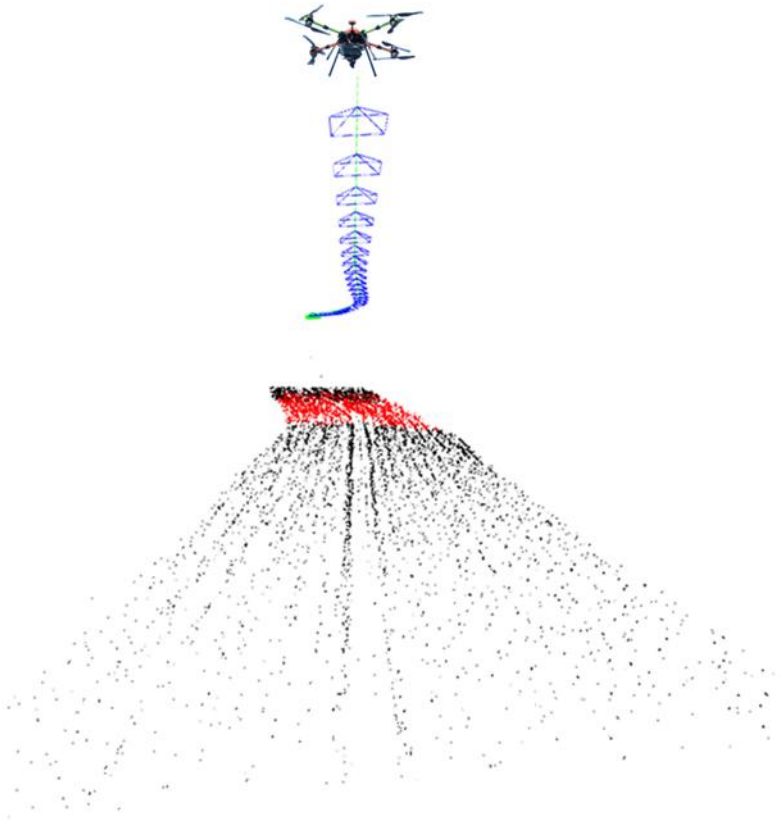
UiO : Department of Technology Systems
University of Oslo

Introduction to TEK5030 – Computer Vision

21.01.2026

Trym Vegard Haavardsholm (trym.haavardsholm@its.uio.no)
Idar Dyrdal (idar.dyrdal@its.uio.no)
Thomas Opsahl (thomasoo@its.uio.no)



Today

- A quick introduction to computer vision
- About the course

- Lab: An introduction to OpenCV
  - Processing live video streams with OpenCV!



Illustration: Generated by Midjourney

«An image is worth more than a thousand words»



«An image is worth more than a thousand words»

How can we extract
this information?



Images and pixels



Images and pixels



Images and pixels



Images and pixels



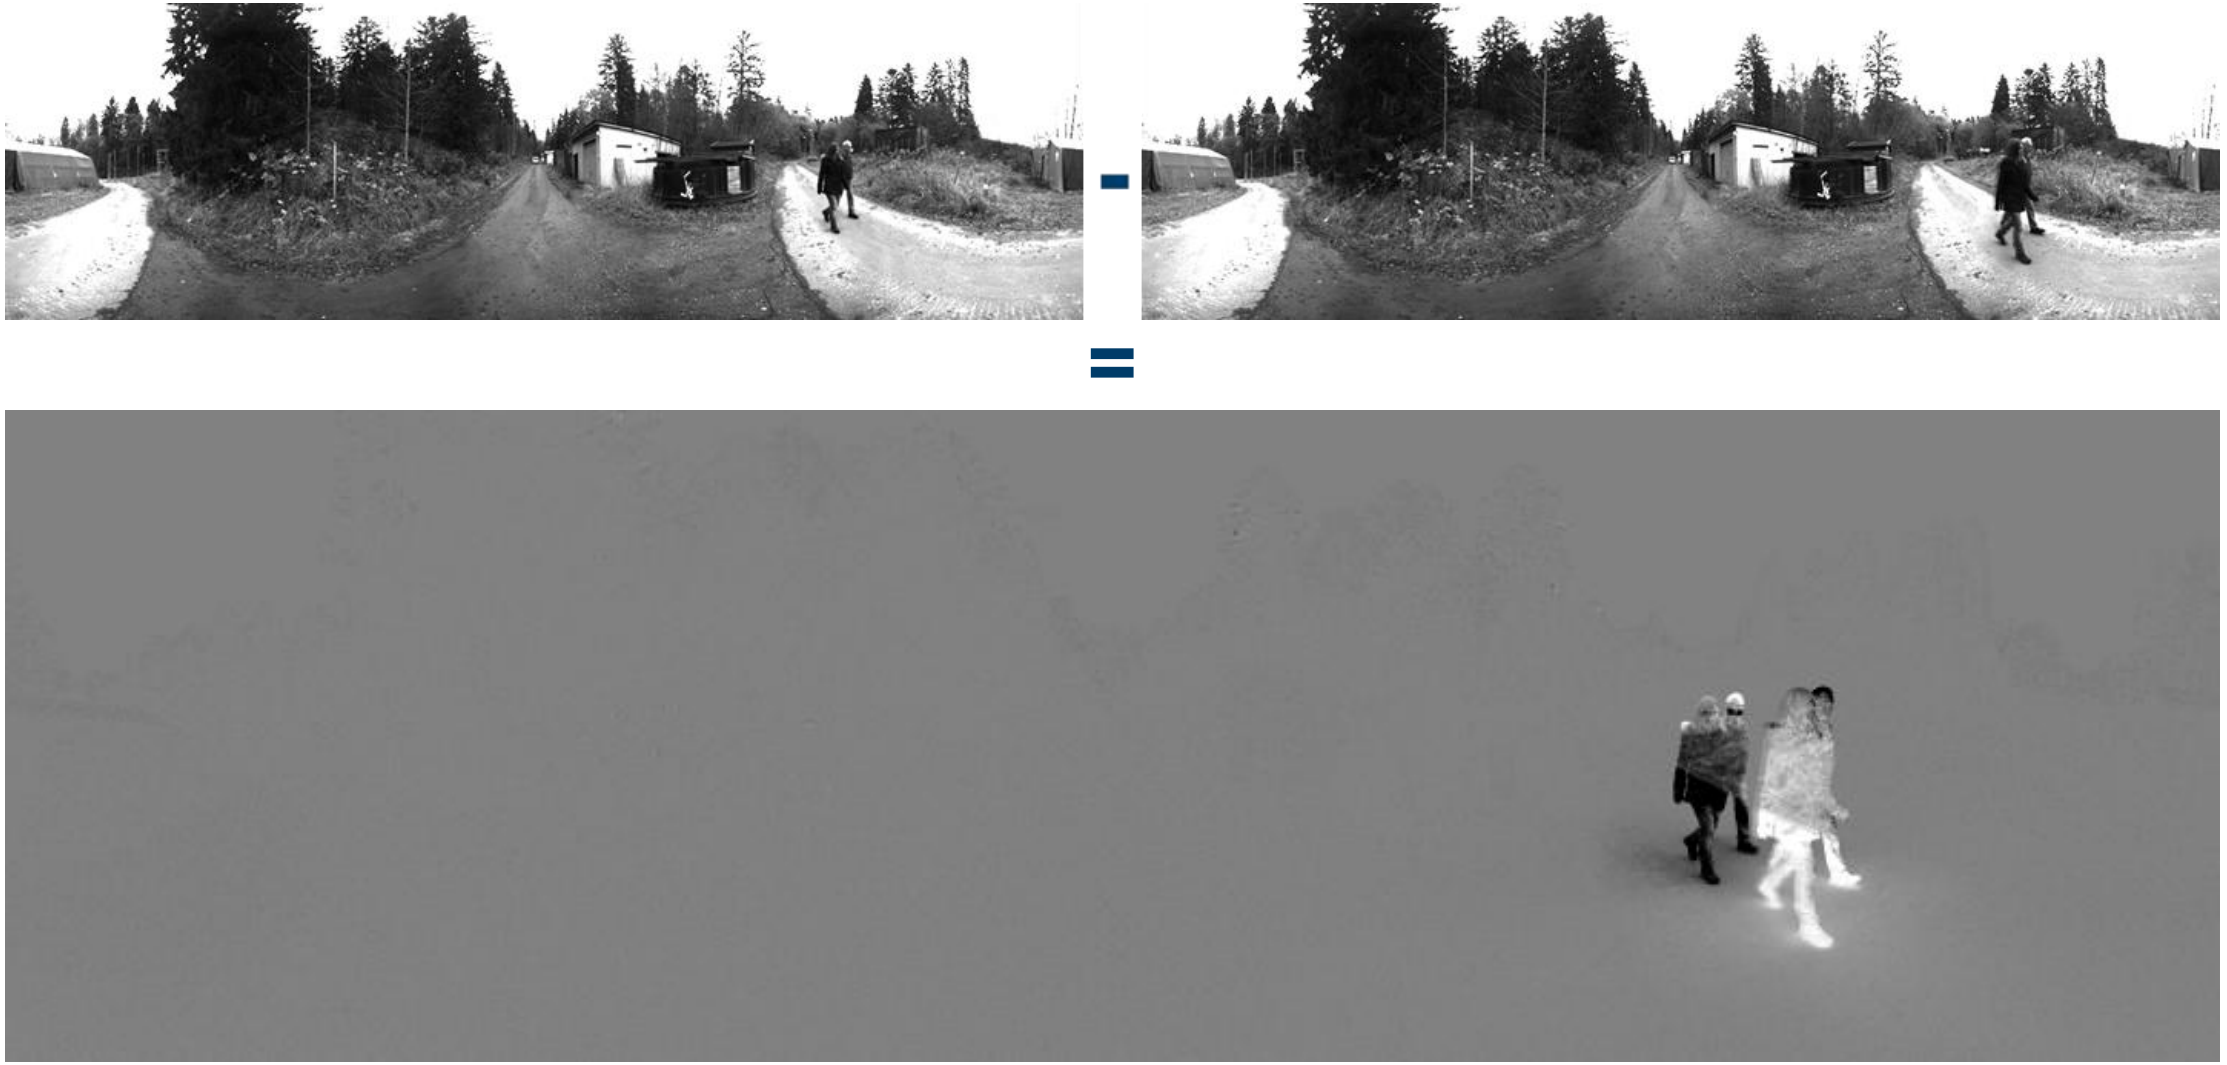
Images and pixels



It is easy to compute with images!



It is easy to compute with images!



How fast can a machine crunch pixel values?



Illustration: Generated by Midjourney

How fast can a machine crunch pixel values?

How many operations per second?

- Thousands?
- Millions?
- Billions?



Illustration: Generated by Midjourney

How fast can a machine crunch pixel values?

How many operations per second?

- 50 billion?
- 200 billion?
- 1000 billion?
- 3000 billion?
- 1000000000 billion?



Illustration: Generated by Midjourney

Floating point operations (FP32) per second (FLOPS)



Photo: Miiicihiaieil Hieinizilieir (CC BY-SA 4.0)

Raspberry PI 4B
48 GFLOPS

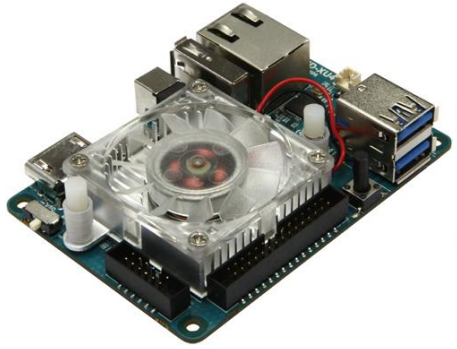


Photo: Alexandre.willame (CC BY-SA 4.0)

Odroid XU-4
214 GFLOPS

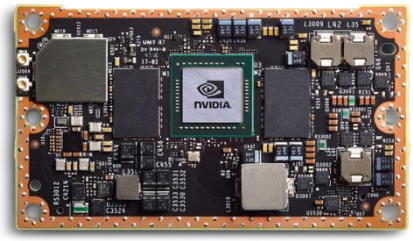


Photo: NVIDIA

Jetson TX2
782 GFLOPS



Photo: HP

My laptop
3 TFLOPS

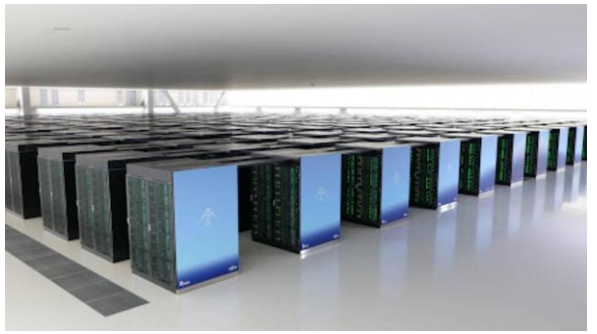


Photo: RIKEN

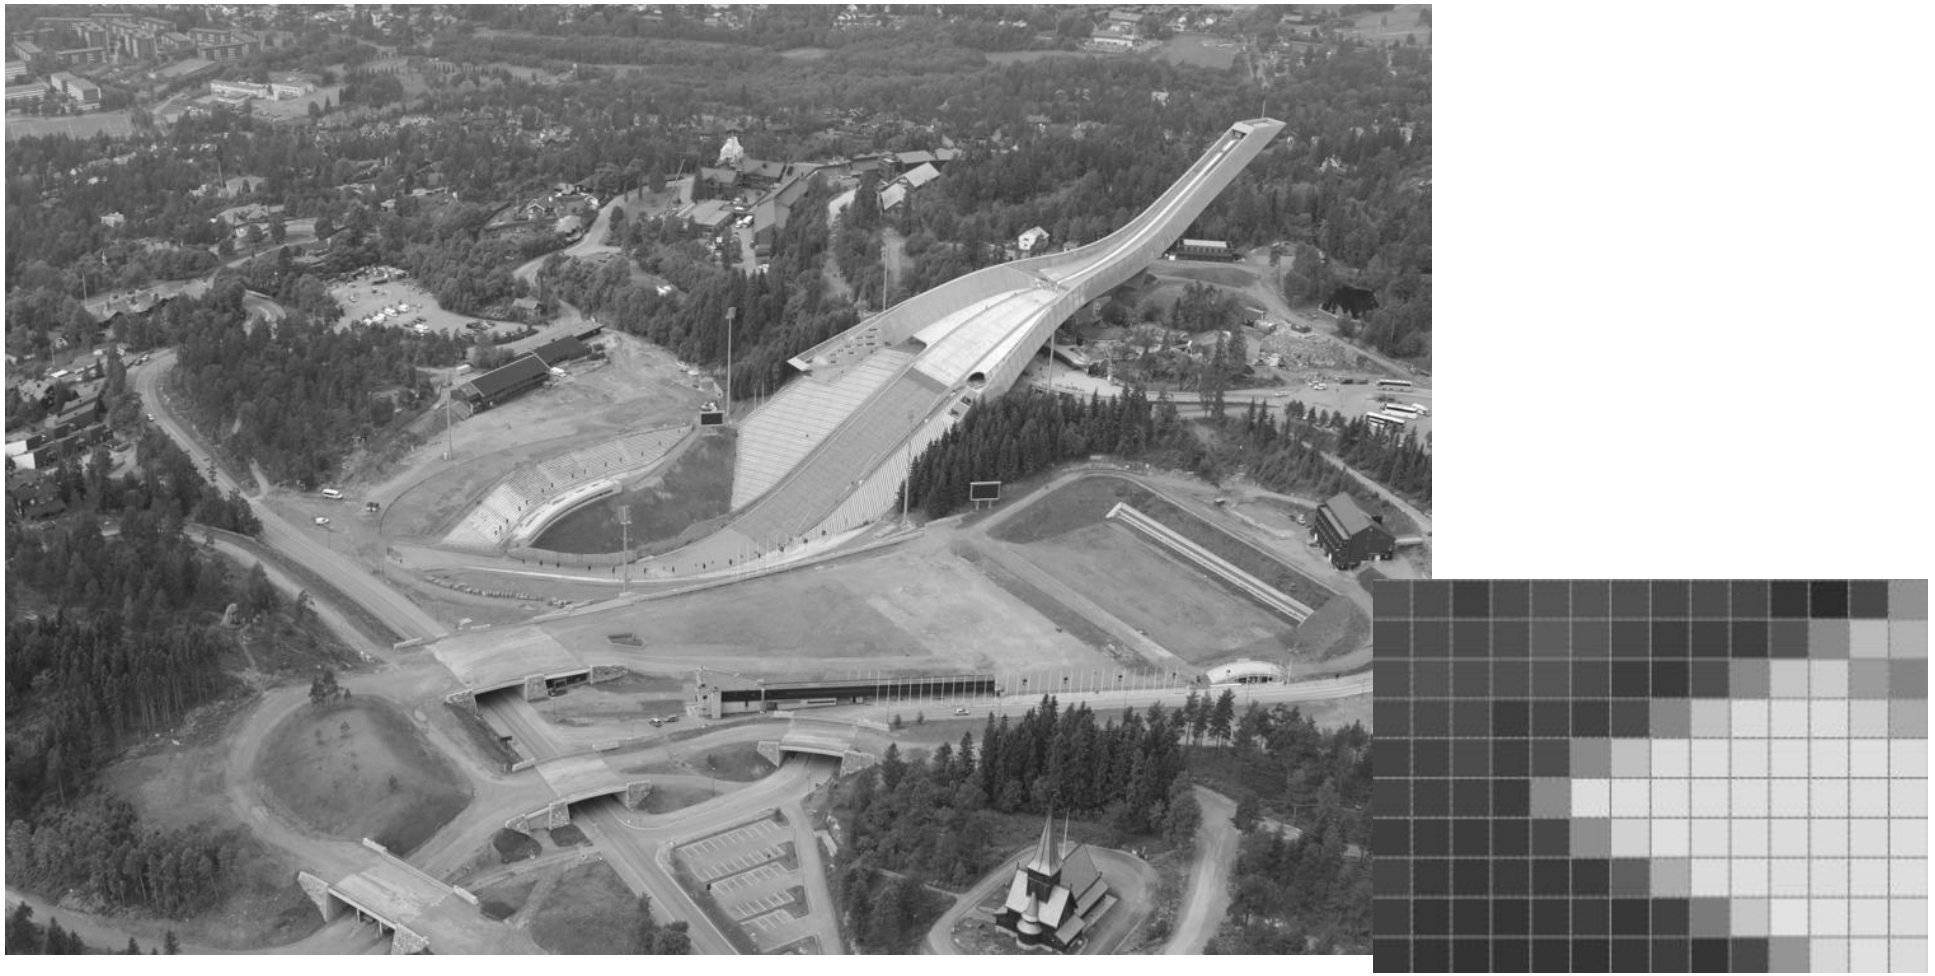
Fugaku supercomputer
>1000 PFLOPS

(theoretical *peak*)

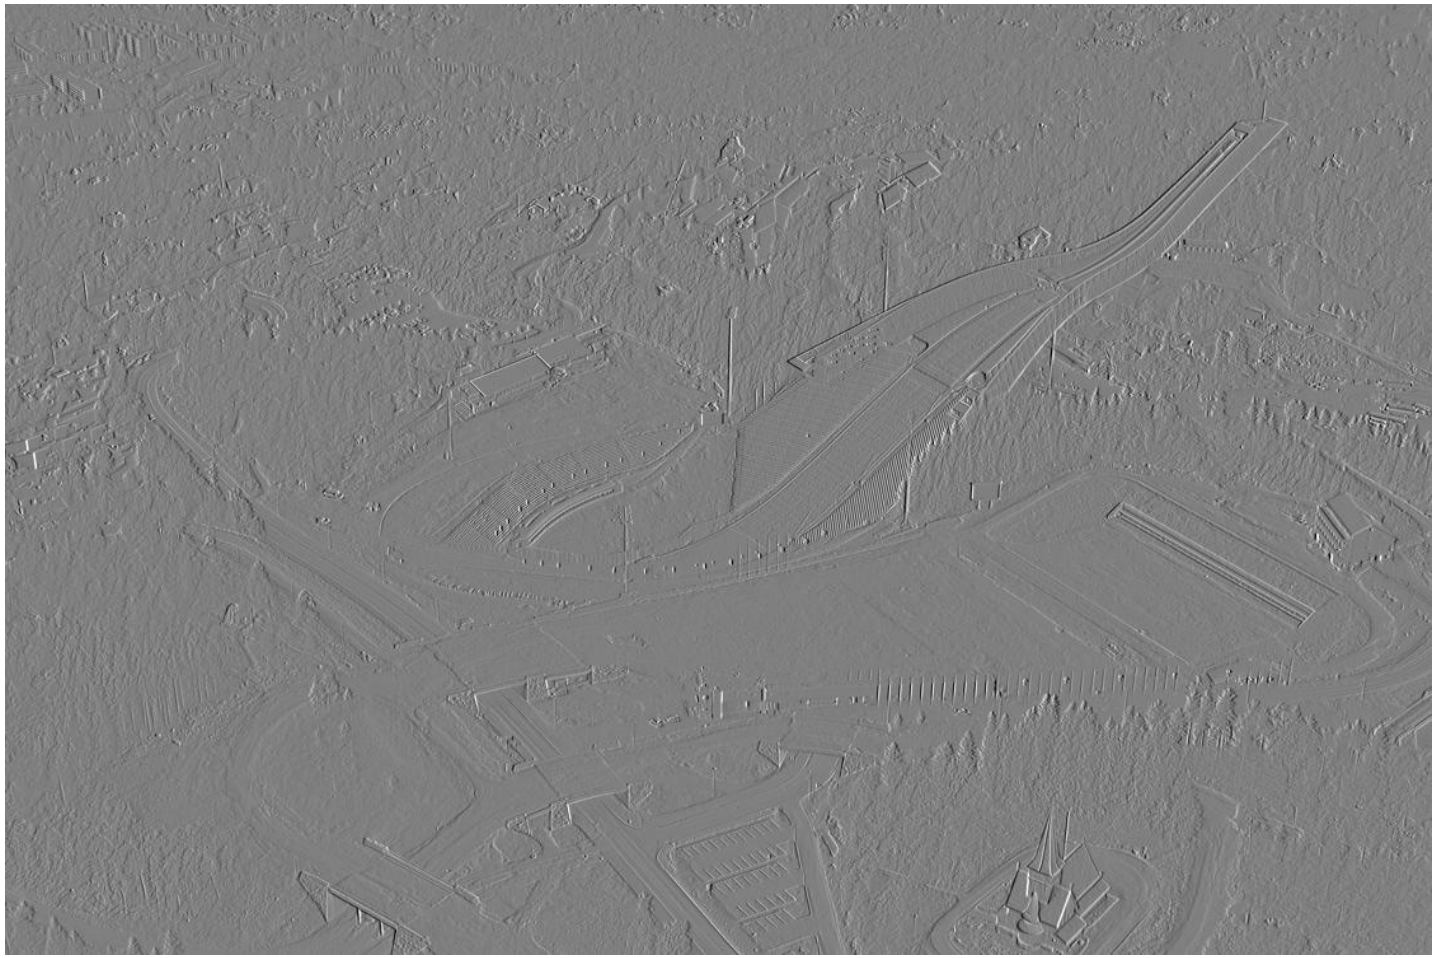
The difference between neighbouring pixels



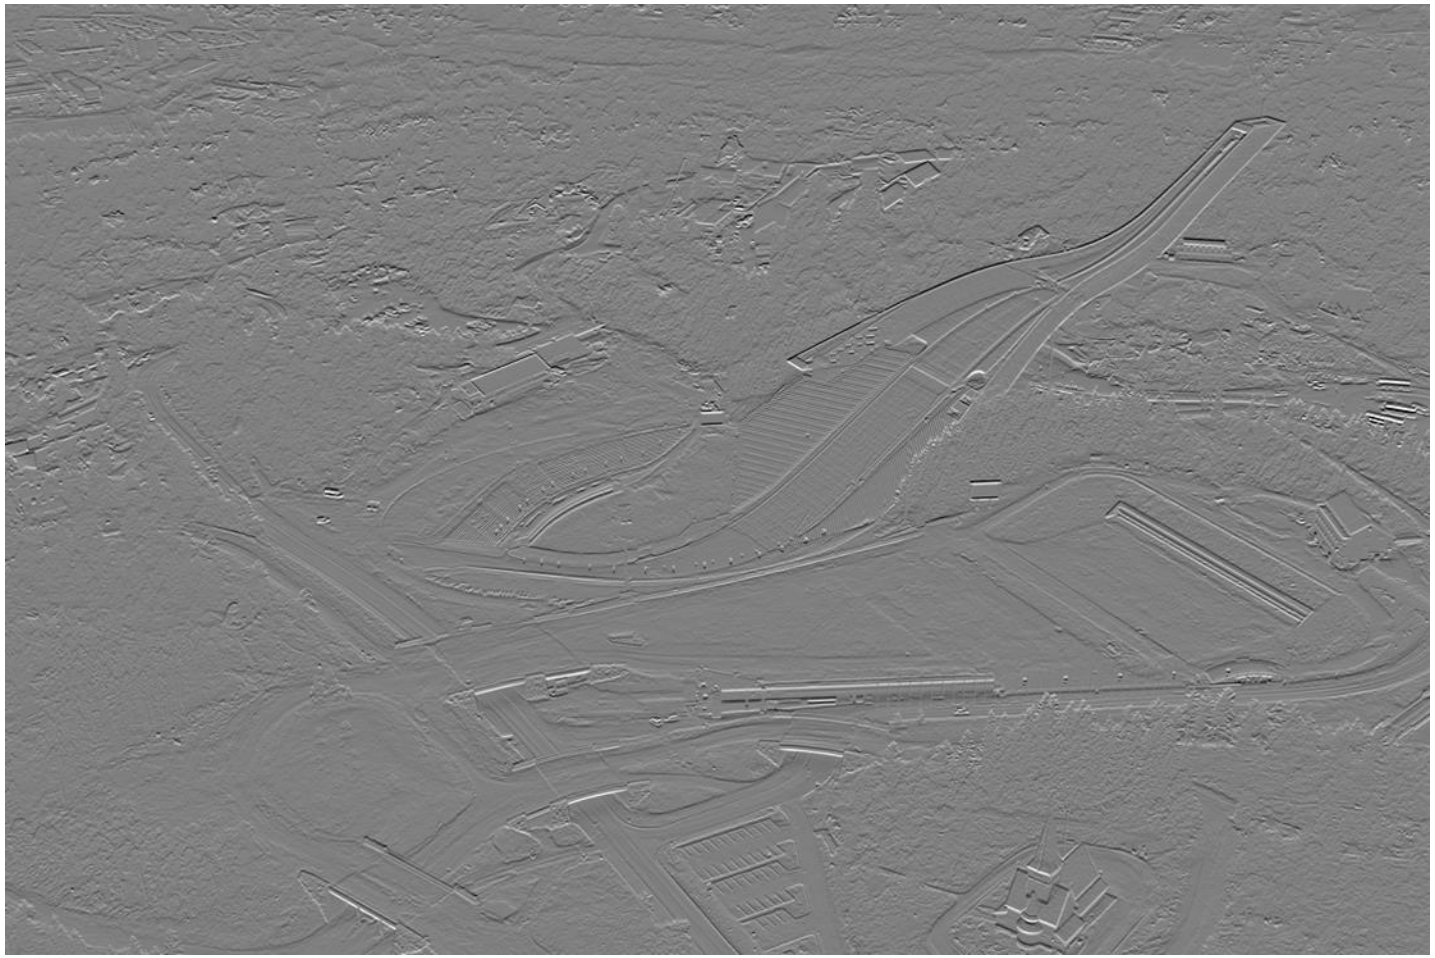
The difference between neighbouring pixels



Horizontal differences



Vertical differences



Edges and corners



Significant corners



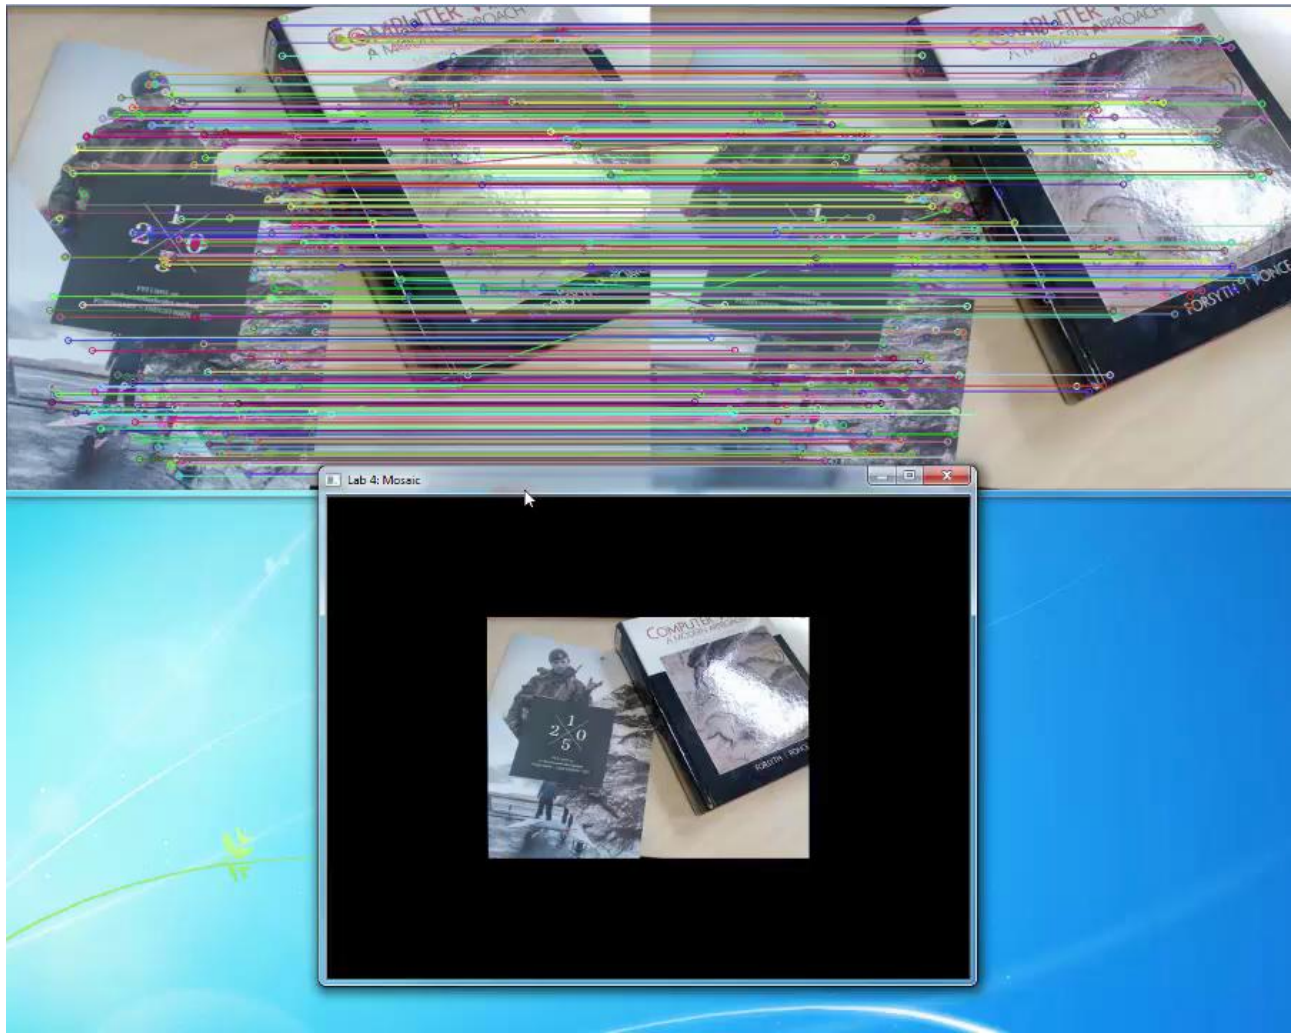
Significant corners



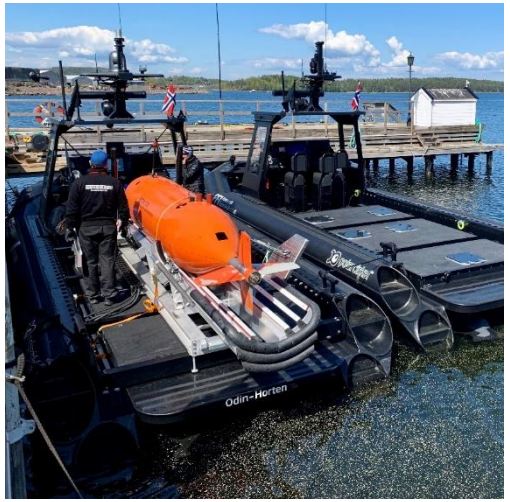
Significant corners

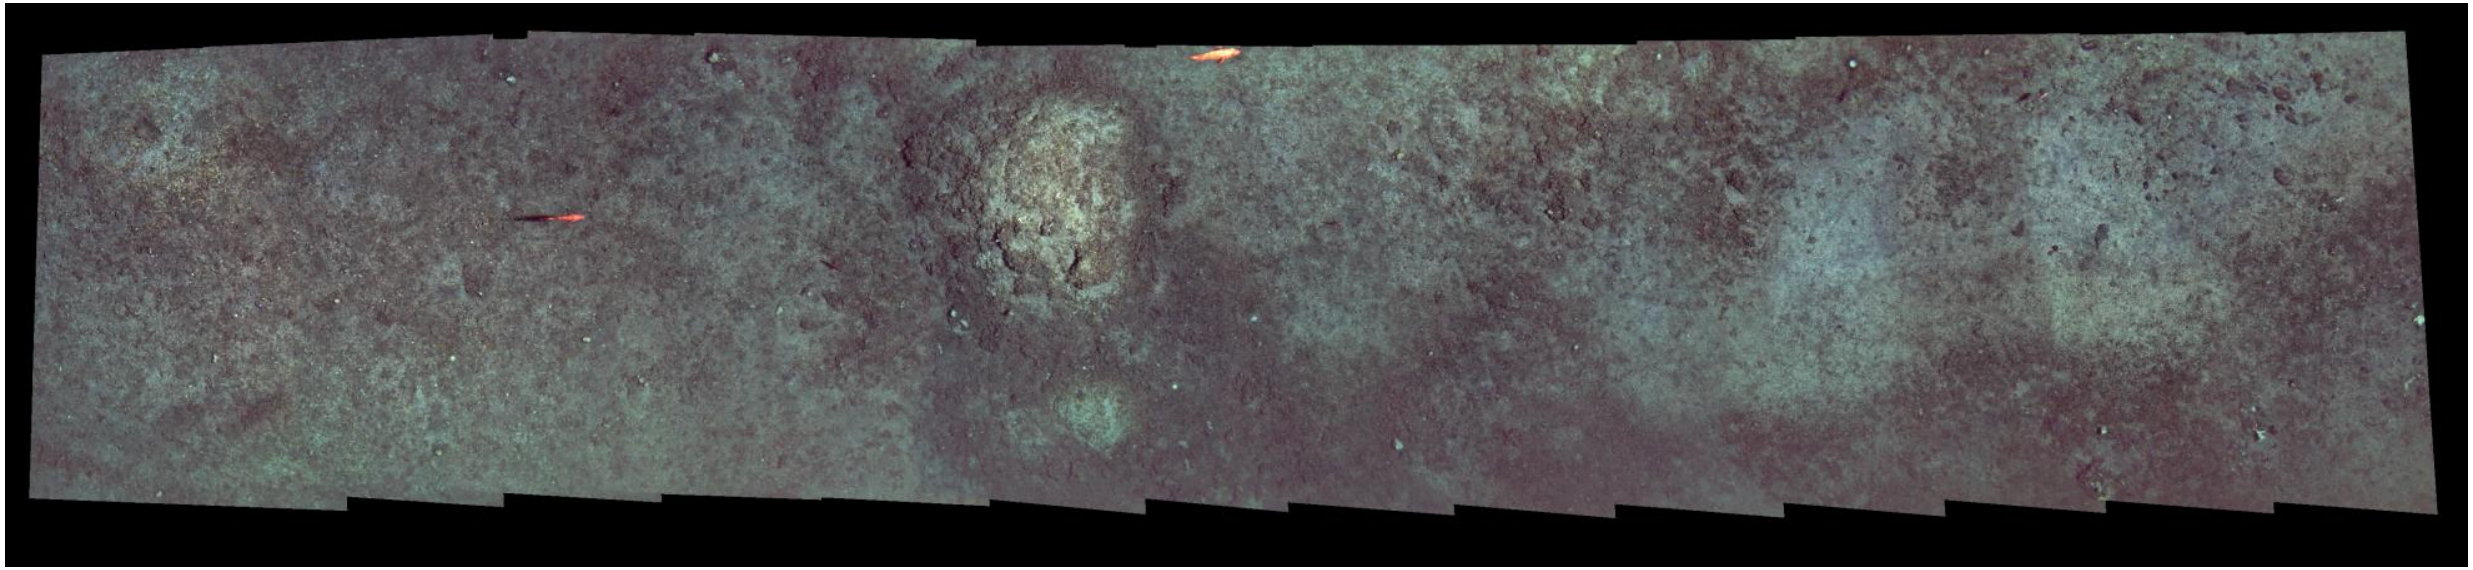


Example: Coregistering images

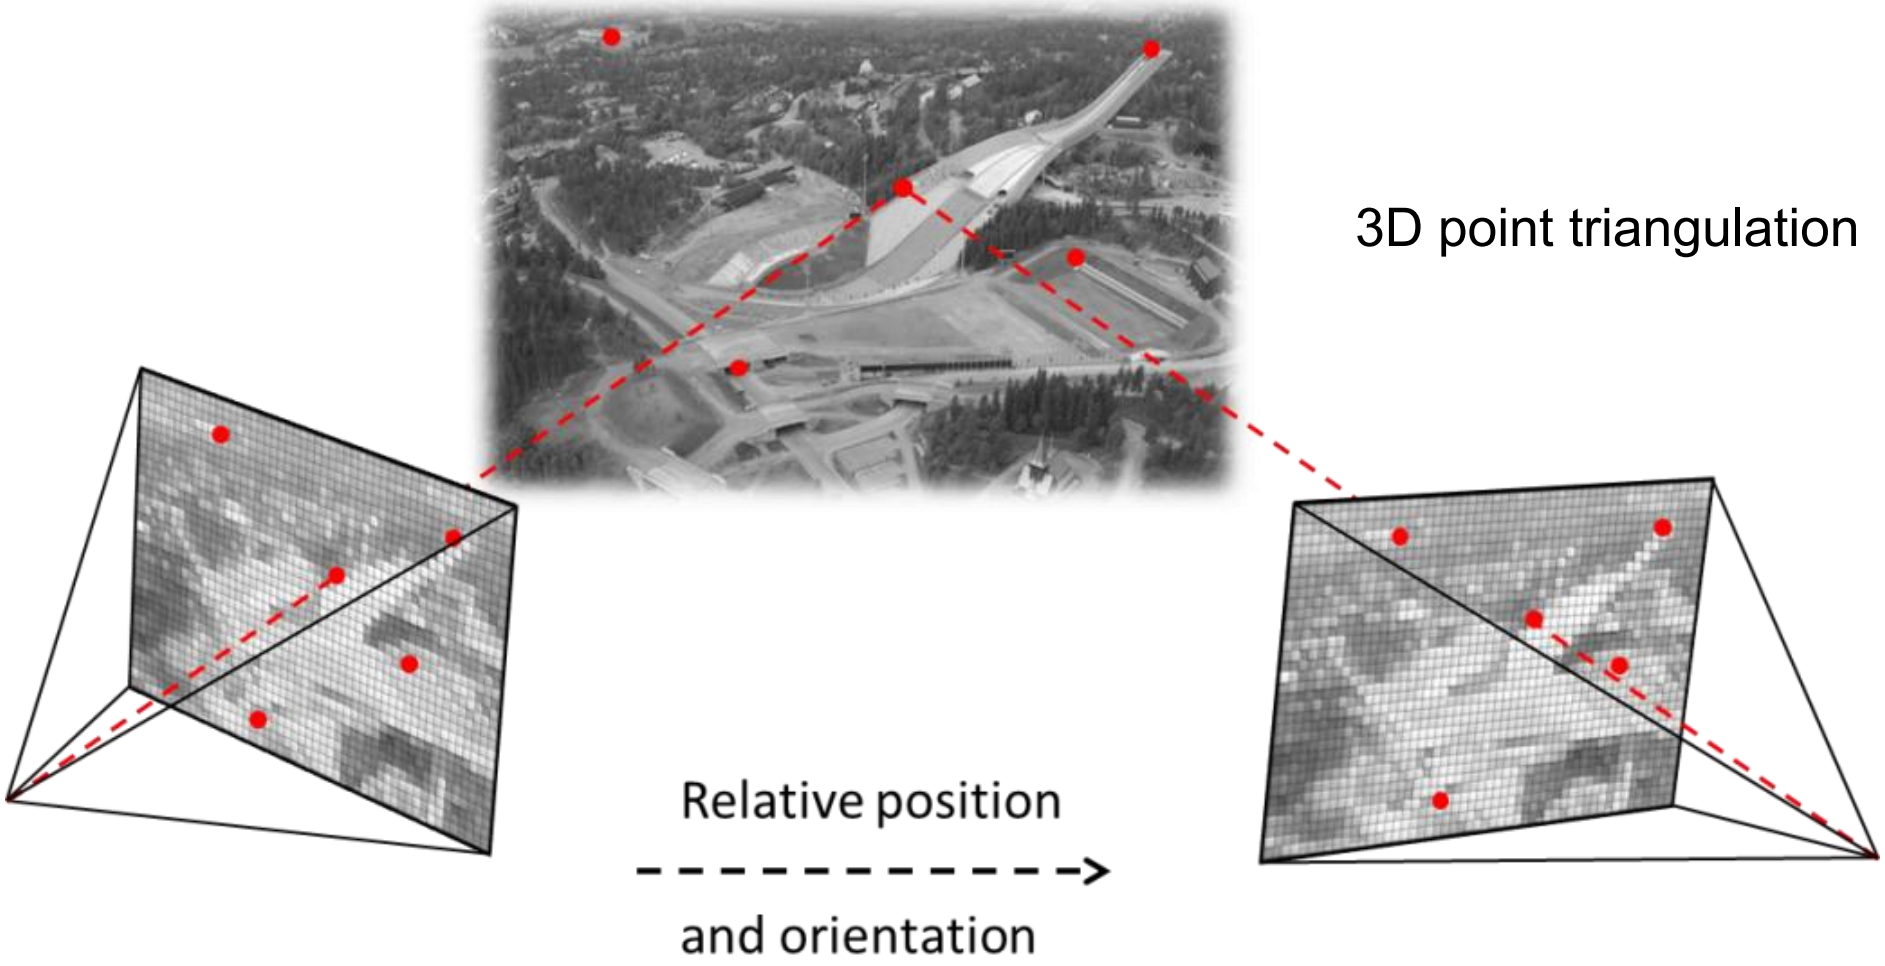


Application: Underwater visual mapping from HUGIN AUV

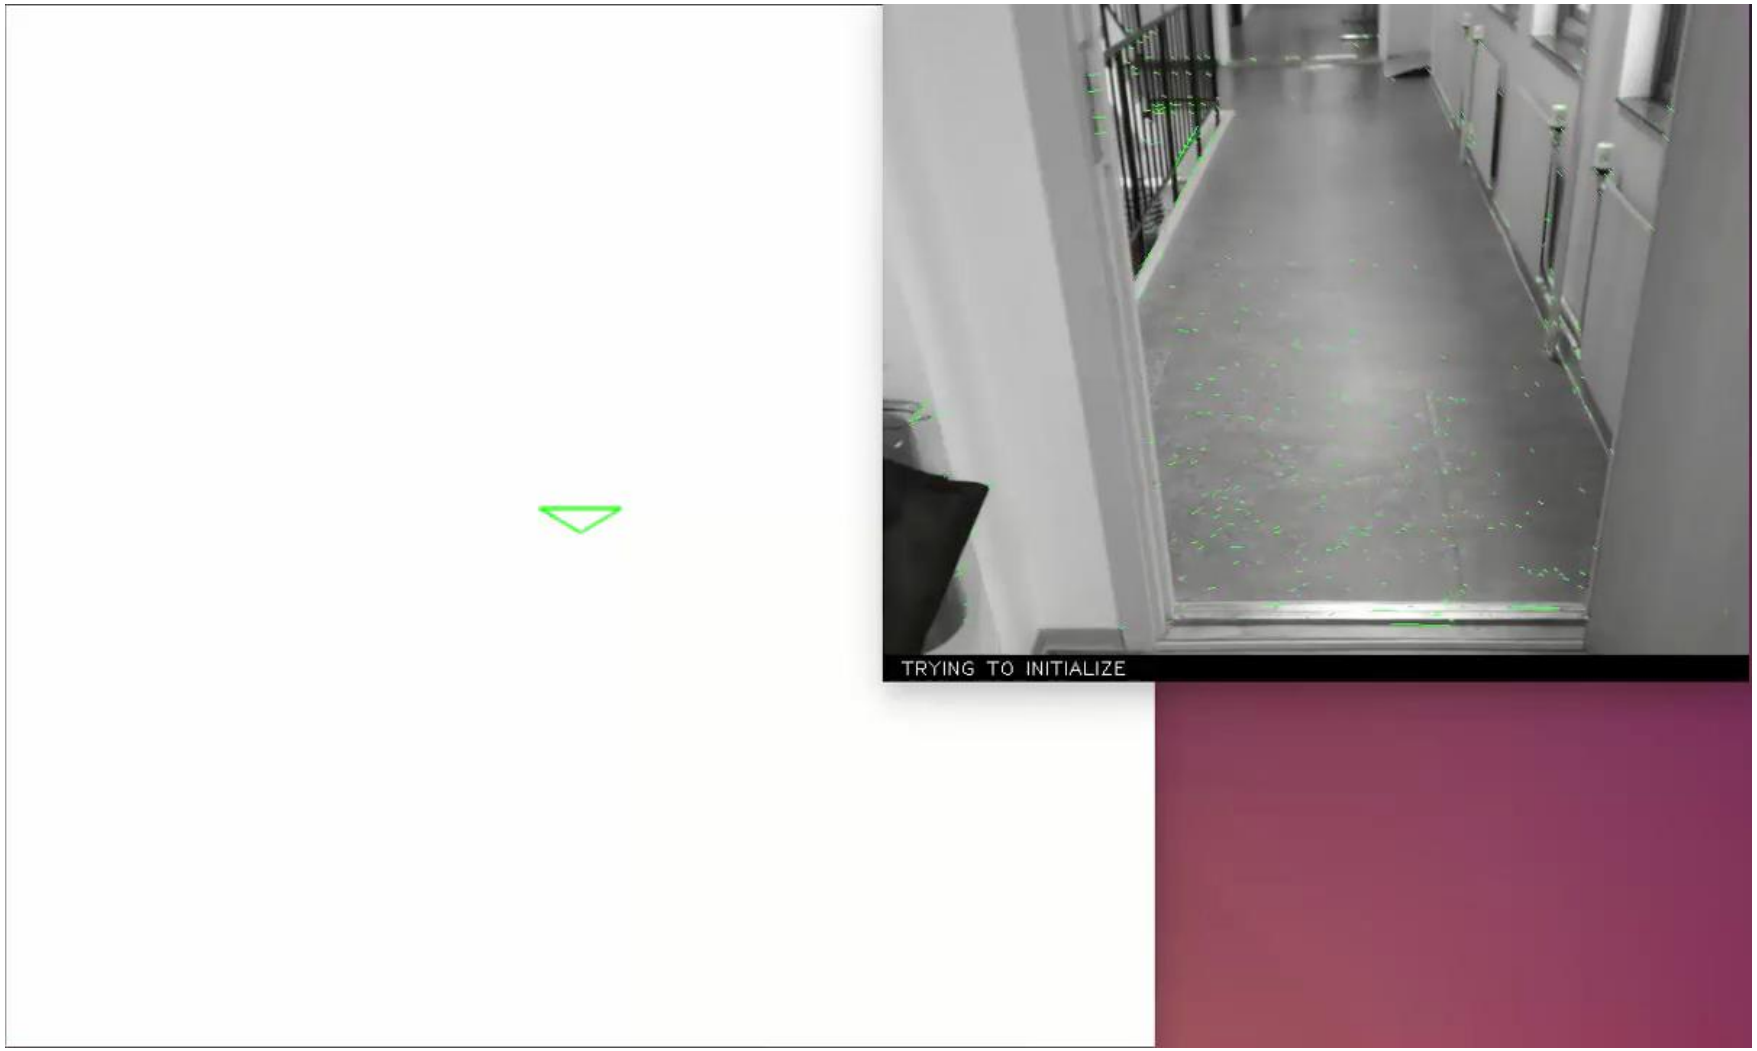




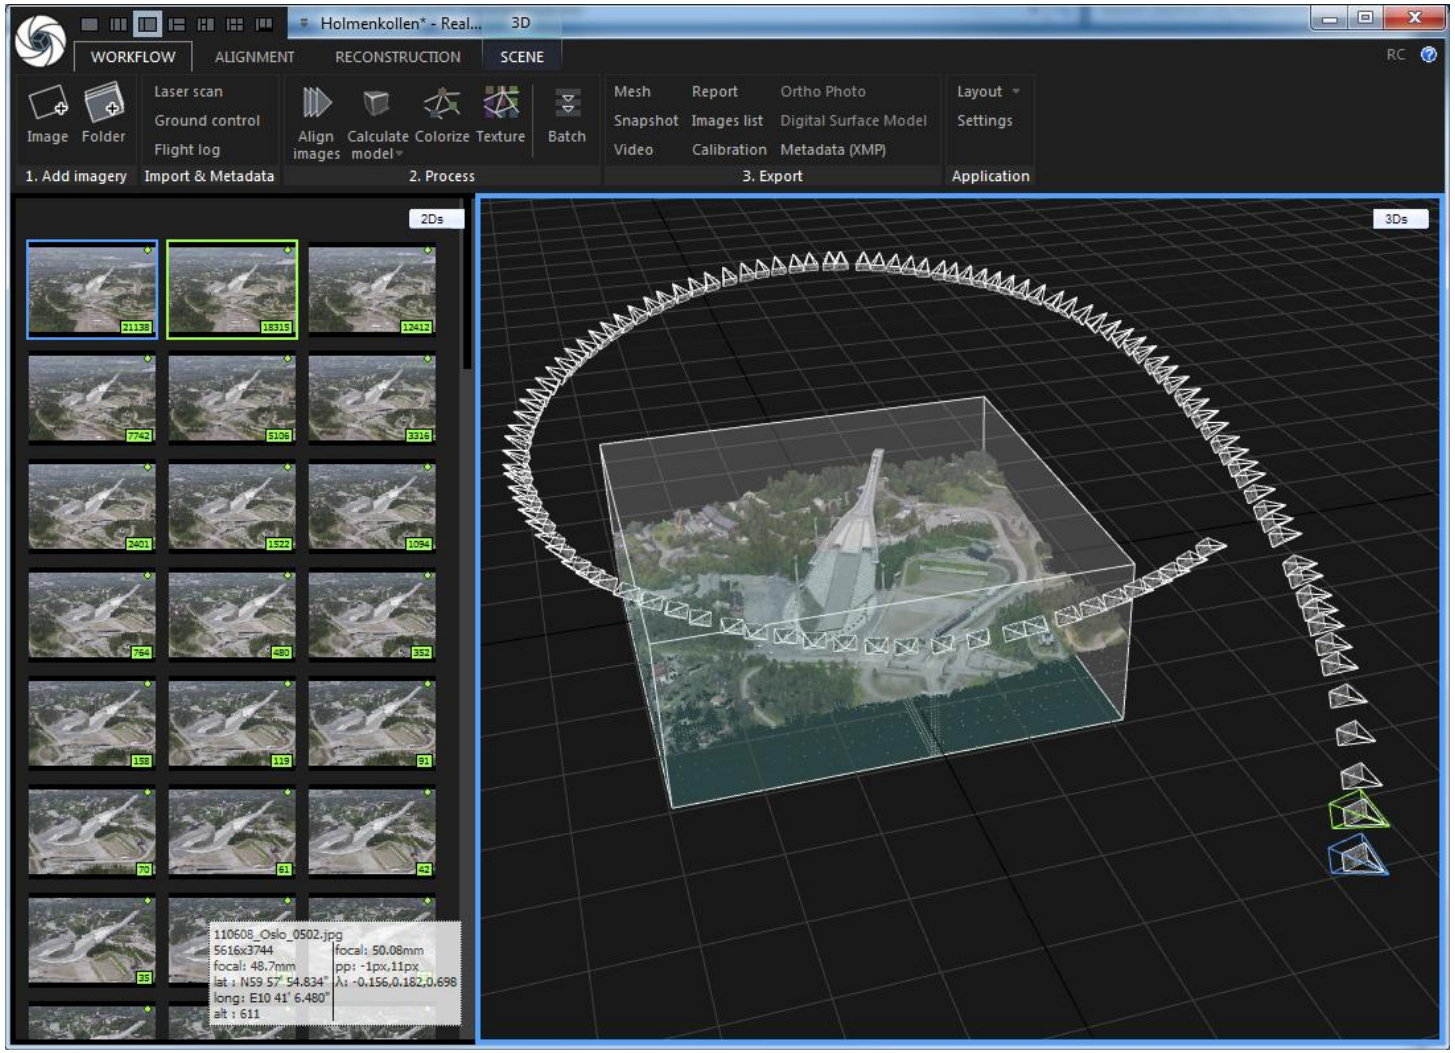
Extracting geometric information from images



Visual navigation



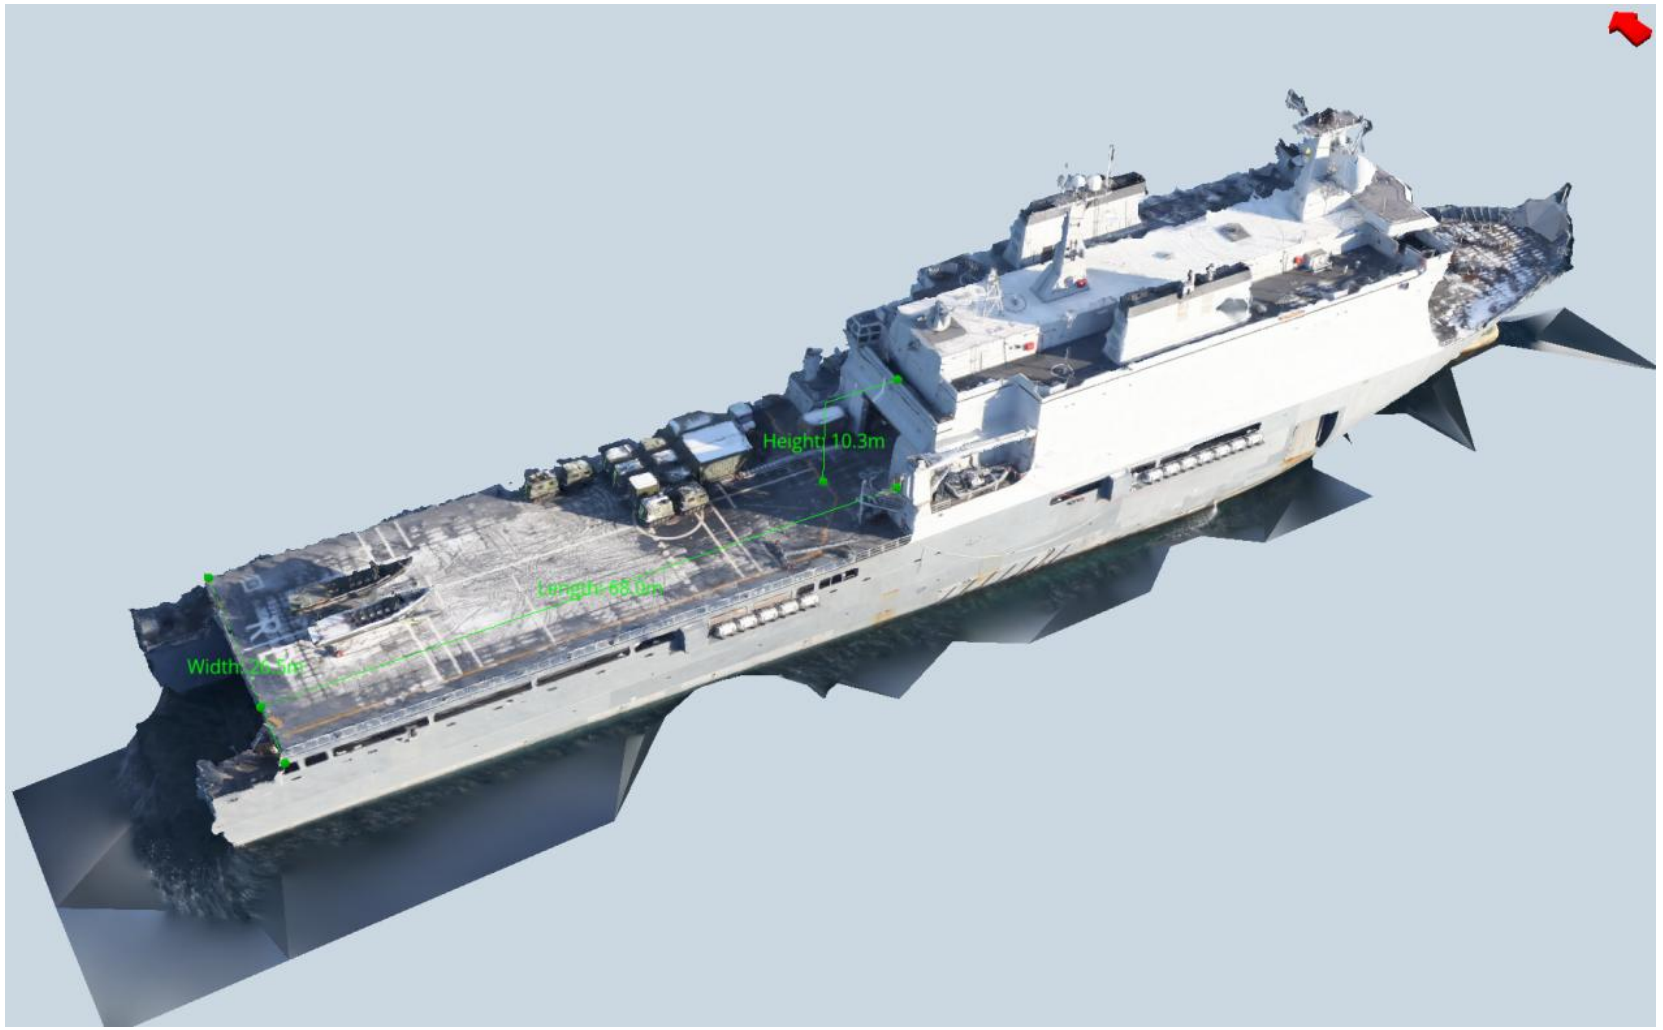
3D reconstruction from images



A detailed 3D surface in colours!

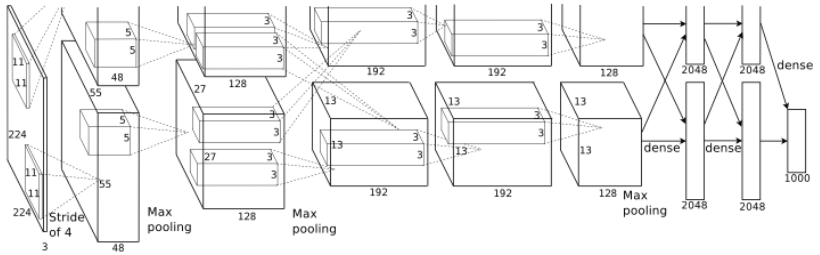


Recognise the shape of a ship



Let the machine learn a better representation itself!

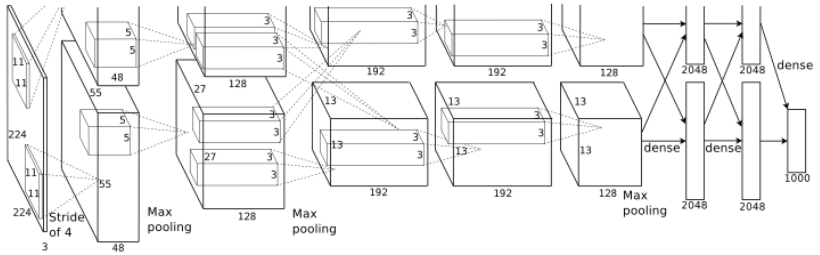




«Warship»

Let the machine learn a better representation itself!

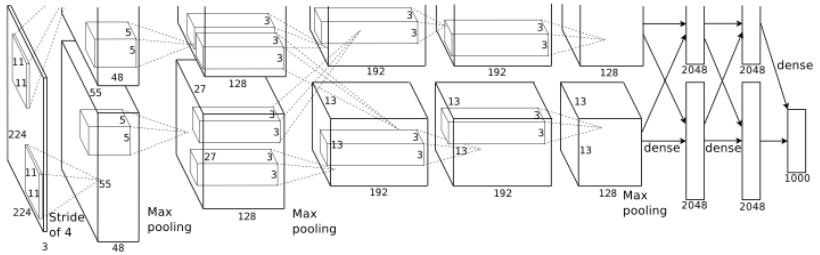




«Warship»

Let the machine learn a better representation itself!



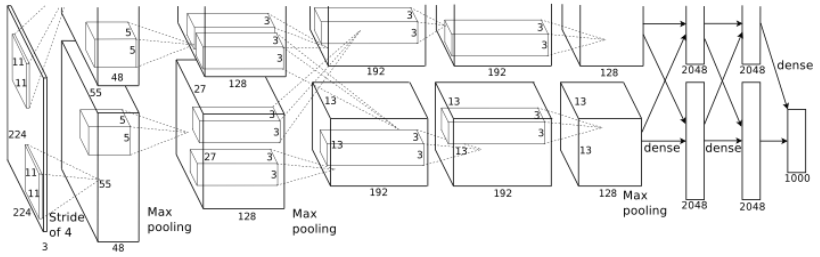


«CV90»

Let the machine learn a better representation itself!

Many millions of
learned parameters

Millions
of images

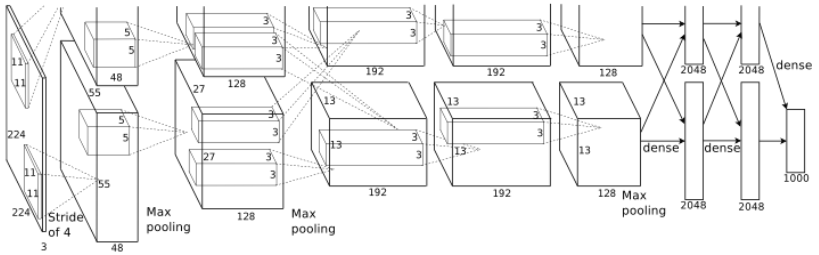


Millions
of labels

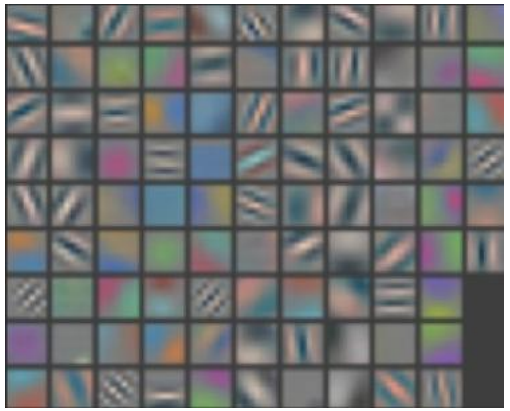
Let the machine learn a better representation itself!

Many millions of
learned parameters

Millions
of images



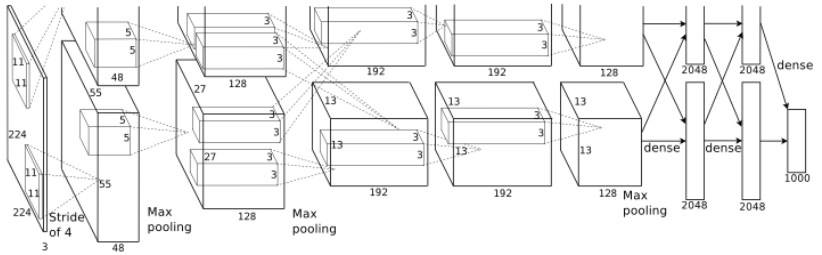
Millions
of labels



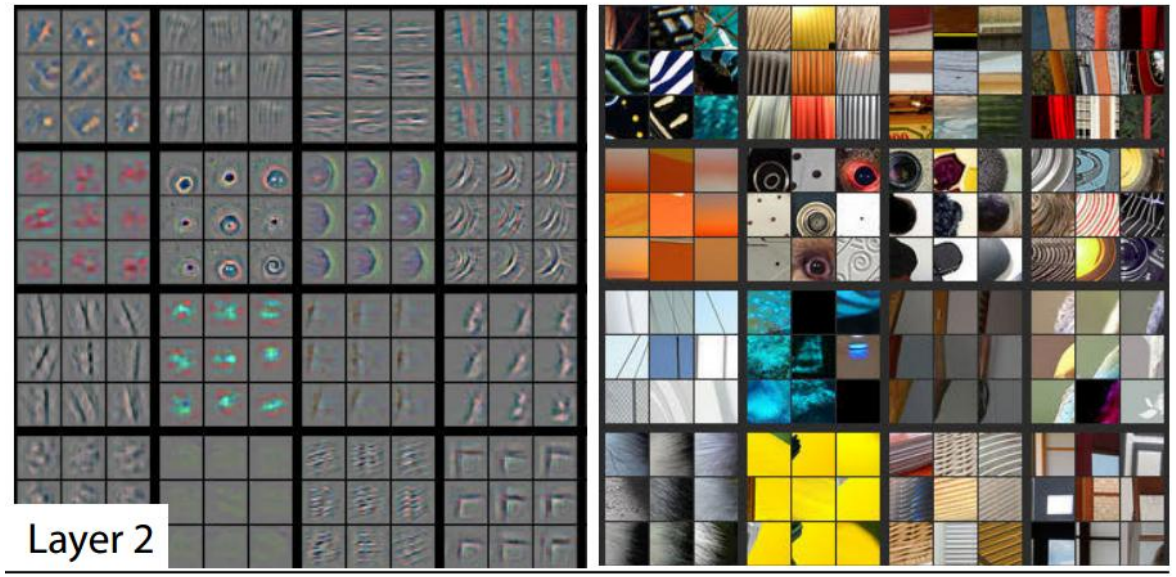
Let the machine learn a better representation itself!

Many millions of
learned parameters

Millions
of images



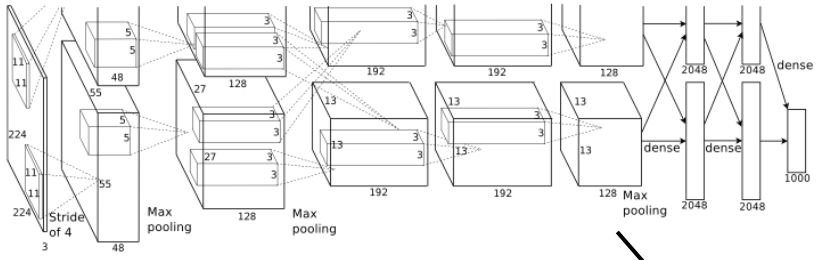
Millions
of labels



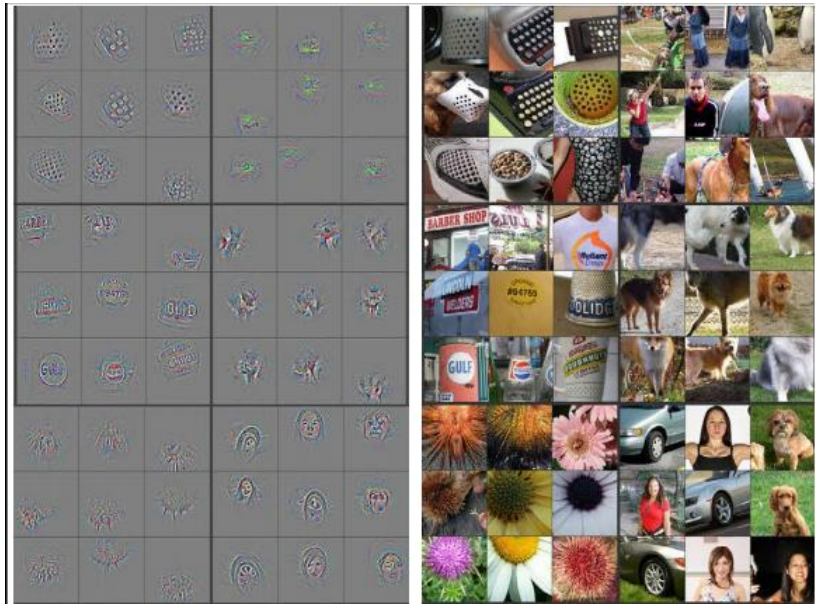
Let the machine learn a better representation itself!

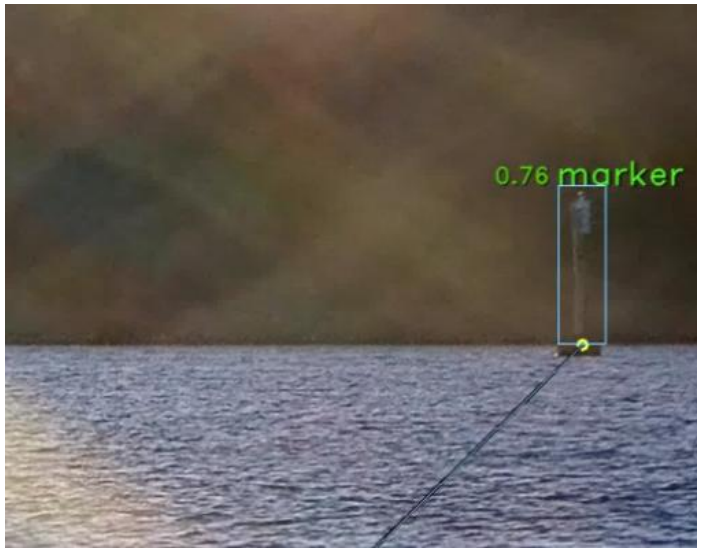
Many millions of
learned parameters

Millions
of images



Millions
of labels




0.76 marker


boat
3
0

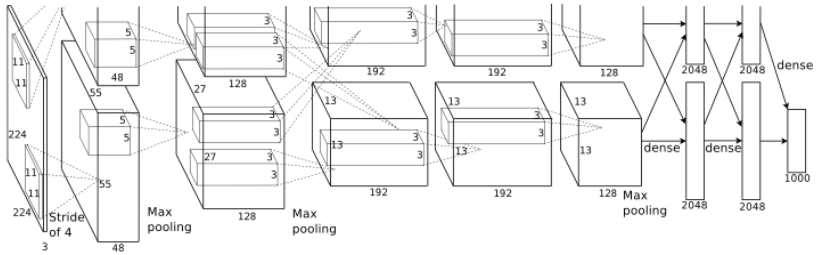



sky
mountain
tree
grass
person
person
road

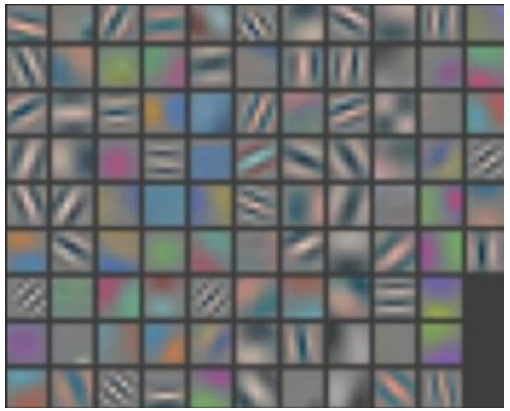
Convolutional Neural Networks (CNNs)

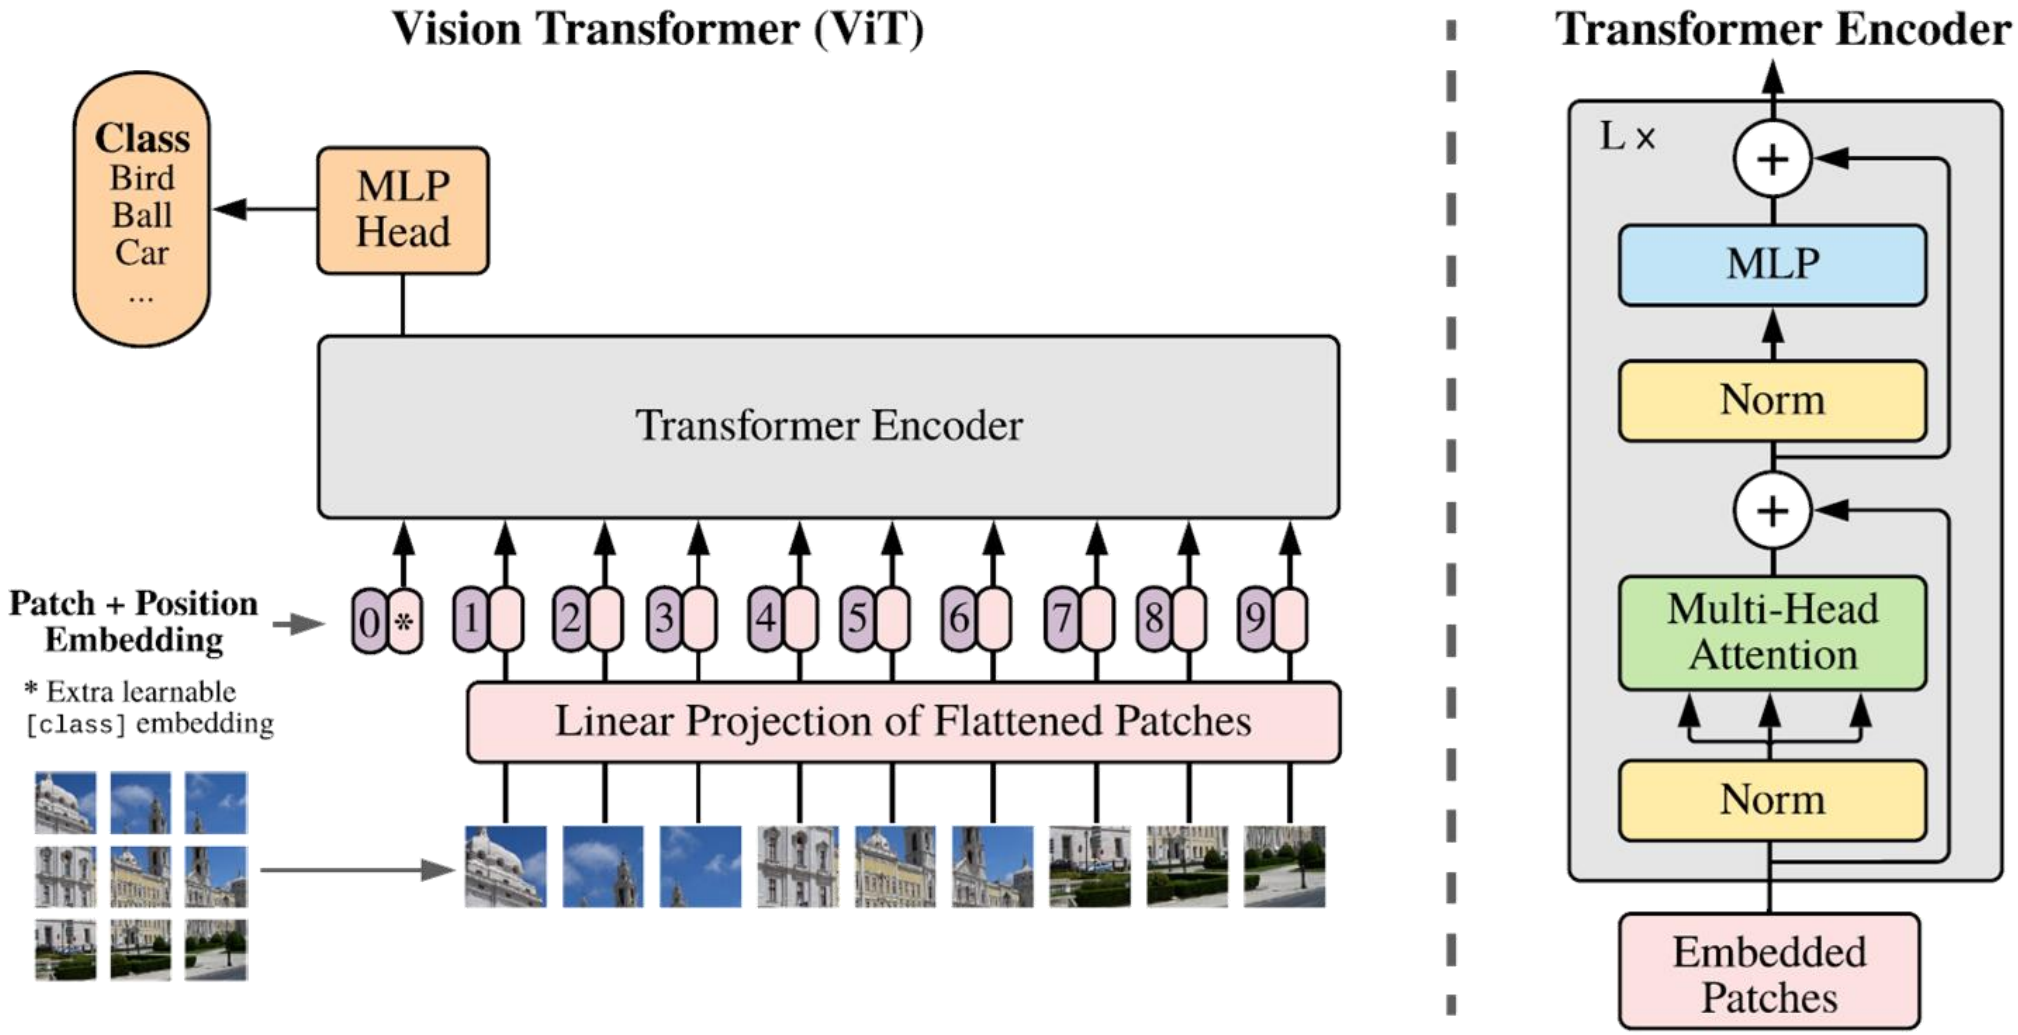
Many millions of
learned parameters

Millions
of images

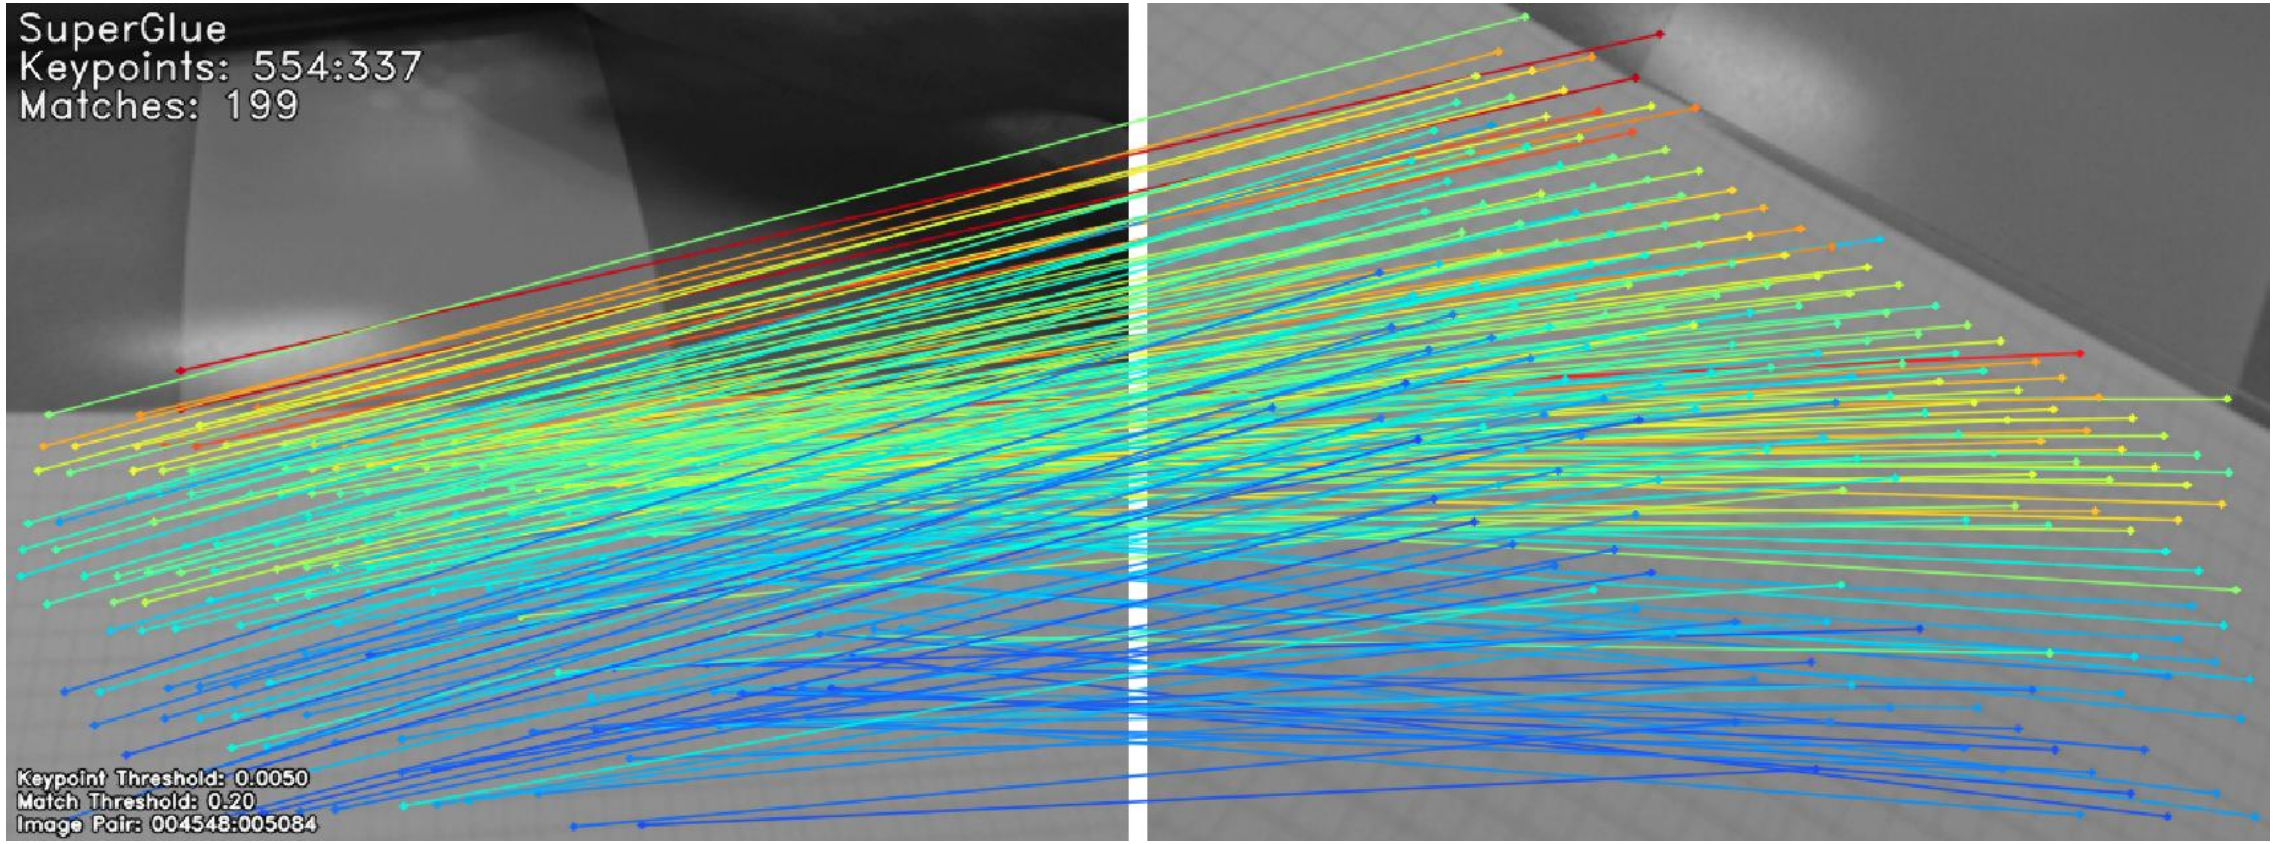


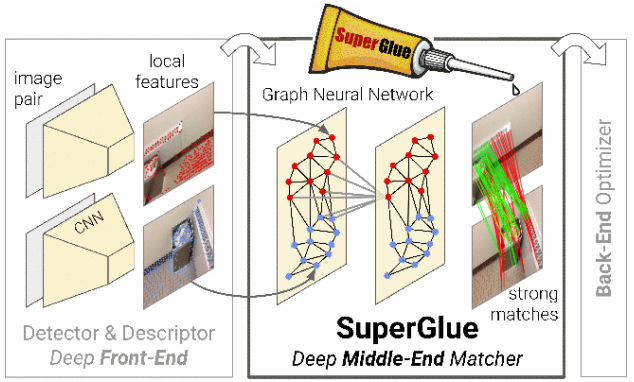
Millions
of labels





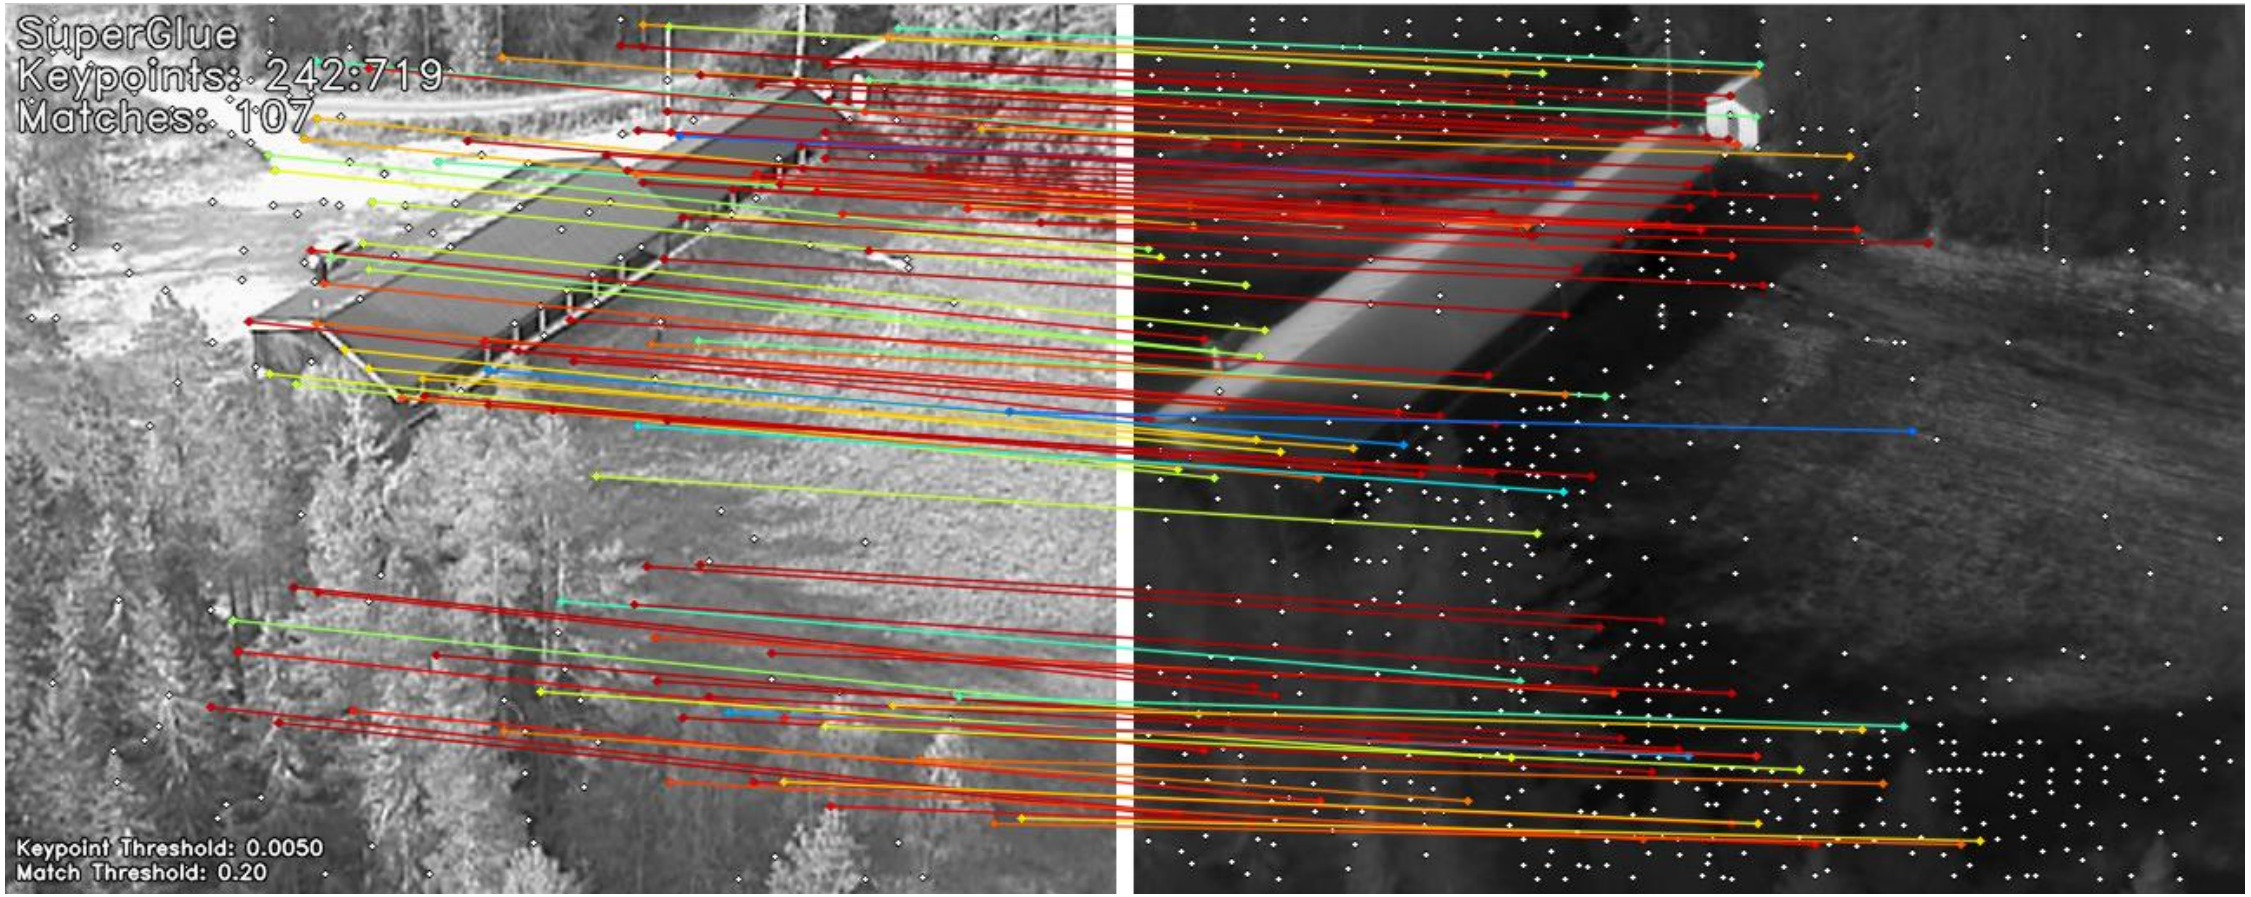
A. Dosovitskiy et al., 'An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale', Jun. 03, 2021, arXiv: arXiv:2010.11929. doi: 10.48550/arXiv.2010.11929.

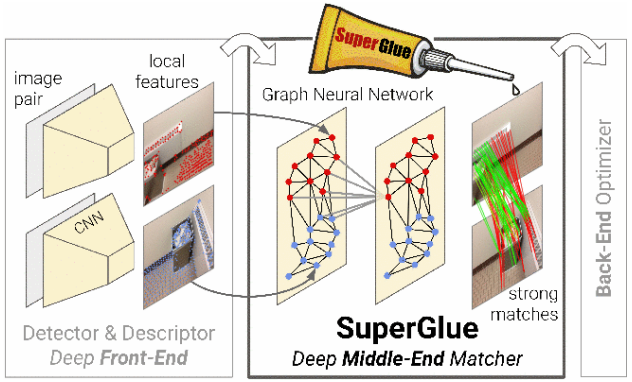




Sarlin, P. E., Detone, D., Malisiewicz, T., & Rabinovich, A. (2020).
SuperGlue: Learning Feature Matching with Graph Neural Networks.
Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 4937–4946.

https://psarlin.com/superglue/





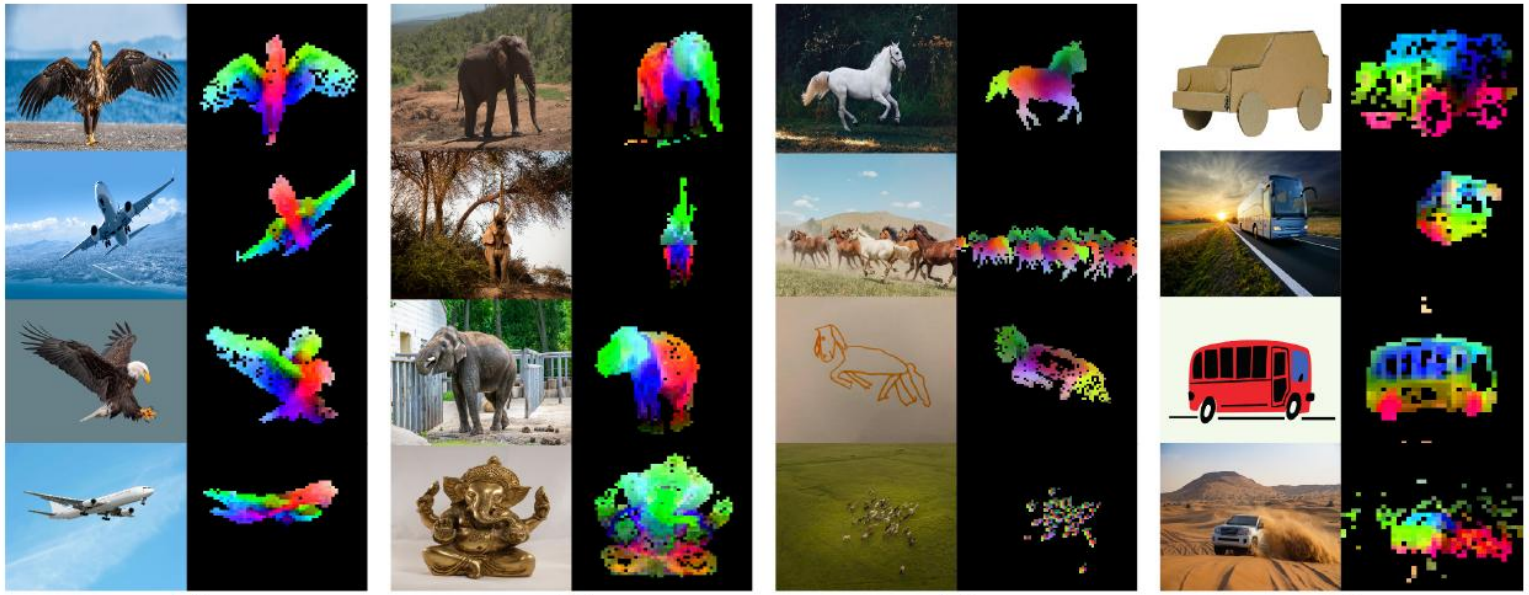
Sarlin, P. E., Detone, D., Malisiewicz, T., & Rabinovich, A. (2020).
SuperGlue: Learning Feature Matching with Graph Neural Networks.
Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 4937–4946.

https://psarlin.com/superglue/

DINOv2: Learning Robust Visual Features without Supervision

Maxime Oquab**, Timothée Darcet**, Théo Moutakanni**,
Huy V. Vo*, Marc Szafraniec*, Vasil Khalidov*, Pierre Fernandez, Daniel Haziza,
Francisco Massa, Alaaeldin El-Nouby, Mahmoud Assran, Nicolas Ballas, Wojciech Galuba,
Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat,
Vasu Sharma, Gabriel Synnaeve, Hu Xu, Hervé Jegou, Julien Mairal[1],
Patrick Labatut*, Armand Joulin*, Piotr Bojanowski*

Meta AI Research [1]*Inria*



(a) (b) (c) (d)

M. Oquab et al., 'DINOv2: Learning Robust Visual Features without Supervision', 2023, arXiv. doi: 10.48550/ARXIV.2304.07193.

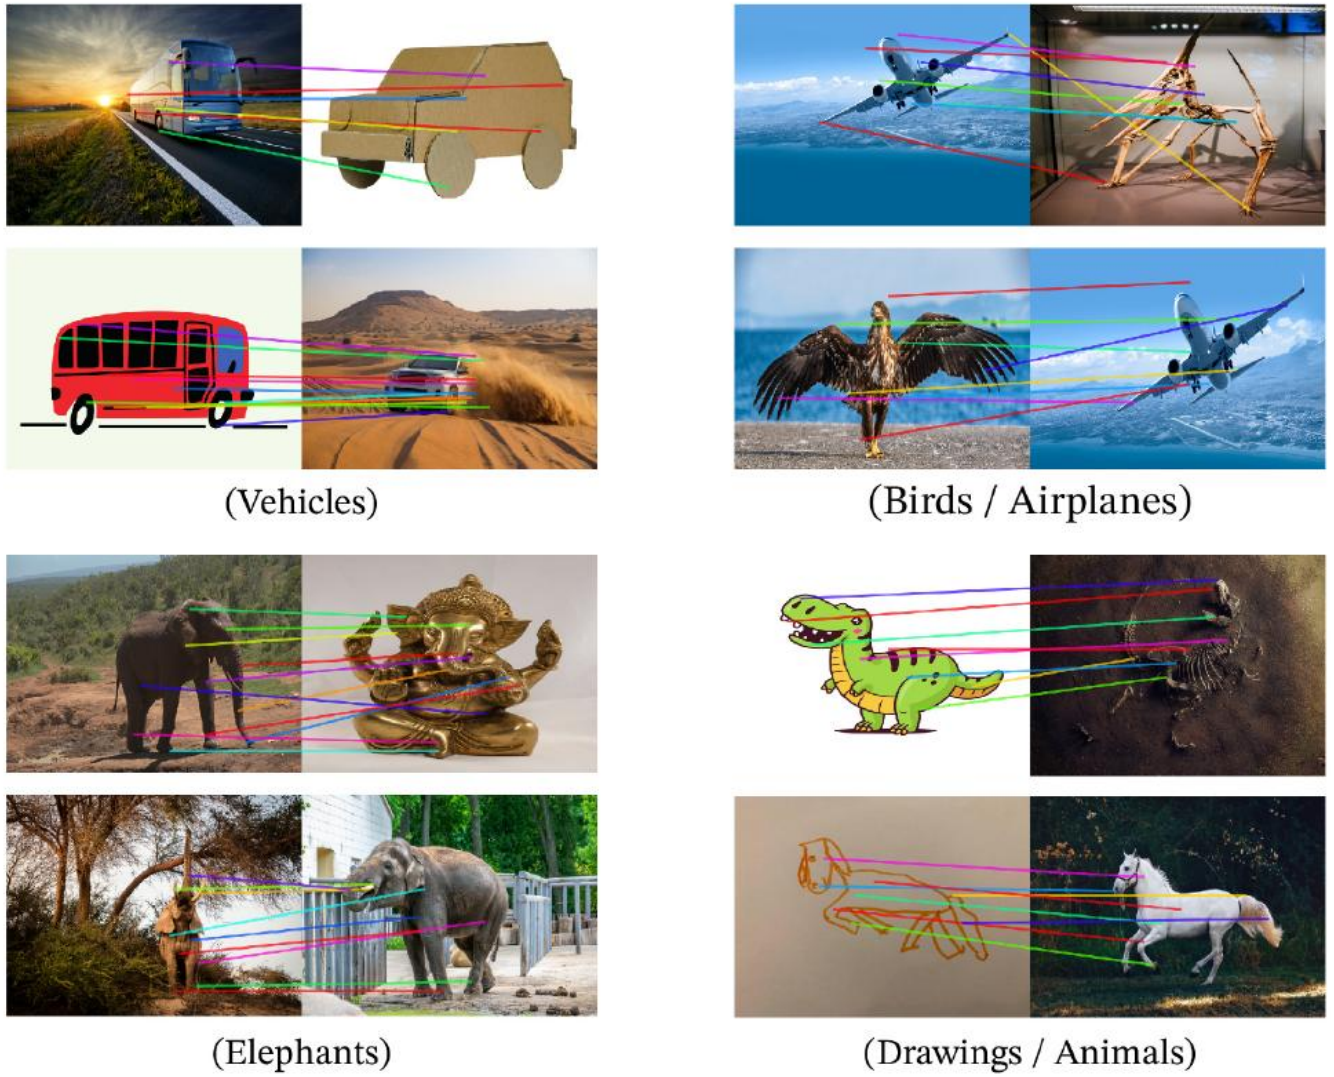
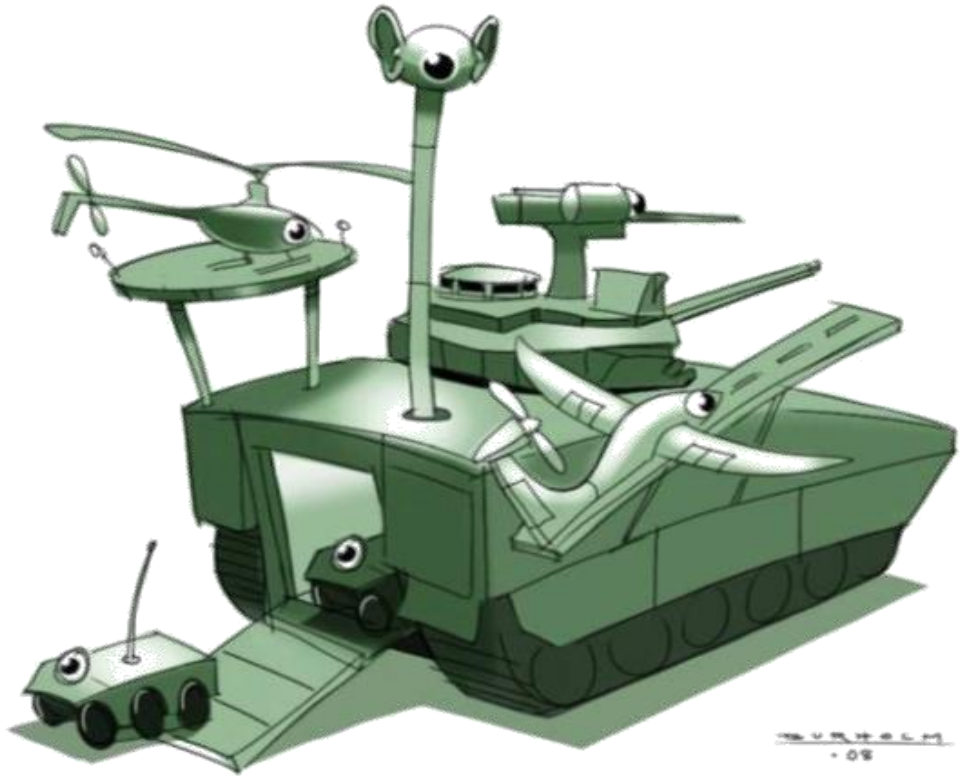


Figure 10: **Matching across images.** We match patch-level features between images from different domains, poses and even objects that share similar semantic information. This exhibits the ability of our model to transfer across domains and understand relations between similar parts of different objects.

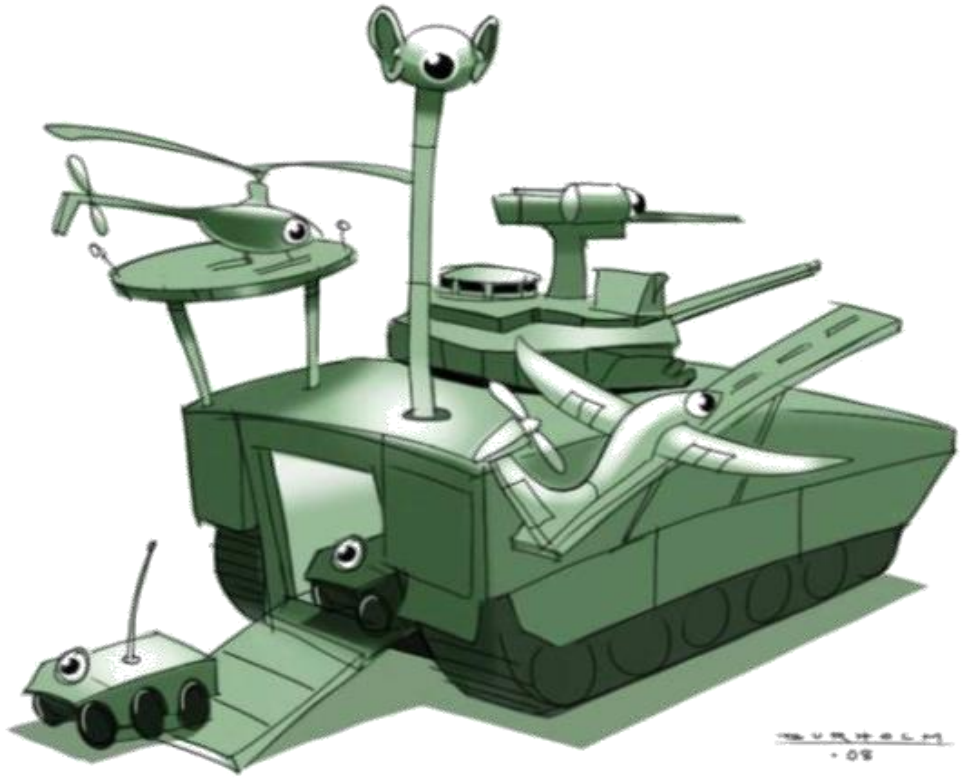
M. Oquab et al., 'DINOv2: Learning Robust Visual Features without Supervision', 2023, arXiv. doi: 10.48550/ARXIV.2304.07193.

What is computer vision?



What is computer vision?

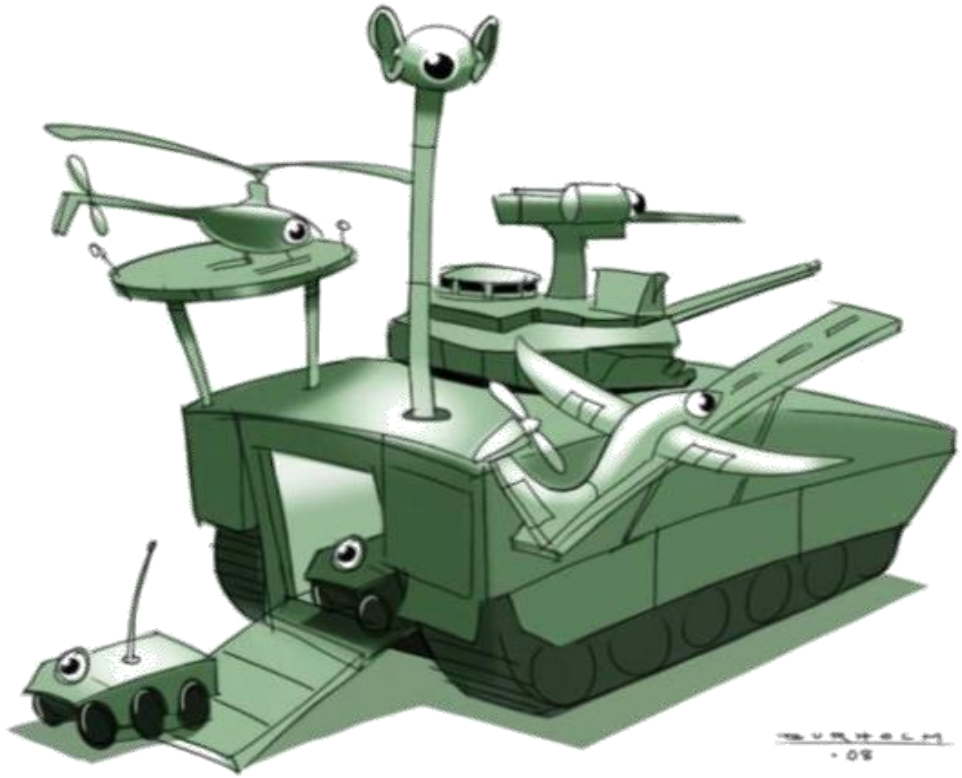
The study of how a machine
can interpret and understand its surroundings
from images



What is computer vision?

The study of how a machine
can interpret and understand its surroundings
from images

- "Enabling computers to see"

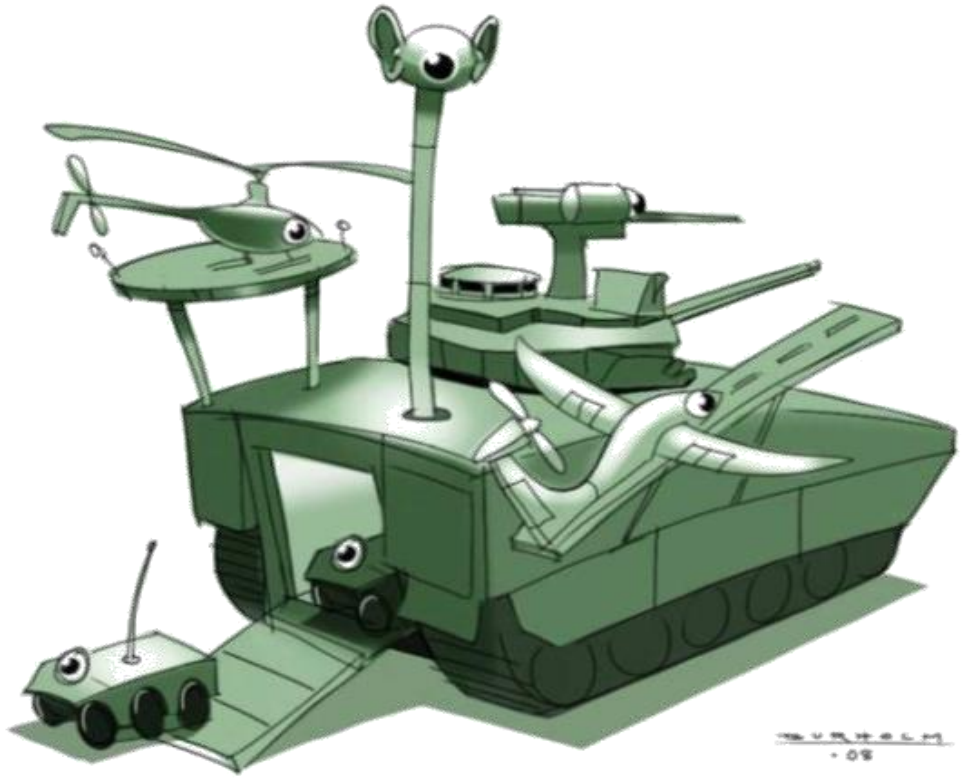


What is computer vision?

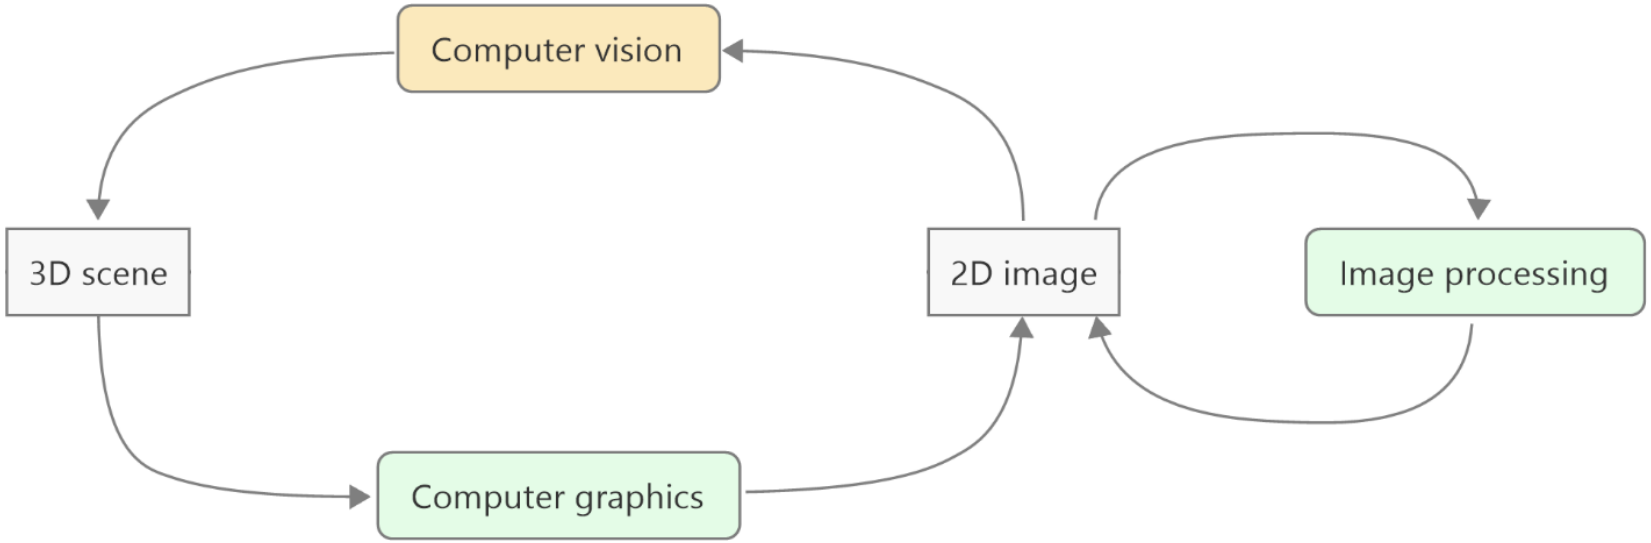
The study of how a machine
can interpret and understand its surroundings
from images

- “Enabling computers to see”

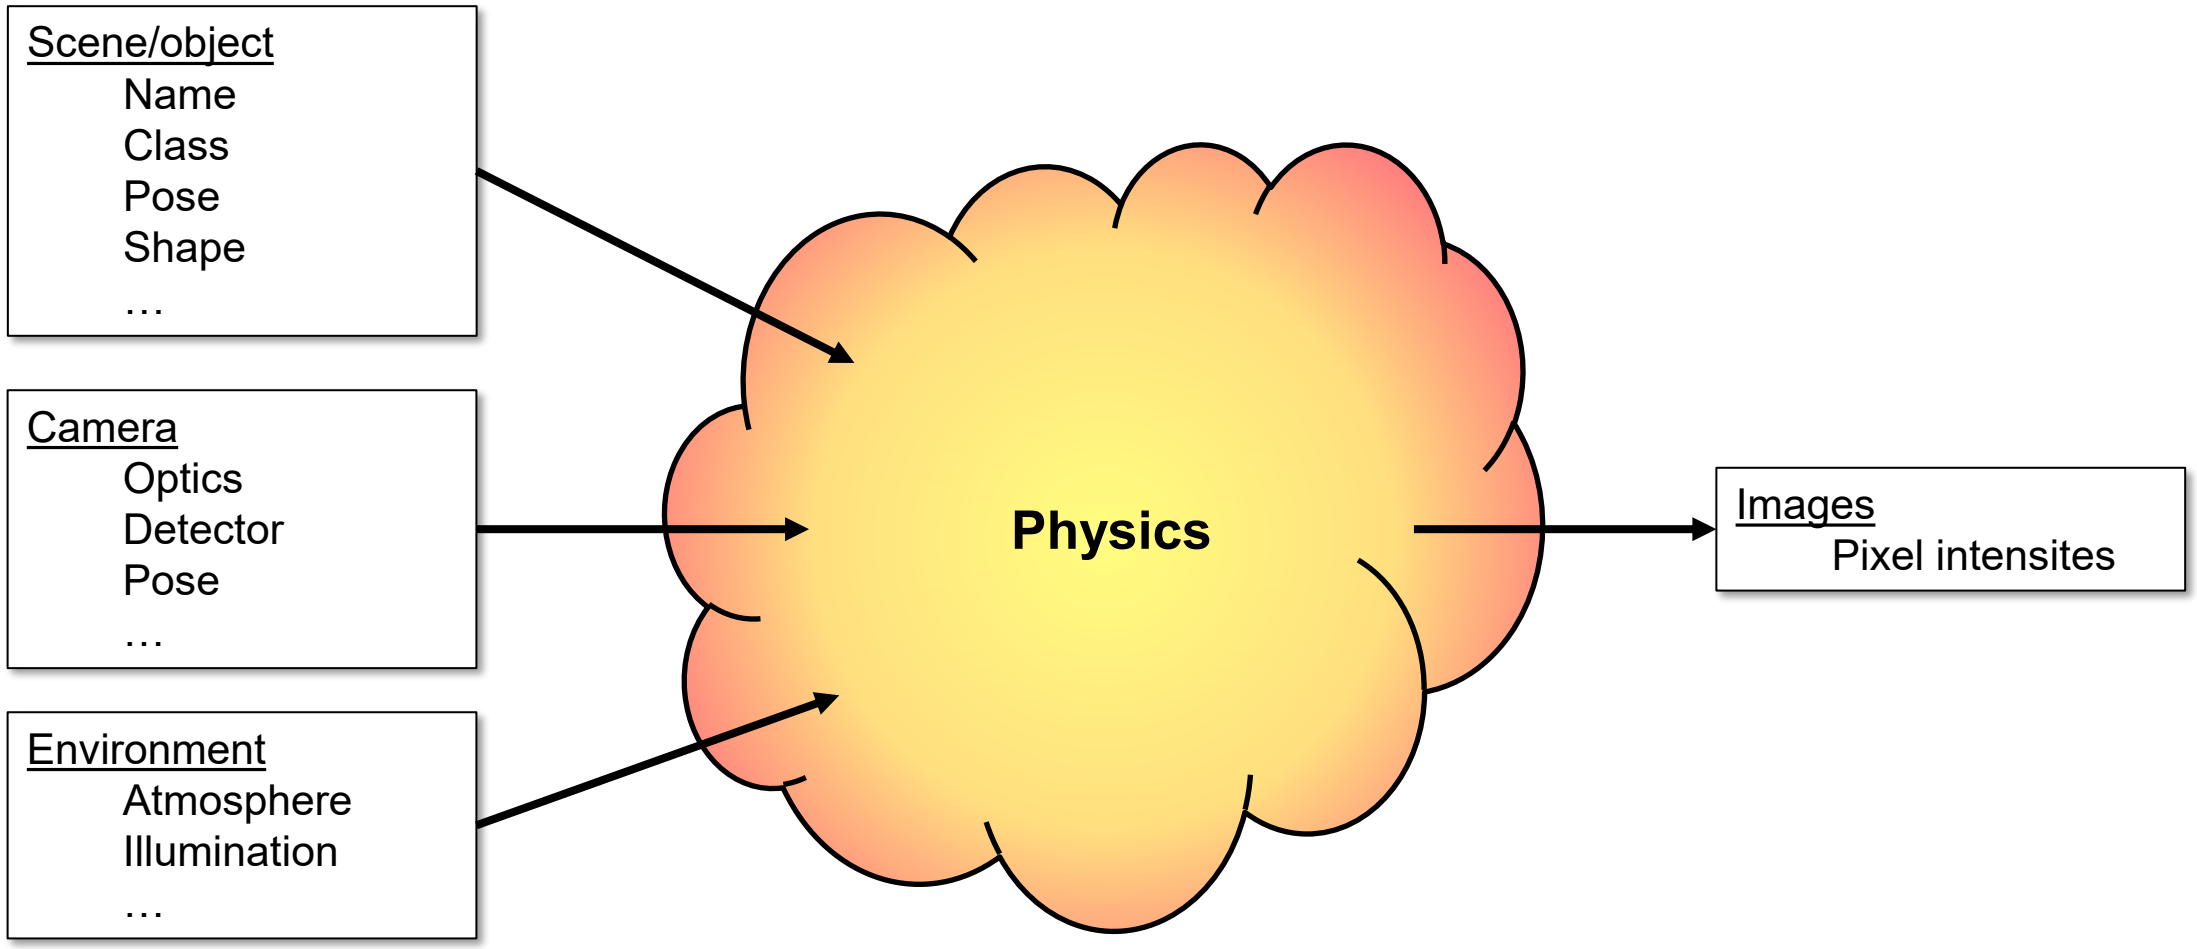
- Image analysis
- Robotic vision
- Perception
- Spatial AI



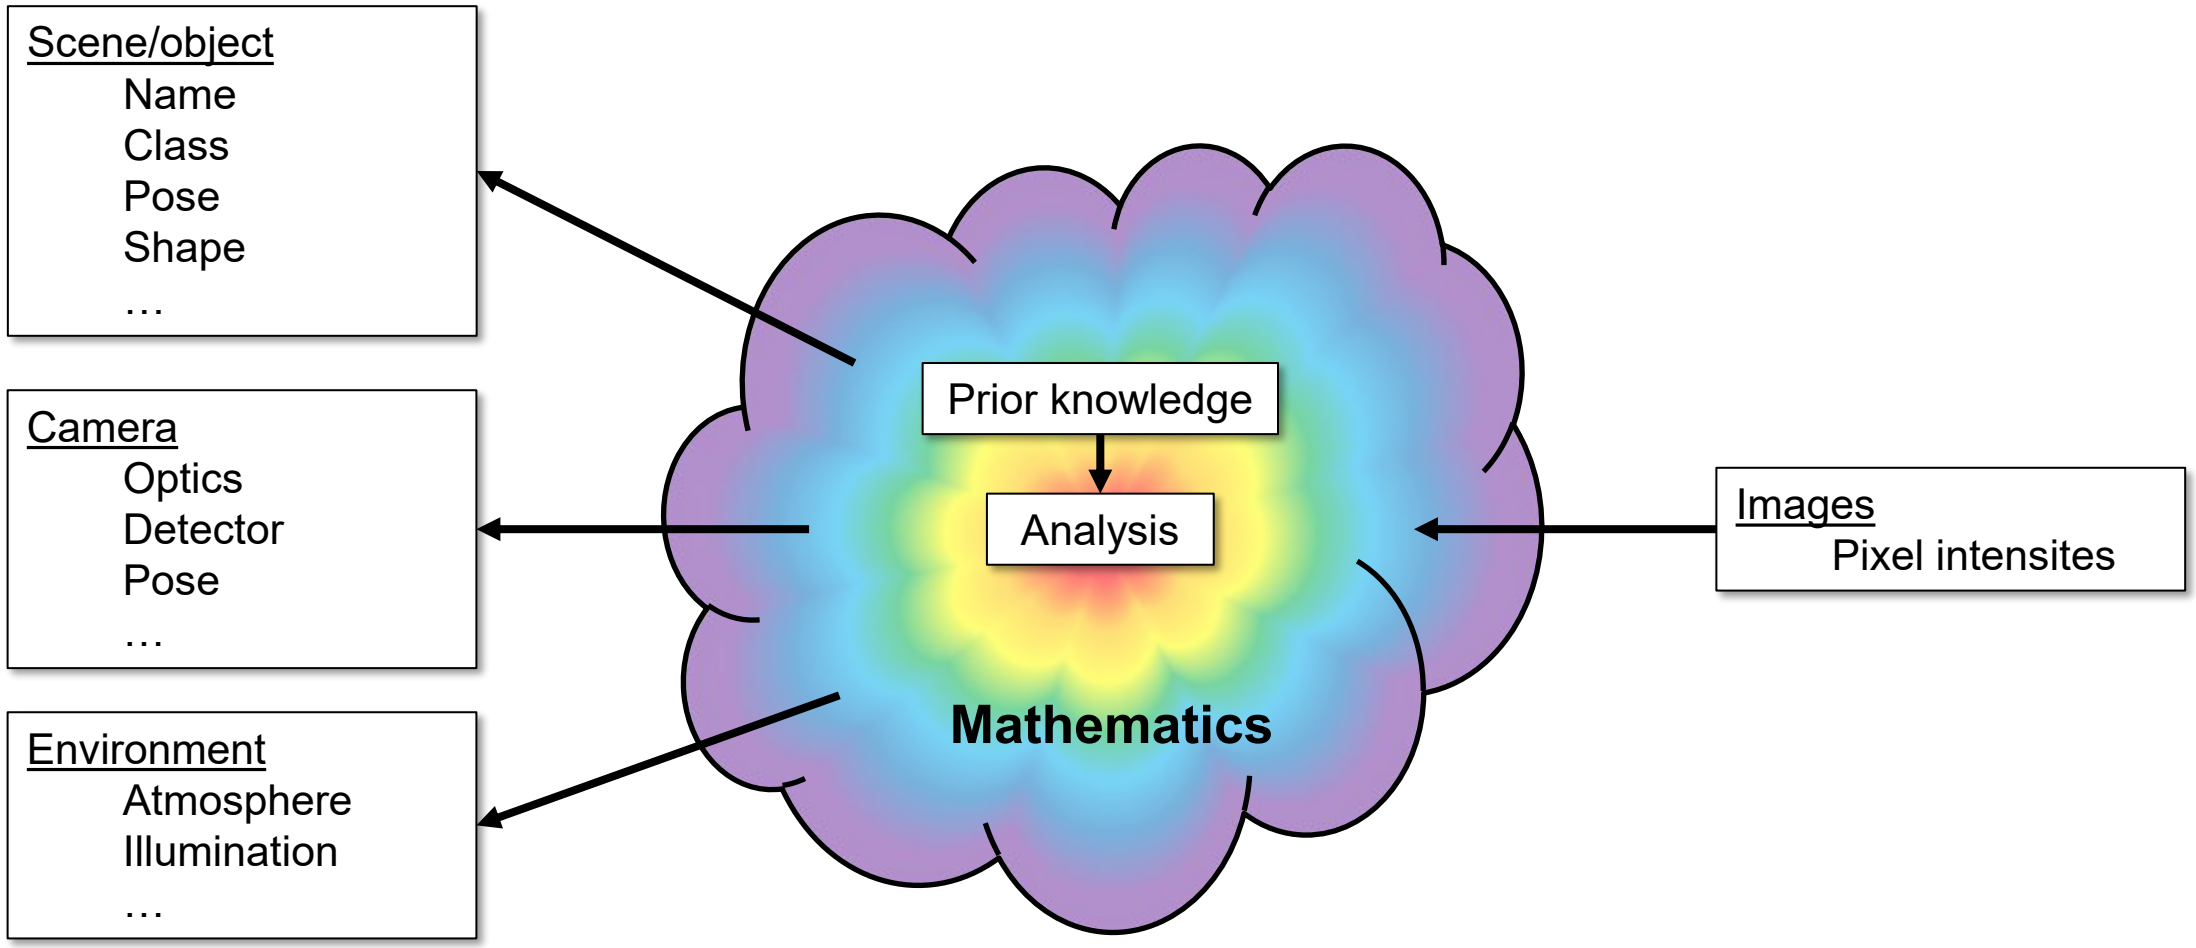
Computer vision is an *inverse problem*!



The *forward* imaging process



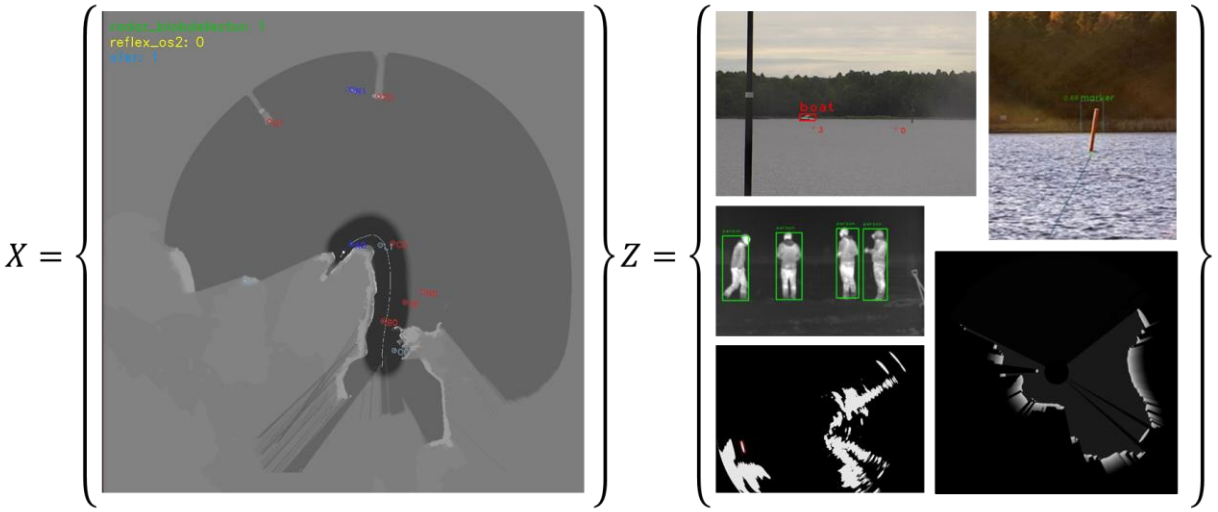
The *inverse* analysis process

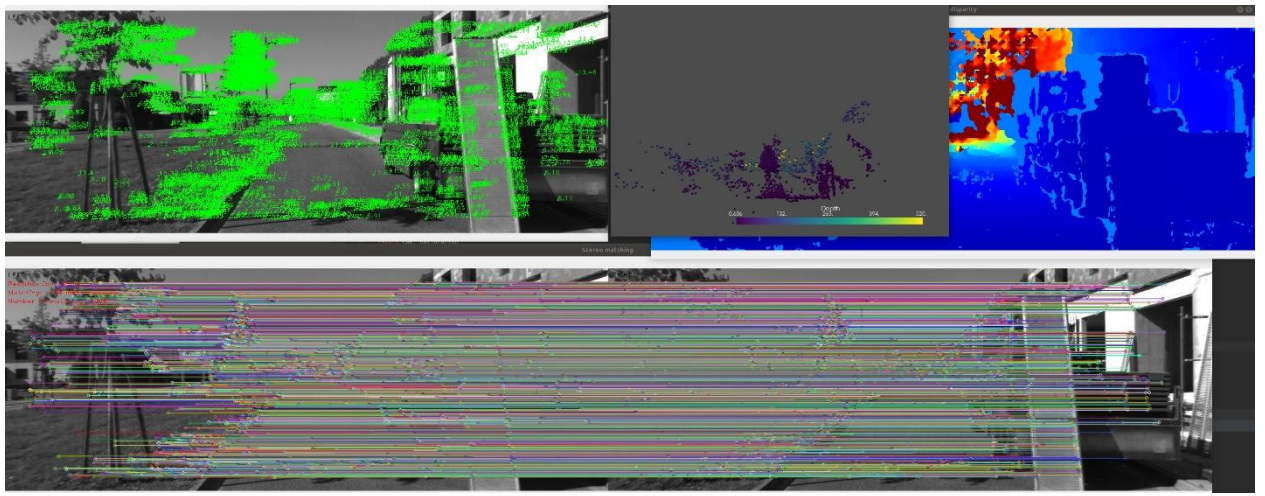


About the course

TEK5030 – Computer vision in practice!

- From pixel intensities to state estimation!
- Real-time video processing
- Supervised programming labs
- One large project





«Flipped classroom»

- Purpose
  - Get as much as possible out of a day together
- Online
  - Pre-recorded lectures most weeks
- Wednesdays 09:15-12:00
  - ~20 min lecture summary and questions
  - ~2.5 hours programming lab




```cpp
#include "opencv2/highgui.hpp"
#include <iostream>

int main()
{
    cv::VideoCapture input_stream(0);

    if (!input_stream.isOpened())
    {
        std::cerr << "Could not open camera\n";
        return EXIT_FAILURE;
    }

    const std::string window_title = "Lab 0: Introduction to OpenCV";
    cv::namedWindow(window_title, cv::WINDOW_NORMAL);

    cv::Mat frame;

    while(true)
    {
        input_stream >> frame;

        if (frame.empty())
        { break; }

        cv::imshow("cam", frame);

        if (cv::waitKey(15) >= 0)
        { break; }
    }

    return EXIT_SUCCESS;
}
```

Reflection portfolio assessment – The task

You must select **two different lecture weeks** (from weeks 2-12) to reflect upon.
For each of the two weeks, you will produce a “deliverable” that answers the following:

- **Theory:** What is the core mathematical concept?
  (Explain it in your own words, don't just copy the slide bullets).

- **Application:** How did you implement this? What were the critical steps?

- **Results & Reflection:** Show us what happened. Did it work? Why/why not? What specific parameters changed the outcome? What are the practical applications?

Reflection portfolio assessment – Deliveries

You may choose the format that best suits your learning style. You can even switch formats.

Option A: The written report/essay

- **Length:** Approx. 3–6 pages per week (including figures).
- **Format:** PDF.
- **Requirements:** Must include visual results, code snippets should be kept short and relevant.

Option B: The screencast/video

- **Length:** 5–10 minutes per week.
- **Format:** Upload to Panopto/YouTube (unlisted) and submit the link or upload a compressed video file.
- **Requirements:** Talk us through the theory and your solution. Show the code running live.

Reflection portfolio assessment – Assessment

This is a **Pass/Fail** activity.

- ✅ **Approved:** The submission demonstrates that you have engaged with the material, written your own code (or understood the provided skeleton), and can explain the link between the theory and the result.

- ❌ **Not Approved:** The submission is a copy-paste of the lecture slides, lacks visual results, shows code without explanation or involves no reflection.

Reflection portfolio assessment – Submission

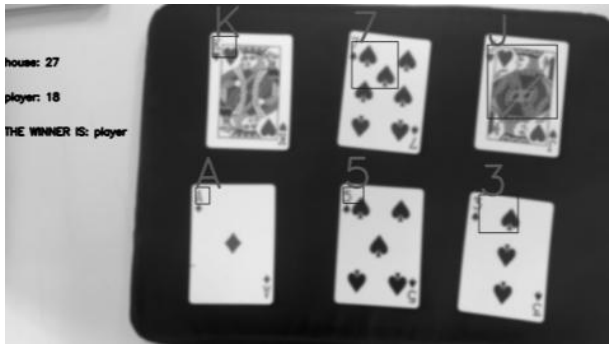
- **Deadline**: You can submit anytime during the semester, but no later than 21 April 23:59.
- **Where**: Submit via the “Reflection portfolio" assignment folder on Canvas.
- **Feedback**: You will receive brief feedback from the lecturers.

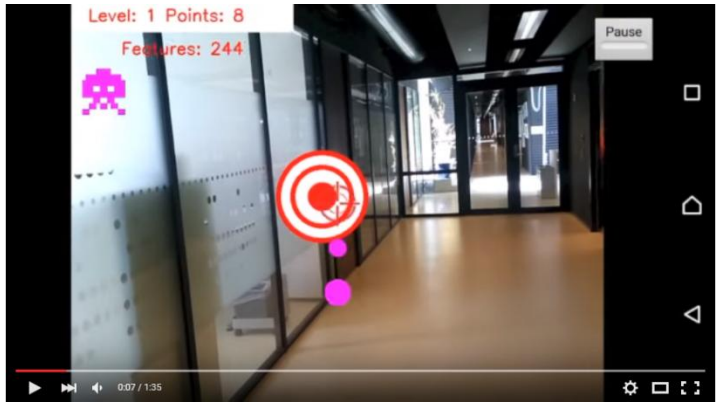
- The submissions are individual.
- If a submission is not approved, you will have to resubmit.

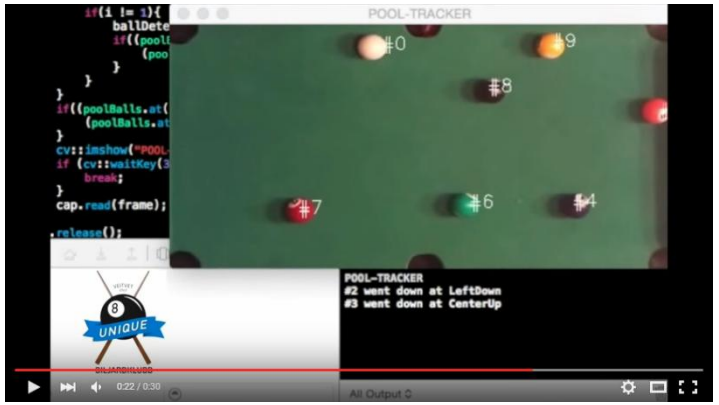
Student projects

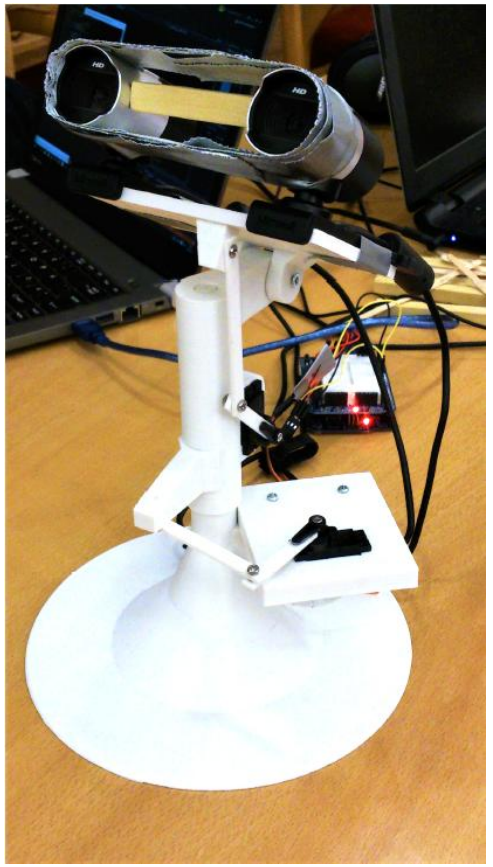
- Develop a working computer vision system that does something interesting
  - Big: More than a month
  - Approved/Not approved

- Project topic of your own choice
- Preferably in groups of up to 3 persons





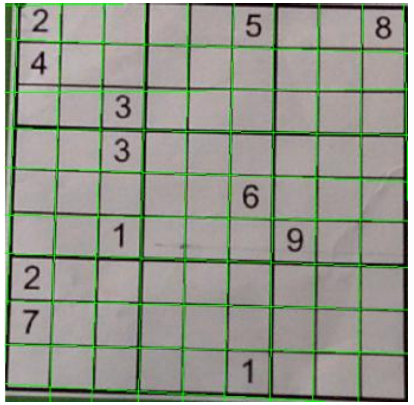


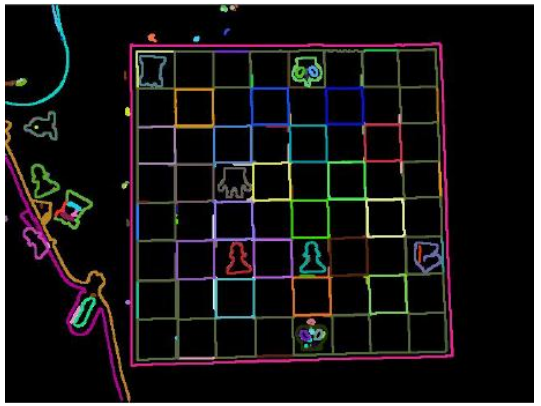


Student projects

- Develop a working computer vision system that does something interesting
  - Big: More than a month
  - Approved/Not approved

- Project topic of your own choice
- Preferably in groups of up to 3 persons



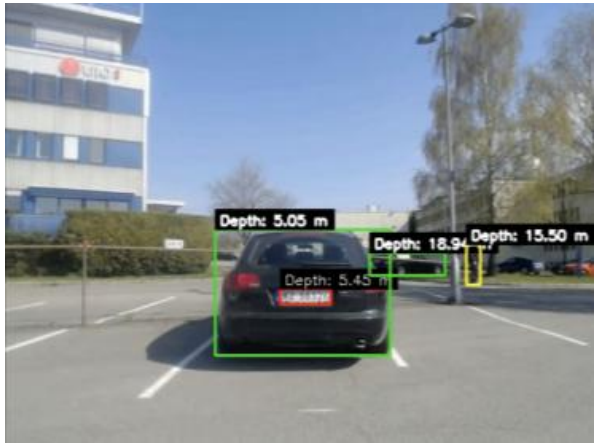






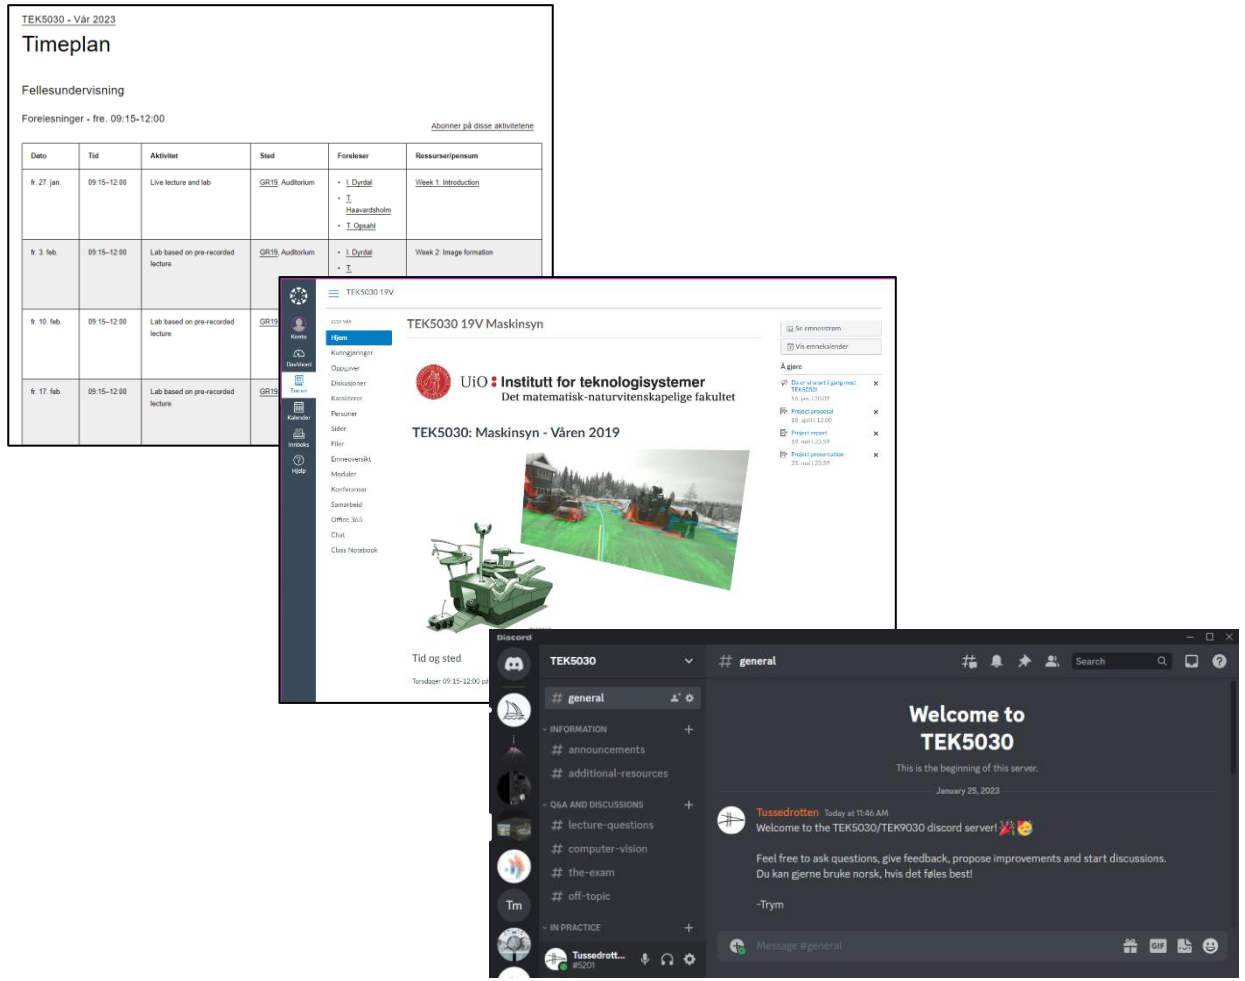






Course sites

- The semester page
  - Course plan
  - Lecture slides and videos
  - Lab exercises
- Canvas
  - Announcements
  - Course overview
  - Zoom links
  - Project deliveries
- Discussion forum?
  - Announcements
  - Questions
  - Discussions



Exam

- Written or oral exam
- First half of June



Feedback

- We encourage feedback during the course
  - We are open to adjusting it!
- Please fill out and deliver
  the course evaluation form after the course!
- Any feedback or questions now?

