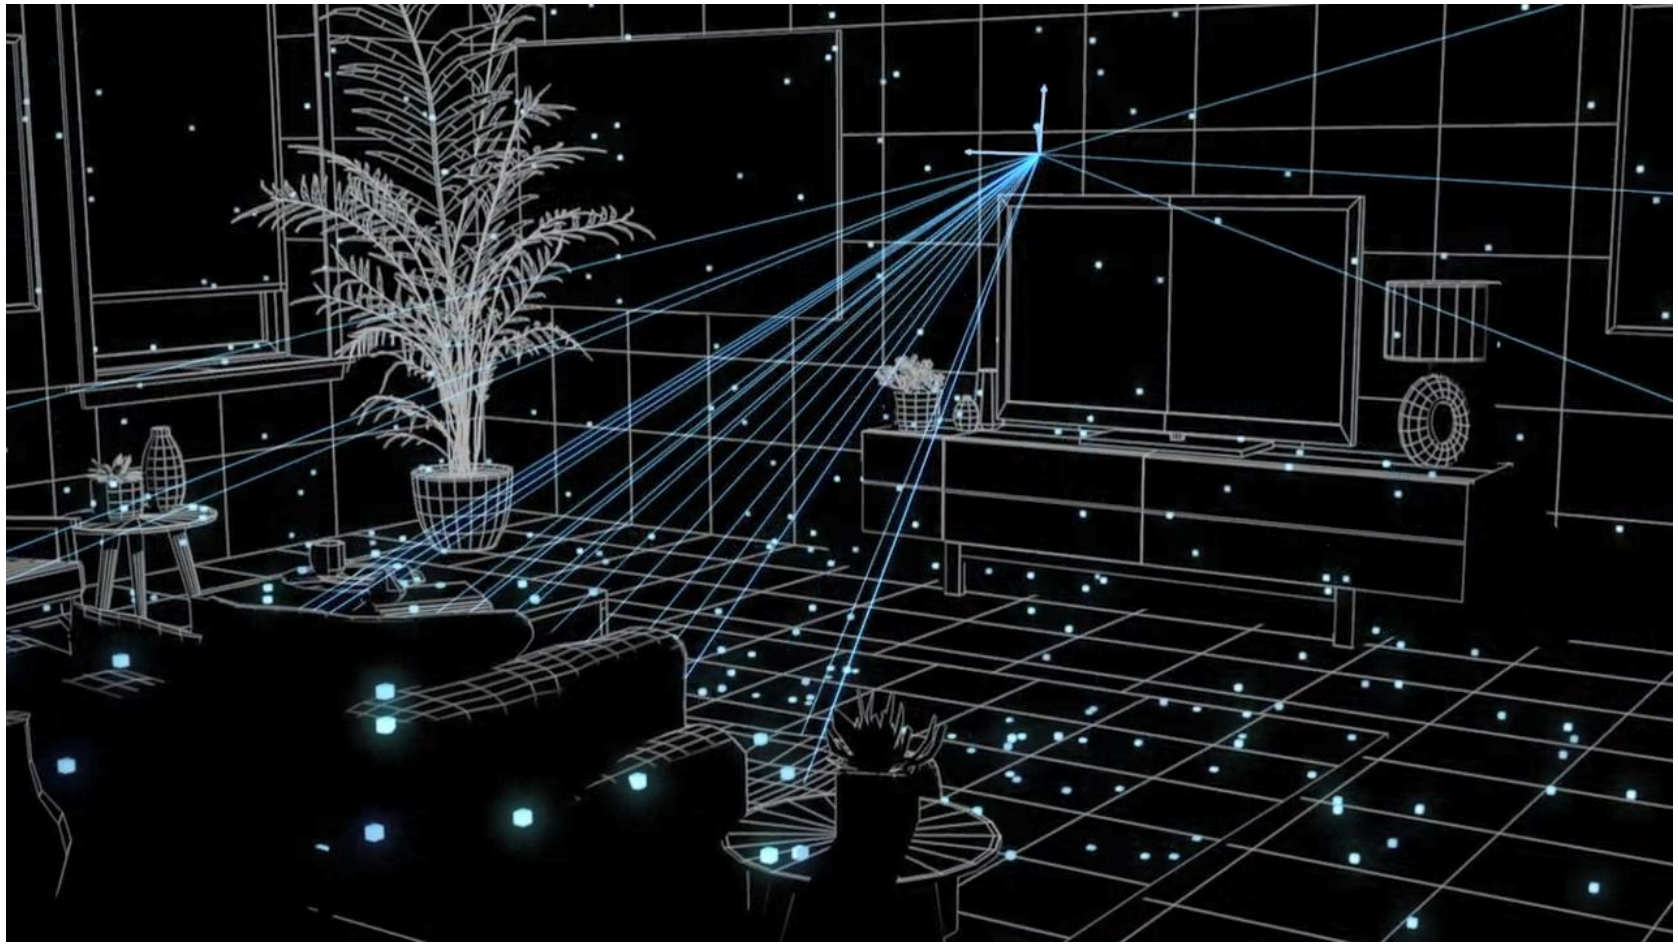
Watch Boston Dynamics keynote at the MIT Robotics Today Seminar Series: https://youtu.be/EGABAx52GKI

Boston Dynamics ATLAS Robot - Since 2018



https://blog.bostondynamics.com/flipping-the-script-with-atlas

VR/AR - Since 2019

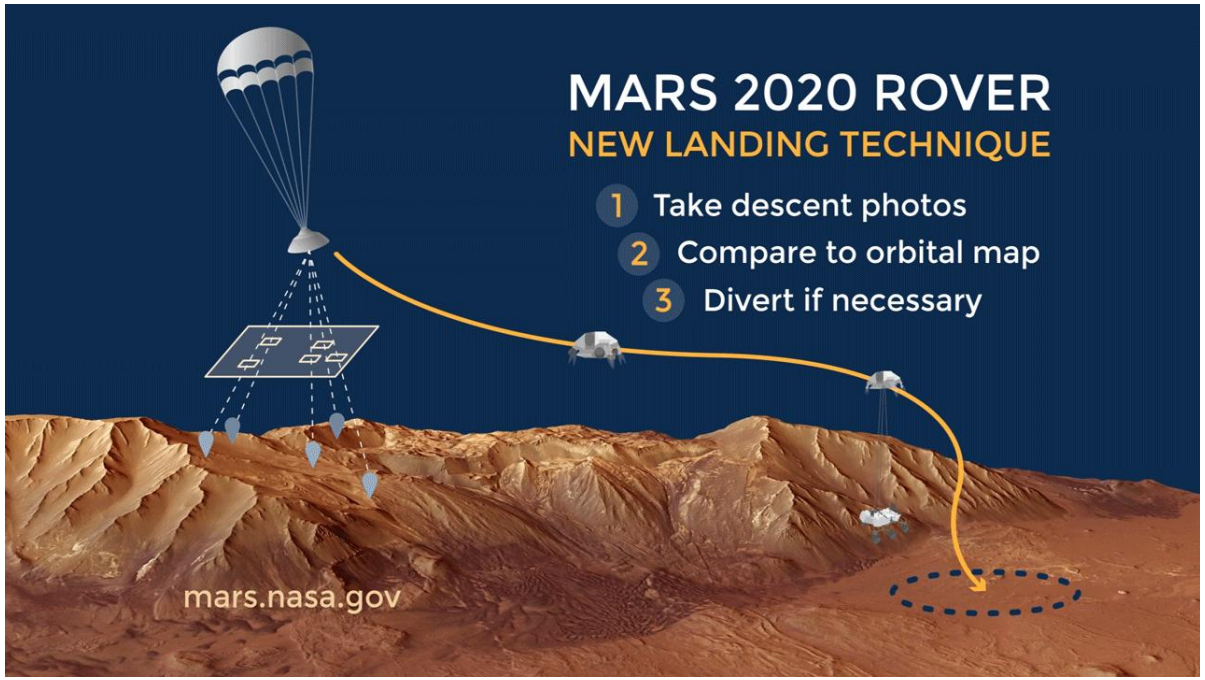


Oculus Quest uses four cameras to track the pose of the head and the controllers

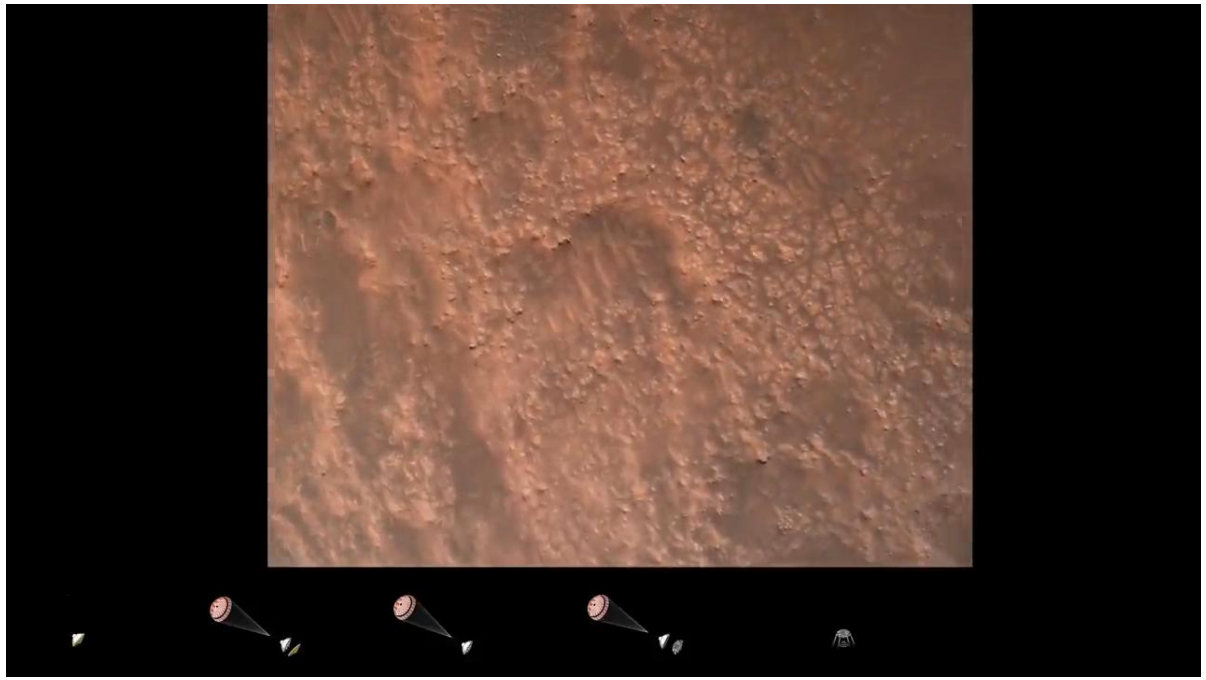
Perseverance Descent via Vision-based Terrain Relative Navigation (TRN) – Since 2021

Landing accuracy: **40 meters** (TRN improved landing uncertainty from 20km in 2012 and 7km in 2021)

"If we didn't have TRN, the probability of landing safely at Jezero Crater was about 80 to 85%. With TRN, the probability increased to 99%." [Link]



Link 1, Link 2



Real footage recorded by Perseverance during descent https://youtu.be/4czjS9h4Fpg

Vision-based Flight on Mars – Since 2021



The NASA Ingenuity helicopter performs the first autonomous vision-based flight on April 19, 2021, on Mars.
https://youtu.be/p1KolyCqICI?t=2502
https://mars.nasa.gov/technology/helicopter/#

On December 3, we will have a lecture by Jeff Delaune, from NASA/JPL, who developed the visual navigation of Ingenuity

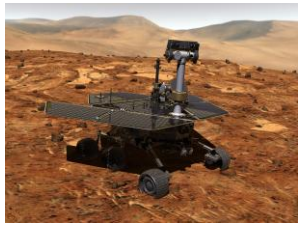
Humanoid robotics 2024



“Mobile ALOHA” first bimanual dexterous manipulation of domestic tools. Each arm has a wrist camera


Cook Shrimp
(autonomous)
2024
3x speed
Stanford
University

Computer vision applications: selected milestones



Mars rovers

- Spirit and Opportunity
- Stereo vision and odometry

2004



Roomba 980

- Visual SLAM maps the home
- Systematic cleaning paths

2015



ATLAS parkour

- Vision guides agile motion
- Terrain and foot placement

2018



Perseverance landing

- Terrain-relative navigation
- Match images to a map

FEB 2021

2007



Driver assistance

- Mobileye EyeQ deployment
- Lane and collision warnings

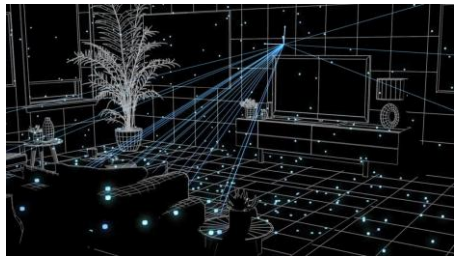
2016–2019



DJI / Skydio drones

- DJI 2016, Skydio 2018/19
- Avoid obstacles, follow people

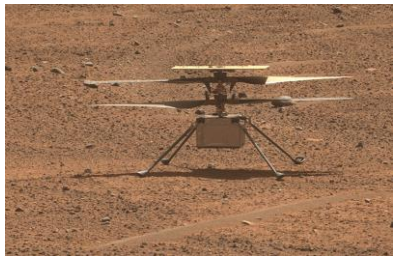
2019



Oculus Quest

- Track head and controllers
- Four cameras support VR

APR 2021

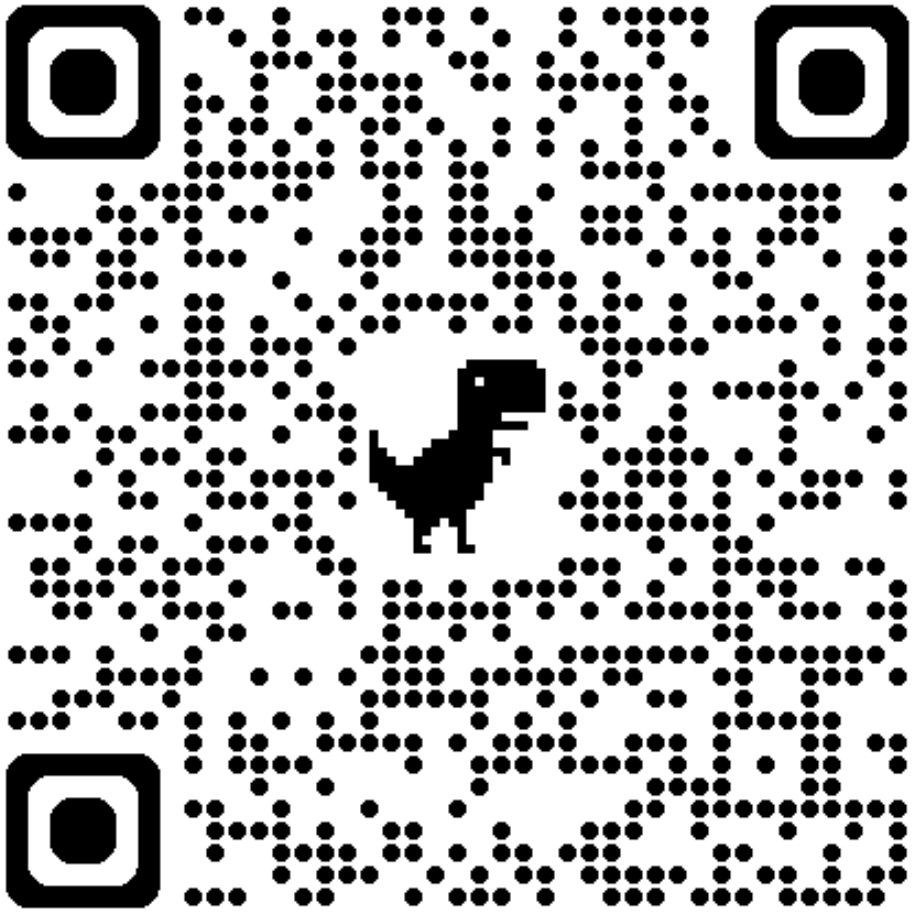


Ingenuity flight

- First powered flight on Mars
- Onboard vision guides flight

Before the Break

Please fill this quick survey: https://tinyurl.com/VAMR2026



Today's Outline

- About me and my research lab
- What is Computer Vision?
- Why study computer vision?
- Examples of Vision Applications
- Organization of the course
- Start: Visual Odometry overview

Instructors



Lecturer: Davide Scaramuzza

- **Contact**: sdavide (at) ifi (dot) uzh (dot) ch
- **Office hours**: every **Thursday from 16:00 to 18:00**
  both in person or via ZOOM possible (**please announce yourself by email**)

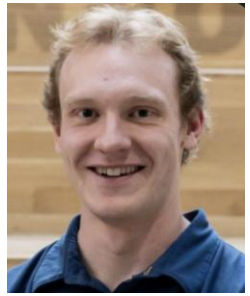
- **Teaching Assistants**: Rong Zou and Roberto Pellerito
- **Tutor**: David Jensen



Rong Zou
zou (at) ifi (dot) uzh (dot) ch



Roberto Pellerito
rpellerito (at) ifi (dot) uzh (dot) ch



David Jensen
Djensen (at) student (dot) ethz (dot) ch

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

Lectures and Exercises

Lectures:

- **08:00 to 09:45** every week. **After class**, **I usually stay for 10 more minutes** for questions
- **Breaks**: always but can vary between 5-15 min
- Room: SOC-F-106, Rämistrasse 69, 8001 Zurich

Exercises:

- **12:15 to 13:45** every week: starting today with a tutorial on camera notation
- Room: same as above

Course & Exam Registration and Cancelation

- Registration and exam cancelation deadline for the course is **October 6, 23:59 hrs**
- NB: at UZH, when you register for a course, you are also automatically registered for the exam. If you want to cancel the exam, you must unbook the course by October 6. Afterward, it will no longer be possible to cancel the exam. No show at the exam will be graded as 1.0. If you cannot take the exam because you fell ill, you must submit a petition with medical certificate no later than five business days after the examination date. If you are a student with a disability, you must request assistance in due time. Further info here

Tentative Course Schedule

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

Two industry seminars:
- Skydio
- NASA-JPL

17.09.2026	**Lecture 01 - Introduction to Computer Vision and Visual Odometry** **Exercise: Camera Notation Tutorial**	Scaramuzza Rong, Roberto
24.09.2026	**Lecture 02 - Image Formation: perspective projection and camera models** **Exercise 01- Augmented reality wireframe cube**	Scaramuzza Rong, Roberto
01.10.2026	**Lecture 03 - Camera Calibration** **Exercise 02 - PnP problem**	Scaramuzza Rong, Roberto
08.10.2026	**Lecture 03 continued** **Lecture 04 - Filtering & Edge detection** **Exercise session replaced by continuation of Lecture 4**	Scaramuzza
15.10.2026	**Lecture 05 - Point Feature Detectors, Part 1** **Exercise 03 - Harris detector + descriptor + matching**	Scaramuzza Rong, Roberto
22.10.2026	**Lecture 06 - Point Feature Detectors, Part 2** **Exercise 04 - SIFT detector + descriptor + matching**	Scaramuzza Ismail, Roberto
29.10.2026	**Lecture 07 - Multiple-view Geometry 1** **Exercise 05 - Stereo vision: rectification, epipolar matching, disparity, triangulation**	Scaramuzza Rong, Roberto
05.11.2026	**Lecture 08 - Multiple-view Geometry 2** **Exercise 06 - Eight-Point Algorithm**	Scaramuzza Rong, Roberto
12.11.2026	**Lecture 09 - Multiple-view Geometry 3** **Exercise 07 - P3P algorithm and RANSAC** **2nd hour of exercise session: seminar by Skydio Zurich (Dr. Davide Falanga)**	Leonard Bauersfeld Rong, Roberto
19.11.2026	**Lecture 10 - Multiple-view Geometry 4** **Continuation of Lecture 10 + Exercise 08 on Bundle Adjustment**	Scaramuzza Rong, Roberto
26.11.2026	**Lecture 11 - Optical Flow and KLT Tracking** **Exercise 09 - Lucas-Kanade tracker**	Scaramuzza Ismail, Roberto
03.12.2026	**1st lecture hour: seminar by Dr. Jeff Delaune from NASA-JPL: "Vision-Based Navigation for Mars Helicopters."** **Lecture 12a (2nd hour) - Place Recognition** **Lecture 12b (3rd and 4th hour, replaces exercise) - Deep Learning Tutorial** **Optional Exercise on Place Recognition**	Scaramuzza Rong, Roberto
10.12.2026	**Lecture 13 - Visual inertial fusion + Lab Open Doors Day at 5pm** **Exercise 10 – Visual Inertial fusion**	Scaramuzza Rong, Roberto
17.12.2026	**Lecture 14 - Event-based vision** **Exercise 11: Contrast Maximization for Event Cameras**	Scaramuzza Rong, Roberto

Study Material

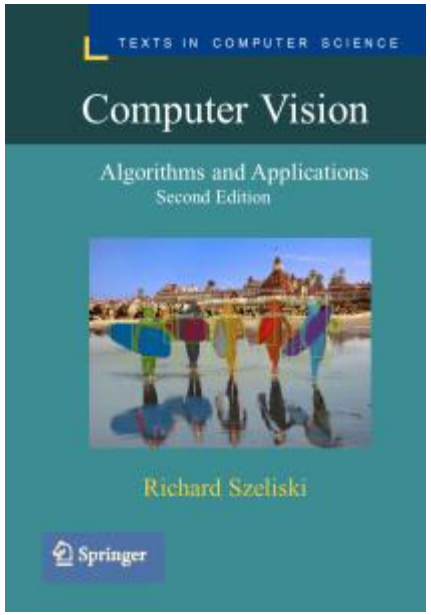
- **Schedule**, **lecture slides, exercise download, mini projects, course info** on the official course website**:** https://rpg.ifi.uzh.ch/teaching.html

- **Video Recordings** of lectures and exercises will be uploaded to **OLAT**: https://lms.uzh.ch/auth/RepositoryEntry/17913872651/CourseNode/112478900774303

- Post any **questions** related to lectures or exercises in the **OLAT Forum**

Reference Textbooks

- **Computer Vision: Algorithms and Applications**, **2nd Edition**, by Richard Szeliski. **Freely downloadable** from the author webpage: https://szeliski.org/Book/
- **Chapter 4 of Autonomous Mobile Robots,** 2nd edition, 2011, by R. Siegwart, I.R. Nourbakhsh, D. Scaramuzza. PDF
- **SLAM Handbook**: http://asrl.utias.utoronto.ca/~tdb/slam/
- **Additional readings** (i.e., **optional and not requested at the exam**) for interested students will be provided along with the slides and linked directly from the course website
- Further readings:
  - *Robotics, Vision and Control: Fundamental Algorithms*, 2nd edition, by Peter Corke
  - *An Invitation to 3D Vision*: Y. Ma, S. Soatto, J. Kosecka, S.S. Sastry
  - *Multiple view Geometry*: R. Hartley and A. Zisserman



Prerequisites

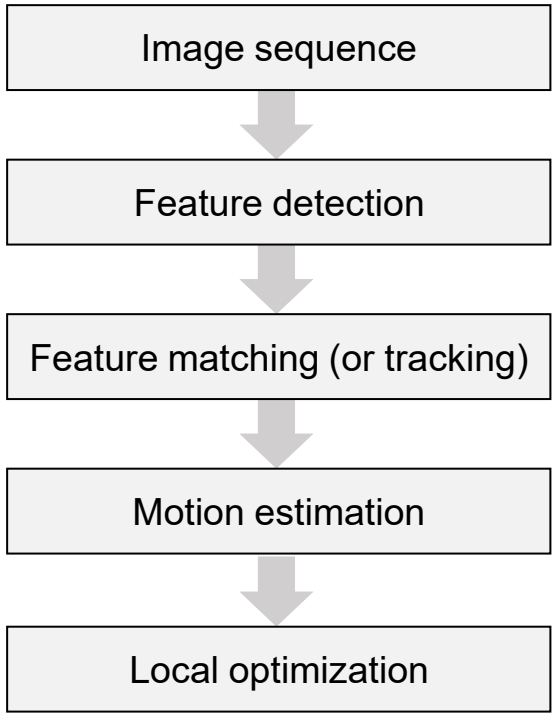
- **Linear algebra**
- **Matrix calculus**: matrix multiplication, inversion, singular value decomposition
  - Check out this Linear Algebra Primer from Stanford University
  - Check out this Immersive Linear Algebra interactive tool by Ström, Åström, and Akenine-Möller
  - Check out this tutorial on camera pose notation by Paul Furgale
- **No prior knowledge** of **computer vision** and **image processing** is required

Learning Objectives

- **High-level goal:** learn to implement the visual-inertial odometry algorithms used in current mobile robots (drones, cars, planetary robots), AR/VR products (Meta Quest and Aria glasses, Microsoft HoloLens, Magic Leap), and Google Visual Positioning Service (e.g., Google Map Live View).
- You will also learn **to implement the fundamental computer vision algorithms** used in mobile robotics, in particular:
  - image formation,
  - filtering,
  - feature extraction,
  - multiple view geometry,
  - dense reconstruction,
  - feature and template tracking,
  - image retrieval,
  - event-based vision,
  - visual-inertial odometry, Simultaneous Localization And Mapping (SLAM),
  - and some basics of deep learning.

Exercises

- **Learning Goal** of the exercises: **Implement a full visual odometry pipeline** (like the one running on Mars rovers).
- **Each week** you will learn how to implement a **building block** of visual odometry, theory (lecture) and practice (exercise).
- **The lectures and exercises are synchronized**
- Two exercises will be dedicated to **system integration**.
- **NB: Questions about the implementation details of each exercise may be asked at the exam.**

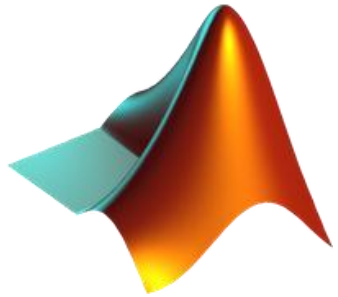


Building blocks of visual odometry along with information flow

Exercises

- **Bring your own laptop**
- Exercises in **Python or Matlab**. You will need to have Matlab or Python already pre-installed on your machine for the exercises.
- Python can be downloaded from here.
- You can download Matlab from:
  - **ETH:** Download: https://itshop.ethz.ch/EndUser/Items/Home
  - **UZH:** Download: https://www.zi.uzh.ch/de/students/software-elearning/campussoft.html
  - An introductory tutorial on Matlab can be found here: http://rpg.ifi.uzh.ch/docs/teaching/2026/MatlabPrimer.pdf
  - **Please install all the toolboxes included in the license.** If you don't have enough space in your PC, then install at least the Image Processing, Computer Vision, and Optimization toolboxes





Grading and Exam

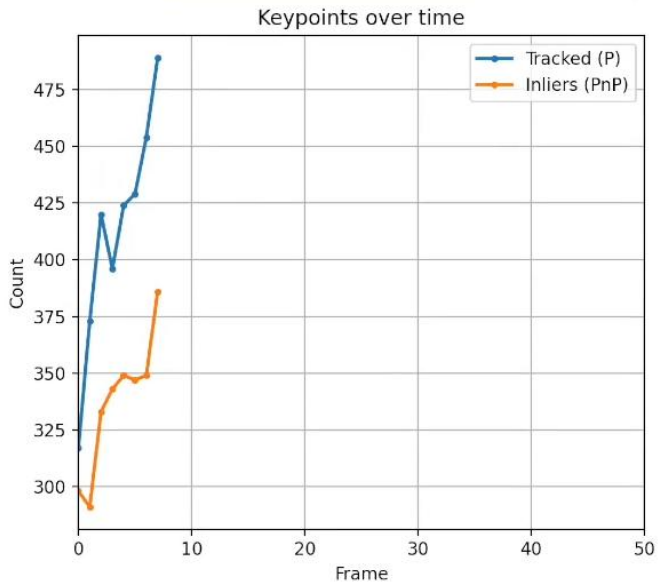
- The **final grade is based on a written exam** (90 minutes)
  - Exam date: **January 7, 2027, from 08:00 to 10:00 on site**
  - **Closed-book exam**
  - **All the exams of the past years can be found on OLAT**

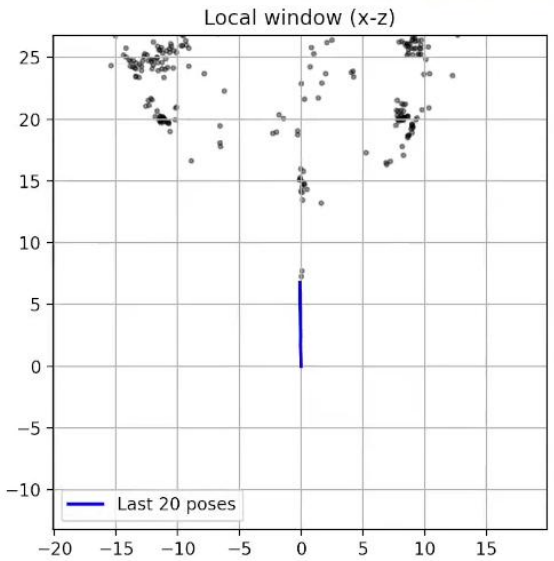
- **Optional mini project:**
  - You have the **option** (i.e., not mandatory) to do a **mini project**, which consists of implementing a working visual odometry algorithm in **Matlab** or **Python** (but C++ is also accepted)
  - If the algorithm runs smoothly, producing a reasonable result, you will be rewarded with an **up to 0.5 grade increase on the final grade**. However, notice that the mini project can be very time-consuming!
  - The **deadline** to hand in the mini project is **04.01.2027 at 23:59 hrs**.
  - **Group work: minimum 2, max 4 people.**
  - **The best mini project team(s) will get an award certificate** (useful for job applications)

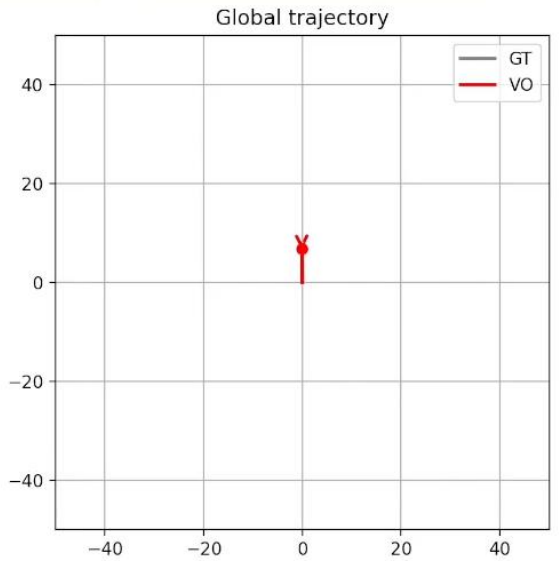


Outcome of last year miniprojects









Class Participation

- **Strong class participation is encouraged!**

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

Today's Outline

- About me and my research lab
- What is Computer Vision?
- Why study computer vision?
- Examples of Vision Applications
- Organization of the course
- Start: Visual Odometry overview

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



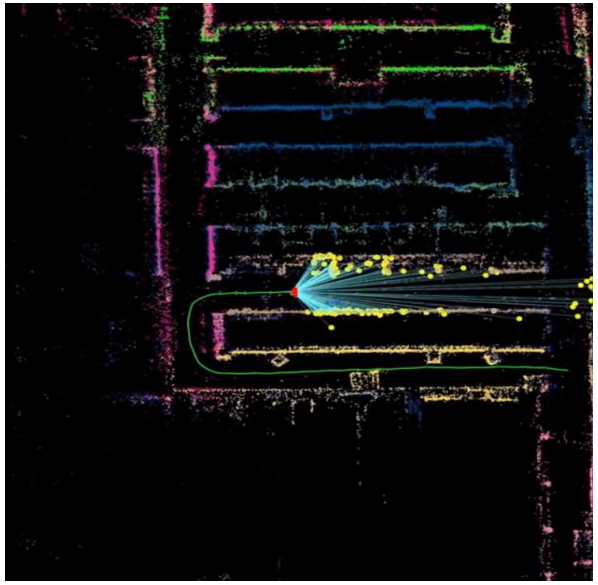


input



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

output



Camera trajectory (3D structure is a plus)

$R_0, R_1, \ldots, R_i$

$t_0, t_1, \ldots, t_i$

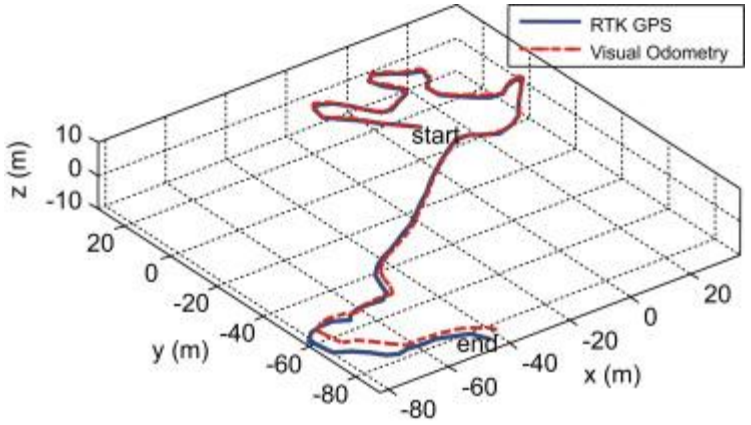
Why VO?

- VO is crucial for **flying**, **walking**, and **underwater** robots
- Contrary to wheel odometry, VO is **not affected by wheel slippage** (e.g., on sand or wet floor)
- Very accurate:
  relative position error is 0.1% – 2% of the travelled distance
- VO can be used as a complement to
  - wheel encoders (wheel odometry)
  - GPS (when GPS is degraded)
  - Inertial Measurement Units (IMUs)
  - laser odometry









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?



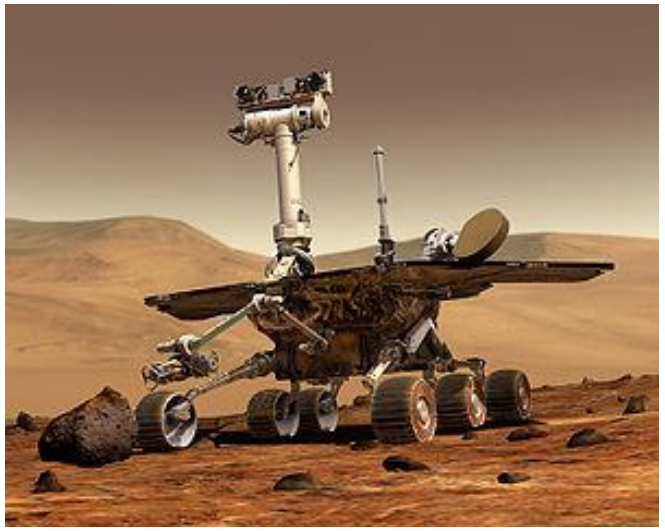
A Brief history of VO

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

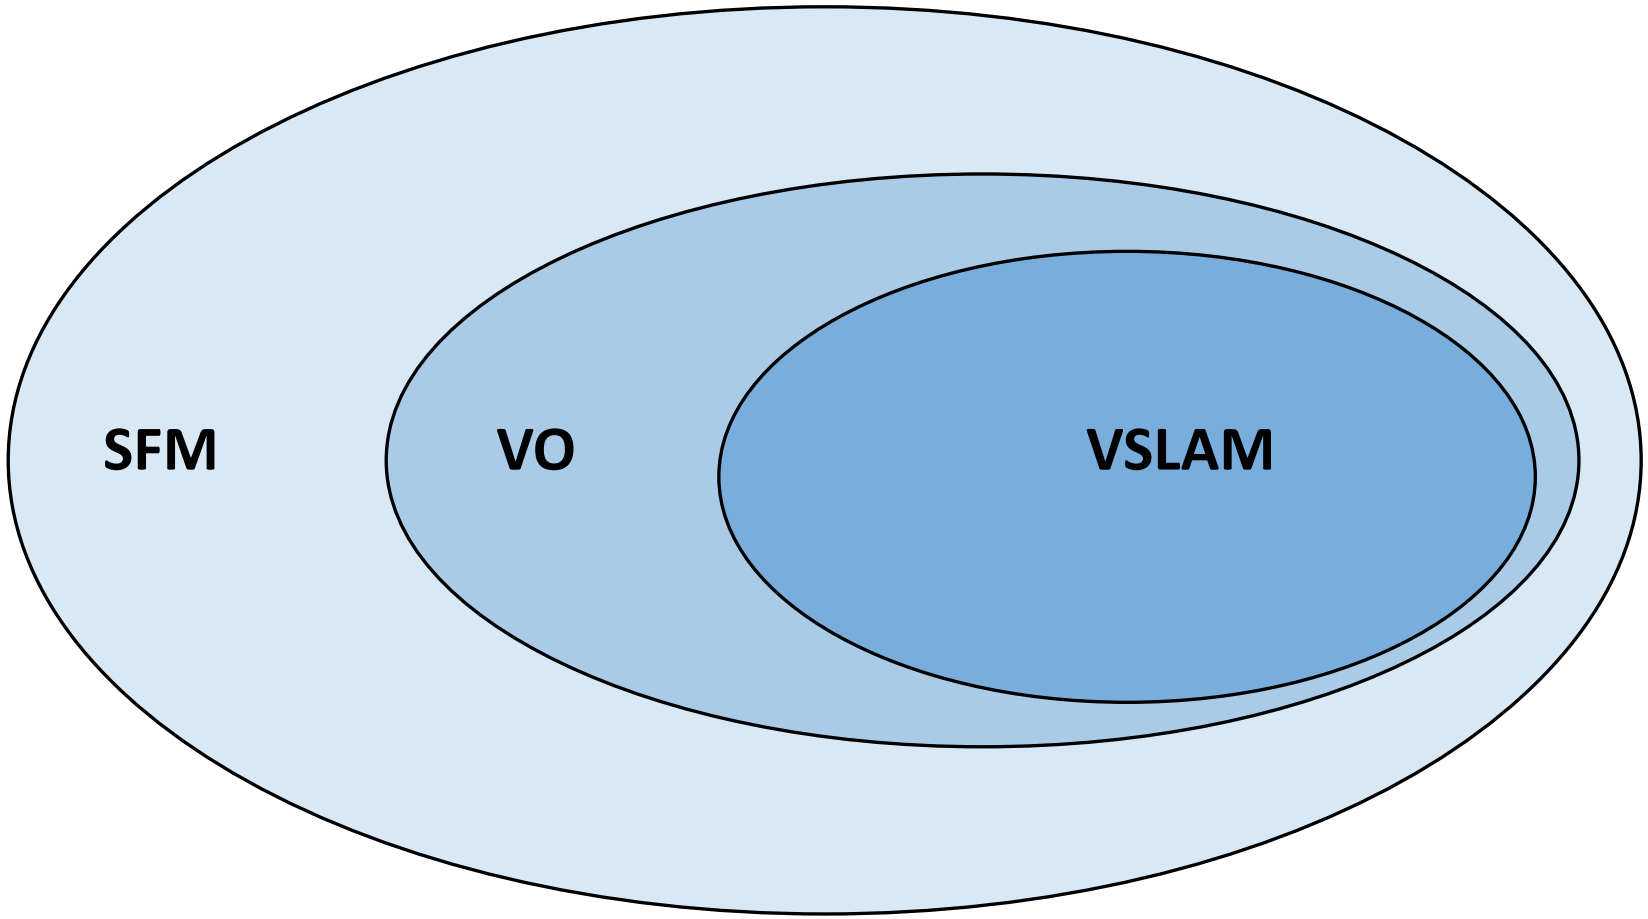


A Brief history of VO

- **1980**: First known VO real-time implementation on a robot by **Hans Moravec**'s 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.
- **2015-today**: VO becomes a **fundamental tool of several products**: VR/AR, drones, smartphones
- **2021**: VO is used on the **Mars helicopter**

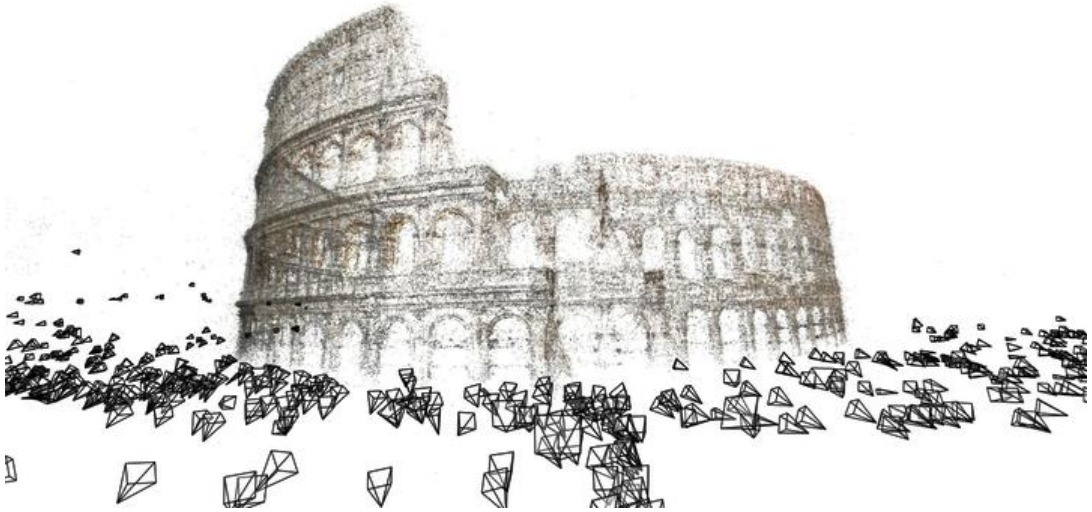


VO vs VSLAM vs SFM



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 on a cluster of 250 computers, 24 hours of computation
Paper: "Building Rome in a Day", ICCV'09.
State of the art software: COLMAP

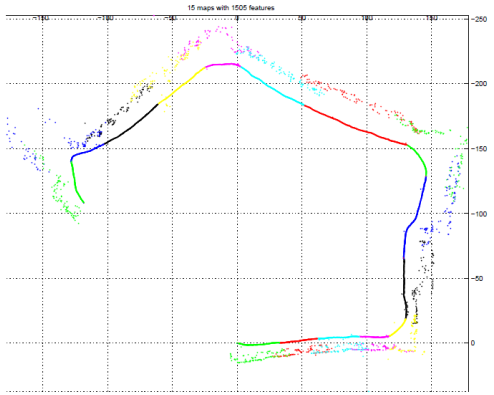
VO vs SFM

- 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**

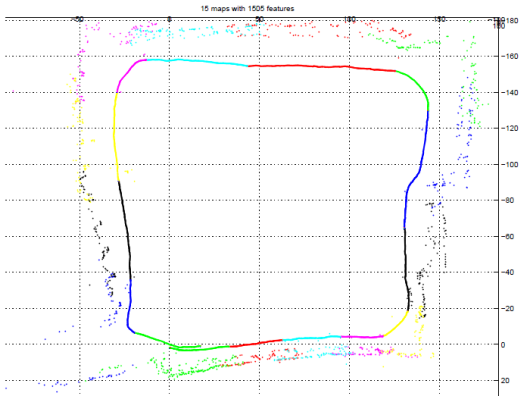
VO vs. Visual SLAM

- **Visual Odometry**
  - Focuses on incremental estimation
  - **Guarantees local consistency** (i.e., estimated trajectory is locally correct, but not globally, i.e. from the start to the end)

- **Visual SLAM** (Simultaneous Localization And Mapping)
  - **It's a particular case of VO**
  - **SLAM = visual odometry + loop detection & closure**
  - **Aims to maintain global consistency** (the estimated trajectory is globally correct, i.e. from the start to the end)



Visual odometry



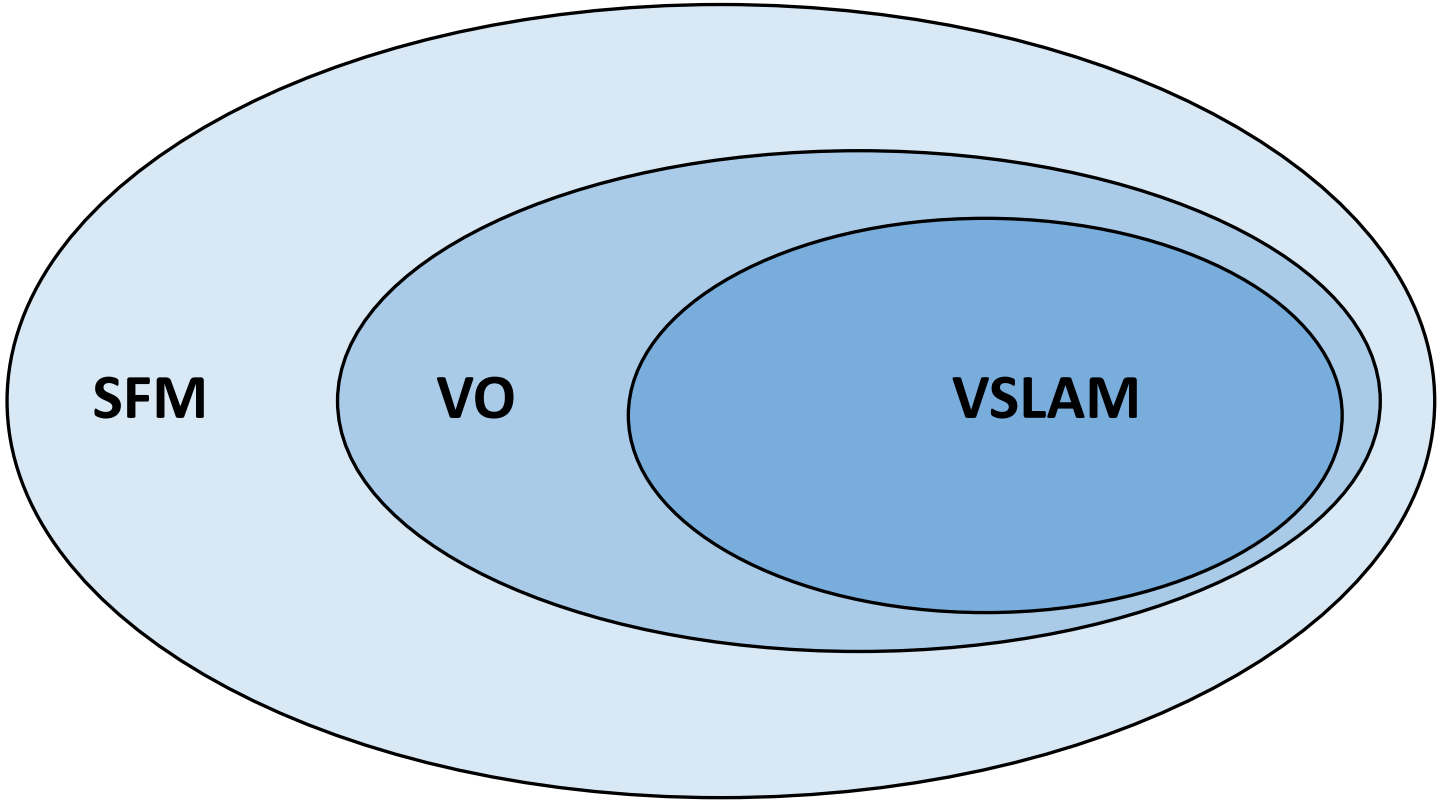
Visual SLAM

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

VSLAM ⊆ VO ⊆ SFM

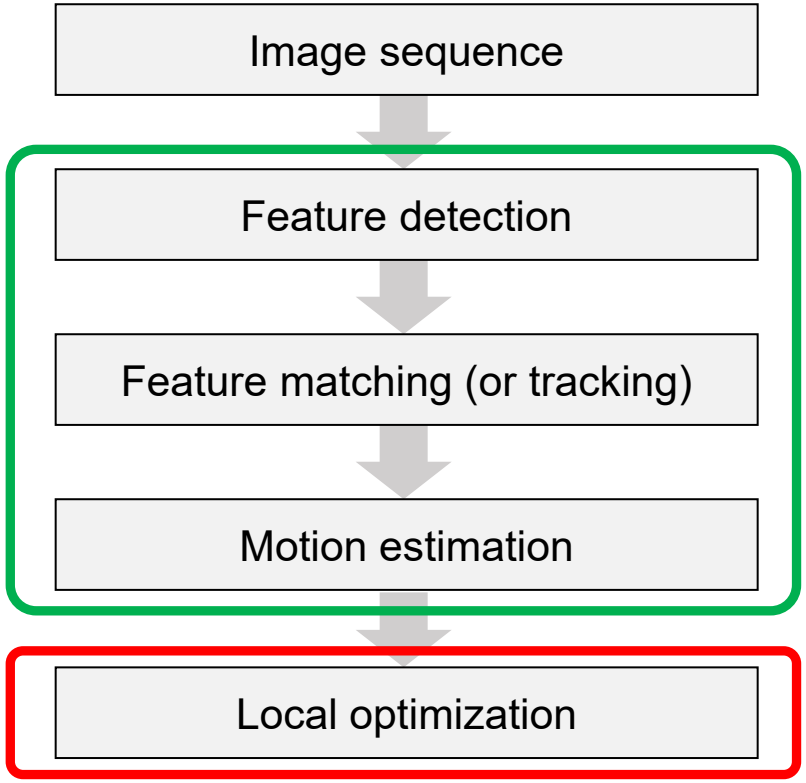
Why?

- **because every VSLAM and VO are SFM**, but not every SFM is VO or SLAM
- **because every VSLAM is a VO**, but not every VO is a SLAM
  - VSLAM applies more stringent requirements, such as loop detection and closure, than VO, making it a particular case. Every VSLAM functions as a VO, given that VSLAM, like VO, incrementally estimates poses (although Bundle Adjustment may further refine these estimations). Moreover, if VSLAM achieves global consistency, it inherently ensures local consistency as well.



VO Flow Chart

VO computes the camera path incrementally (pose after pose)

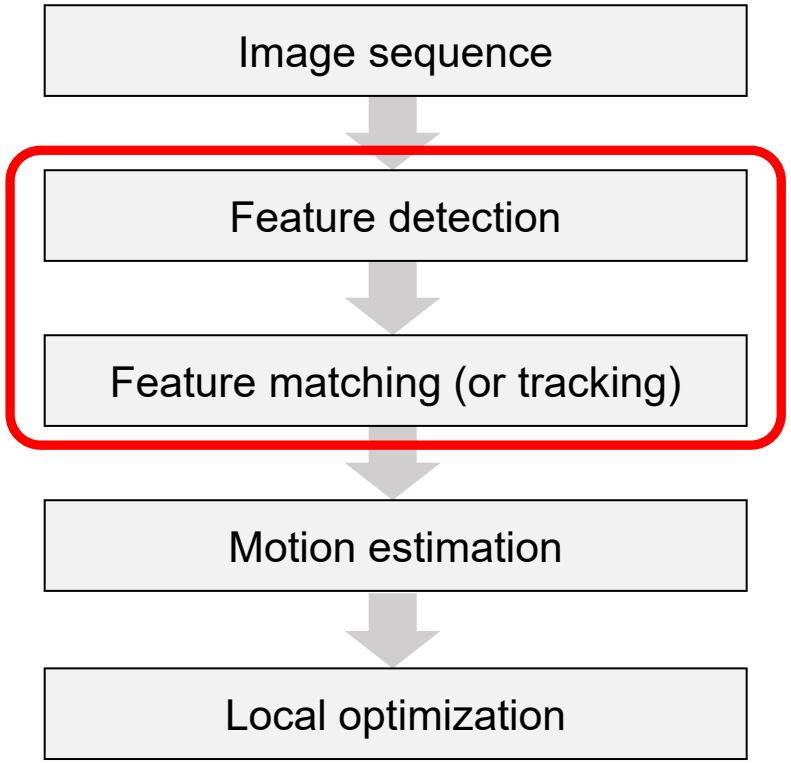


Front-end: outputs the *relative pose* between the *last two frames*

Back-end: "*adjusts*" the relative poses among *multiple recent frames*

VO Flow Chart

VO computes the camera path incrementally (pose after pose)

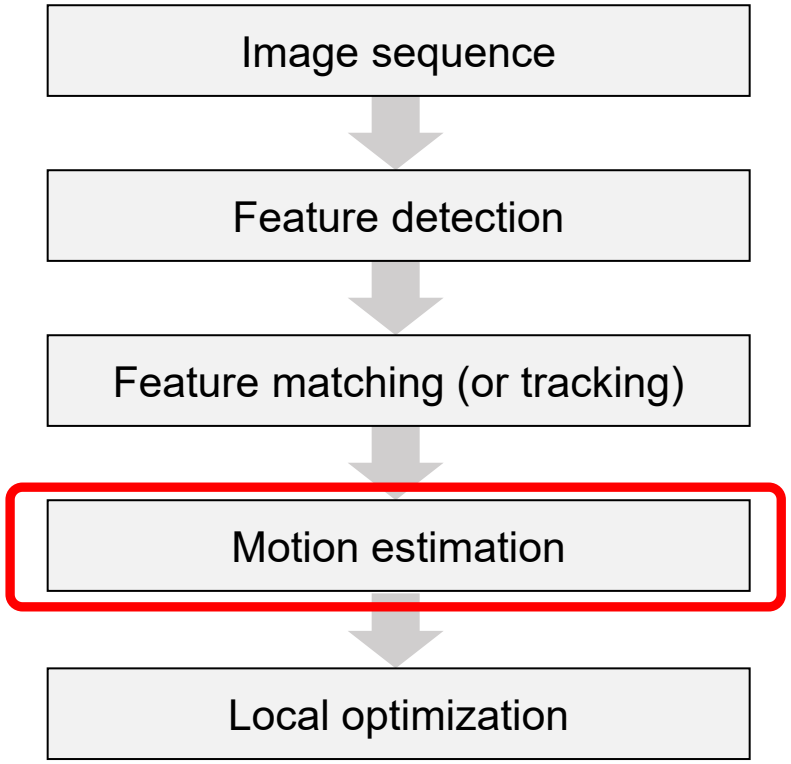


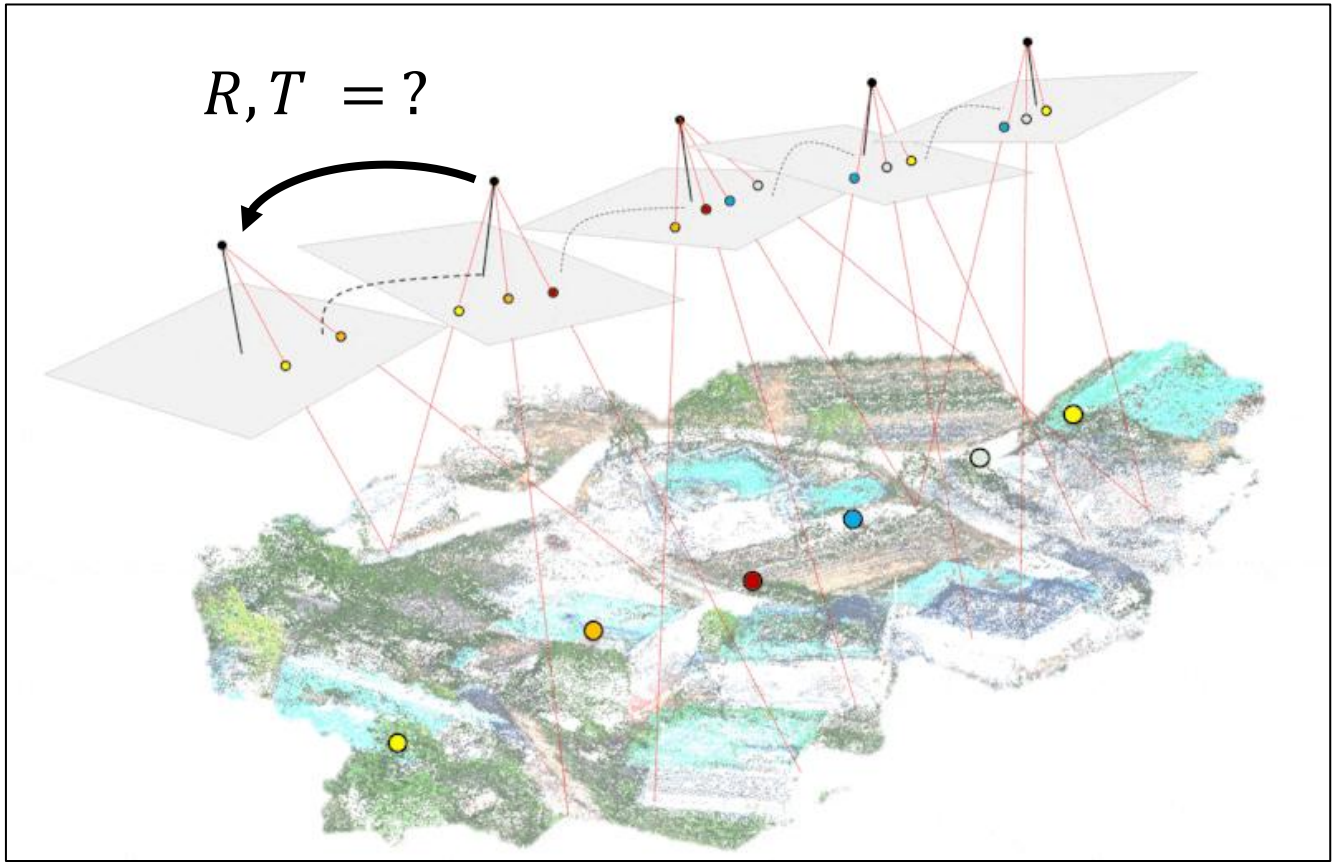


Features tracked over multiple recent frames overlaid on the last frame

VO Flow Chart

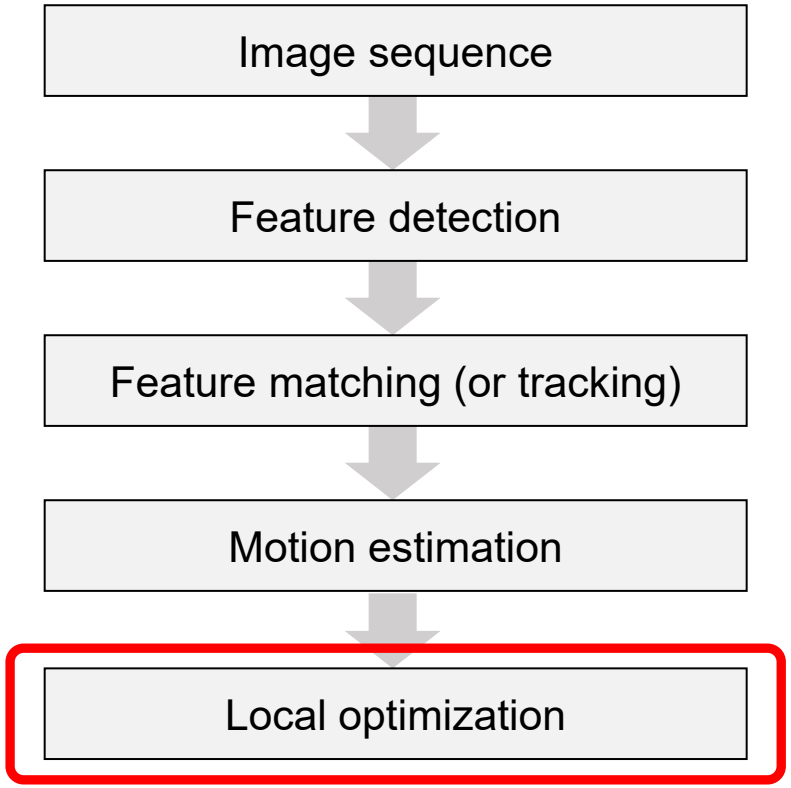
VO computes the camera path incrementally (pose after pose)

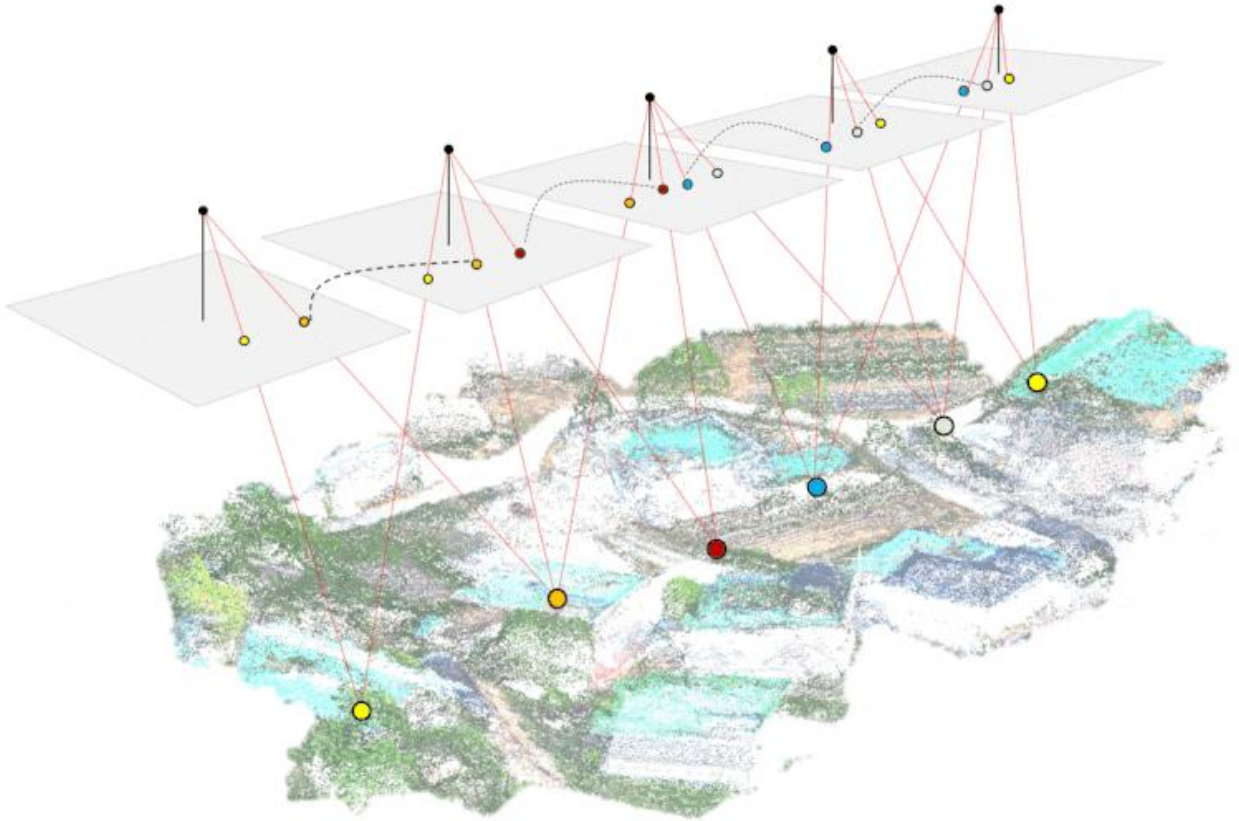




VO Flow Chart

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
- Deep learning
- Visual inertial fusion
- Event-based Vision

Course Topics

- Principles of image formation
  - Perspective projection
  - Camera calibration

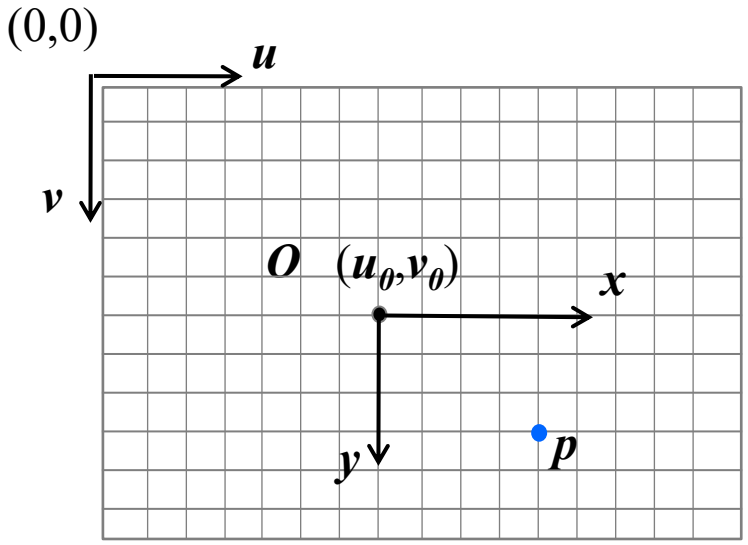


Image plane

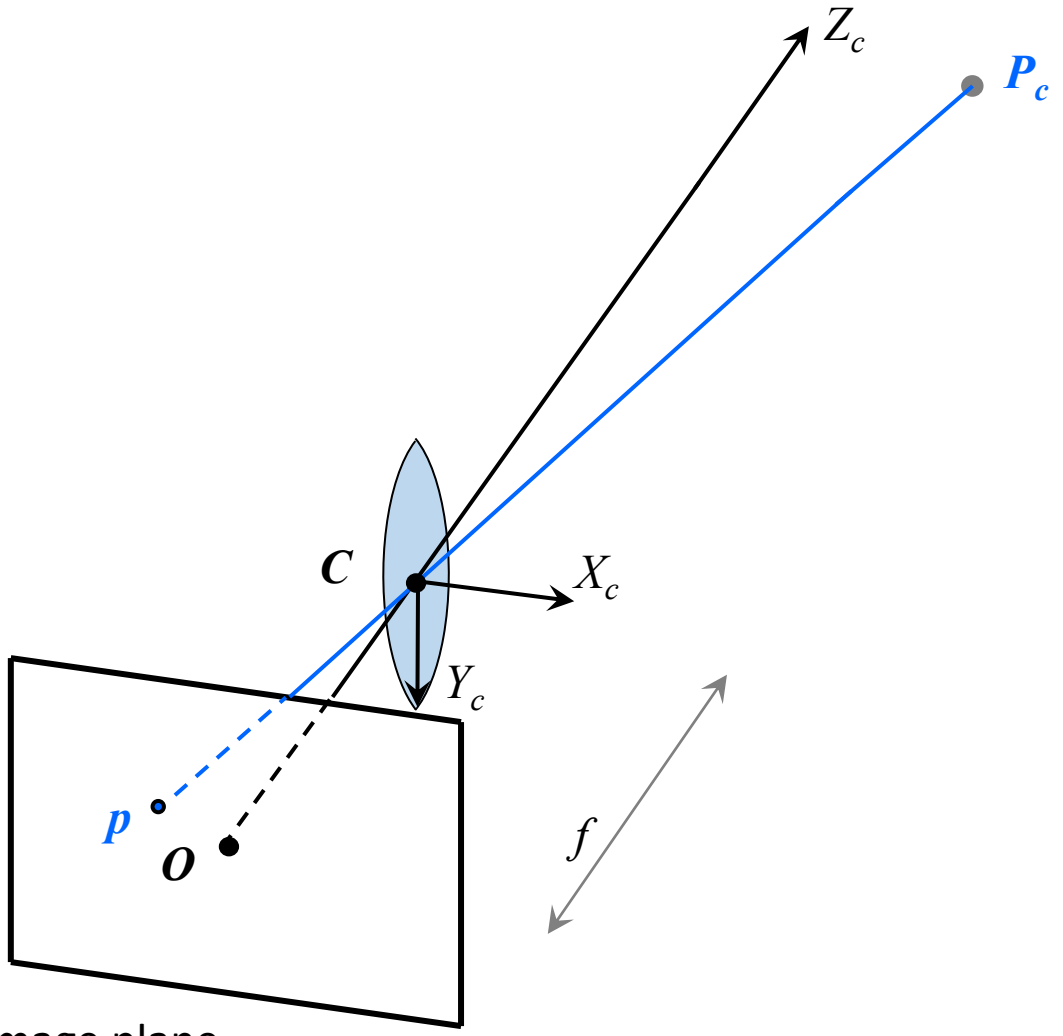
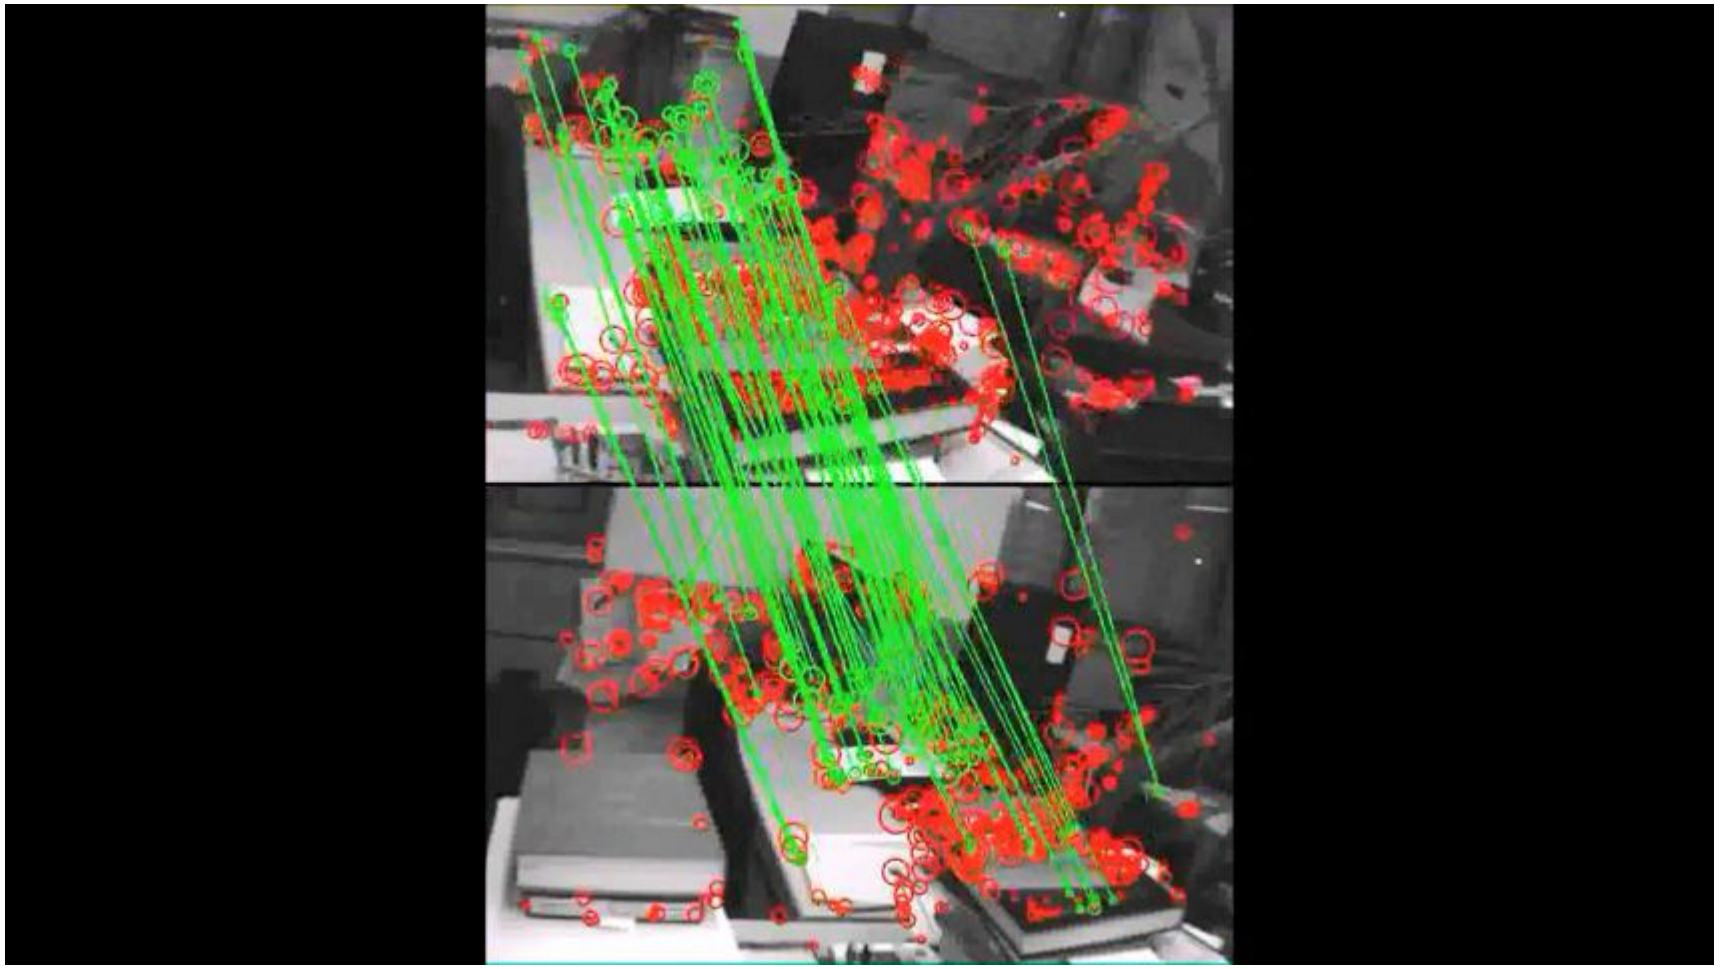


Image plane

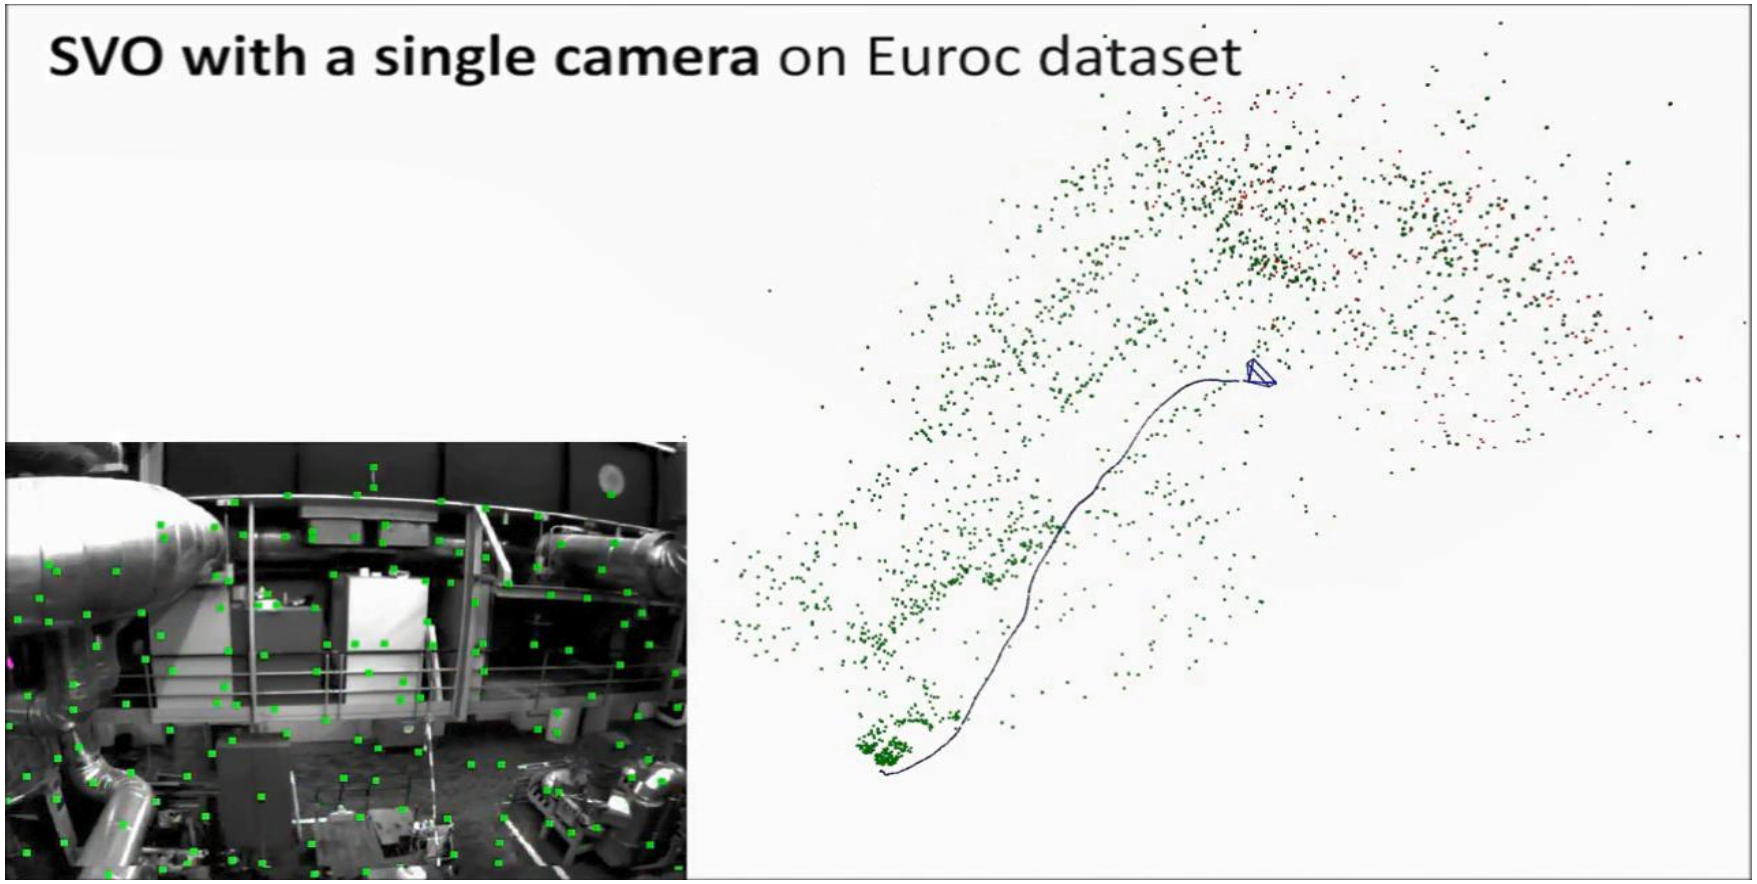
Course Topics

- Feature detection and matching



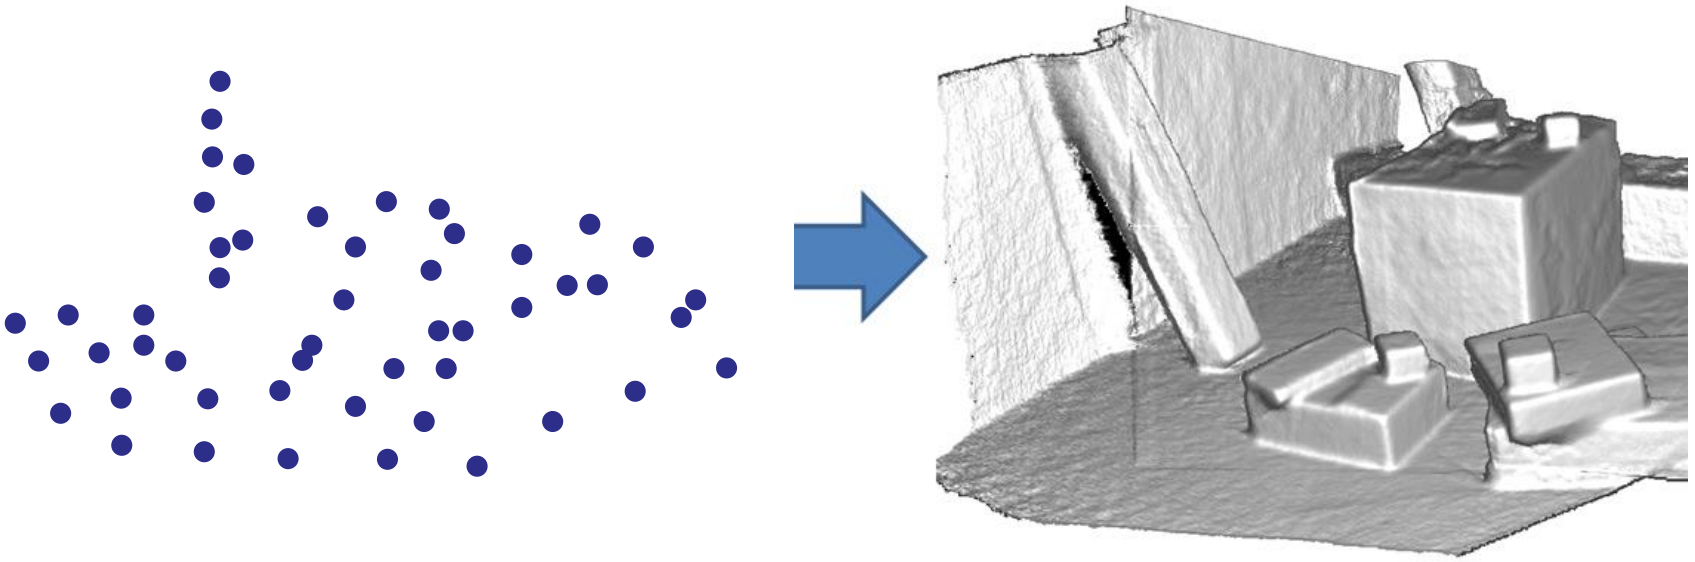
Course Topics

- Multi-view geometry and sparse 3D reconstruction



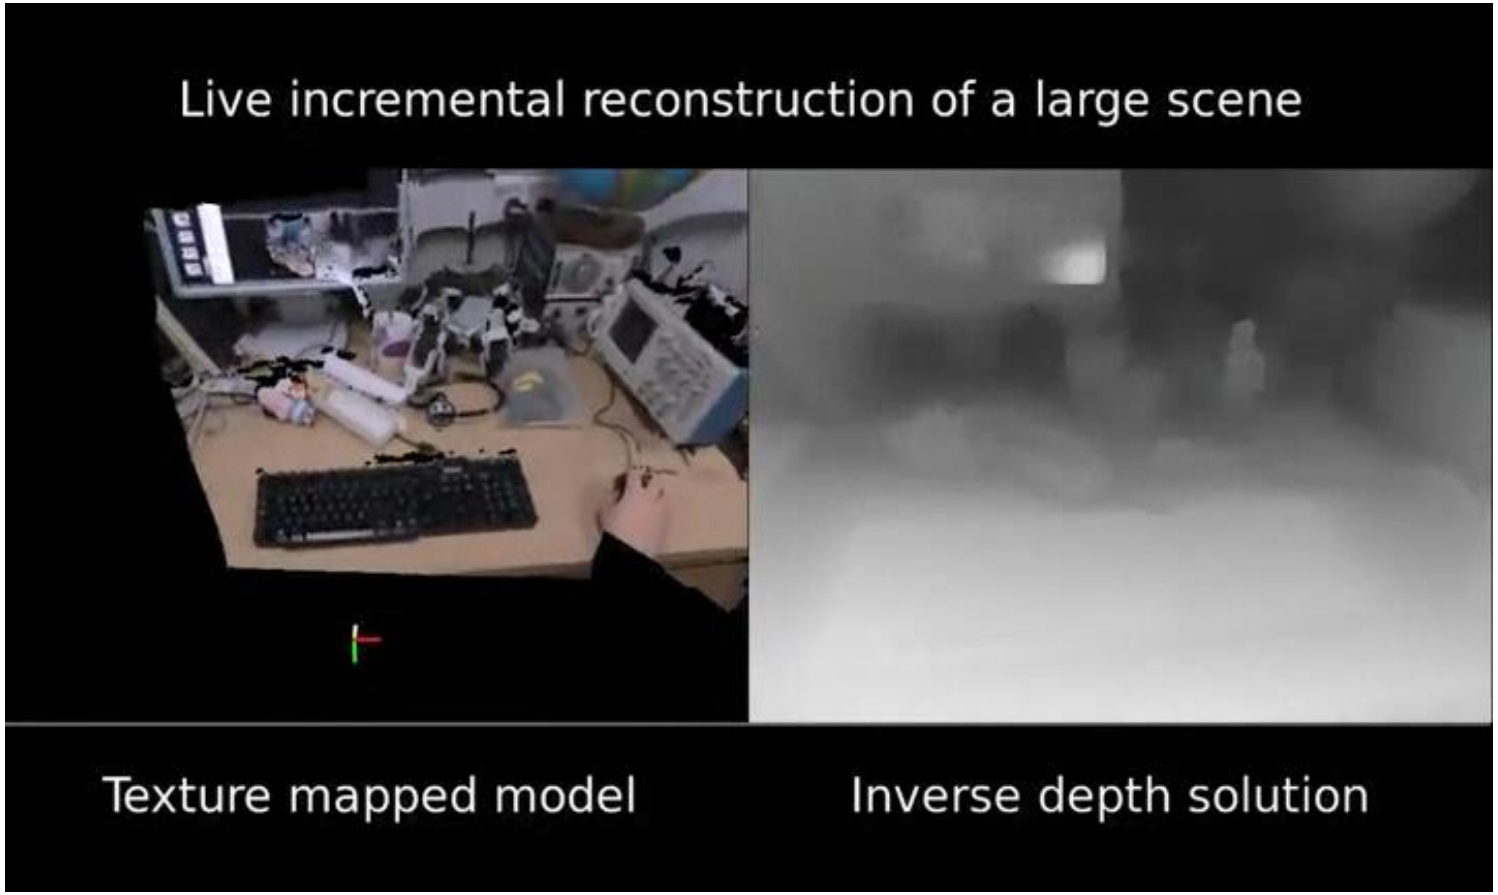
Course Topics

- Dense 3D reconstruction



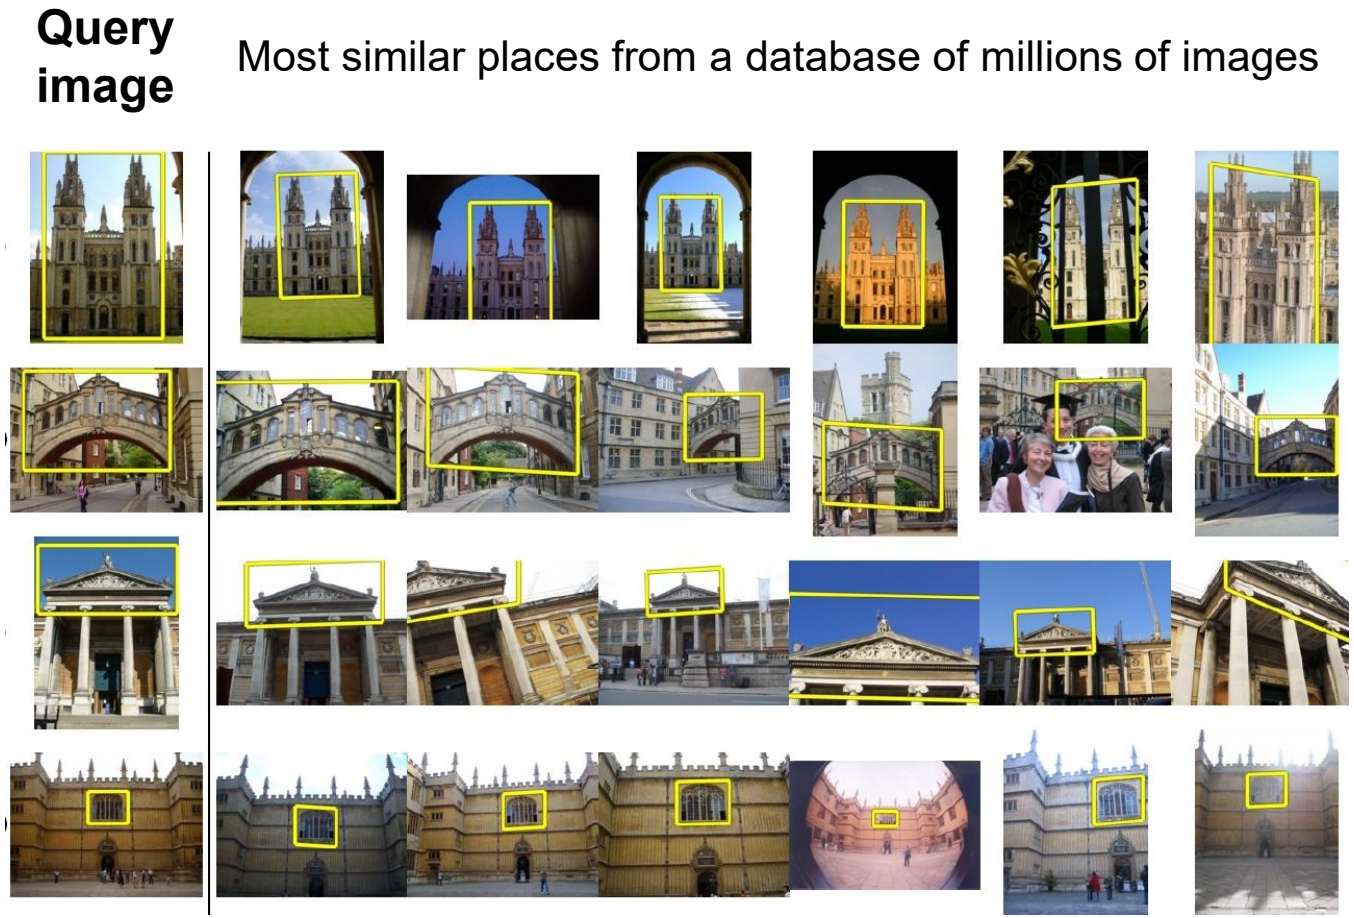
Course Topics

- Dense 3D reconstruction



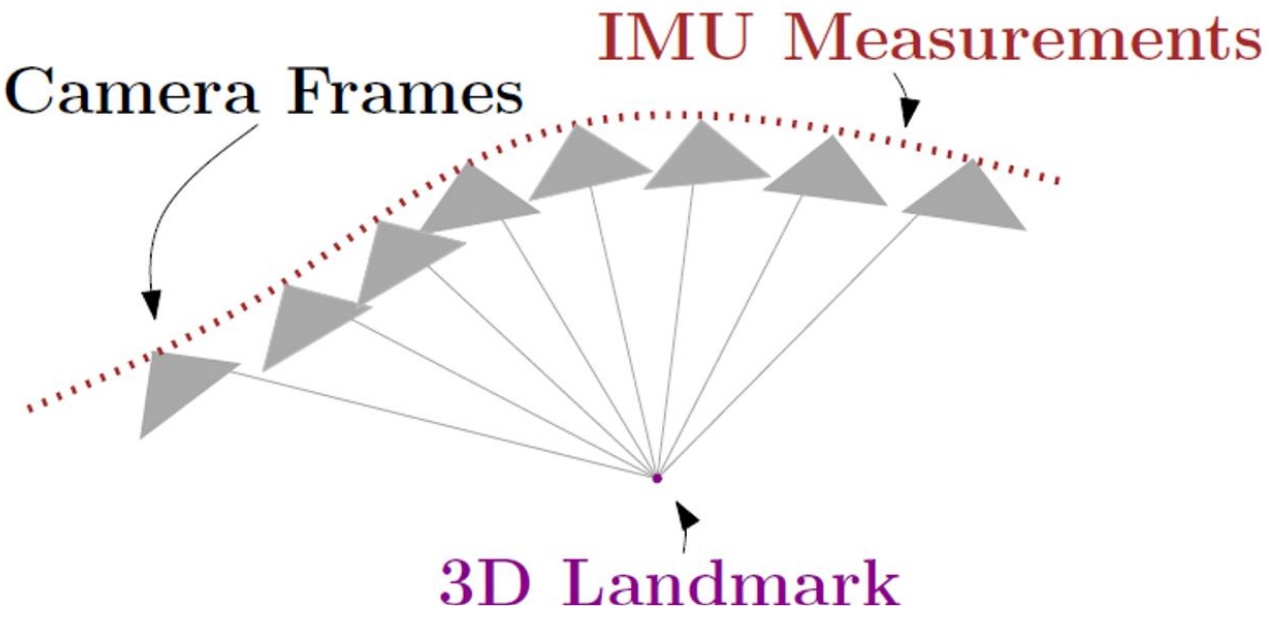
Course Topics

- Place recognition and deep learning



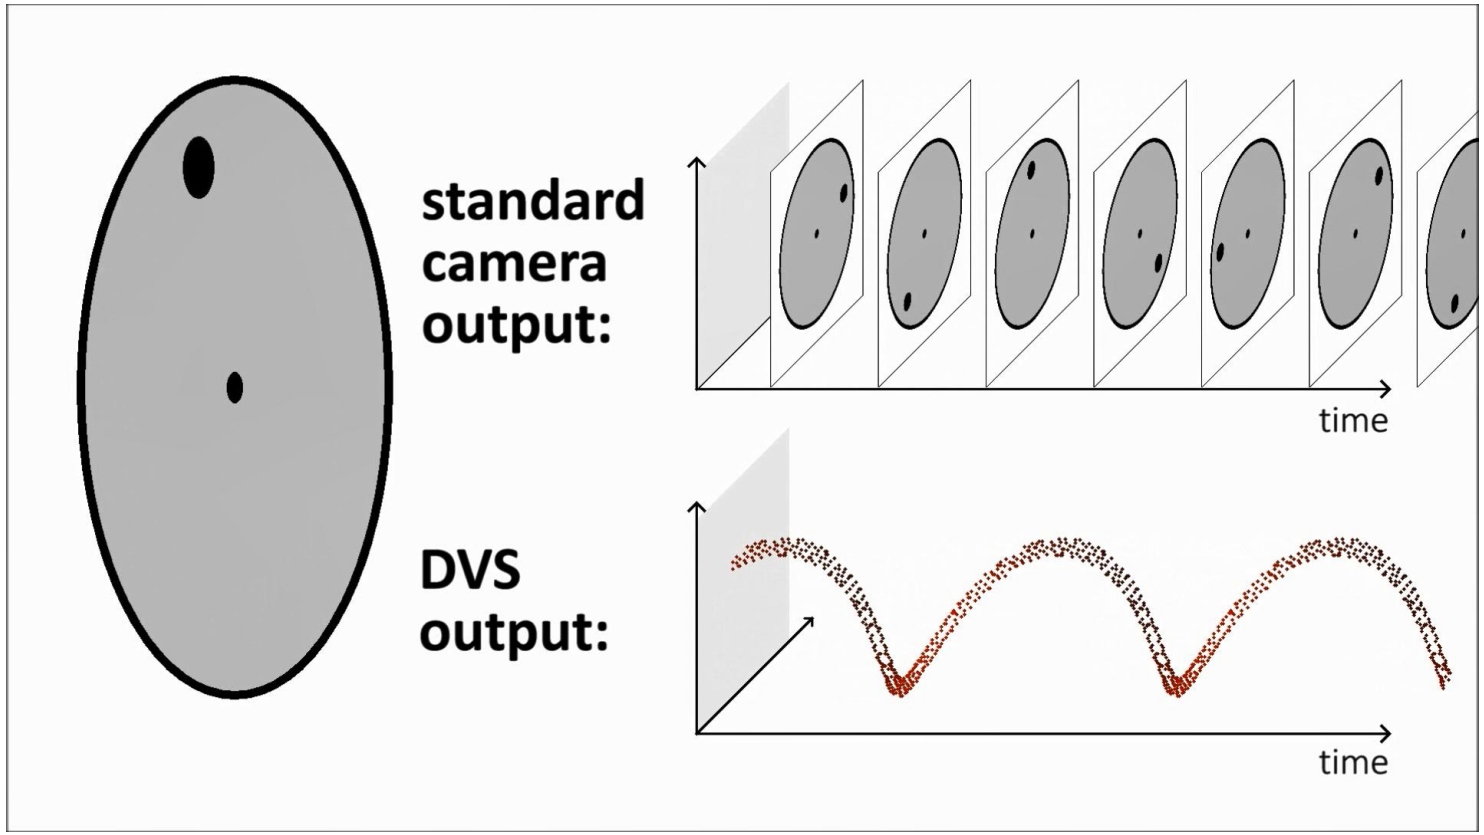
Course Topics

- Visual-inertial fusion

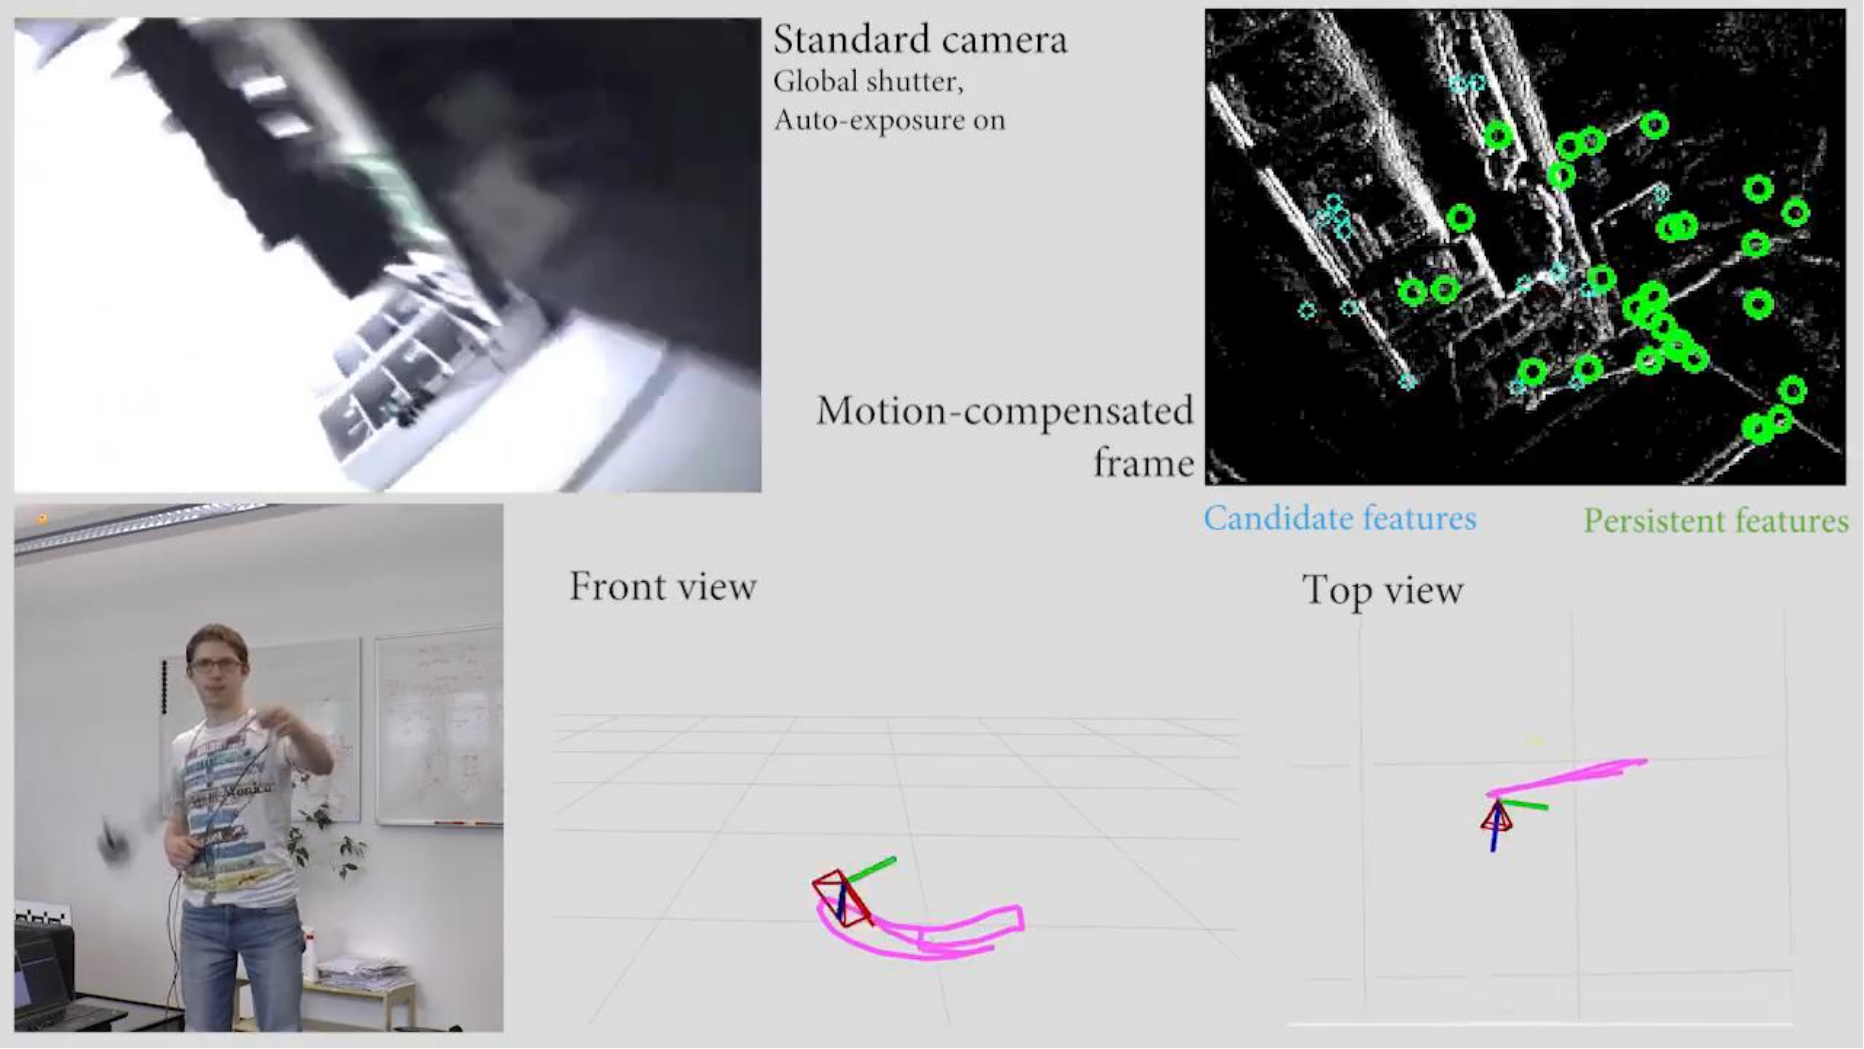


Course Topics

- Event cameras

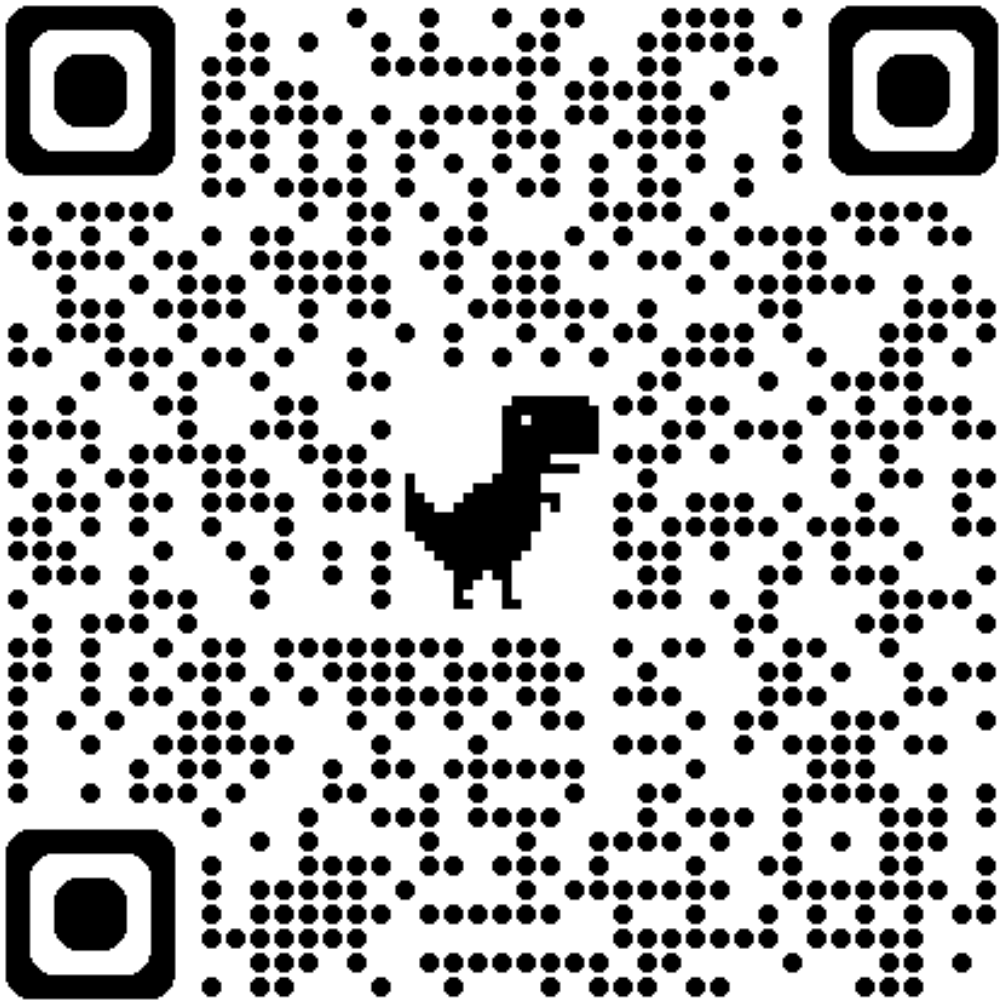


Application of event cameras: high-speed VO



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

1-Minute Understanding Quiz



Reading

- 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

- **Chapter 1 of SLAM Handbook**: PDF

Understanding Check

Are you able to:

- Provide a definition of Visual Odometry?
- Explain the most important differences between VO, VSLAM, and SFM?
- What assumptions does VO rely on?
- Illustrate the flow chart of VO?