Large-Scale 3D Point Cloud Processing Tutorial 2013

Features

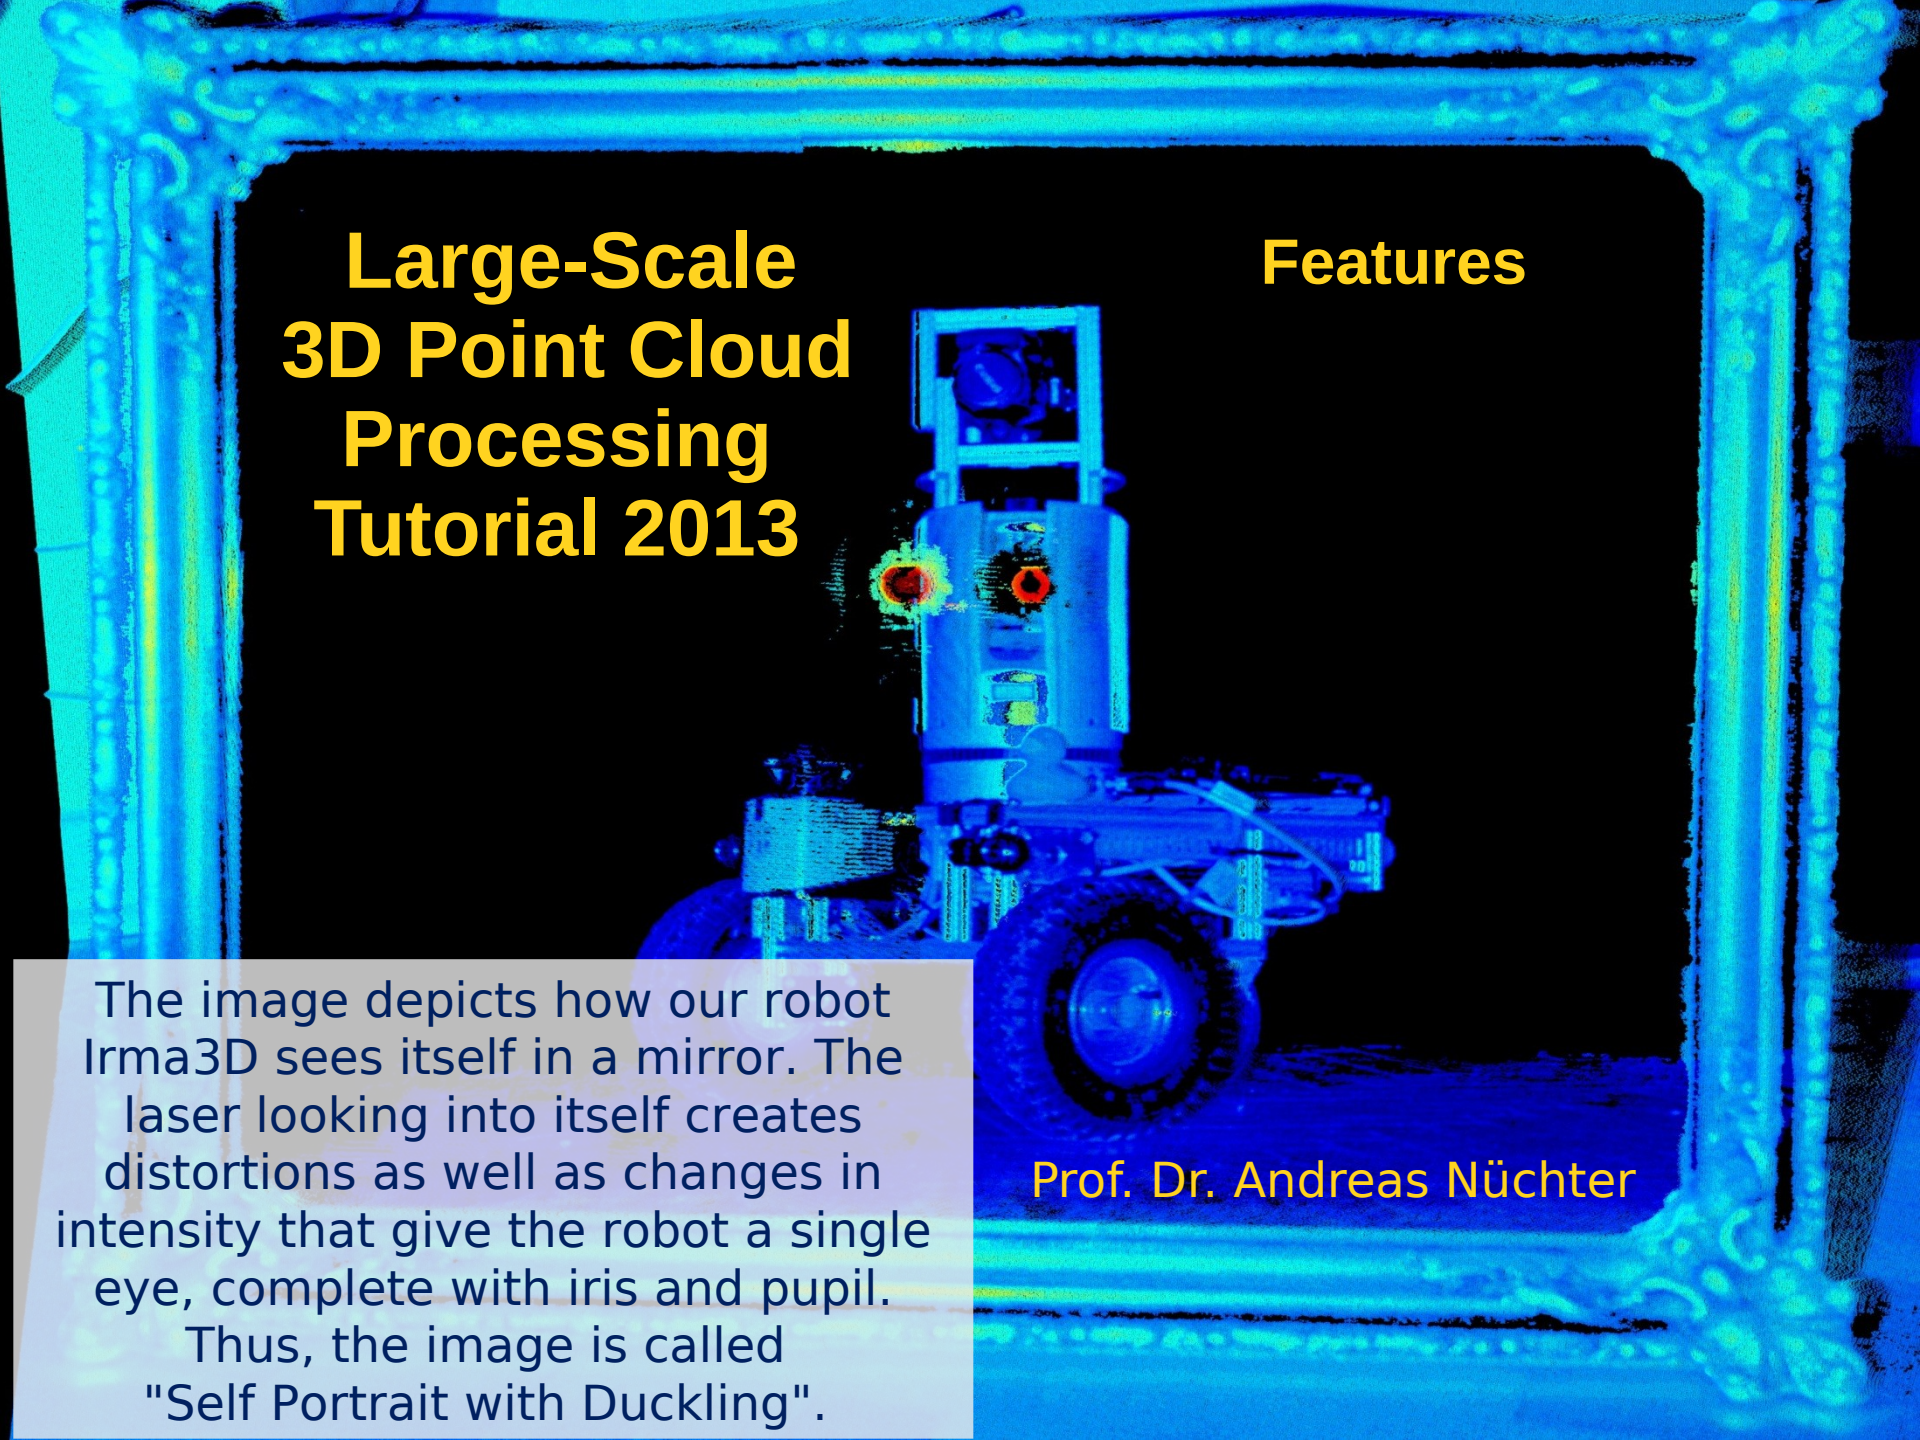

The image depicts how our robot Irma3D sees itself in a mirror. The laser looking into itself creates distortions as well as changes in intensity that give the robot a single eye, complete with iris and pupil. Thus, the image is called "Self Portrait with Duckling".

Prof. Dr. Andreas Nüchter

Features (1)

- What is a feature?

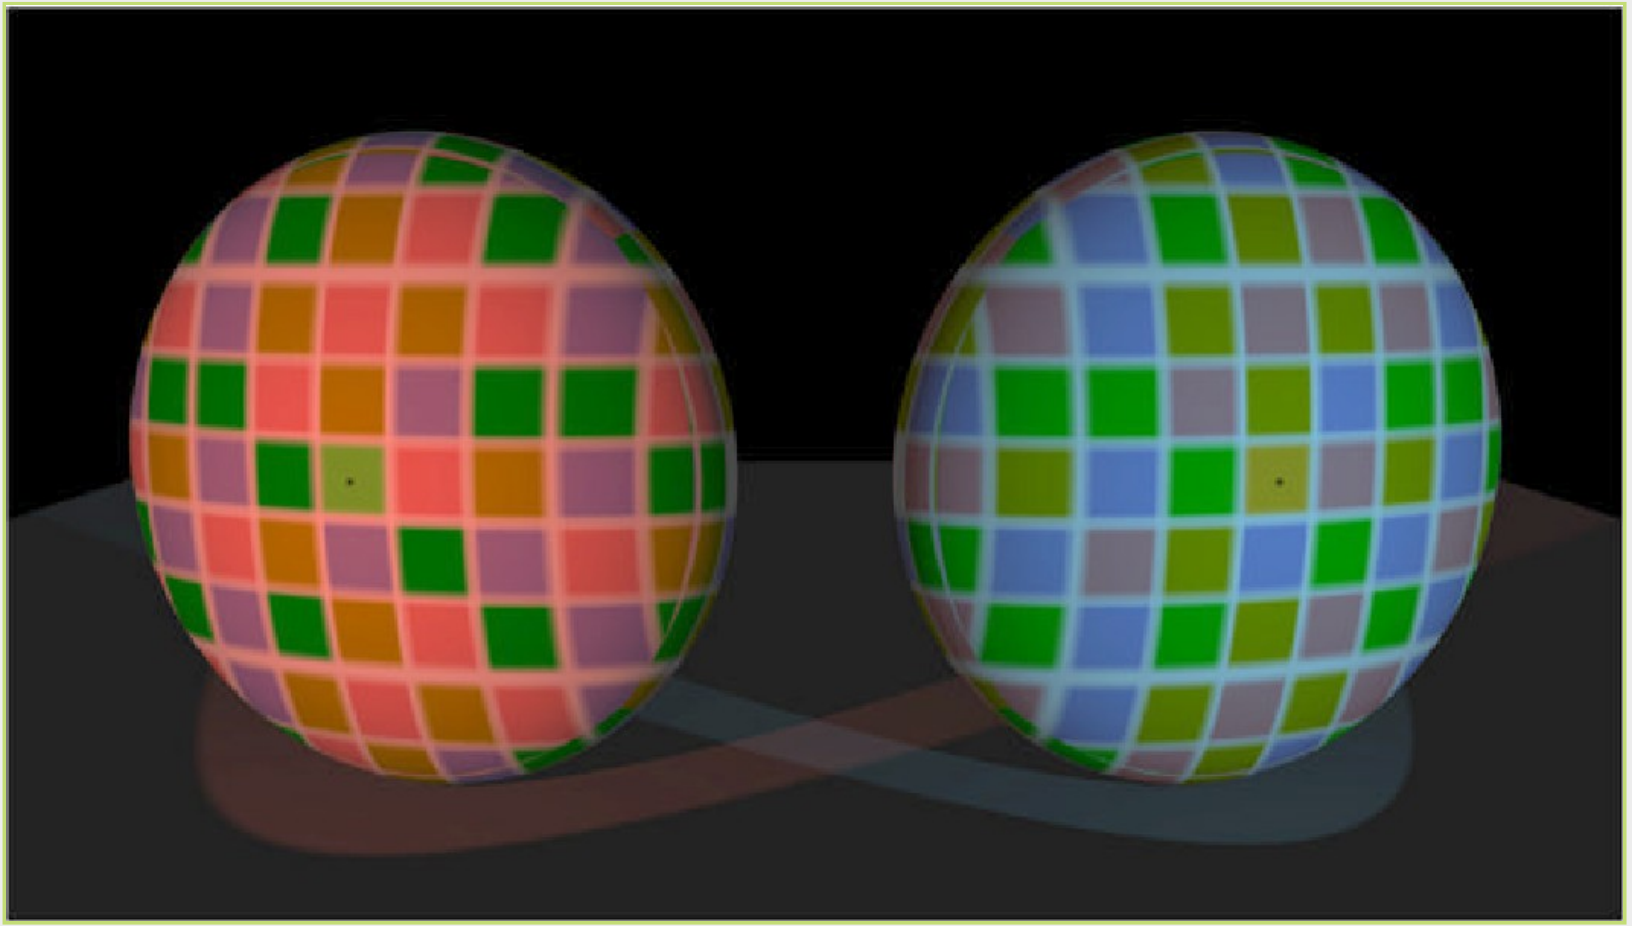

Features (1)

- What is a feature?

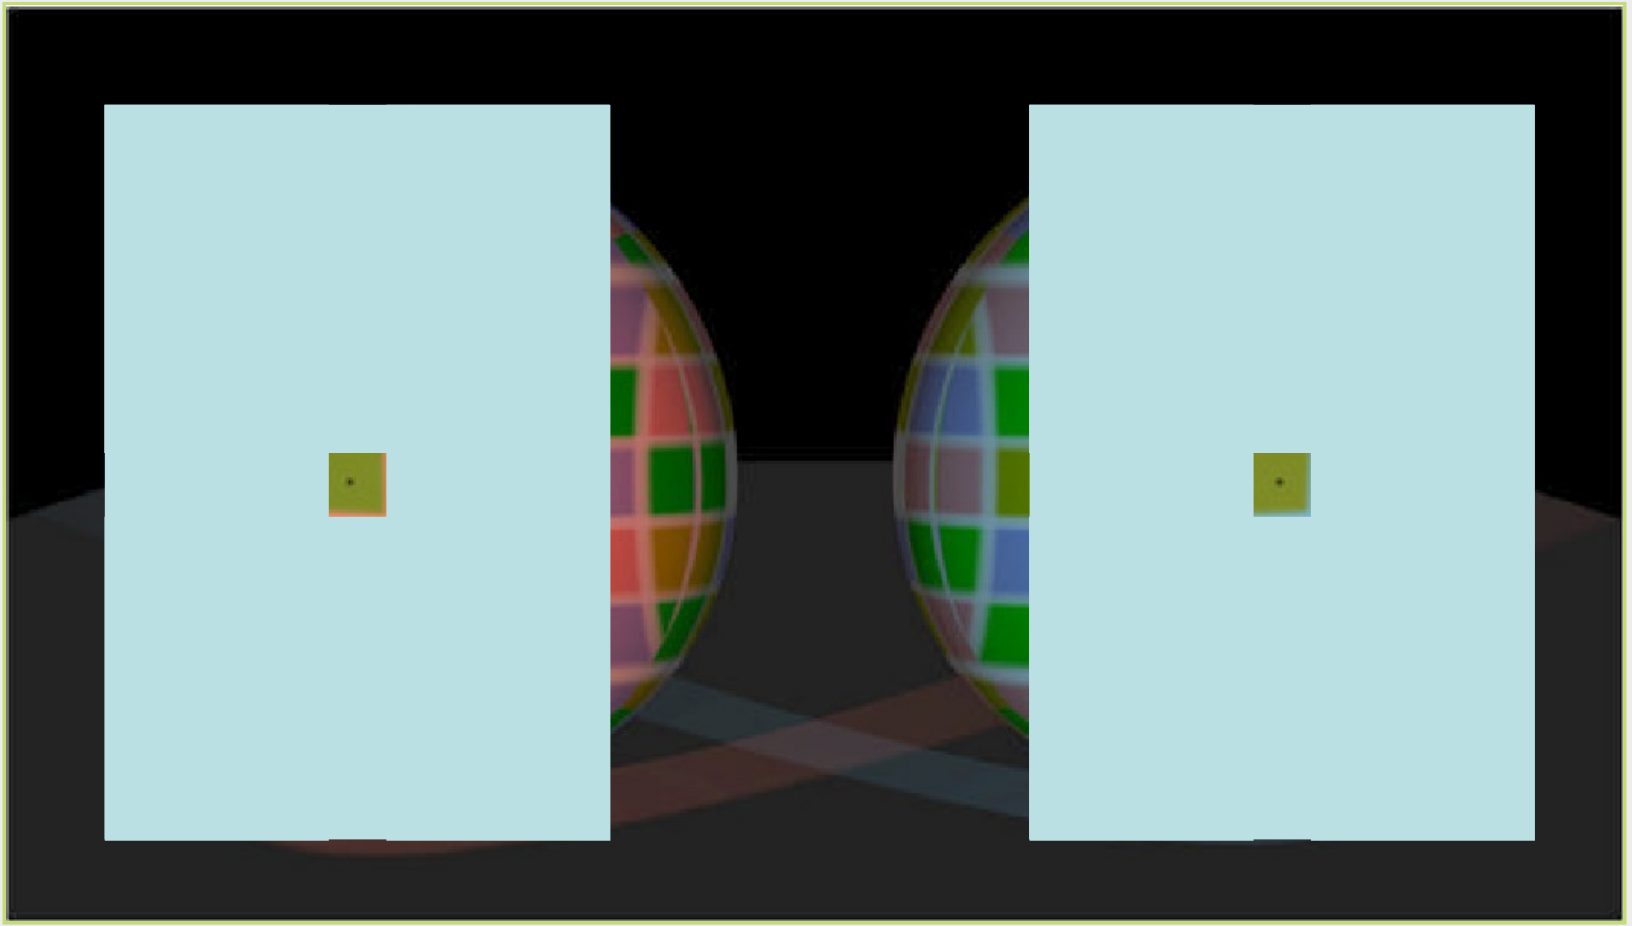

  - Local, meaningful, detectable parts of the image

Features (2)

- What is a feature?
  - Image region of sudden change
  - Shape features, contours texture features

Features (2)

- What is a feature?
  - Image region of sudden change
  - Shape features, contours texture features

Features (2)

- What is a feature?
  - Image region of sudden change
  - Shape features, contours texture features

- Why use features?
  - Information content high
  - Invariant to change of view point, illumination
  - Reduces computational burden

Work flow in Computer Vision

Image 1

Image 2

feature 1
feature 2
:
feature N

feature 1
feature 2
:
feature l N

Computer Vision Algorithm

Where are features used?

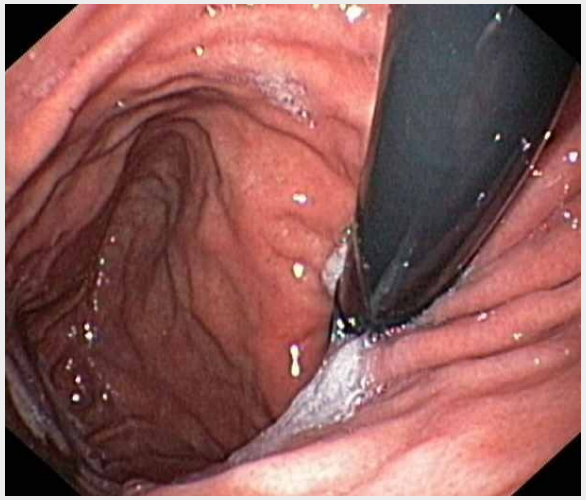

- Calibration (we have already done this)
- Image Segmentation
- Correspondence in multiple images (stereo, structure from motion)
- Object detection, classification

  ⇨ Good features support these processes

Edge- and Corner Detection

- **Goal:** Identify of a sudden change (discontinuity) in an image
- These are the locations, where the most information is “saved”
- **Example:** Sketch of an artist (But the painter uses knowledge about objects)

Already seen: Edge detection

Source: S. Thrun

From Edges to Corners

- Edge detectors perform poorly at corners.
- Corners provide repeatable points for matching, so are worth detecting.

Idea:

- Exactly at a corner, gradient is ill defined.
- However, in the region around a corner, gradient has two or more different values.

The Harris Corner Detector (1)

Form the second-moment matrix:

Sum over a small region around the hypothetical corner

Gradient with respect to x, times gradient with respect to y

$$C = \begin{bmatrix} \sum I_x^2 & \sum I_x I_y \\ \sum I_x I_y & \sum I_y^2 \end{bmatrix}$$

Matrix is symmetric

The Harris Corner Detector (2)

Simple case – first, consider case where:

$$C = \begin{bmatrix} \sum I_x^2 & \sum I_x I_y \\ \sum I_x I_y & \sum I_y^2 \end{bmatrix} = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix}$$

This means dominant gradient directions align with x or y axis

If either λ is close to 0, then this is not a corner, so look for locations where both are large.

The Harris Corner Detector (3)

General case

It can be shown that since C is rotationally symmetric:

$$C = R^{-1} \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} R$$

So every case is like a rotated version of the one on last slide.

The Harris Corner Detector (4)

$$A = \sum_u \sum_v w(u,v) \begin{bmatrix} I_x^2 & I_x I_y \\ I_x I_y & I_y^2 \end{bmatrix} = \begin{bmatrix} \langle I_x^2 \rangle & \langle I_x I_y \rangle \\ \langle I_x I_y \rangle & \langle I_y^2 \rangle \end{bmatrix}$$

- Harris and Stephens noted that exact computation of the eigenvalues is computationally expensive, since it requires the computation of a square root, and instead suggest the following function M_c, where κ is a tunable sensitivity parameter:

$$M_c = \lambda_1 \lambda_2 - \kappa (\lambda_1 + \lambda_2)^2 = \det(A) - \kappa \operatorname{trace}^2(A)$$

- Therefore, the algorithm does not have to actually compute the eigenvalue decomposition of the matrix A and instead it is sufficient to evaluate the determinant and trace of A to find corners, or rather interest points in general.

Problem: Features for Recognition

Want to find

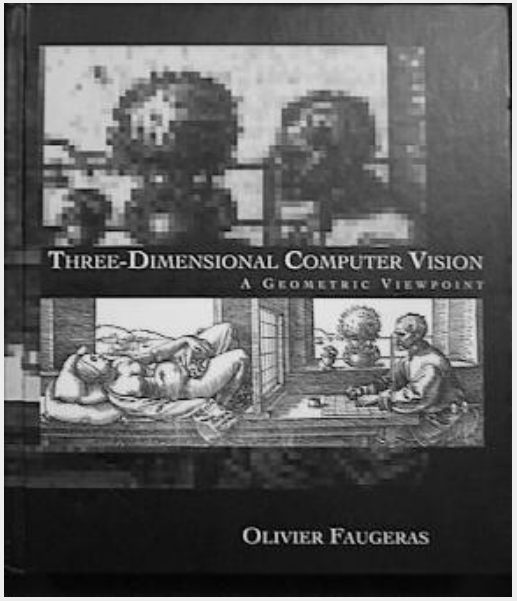

… in here

Solution: SIFT Features

- Invariants:
  - Scaling — Yes
  - Rotation — Yes
  - Illumination — Yes
  - Perspective projection — ???
- additional
  - Good localization — Yes

Switch to Hamid