Towards Semantic 3D Maps

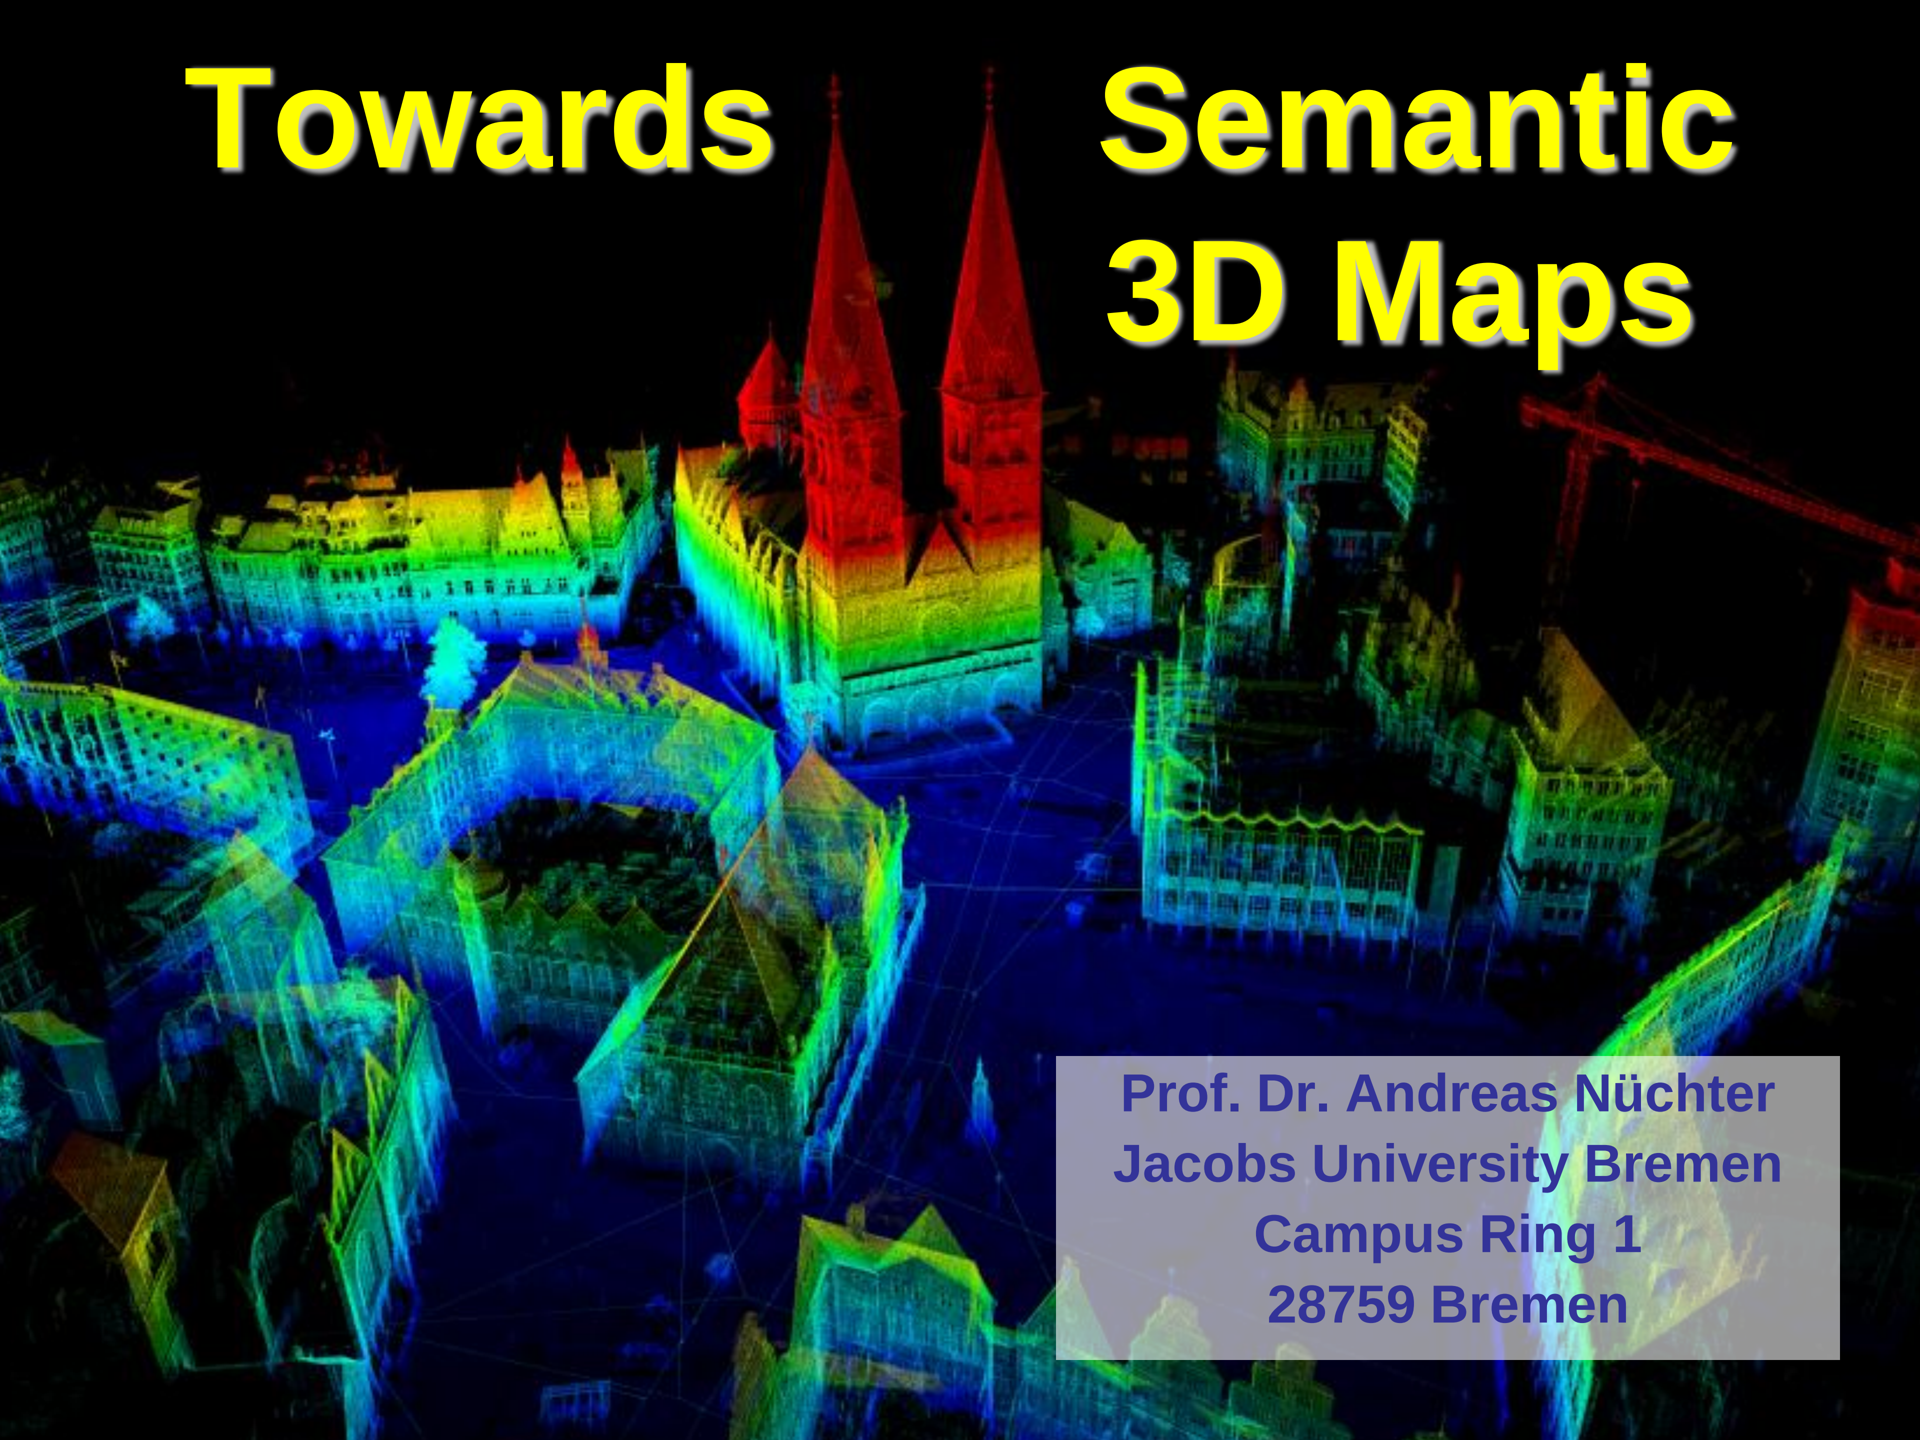

Prof. Dr. Andreas Nüchter
Jacobs University Bremen
Campus Ring 1
28759 Bremen

Acknowledgements

- I would like to thank the following researchers for joint work and inspiration in the past years:
  - Joachim Hertzberg
  - Kai Lingemann
  - Stefan Stiene
  - Hartmut Surmann
  - Oliver Wulf
  - Bernardo Wagner
  - Simone Frintrop
  - Dietrich Paulus
  - Sara Mitri
  - Dorit Borrmann
  - Jan Elseberg
  - Sebastian Thrun
  - Thomas Christaller

Washington State University
World Class. Face to Face.

GMD AiS

AIS
Fraunhofer Institut Autonome Intelligente Systeme

Uni Bonn
Rheinische Friedrich-Wilhelms-Universität

Universität Osnabrück

Universität Koblenz · Landau

Robotics
Jacobs University

About this talk…

(video)

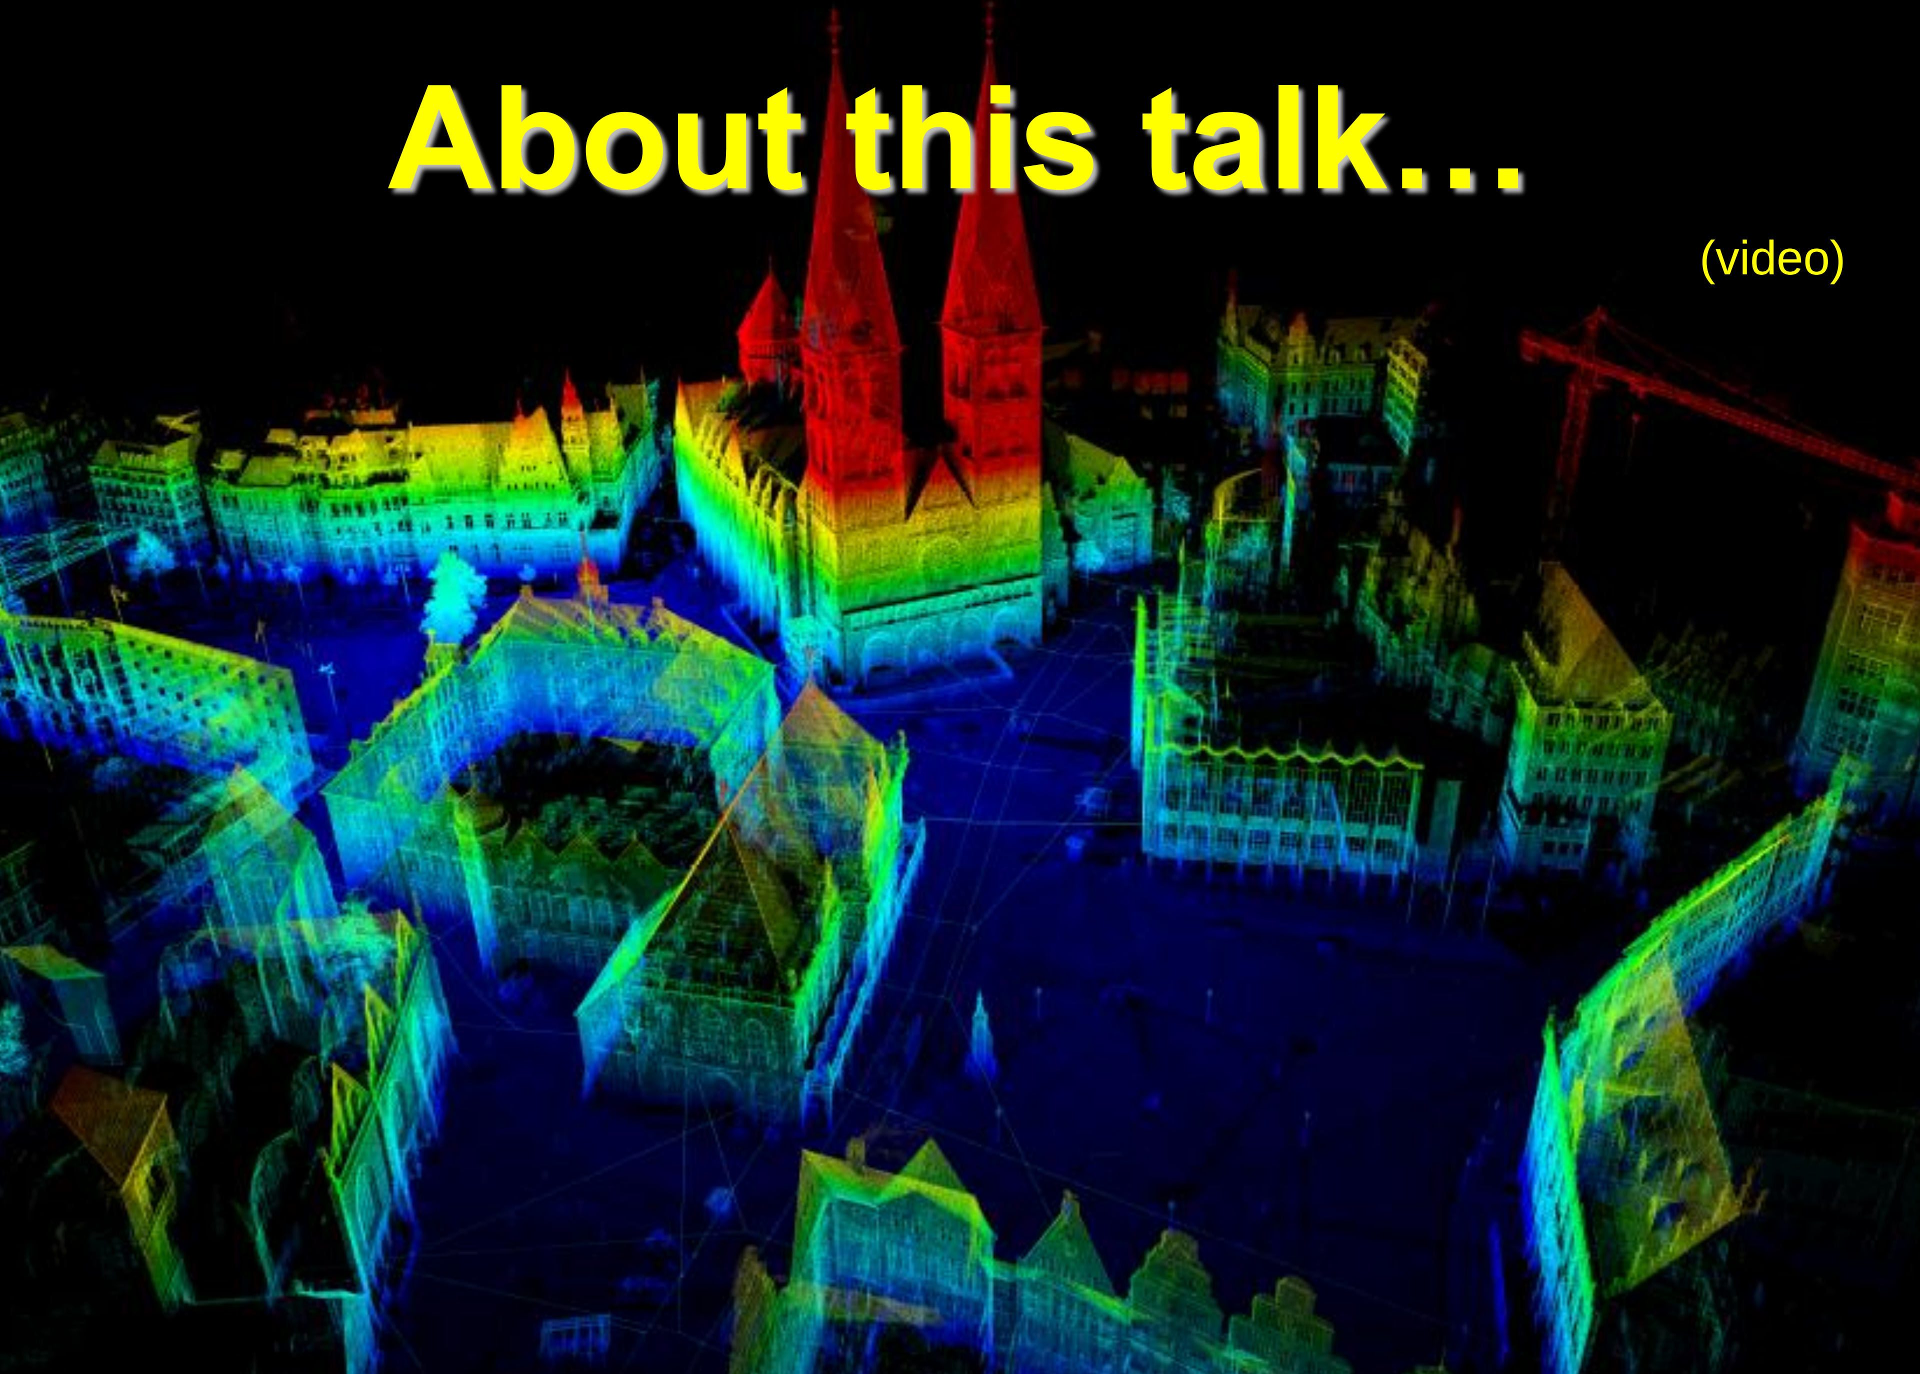

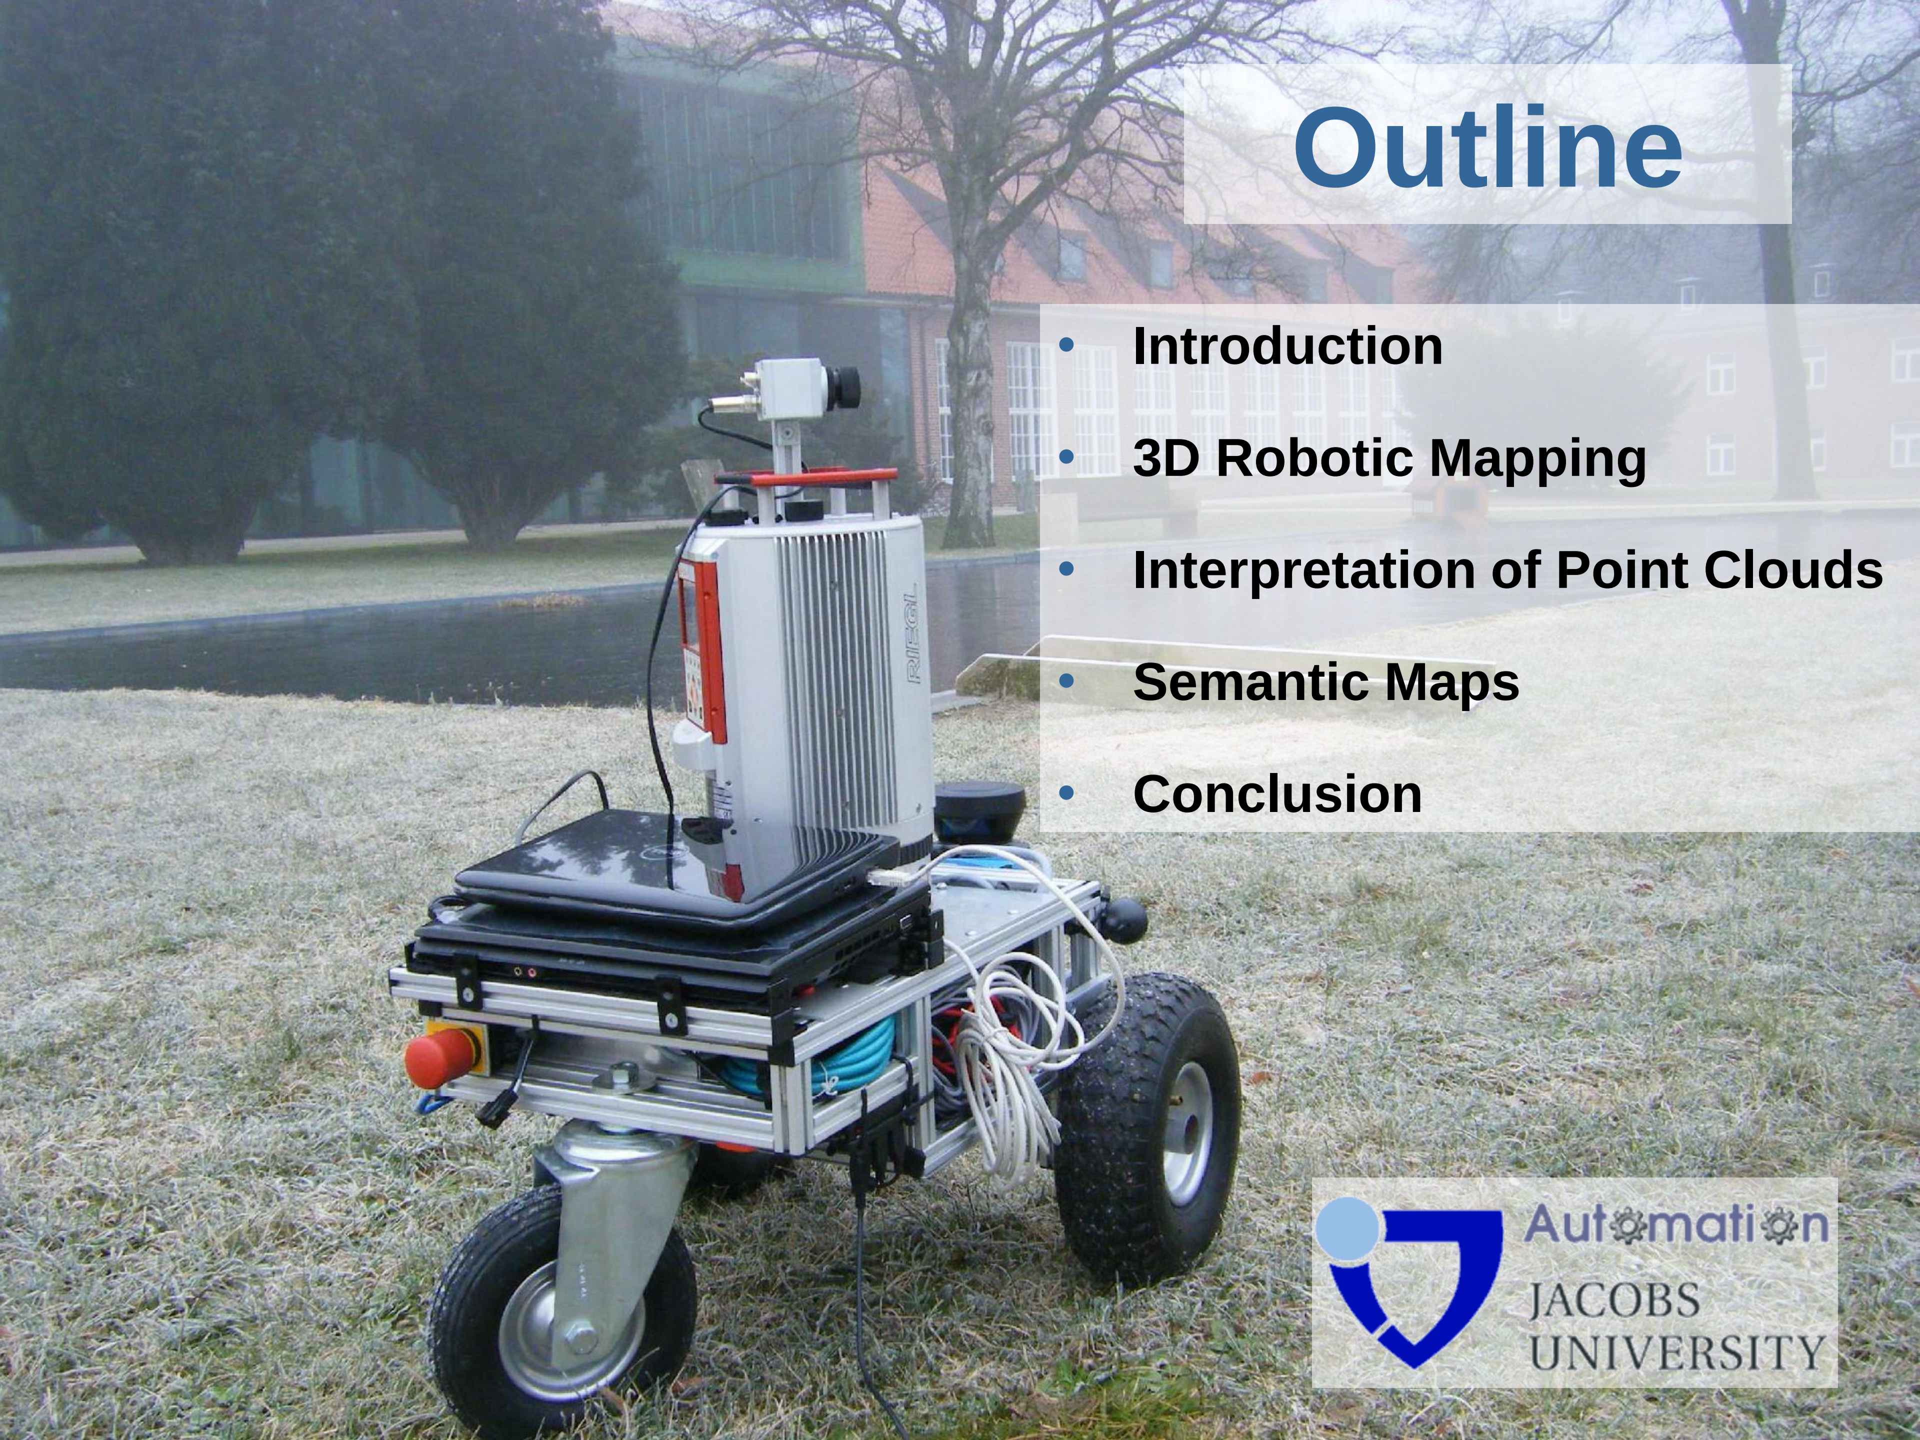

Outline

- Introduction
- 3D Robotic Mapping
- Interpretation of Point Clouds
- Semantic Maps
- Conclusion

Automation
JACOBS UNIVERSITY

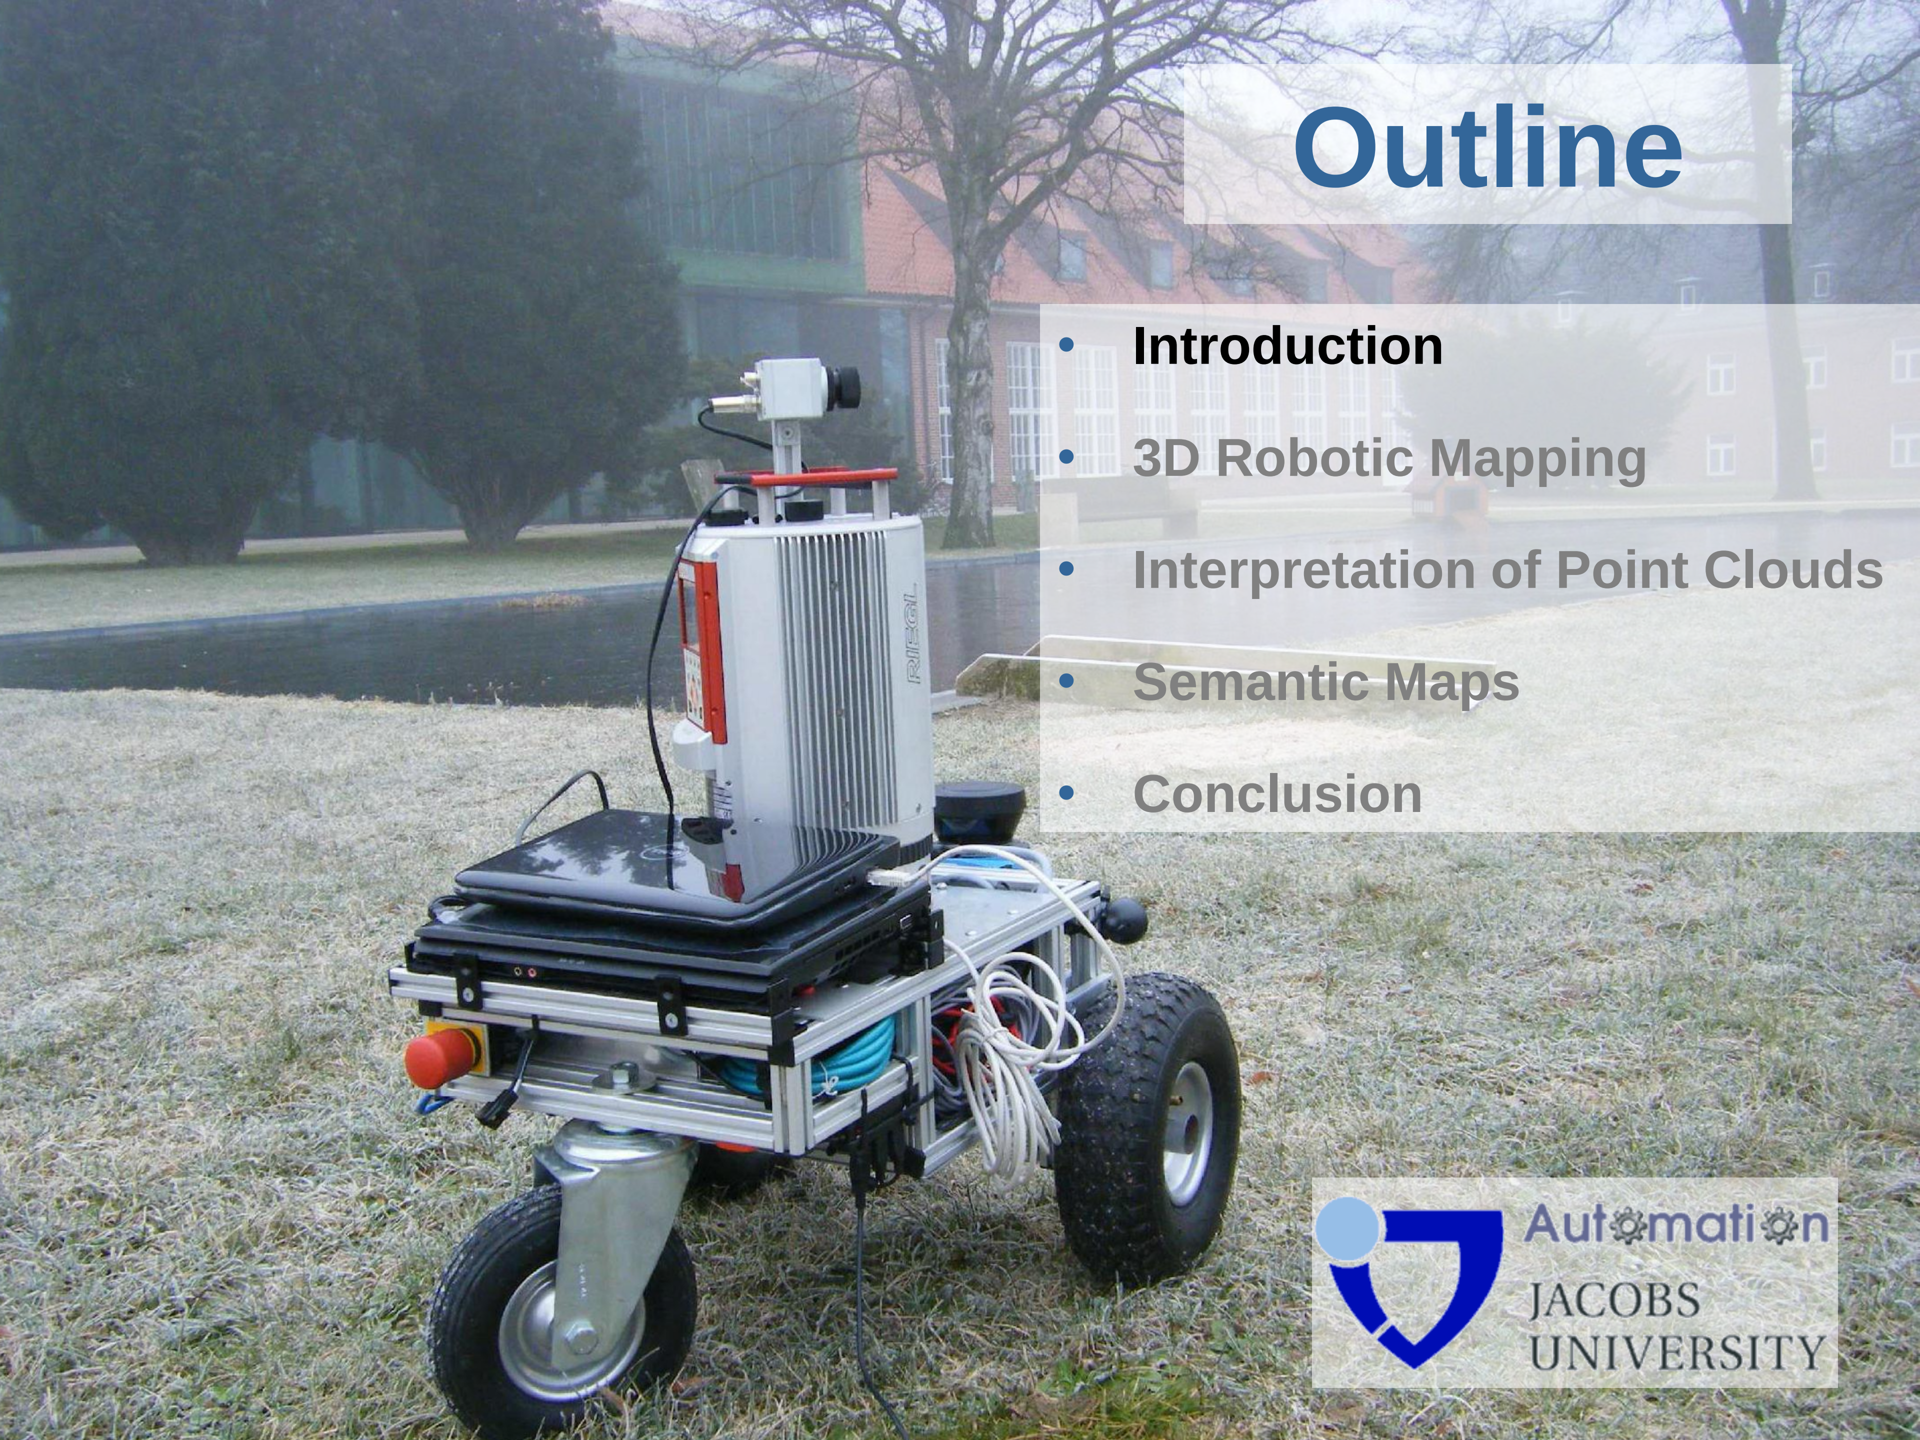

Outline

- **Introduction**
- 3D Robotic Mapping
- Interpretation of Point Clouds
- Semantic Maps
- Conclusion

Automation
JACOBS UNIVERSITY

3D Laser Scanning

- 2D Laser range finder are standard tools in automation and robotics
- 3D laser scanner for mobile robots based on SICK LMS 200

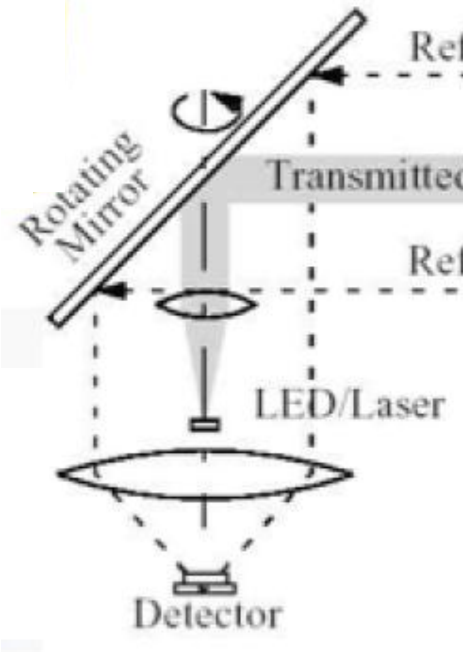

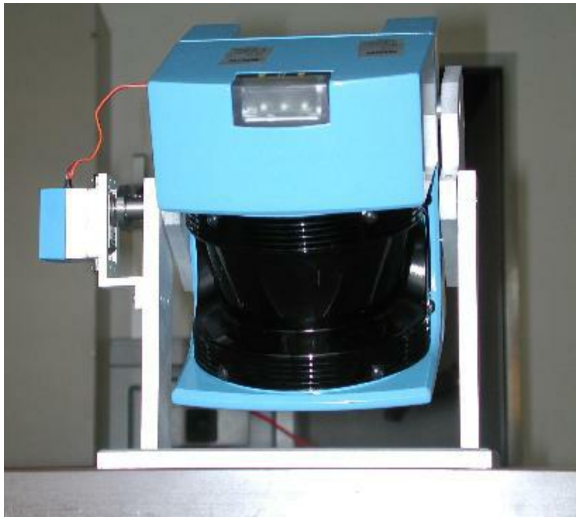

- Based on a laser scanner
- Relatively cheap sensor
- Controlled pitch motion (120° v)
- Various resolutions and modi, e.g., intensity measurement {181, 361, 721} [h] x {128, …, 500} [v] pts
- Fast measurement, e.g., 3.4 sec (181x256 points)

Kinecting

Microsoft Kinect

- Video 30 Hz
- RGB video: 8-bit VGA resolution (640 × 480 Pixel)
- Monochrome Video Stream
  (depth information): 11-bit VGA
  2048 depth values
- Depth: 1,2 – 3,5 m, (enhanced: 0,7 – 6 m)
- FOV: 57° (h) × 43°(vert)
- Tilt unit 27°
- Cost effective

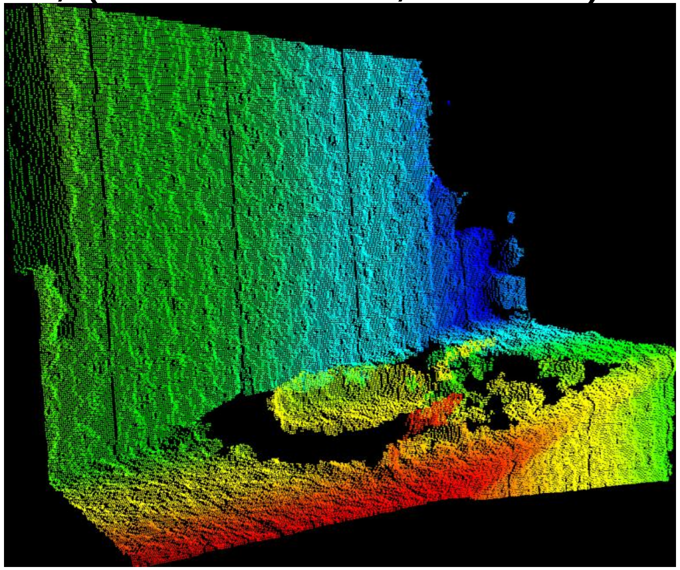

Professional 3D Scanners

- Structured light (close range, trian-gulation)
- Time-of-flight (pulsed laser vs. continuous wave)

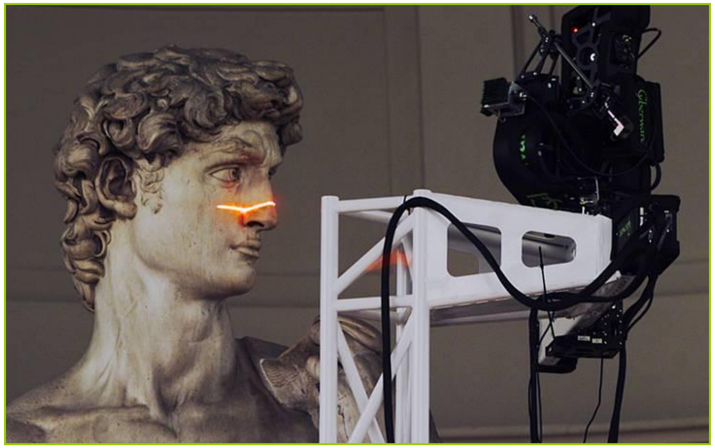

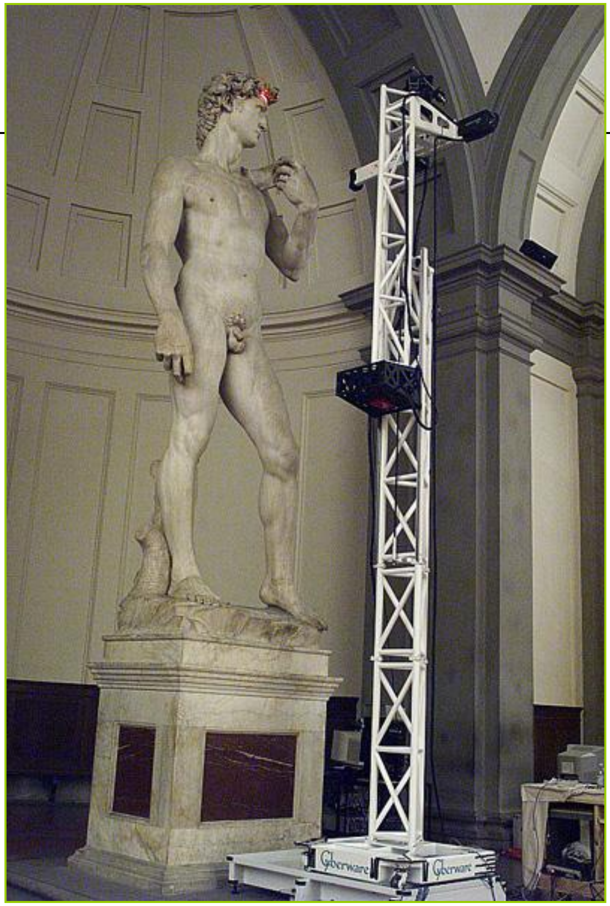

Background Robotic Mapping

- If one knows the pose of a mobile robot precisely, then the sensor readings can be used to build a map.
- The pose of a robot is easy to compute from sensor readings, given a map.
- Simultaneous Localization and Mapping (SLAM)

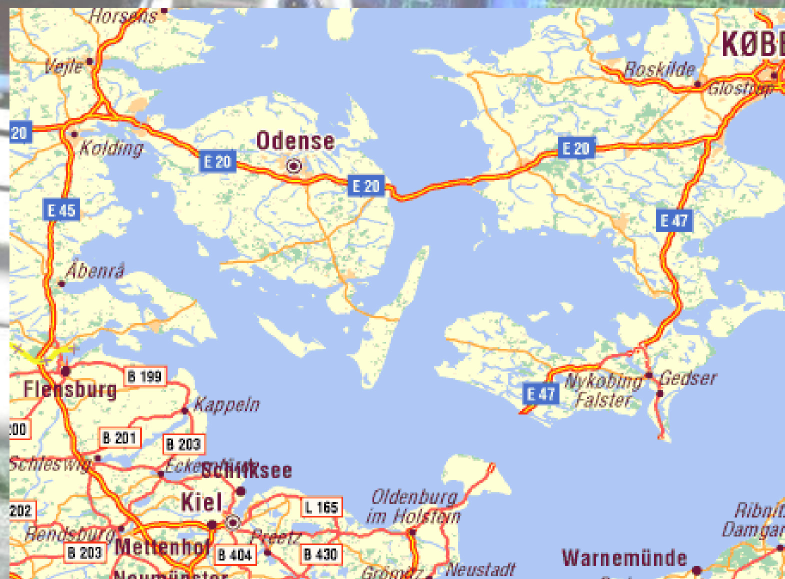

The *chicken and egg* dilemma…

Background Robotic Mapping

- If one knows the pose of a mobile robot precisely, then the sensor readings can be used to build a map.
- The pose of a robot is easy to compute from sensor readings, given a map.
- Simultaneous Localization and Mapping (SLAM)

(Fig. Burgard et al.)

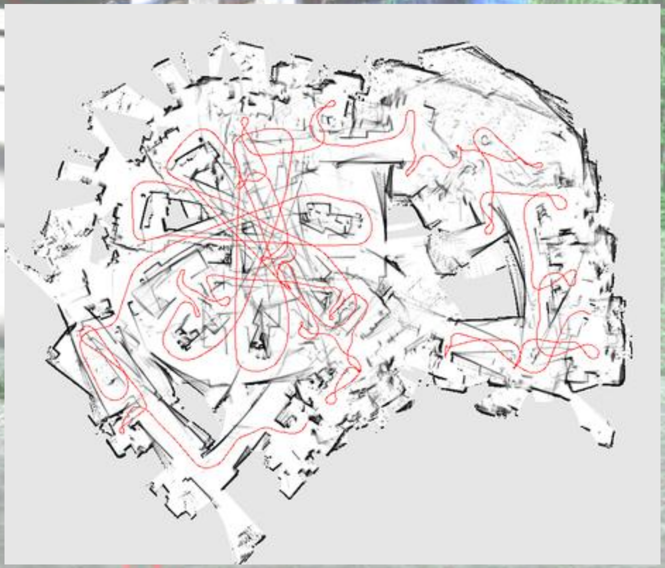

Background Robotic Mapping

- If one knows the pose of a mobile robot precisely, then the sensor readings can be used to build a map.
- The pose of a robot is easy to compute from sensor readings, given a map.
- Simultaneous Localization and Mapping (SLAM)

⇨ 2D Data, 3DoF Poses

(Fig. Burgard et al.)

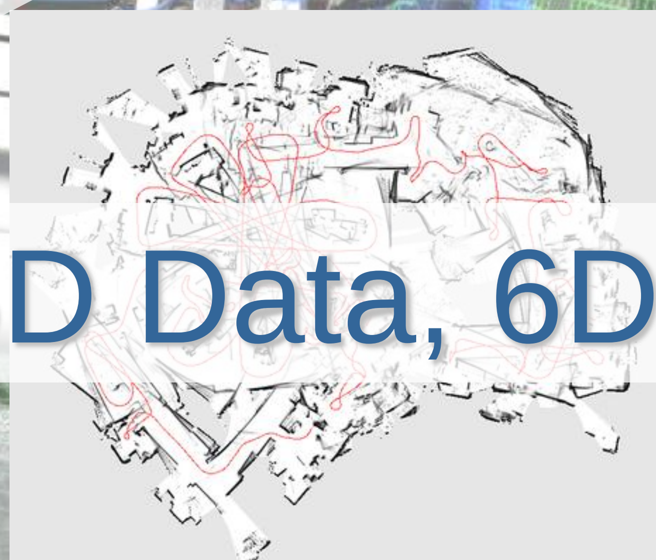

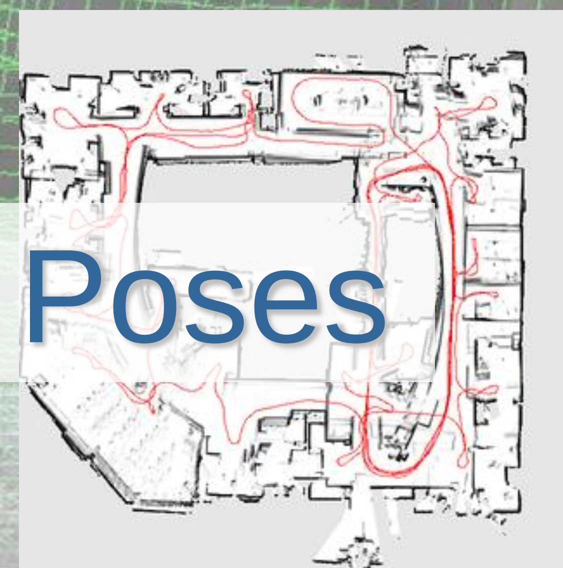

⇨ Here: 3D Data, 6D Poses

The Ariadne Robot (2002/2003)

First, we used the 3D information for obstacle avoidance. Later of we did initial 3D mapping experiments.

The motion of the robot
3 DoF

(Video Crash) (Video NoCrash)


RoboCup
201
5-11 July
Istanbul -
HOSTED BY
RoboCup
ORGANIZAT
Lenar
EVG
SanktAugustin
Heri
EVG
Sankt Augustin
Hendrik
EVG
Sankt Augustin
Morit
EVG
Sankt Augustin

The Mobile Robot Kurt3D (2004 – 2009)

- Kurt3D is a lightweight (25 kg)
- Two 90W (200W) motors, 48 NiMH a 4500mAh, C167 Microcontroller, CAN Controller, Centrino Notebook

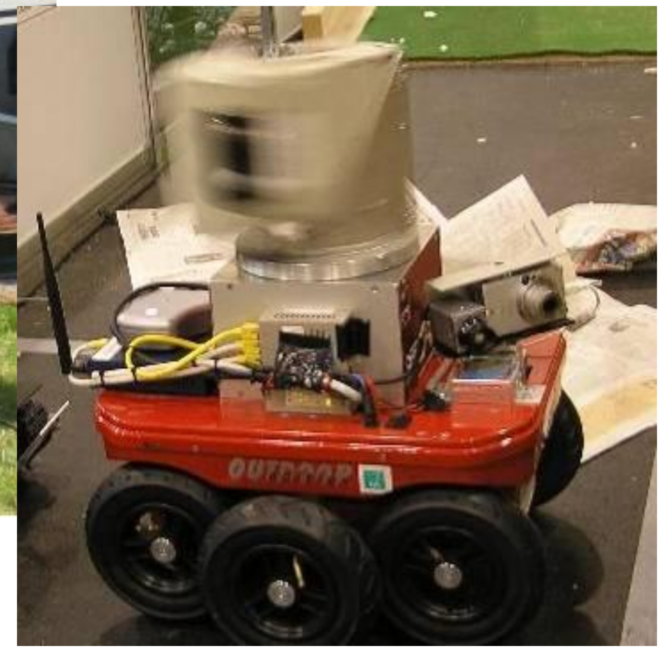

- Indoor/Outdoor versions available
- main Sensor: 3D scanner ⇨ 3D data, 6D poses

(Video Osaka)

The Mobile Robot Irma3D (2010, ongoing)

Technical Data:

(video)

- Base: volksbot
- 2D-Laserscanner: SICK LMS-100
- 3D-Scanner: Riegl VZ-400
- IMU: XSens
- Up to 3 color cameras
- Optris Imager PI
- Canon 1000D
- 12” Notebook for control
  RTLinux + ROS
- Battery capacity
  4×7.2Ah@12V

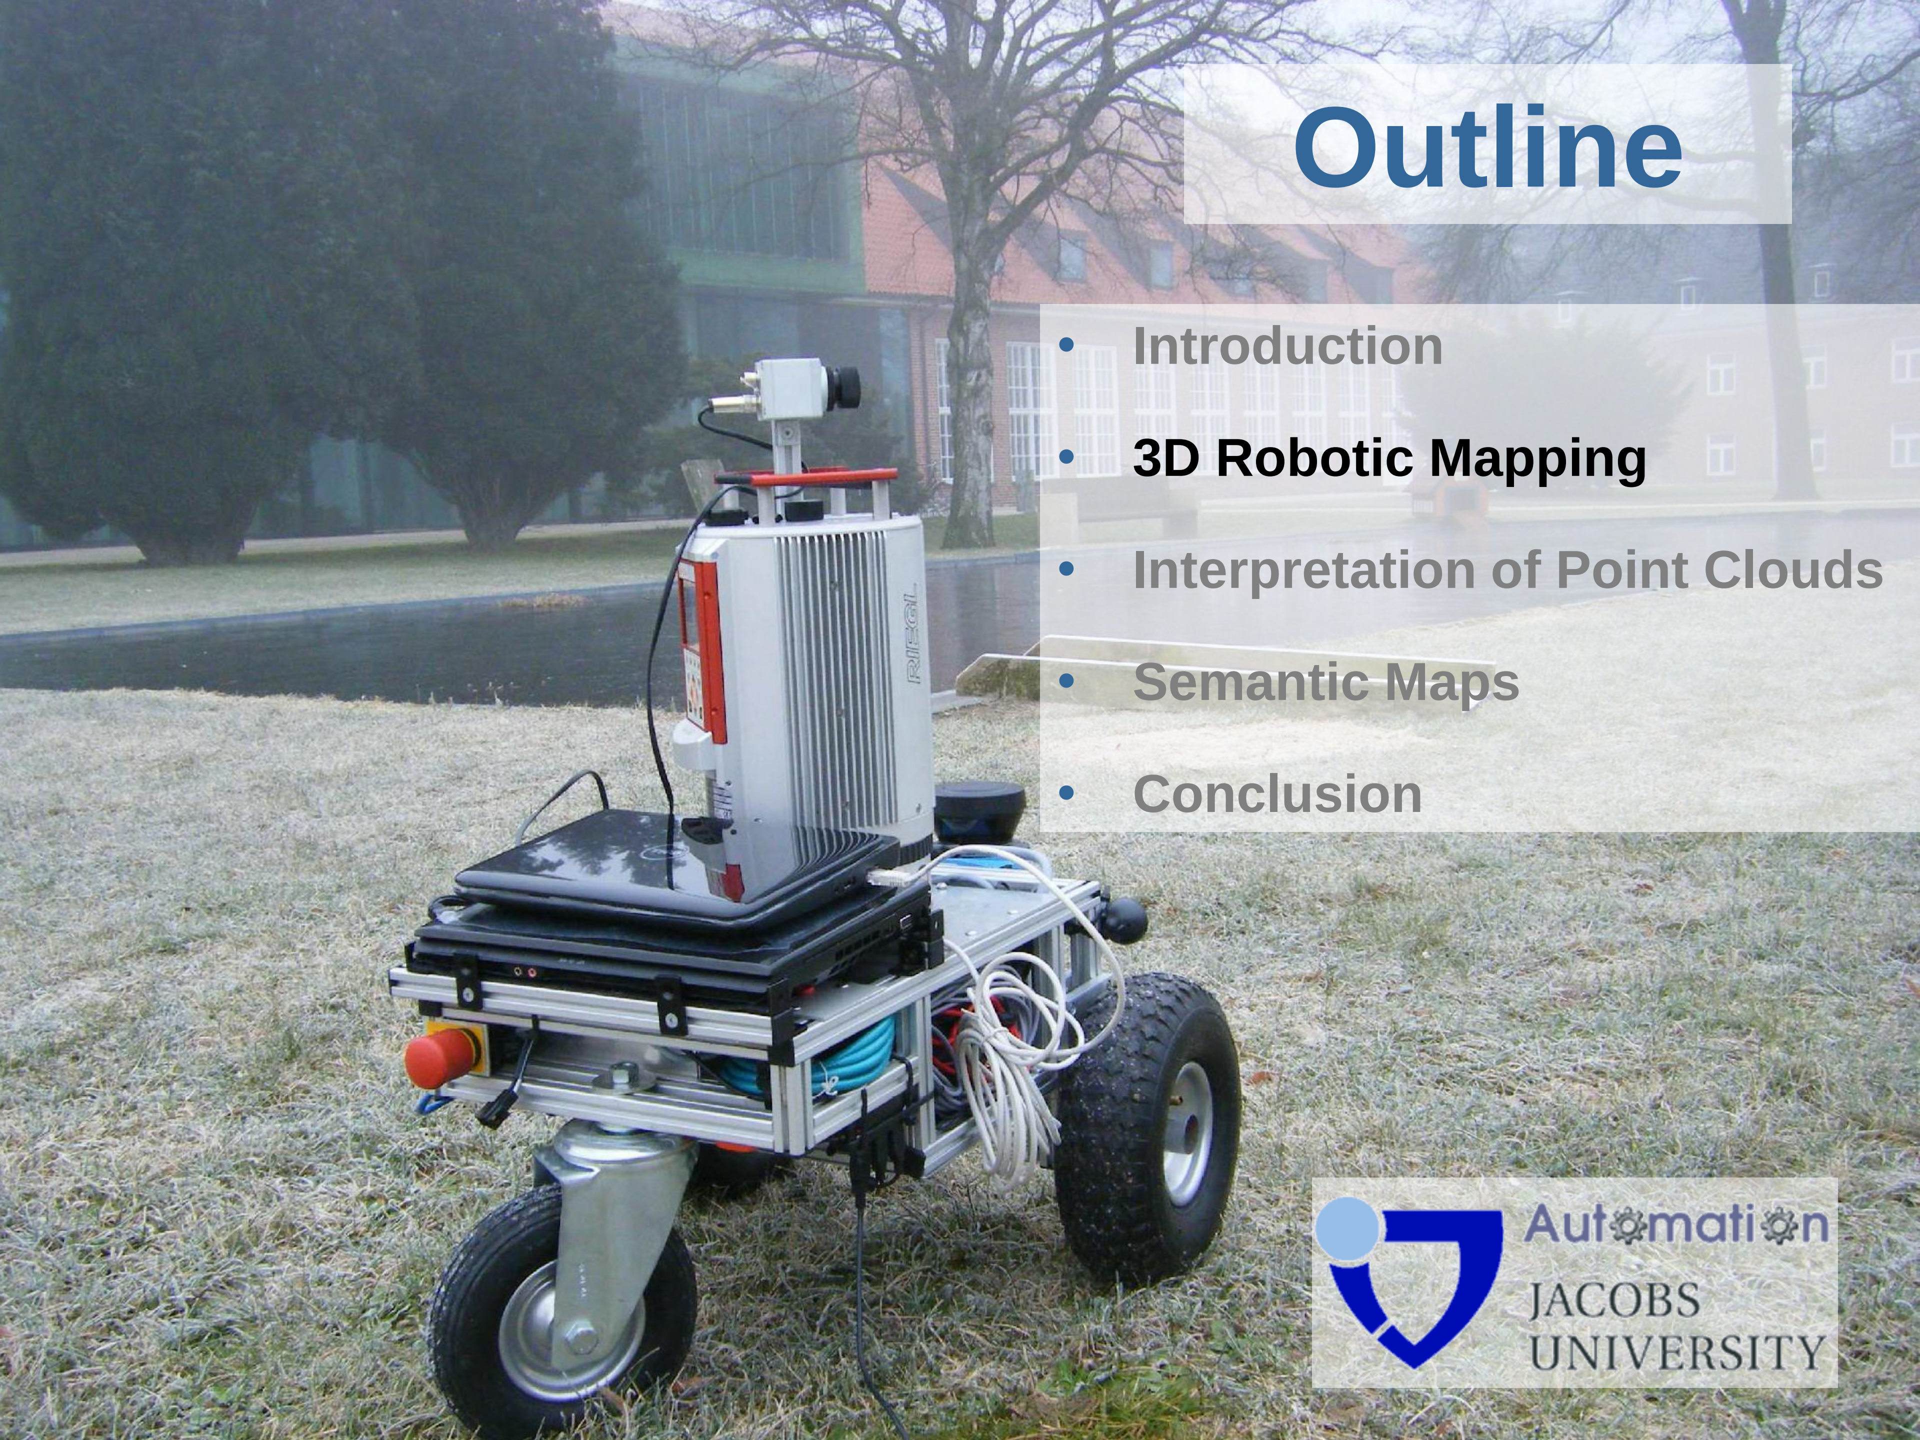

Outline

- Introduction
- **3D Robotic Mapping**
- Interpretation of Point Clouds
- Semantic Maps
- Conclusion

Automation
JACOBS
UNIVERSITY

The ICP Algorithm

Scan registration Put two independent scans into one frame of reference

Iterative Closest Point algorithm [Besl/McKay 1992]

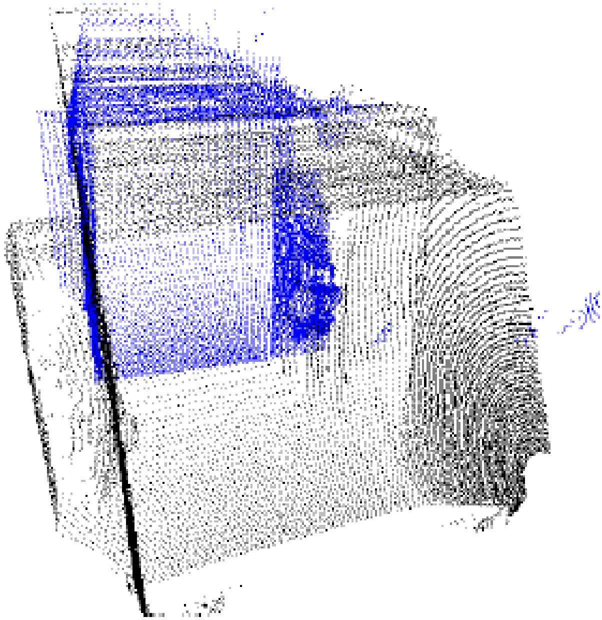

For prior point set M ("model set") and data set D

1. Select point correspondences $w_{i,j}$ in {0,1}
2. Minimize for rotation $\mathbf{R}$, translation $\mathbf{t}$

$$E(\mathbf{R},\mathbf{t}) = \sum_{i=1}^{N_m}\sum_{j=1}^{N_d} w_{i,j}\,||\mathbf{m}_i - (\mathbf{R}\mathbf{d}_j + \mathbf{t})||^2$$

3. Iterate **1.** and **2.**

Four closed form solution for the minimization

- works in 3 translation plus 3 rotation dimensions

3D Mapping Examples

CMU 3D mapping of abandoned mines

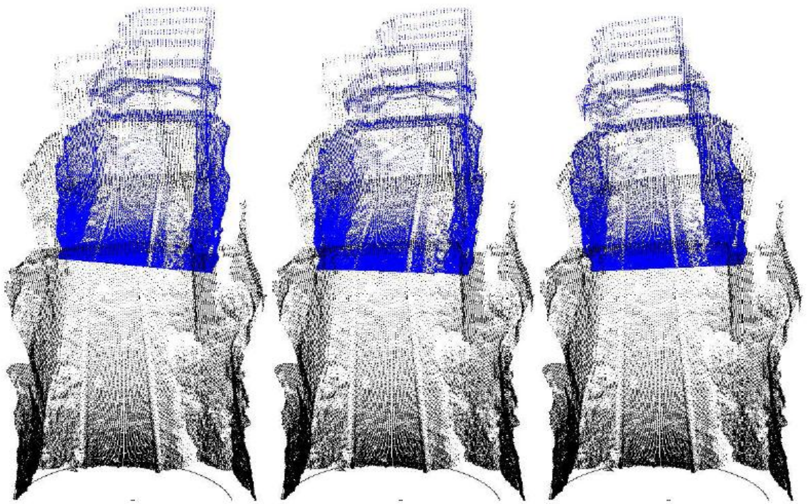

RoboCup Rescue

3D reconstruction in the context of medical imaging

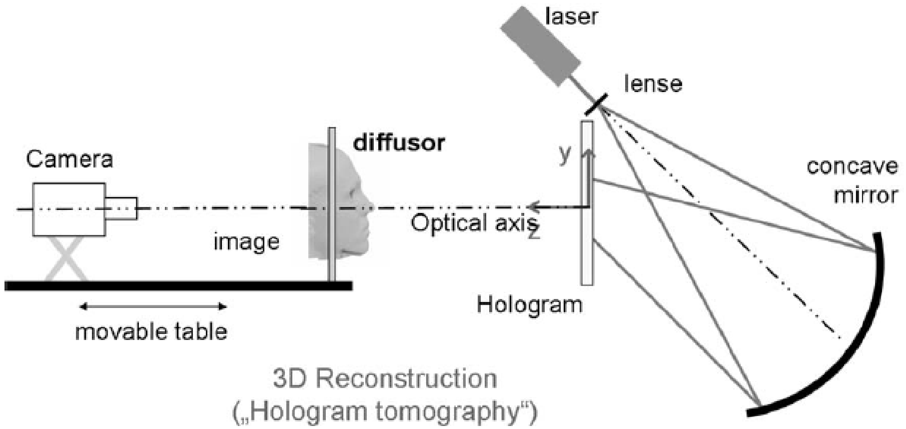

3D Reconstruction
(„Hologram tomography“)

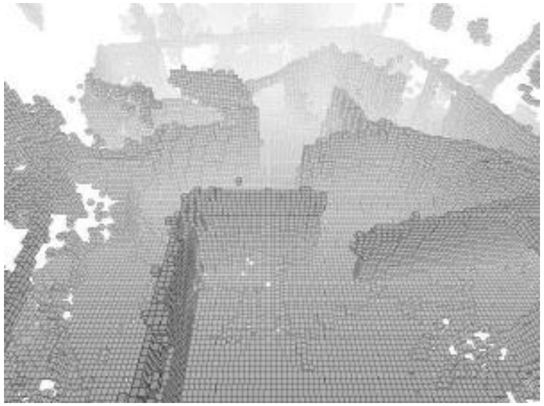

The ICP Algorithm
The global ICP Algorithm

Scan registration Put two independent scans into one frame of reference

Iterative Closest Point algorithm [Besl/McKay 1992]

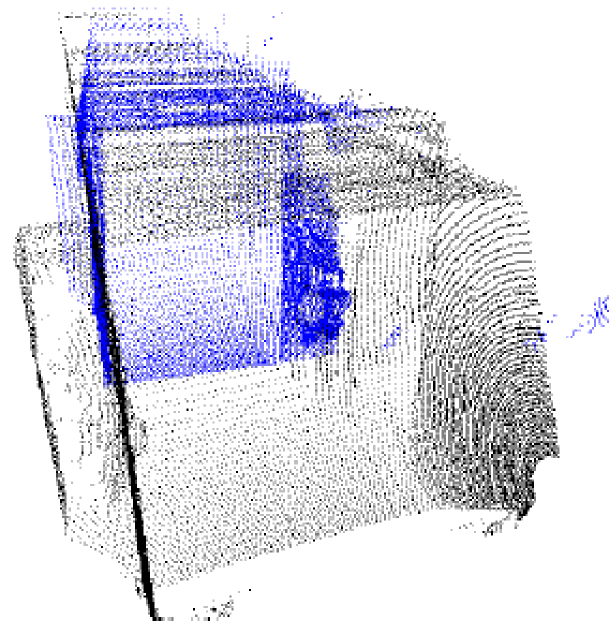

For prior point set M ("model set") and data set D

1. Select point correspondences $w_{i,j}$ in {0,1}
2. Minimize for rotation $\mathbf{R}$, translation $\mathbf{t}$

$$E(\mathbf{R},\mathbf{t}) = \sum_{i=1}^{N_m}\sum_{j=1}^{N_d} w_{i,j}\,||\mathbf{m}_i - (\mathbf{R}\mathbf{d}_j + \mathbf{t})||^2$$

3. Iterate **1.** and **2.**

Four closed form solution for the minimization

Global consistent registration

$$E = \sum_{j\rightarrow k}\sum_{i} |\mathbf{R}_j\mathbf{m}_i + \mathbf{t}_j - (\mathbf{R}_k\mathbf{d}_i + \mathbf{t}_k)|^2$$

Minimize for all rotations $\mathbf{R}$ and translations $\mathbf{t}$ at the same time

Parameterizations for the Rigid Body Transformations

$$E = \sum_{j \to k} \sum_i \left| \mathbf{R}_j \mathbf{m}_i + \mathbf{t}_j - (\mathbf{R}_k \mathbf{d}_i + \mathbf{t}_k) \right|^2$$

- Helix transformation

$$\mathbf{v}(\mathbf{p}) = \bar{\mathbf{x}} + \mathbf{x} \times \mathbf{p}.$$

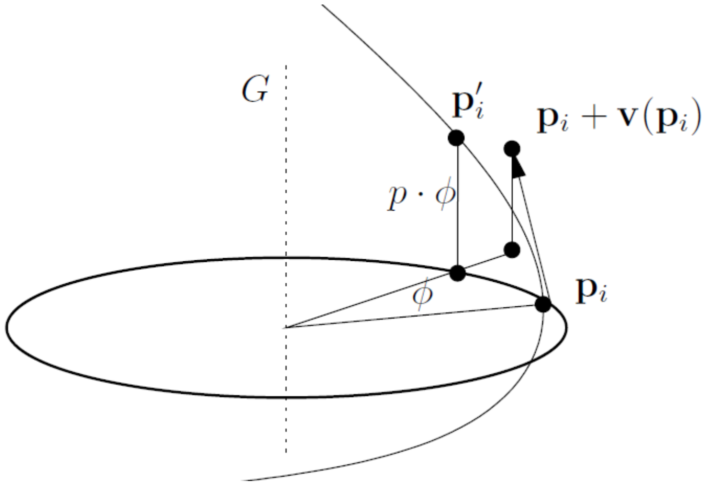

$$E = \sum_{j \to k} \sum_i \left(\mathbf{m}_i - \mathbf{d}_i + (\bar{\mathbf{x}}_j + \mathbf{x}_j \times \mathbf{m}_i) - (\bar{\mathbf{x}}_k + \mathbf{x}_k \times \mathbf{m}_i) \right)^2$$

… solving a system of linear equations

Parameterizations for the Rigid Body Transformations

$$E = \sum_{j \rightarrow k} \sum_{i} \left| \mathbf{R}_j \mathbf{m}_i + \mathbf{t}_j - (\mathbf{R}_k \mathbf{d}_i + \mathbf{t}_k) \right|^2$$

- Small angle approximation

$$\sin\theta \approx \theta - \frac{\theta^3}{3} + \frac{\theta^5}{5} - \cdots$$

$$\cos\theta \approx 1 - \frac{\theta^2}{2} + \frac{\theta^4}{4} - \cdots$$

$$\mathbf{R} \approx \begin{pmatrix} 1 & -\theta_z & \theta_y \\ \theta_x\theta_y + \theta_z & 1 - \theta_x\theta_y\theta_z & -\theta_x \\ \theta_x\theta_z - \theta_y & \theta_x + \theta_y\theta_z & 1 \end{pmatrix}$$

$$\mathbf{R} \approx \begin{pmatrix} 1 & -\theta_z & \theta_y \\ \theta_z & 1 & -\theta_x \\ -\theta_y & \theta_x & 1 \end{pmatrix}$$

… solving a system of linear equations

Parameterizations for the Rigid Body Transformations

$$E = \sum_{j \rightarrow k} \sum_{i} |\mathbf{R}_j \mathbf{m}_i + \mathbf{t}_j - (\mathbf{R}_k \mathbf{d}_i + \mathbf{t}_k)|^2$$

- Explicit modeling of uncertainties
- Assumptions: The unknown error is normally distributed

$$\begin{aligned} W &= \sum_{j \rightarrow k} (\bar{\mathbf{E}}_{j,k} - \mathbf{E}'_{j,k})^T \mathbf{C}_{j,k}^{-1} (\bar{\mathbf{E}}'_{j,k} - \mathbf{E}'_{j,k}) \\ &= \sum_{j \rightarrow k} (\bar{\mathbf{E}}_{j,k} - (\mathbf{X}'_j - \mathbf{X}'_k)) \mathbf{C}_{j,k}^{-1} (\bar{\mathbf{E}}'_{j,k} - (\mathbf{X}'_j - \mathbf{X}'_k)). \end{aligned}$$

$$E_{j,k} = \sum_{i=1}^{m} \|\mathbf{X}_j \oplus \mathbf{d}_i - \mathbf{X}_k \oplus \mathbf{m}_i\|^2 = \sum_{i=1}^{m} \|\mathbf{Z}_i(\mathbf{X}_j, \mathbf{X}_k)\|^2$$

Comparisons of the Parametrizations

Global ICP	Classical Pose GraphSLAM
• Gaussian noise in the „3D Point Cloud“ space	• Gaussian noise in the space of poses
• Locally optimal	• Gradient descent needed
• ICP-like iterations using new point correspondences	• ICP-like iterations using new point correspondences needed as well

- Riegl Laser Measurement GmbH

(Video courtesy of Riegl)

(Video 1) (Video 2) (Video 3)

Closed Loop Detection and Global Relaxation

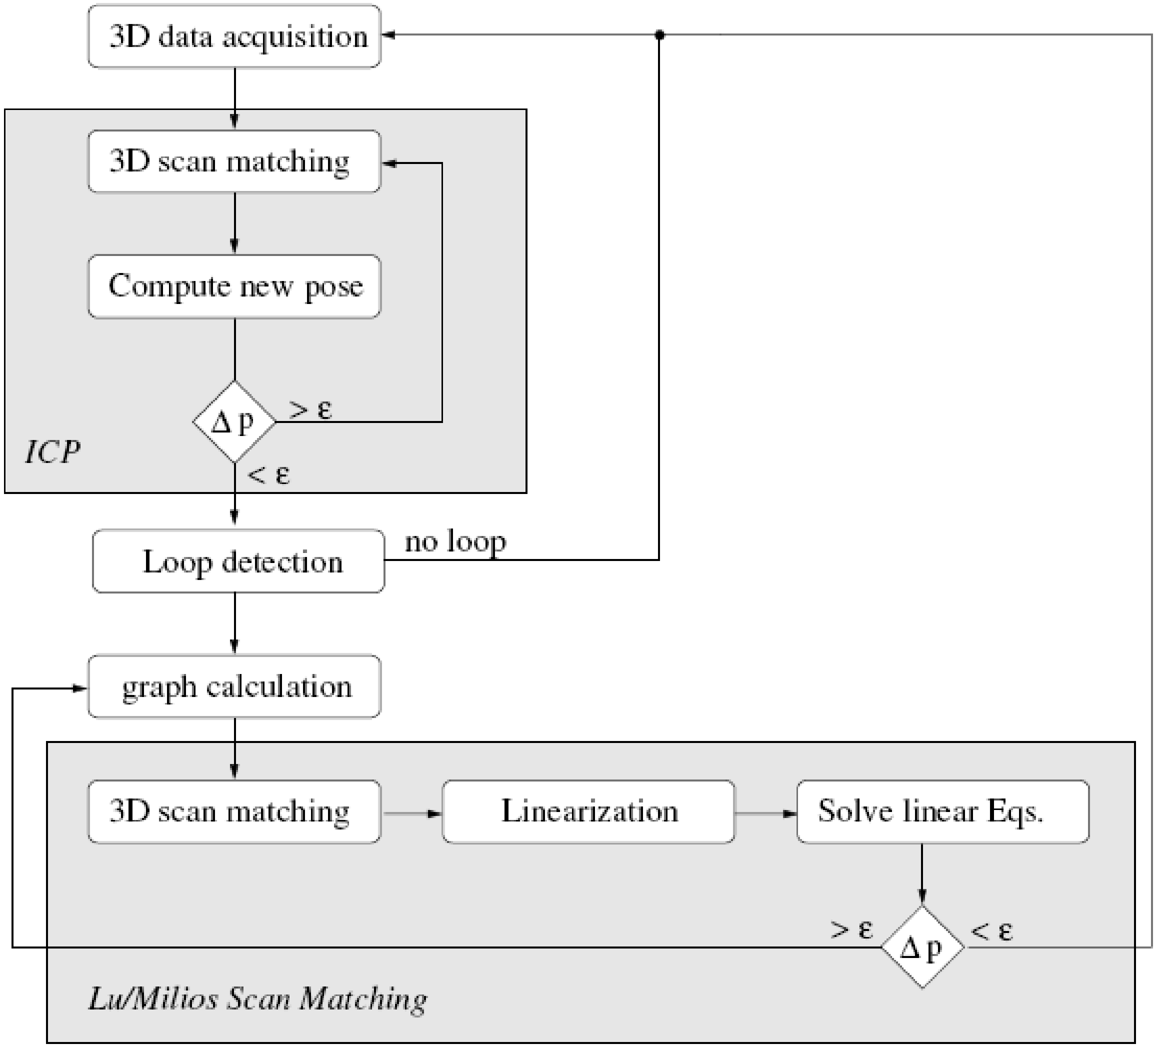

6D SLAM – Full Example

- Leibniz University Hannover (RTS)

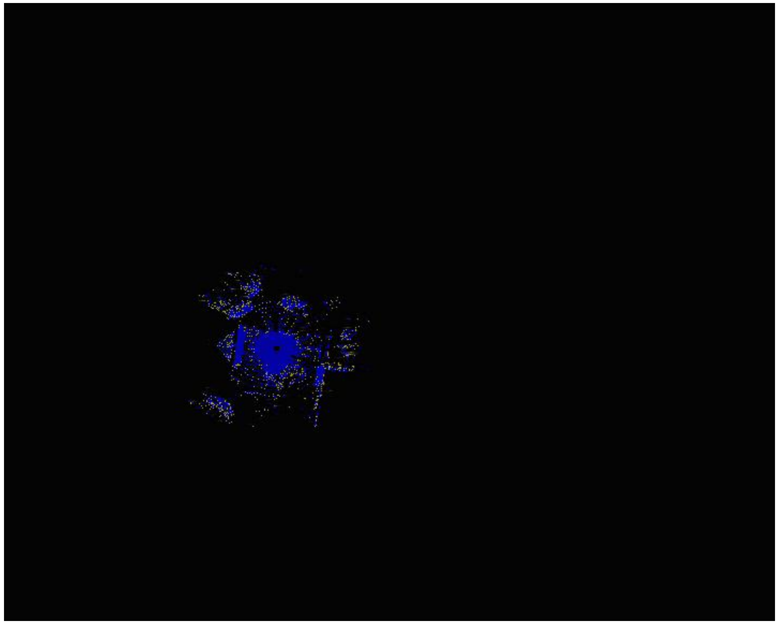

3D Point Cloud Processing – Emerging Trends

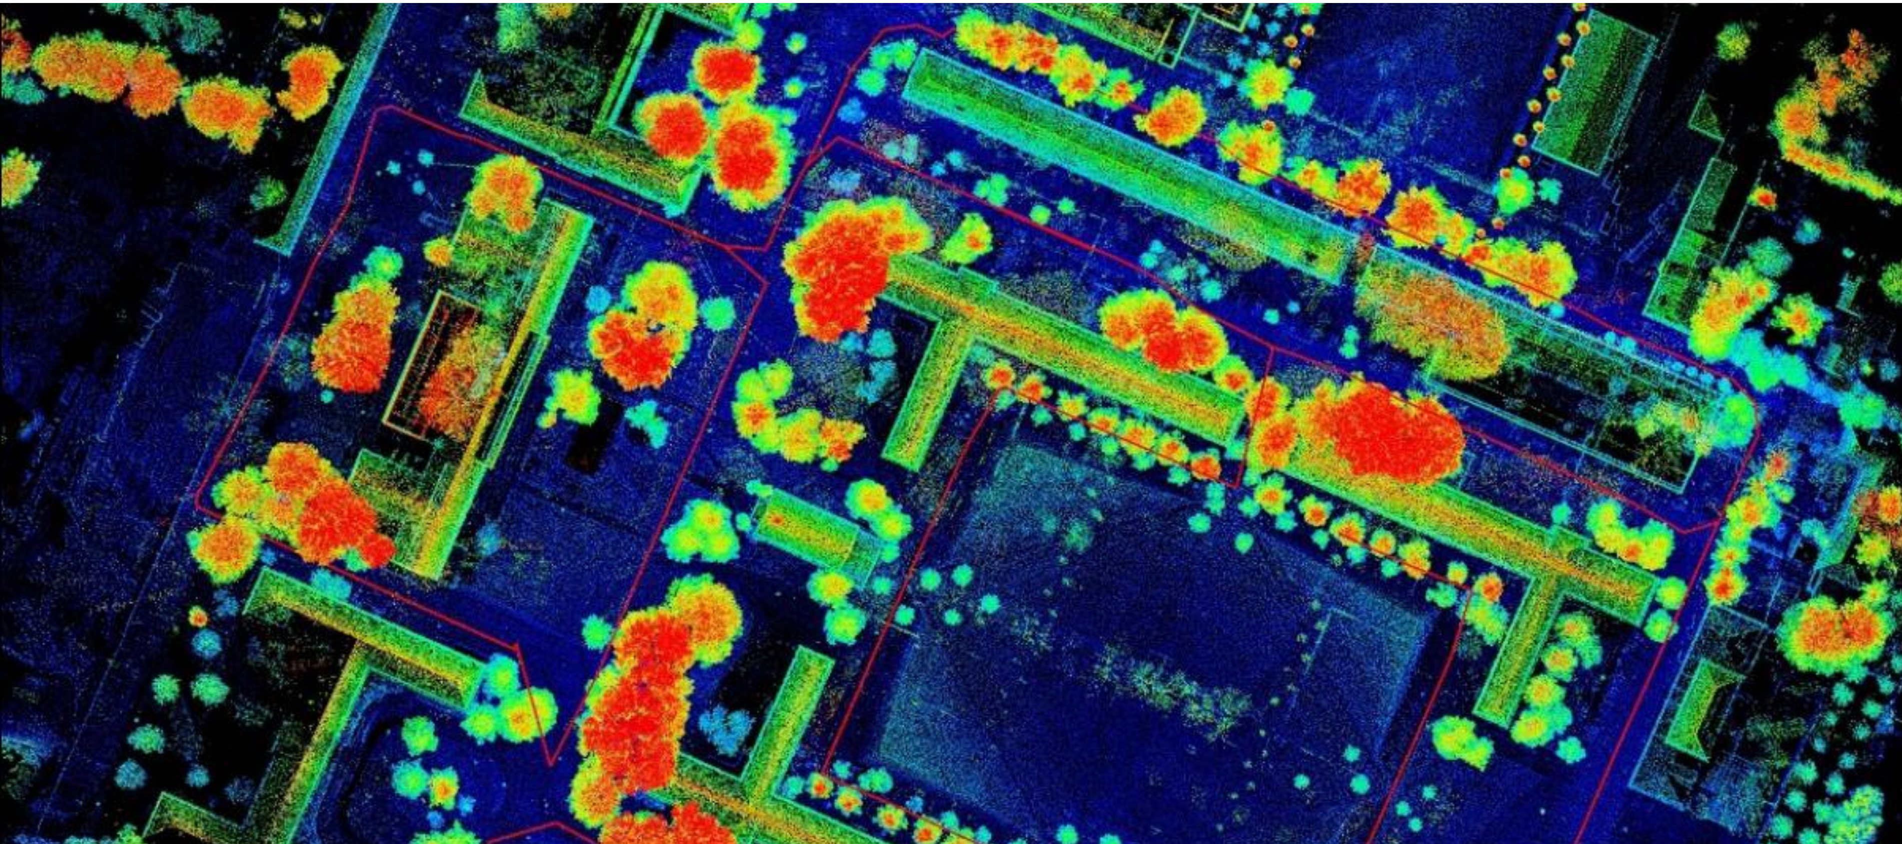

- 122 Scans @ more than 2 Billion points
- 4 coordinates per point, 8 bytes per coordinate => 59.6 Gb
- Compressed only 8.8 Gb @ ~100 micron precision

Automation
JACOBS UNIVERSITY

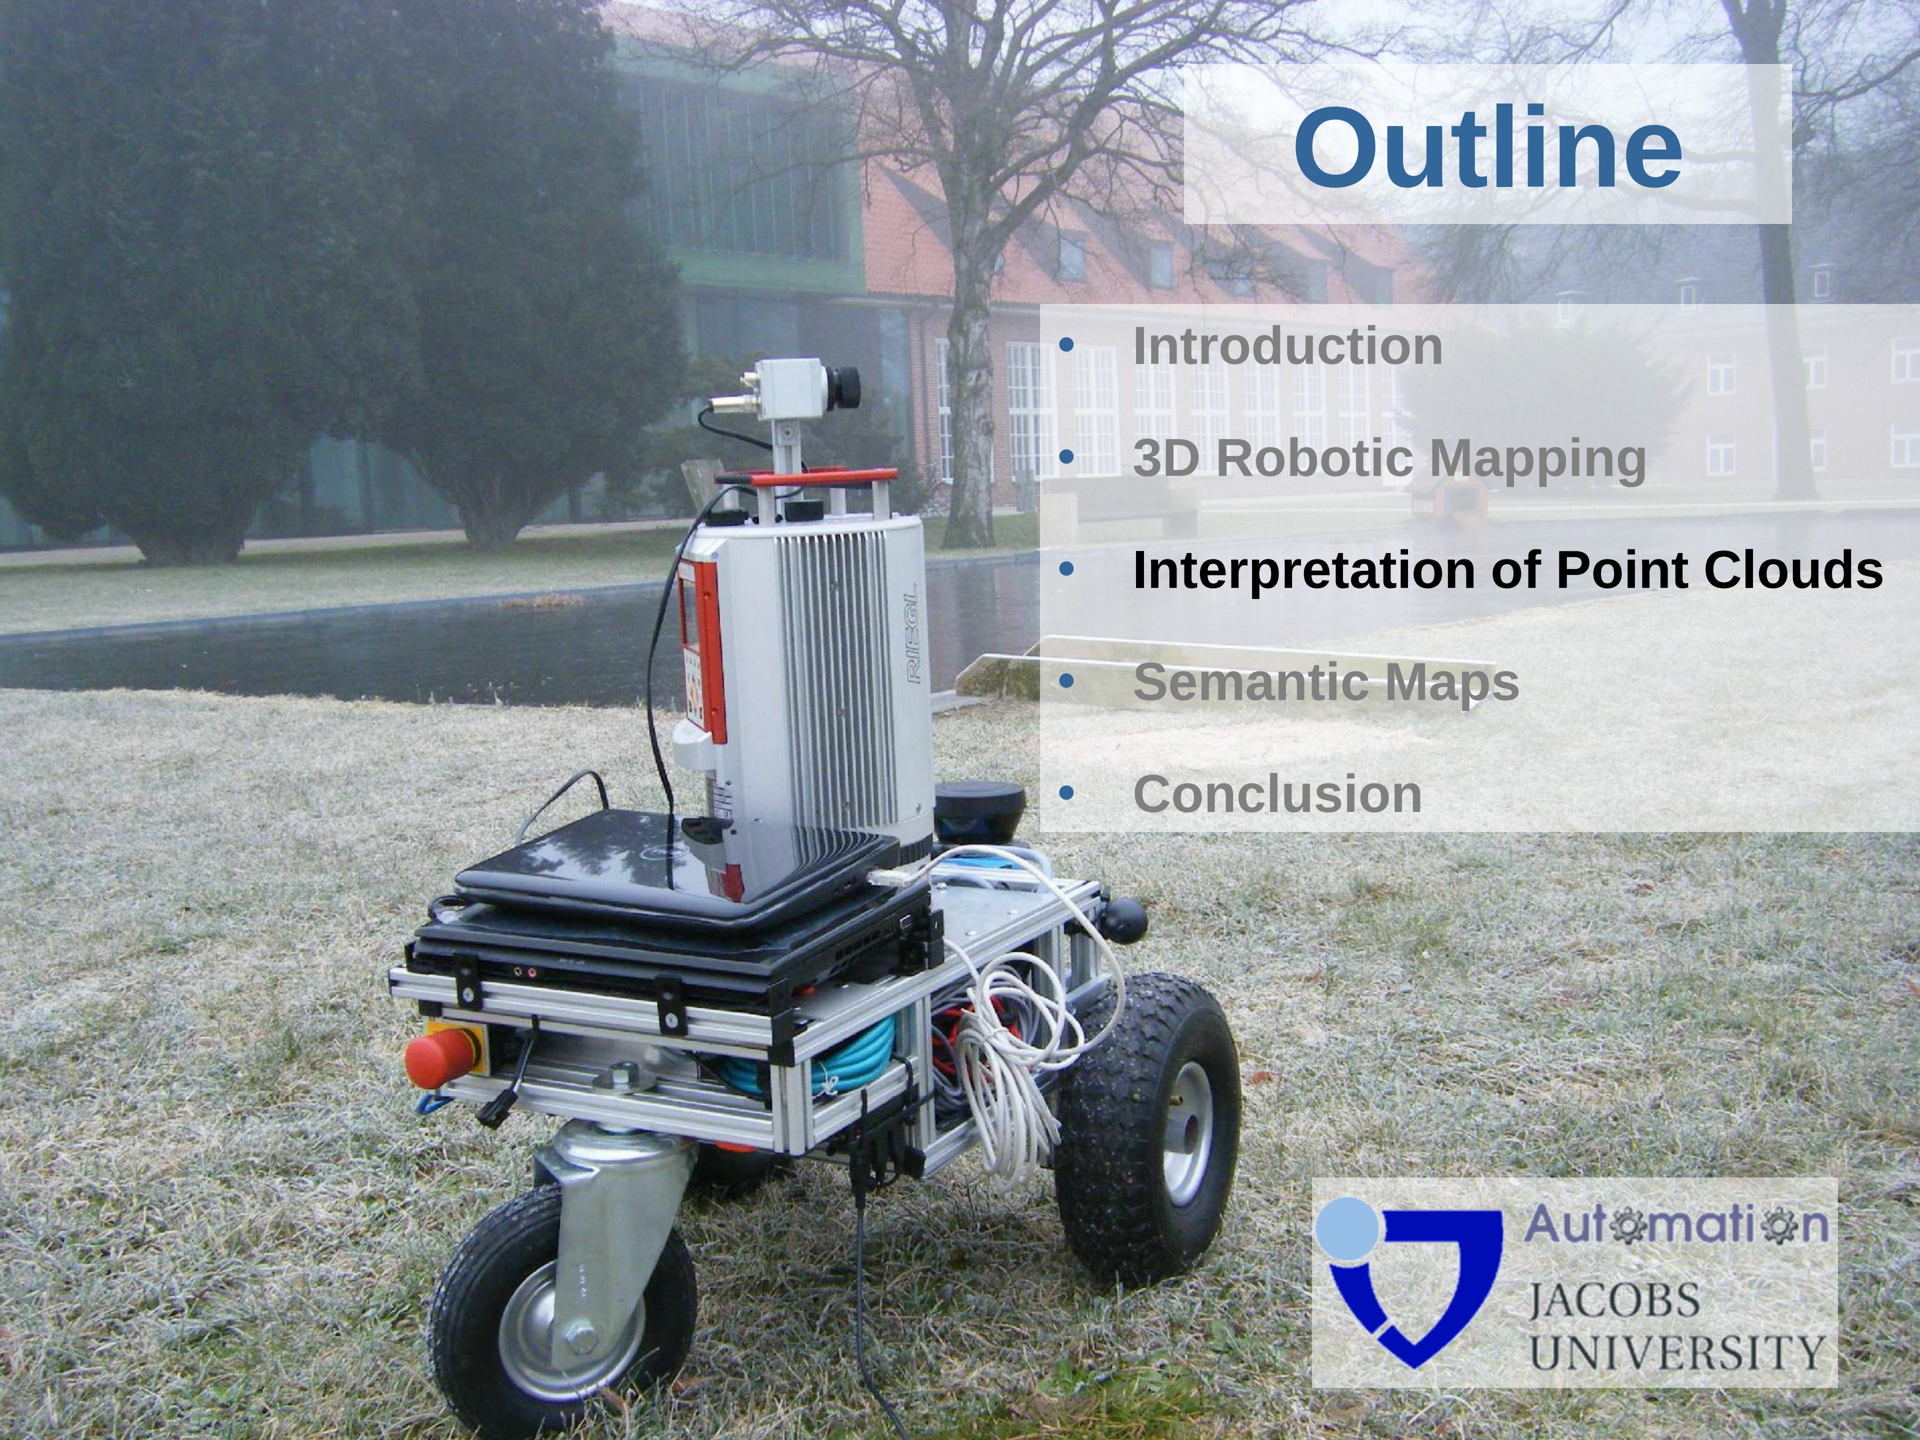

Outline

- Introduction
- 3D Robotic Mapping
- **Interpretation of Point Clouds**
- Semantic Maps
- Conclusion

Automation
JACOBS UNIVERSITY

Scene Interpretation

- Plane extraction using a novel RANSAC+ICP

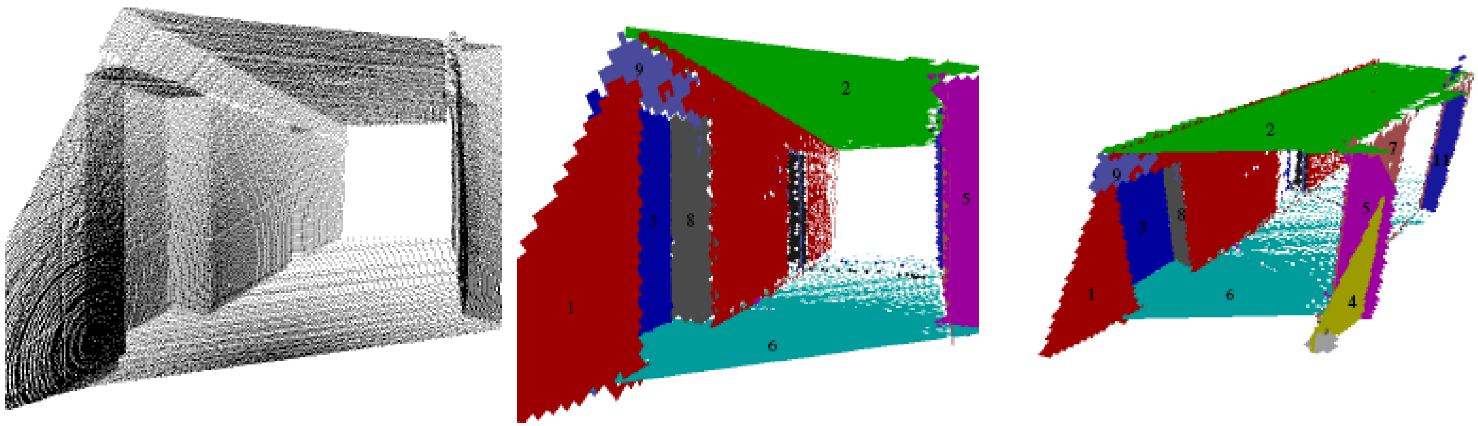

- Semantic net for assigning meaningful labels to the planes

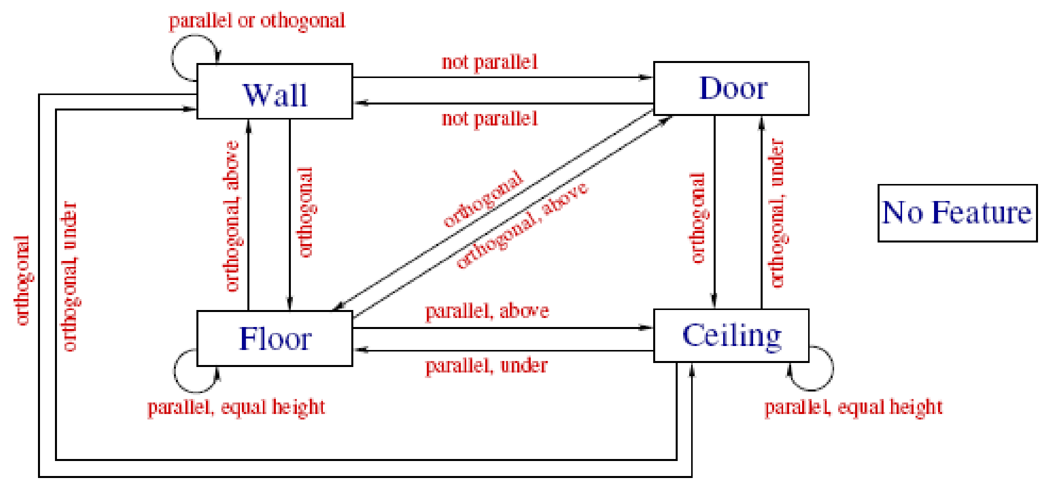

JACOBS UNIVERSITY

Detecting Shapes with RANSAC

- Improve selection of sample points
  - Choose points with higher likelihood if in close proximity
  - Lower number of draws required
- Speed up validation of hypothesis

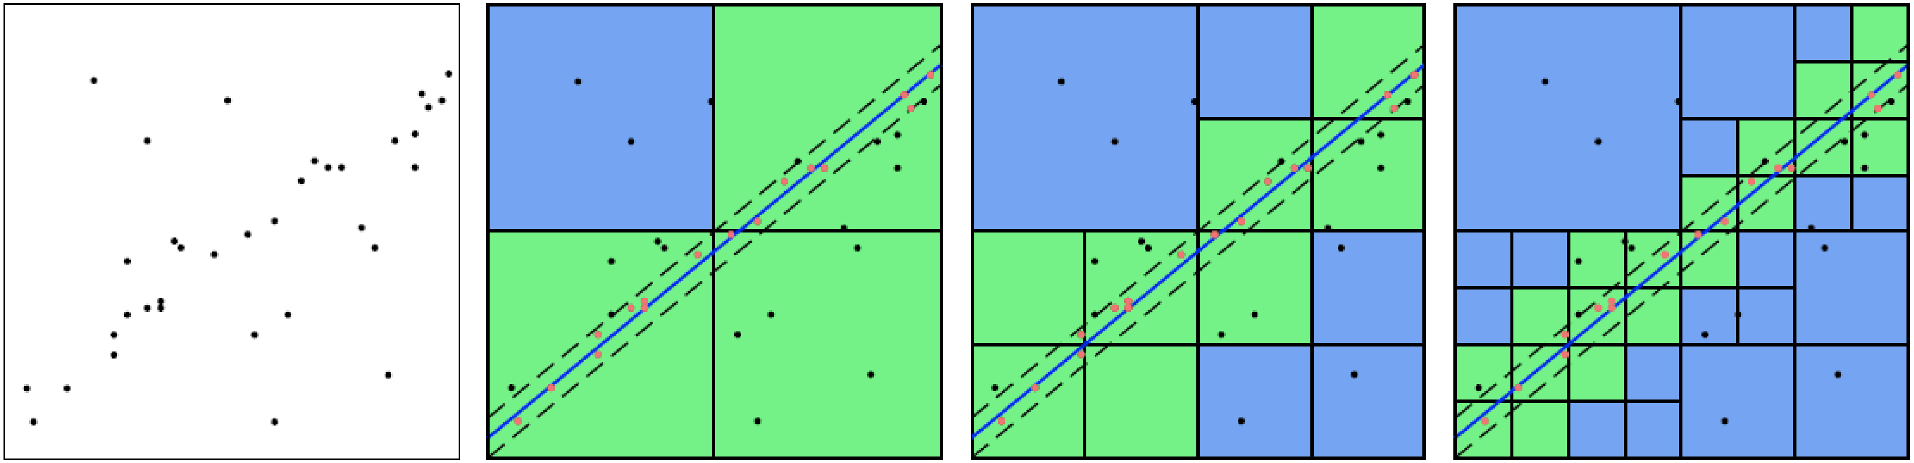

AVERAGE COMPUTING TIME IN *ms* OF RANSAC.

Data set	no octree	octree	speedup
Kurt3D	1666.57	176.69	9.43
Kinect	6905.94	429.32	16.08
city	388551.55	11084.81	35.05

R. Schnabel, R. Wahl, and R. Klein. Efficient RANSAC for Point-Cloud Shape Detection. *Computer Graphics Forum*, 2007.

Semantics by Point Labeling

- Classification of 3D points $p_{i,j} = (\varphi_i, z_{i,j}, y_{i,j})$ is in the i-th vertical scan the j-th point (start counting from the bottom)

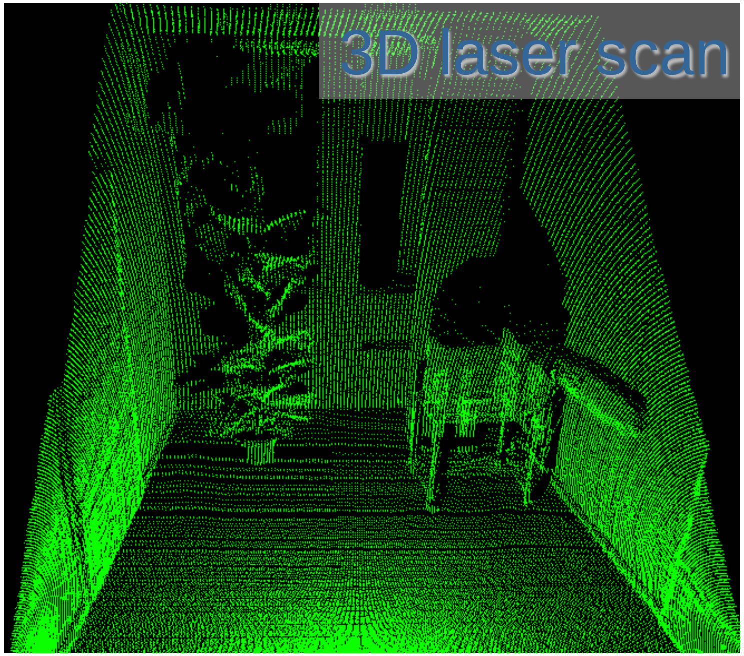

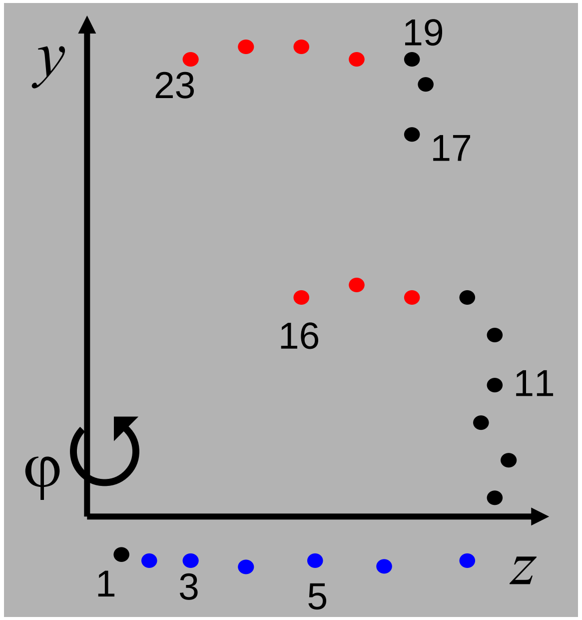

„**drivable points**“
Flat angle in scanning order

$$\left|\alpha_{i,j}\right| < \tau$$

Drivable Surface Classification

Matched Scans, floor blue

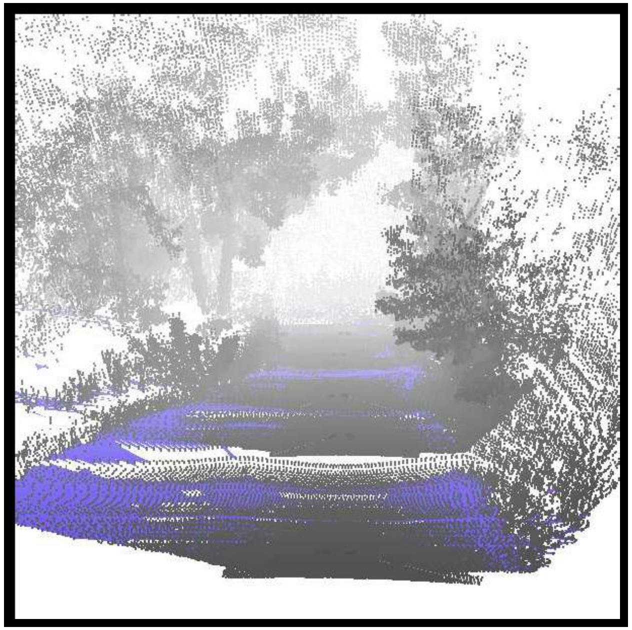

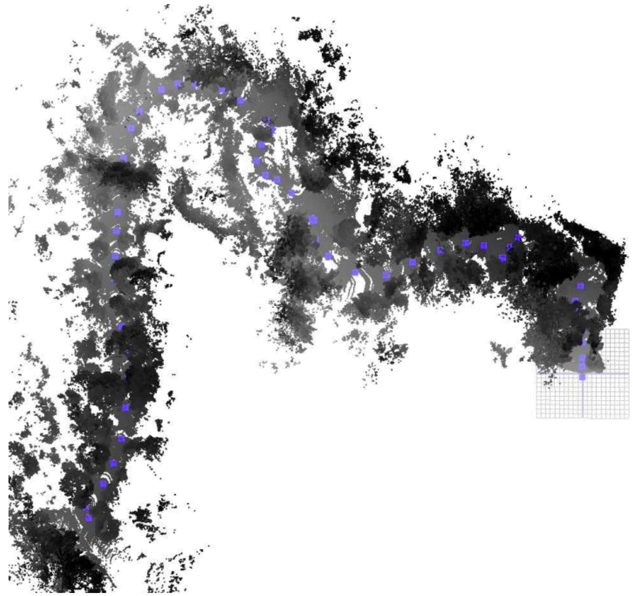

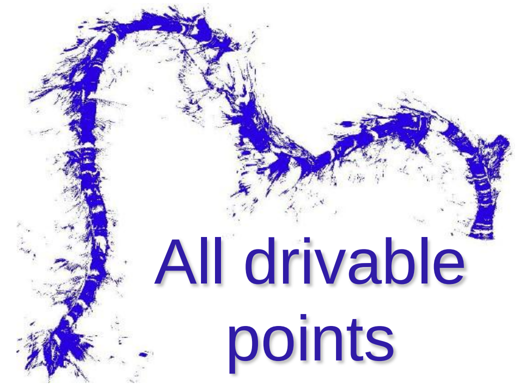

- Unfortunately not all robots are equipped with a 3D scanner

⇨ Classification based on camera images and 2D laser scans for path planning in natural environments

(Video 1) (Video 2) (Video 3) (Video 4)

Finding Objects in 3D Point Clouds

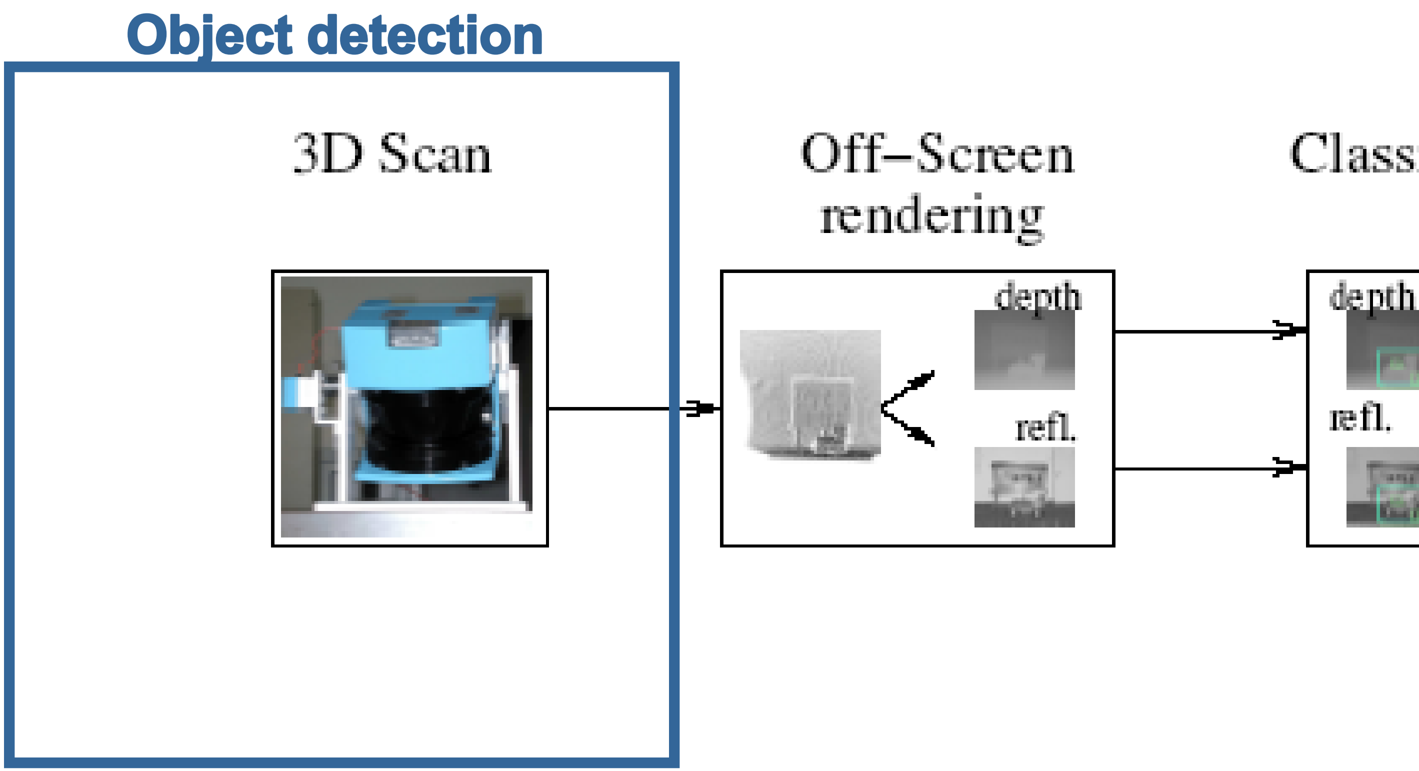

Feature Detection Using Integral Images

- Efficient representation of images using integral images

$$I(x,y) = \sum_{x'=0}^{x} \sum_{y'=0}^{y} N(x',y')$$

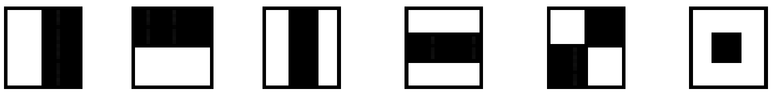

- Calculate features in integral images

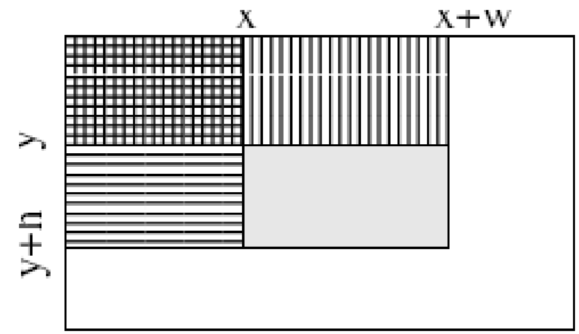

$$\begin{aligned} F(x,y,h,w) \quad &= \quad I(x,y) + I(x+w,y+h) - \\ &\quad\;\; I(x,y+h) - I(x+w,y). \end{aligned}$$

- Rotated features can be calculated with rotated integral images

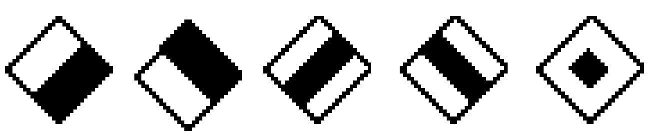

$$I_r(x,y) = \sum_{x'=0}^{x} \sum_{y'=0}^{x-|x'-y|} N(x',y')$$

Automation JACOBS UNIVERSITY

Learning a Classifier (1)

- Learn objects from 3D-Scans
- Create a classification window that contains all possible features
- Learn combination of features using Ada-Boost [Viola und Jones 01, Freund und Schapire 96]

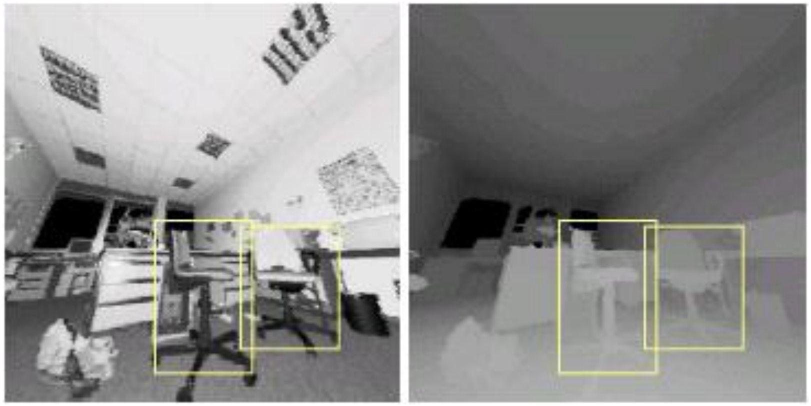

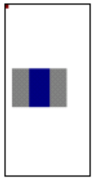

thr. = 0.002739
$\alpha = -0.9544$
$\beta = 0.8265$

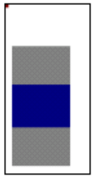

thr. = −0.01696
$\alpha = 0.7638$
$\beta = -0.8637$

- Objects of different sizes are detected by rescaling the classifier window.

Learning a Classifier (2)

- Last but not least use a cascade to lower the false detection rate.

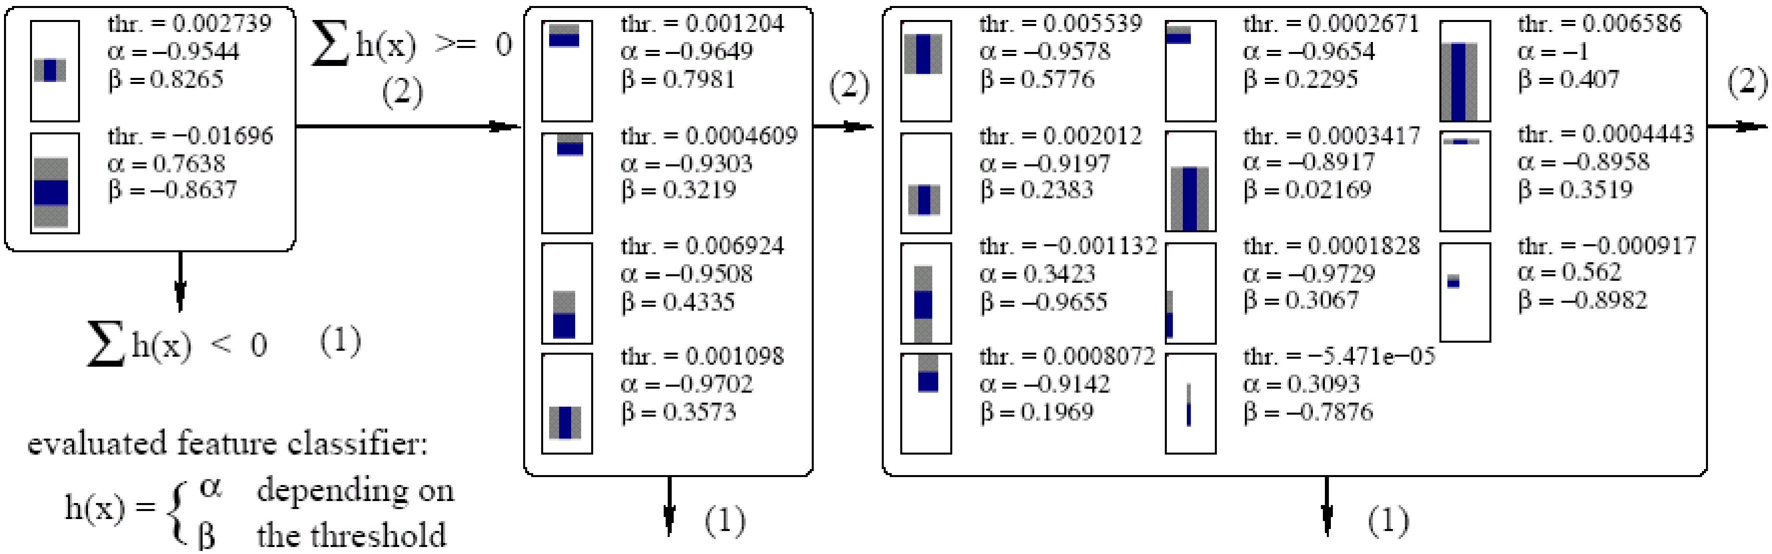

- Improvement: Use Classification and Regression Trees!

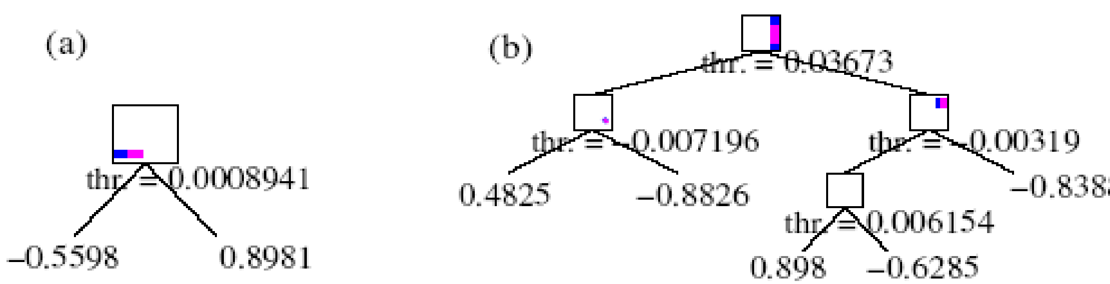

Object Detection

- Use the cascade for detection in the depth and reflectance image
- Logical AND yields reliable detection (false detection ~ 0%)

Localize the Objects (1)

- Get object points via ray tracing

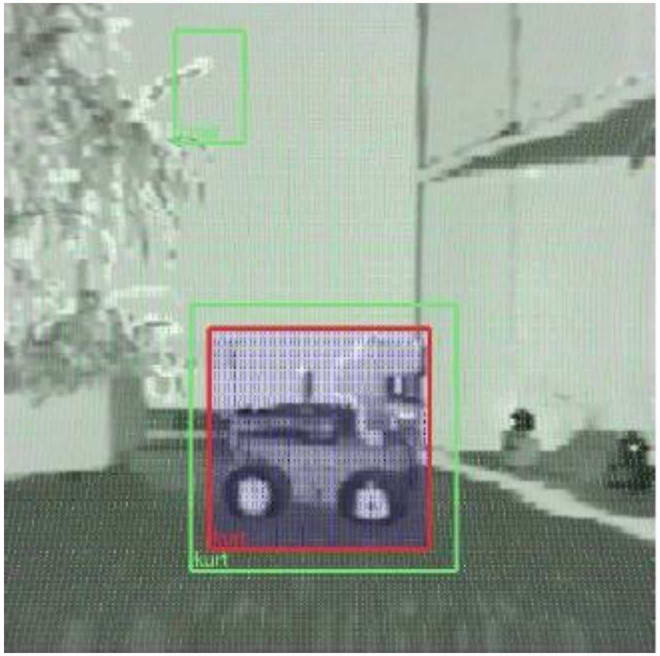

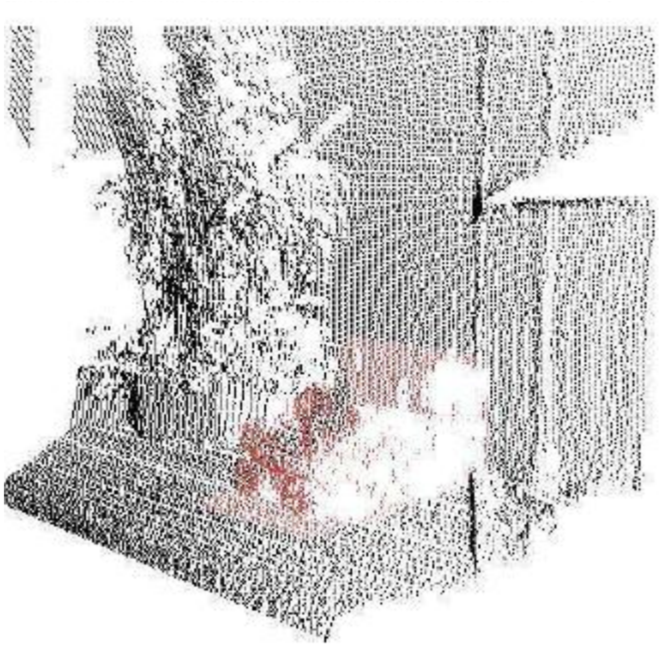

Localize the Objects (2)

- Fit objects in point cloud using an ICP variant

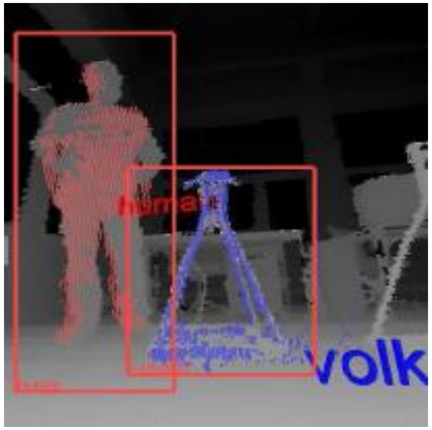

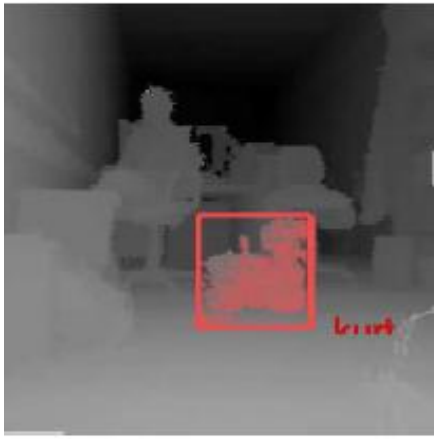

- For prior point set *M* ("model set") and data set *D*

1. Select point correspondences $w_{i,j}$ in {0,1}
2. Minimize for rotation **R**, translation **t**

$$E(\mathbf{R},\mathbf{t}) = \sum_{i=1}^{N_m}\sum_{j=1}^{N_d} w_{i,j}\,||\mathbf{m}_i - (\mathbf{R}\mathbf{d}_j + \mathbf{t})||^2$$

3. Iterate **1.** and **2.**

Point Semantic for Object Detection

- Point labeling removes the ground
- Extract contour features
- Learning
- Detect objects

⇨ Map building with labeled objects

- Task: Detect Objects in depth images

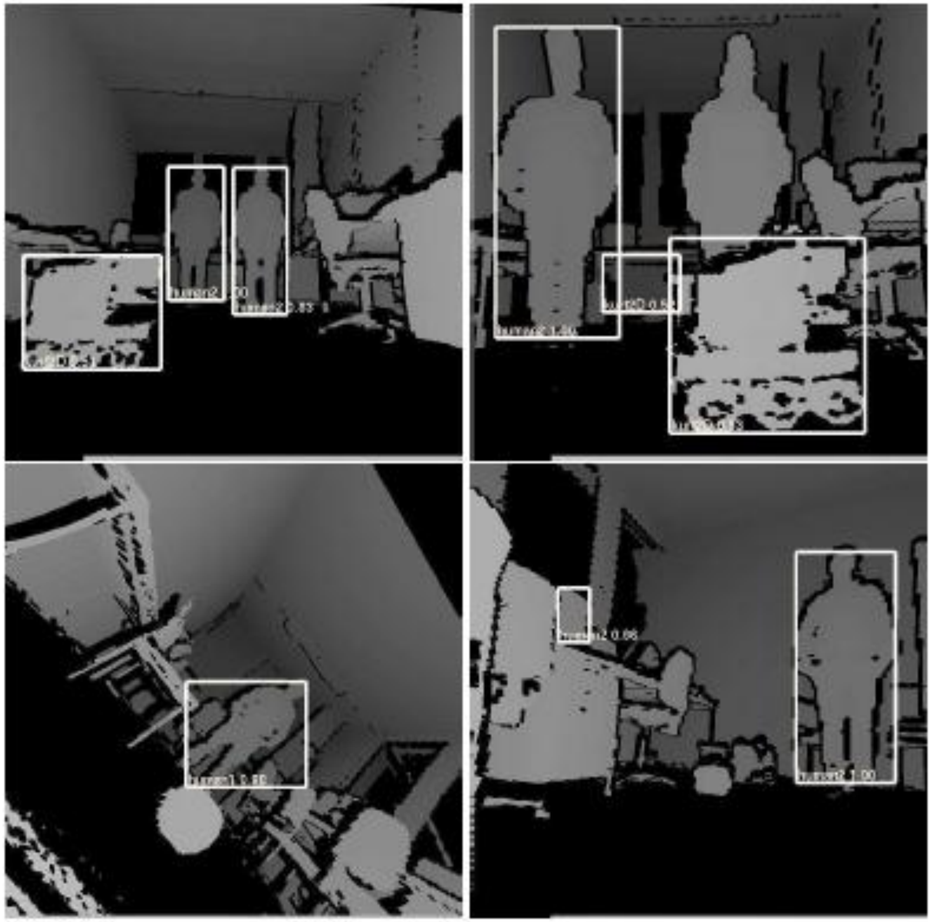

Identifying 3D Google Warehouse Models (1)

Task: Recognize “Audi A4” in laser scan data

- Input: String “Audi A4”, 3D laser scan

- Output: Pose of the object with 6 DoF

Identifying 3D Google Warehouse Models (2)

- Algorithm:
  - Query Google's data base and download all models
  - Convert the models in point cloud data
  - Segment your 3D scan and remove obvious objects
  - Match the models into the segmented scan with a modified version of ICP (scale of the models is unknown!)
  - Design an evaluation function to find the best match

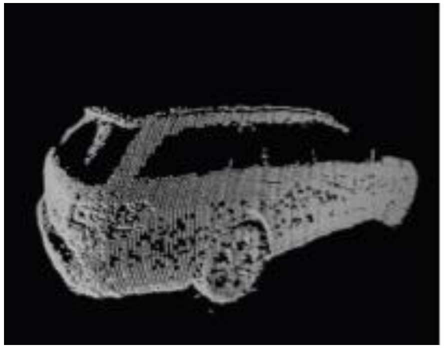

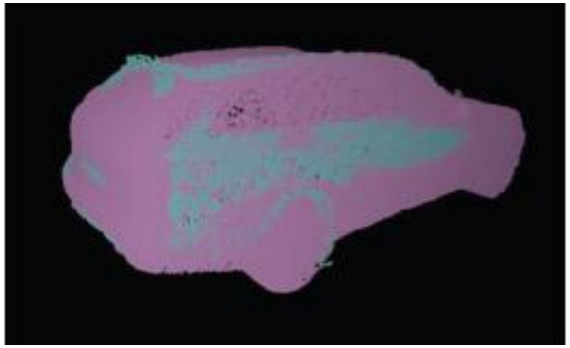

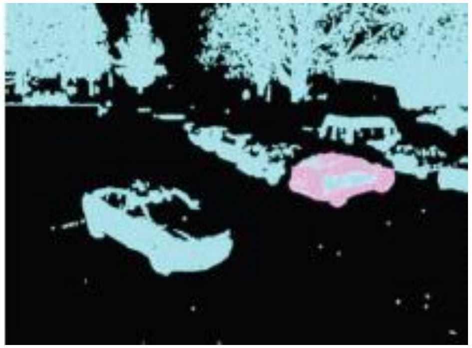

(video)

Full Wave Analysis

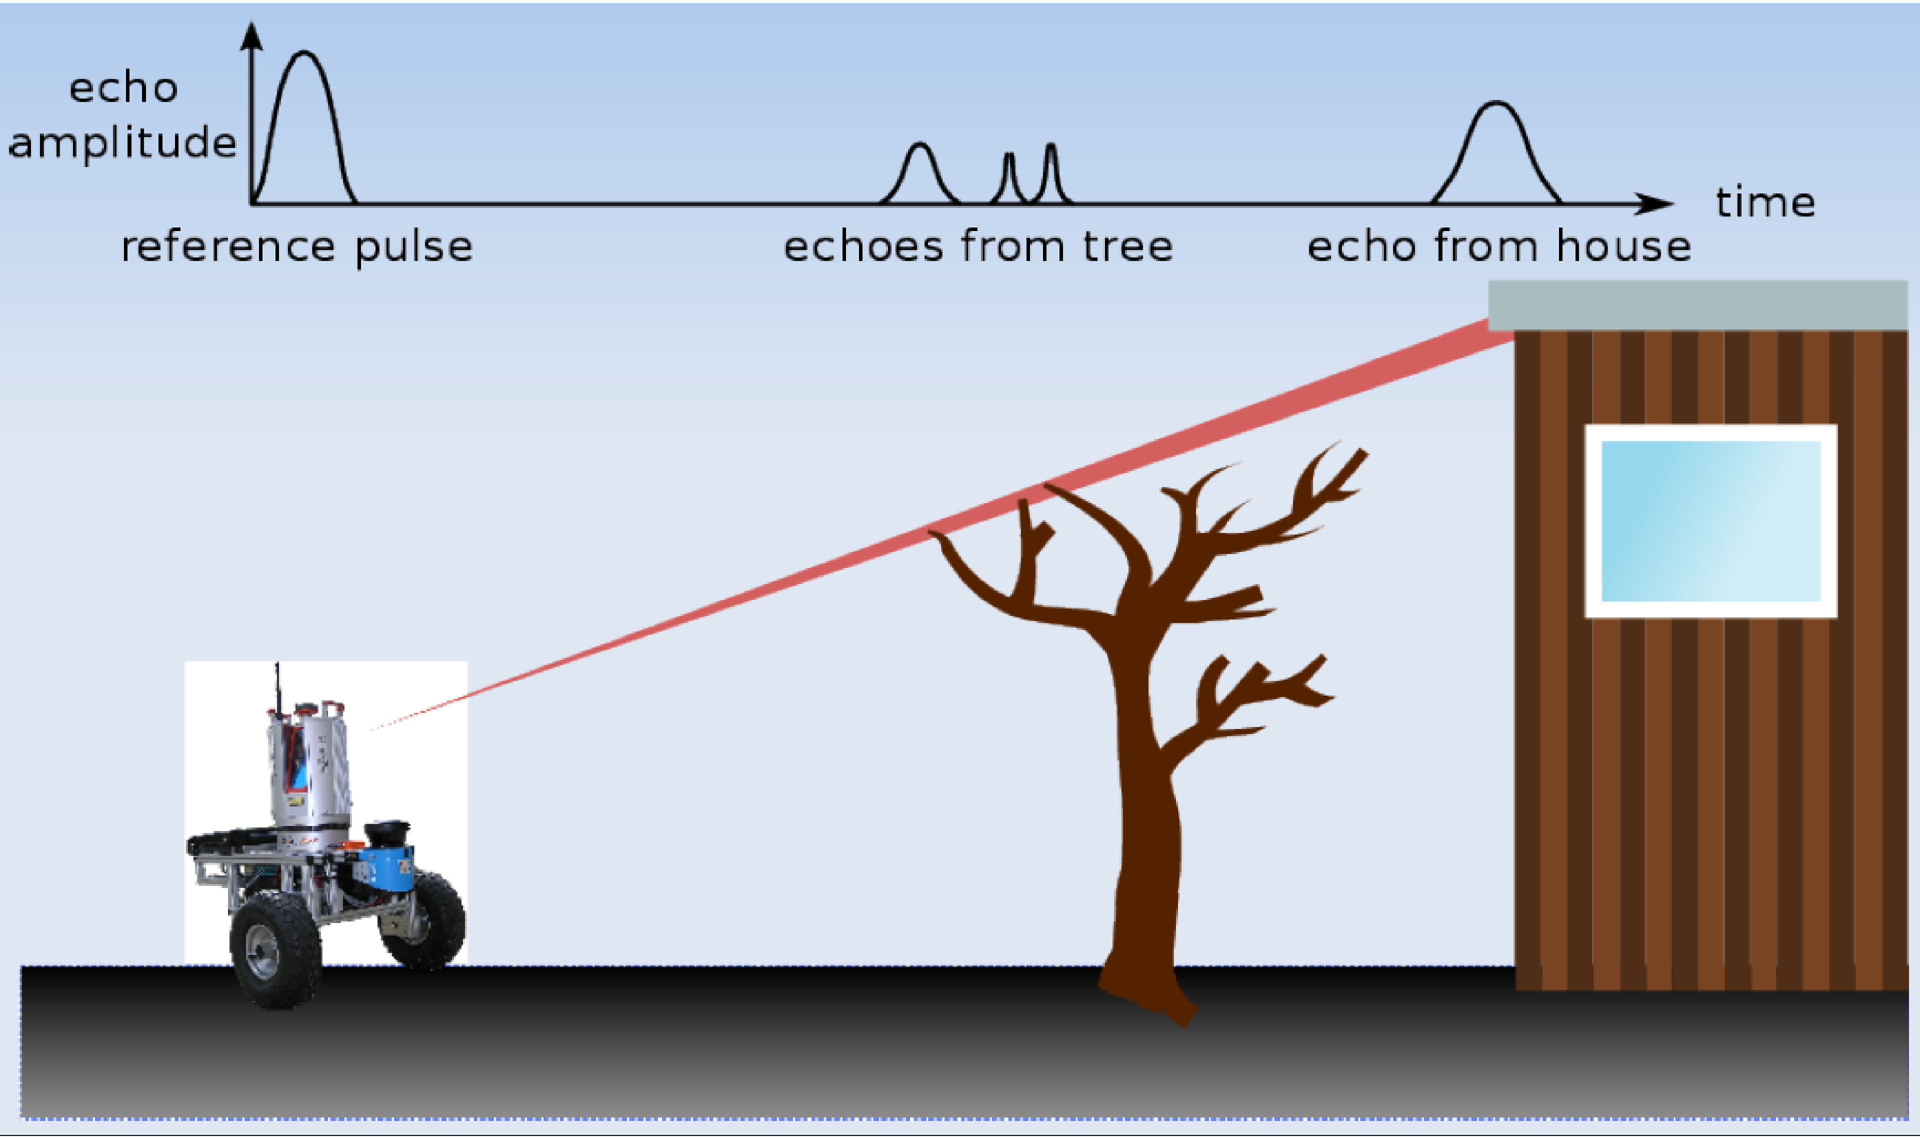

Multiple Echos

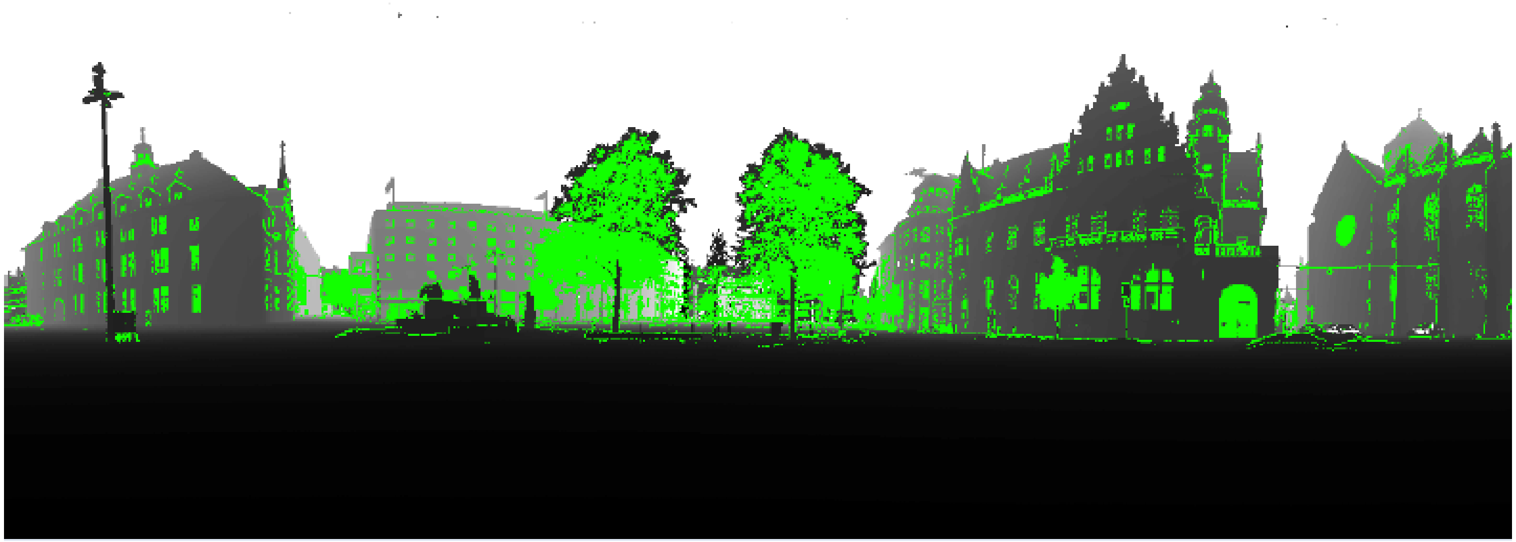

Vegetation Detection

1. Extract inner echos
2. Extract ROI
3. Classify regions

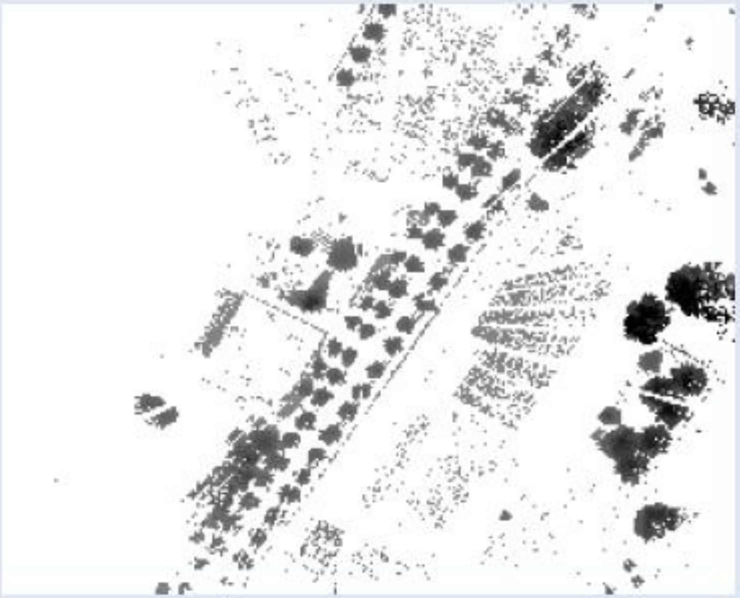

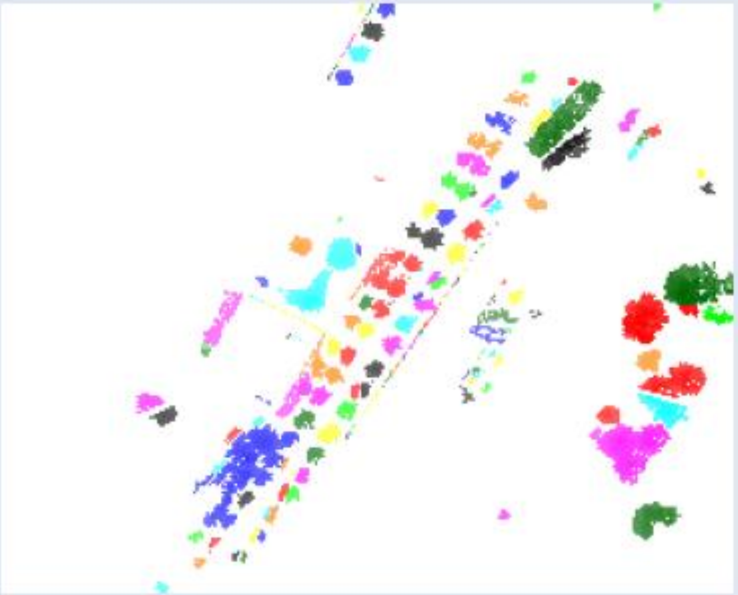

Vegetation Detection

Vegetation Detection

Outline

- Introduction
- 3D Robotic Mapping
- Interpretation of Point Clouds
- **Semantic Maps**
- Conclusion

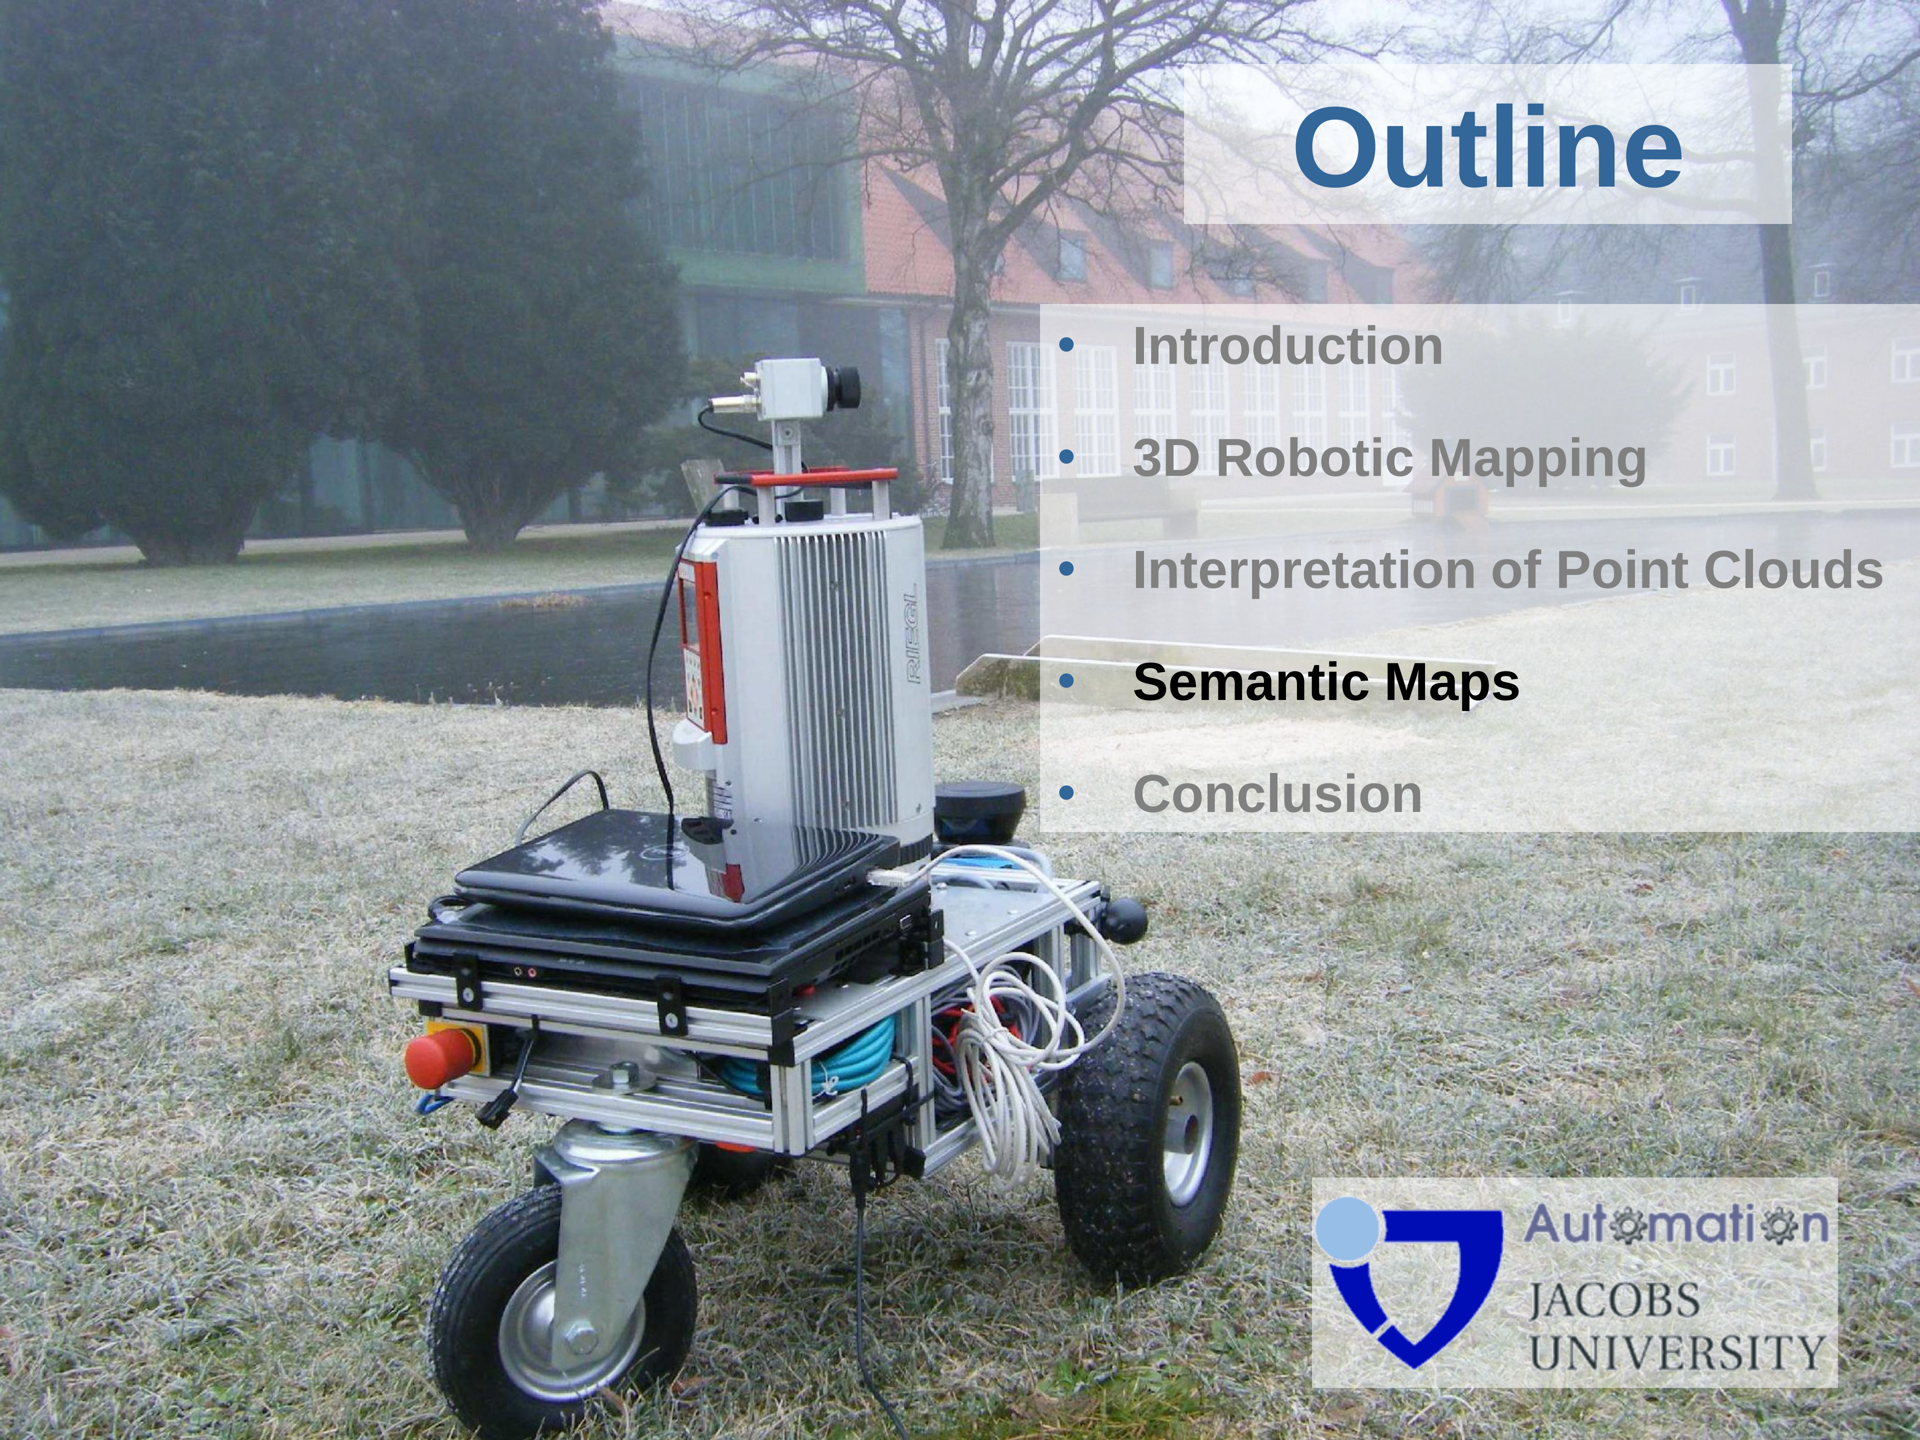

Automation
JACOBS UNIVERSITY

Definition / Discussion

- A semantic 3D map is a metrical map that contains in addition to geometrical information semantic label of the data points.

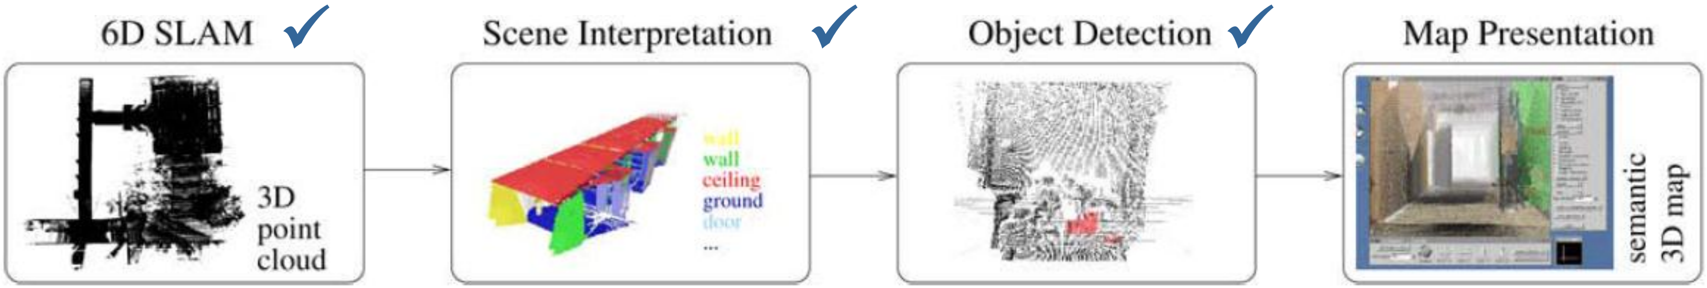

- Map presentation as video

(Video)

Where does knowledge come from? Where does it go?

Knowledge based software system

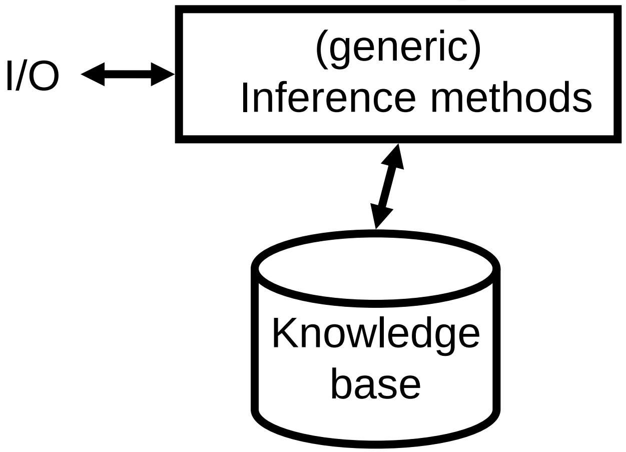

Example expert system:

- *Knowledge Engineer* sets up knowledge base (off-line)
- User causes input and edit the output

Example robot:

- *Knowledge Engineer* creates knowledge base off-line
- Input of KBS comes from sensors
- Output drives the robot

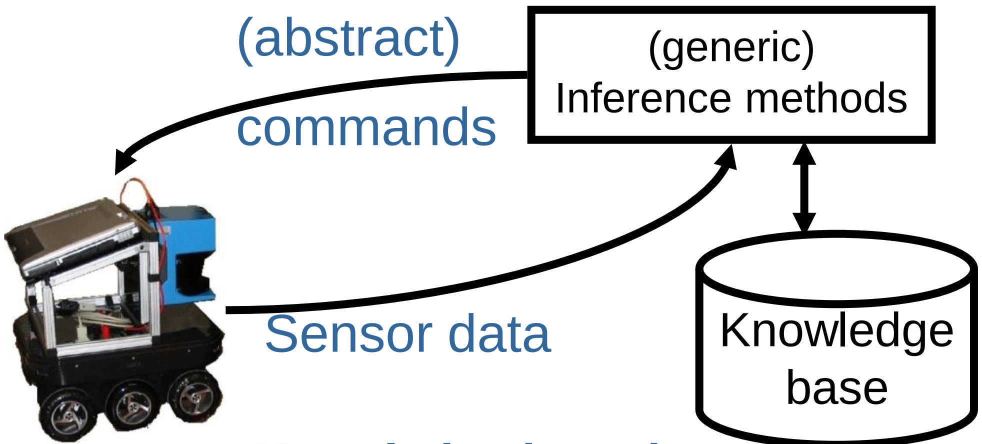

Knowledge based robotic system

Robots in natural environments must translate sensor data into symbols and inferences eventually into control data!

Symbol Grounding

S. Harnad: *The Symbol Grounding Problem*
Physica D 42:335–346, 1990
cogprints.org/3106/01/sgproblem1.html

How is meaning of a symbol to be grounded in something other than just more meaningless symbols?

Is that an important question?

- Some (AI) say: No, at the best that's a technical problem!
- Some (Philosophy, Cog. Sci.) say: That's the point, at which artificial intelligence systems are doomed to fail!
- Some (AI, Cog. Sci., Robotics people, Nüchter) say: That's currently the most interesting point in fundamental research in AI

Specialization: Object anchoring

S. Coradeschi, A. Saffiotti: *An Introduction to the Anchoring Problem*
Robotics & Autonomous Systems 43(2–3):85–96, 2003

www.aass.oru.se/~asaffio/Papers/ras03.html

- Anchoring (object anchoring): *the process of creating and maintaining the correspondence between symbols and sensor data that refer to the same physical objects*
- Anchoring problem: *the problem of how to perform anchoring in an artificial system*
- Specialization with symbol anchoring in general:
  Is related only to physical objects, e.g., no abstract objects like weather or no attributes („red“)

Outline

- Introduction
- 3D Robotic Mapping
- Interpretation of Point Clouds
- Semantic Maps
- **Conclusion**

Conclusions

- Practical (on-line, on-board) variant of ICP for high-resolution point sets
- Generating overall consistent 3D maps with global relaxation
- Tested on various data sets (including borrowed ones, e.g., CMU mine mapping)
- Interpretation of 3D maps resulting in 3D object maps
- Integrated into robot controller for 3D environment mapping
- RoboCup Rescue as evaluation for our mapping approach
  - 2004 second place, SSRR 2005 best paper award, 2005 6th place

- However, there is still no theory about „Semantic Maps in Robotics“

Conclusion – Semantic Maps

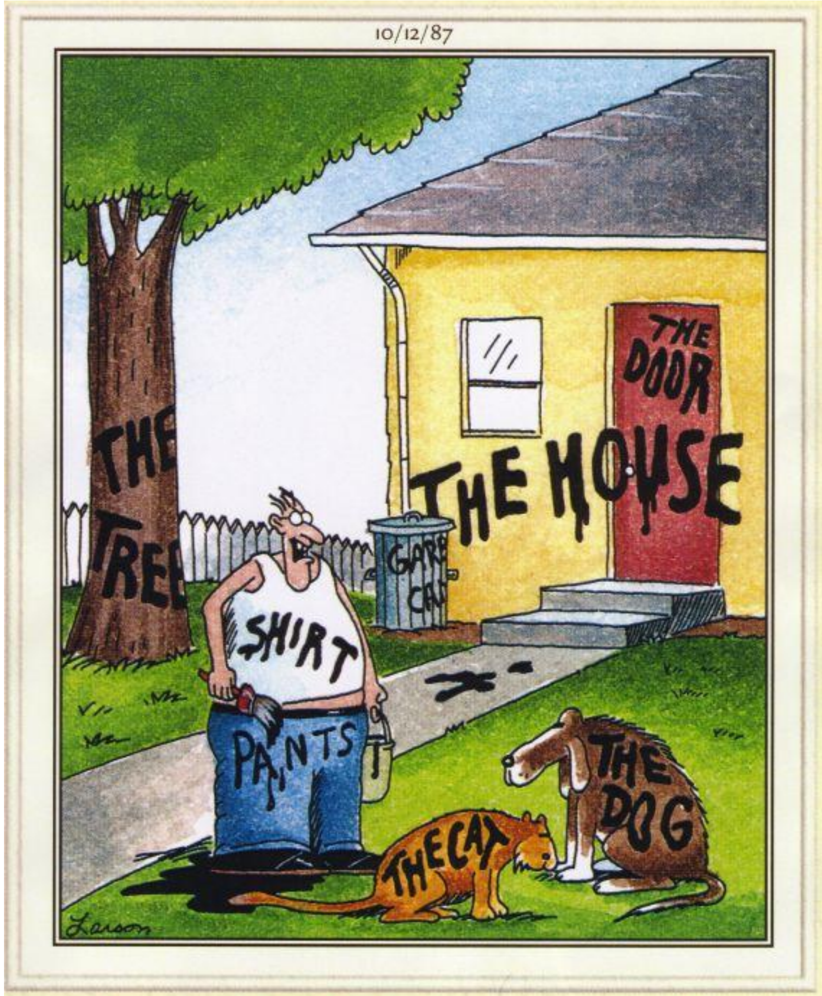

"Now! ... *That* should clear up a few things around here!"

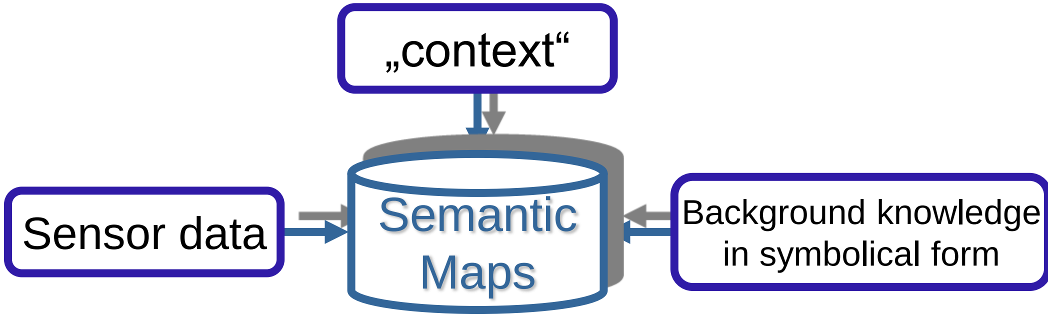

- For constructing semantic maps a robot uses sensor and semantic information to
  - interpret the sensor data (e.g., for disambiguation)
  - conclude additional semantic information (that extrapolates the sensor values, contradicts the sensors values)
  - acquire goal-directed new sensor information (e.g., attention based control, active vision)

Recent Work

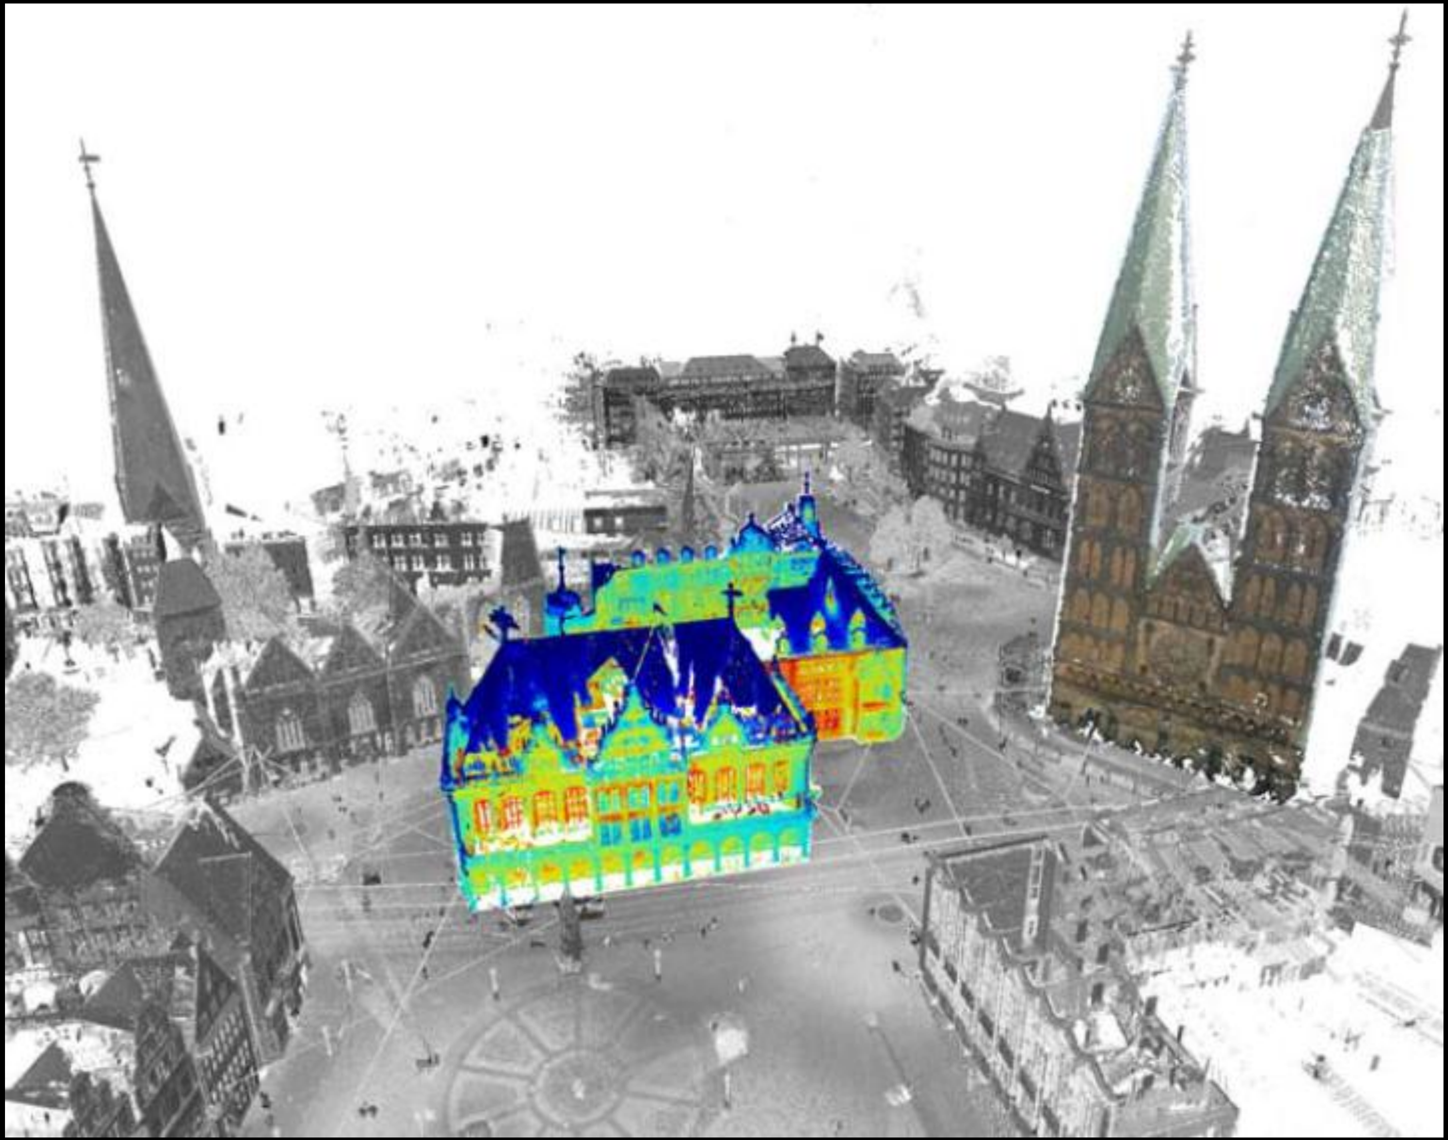

Automation JACOBS UNIVERSITY

(video)

Recent Work

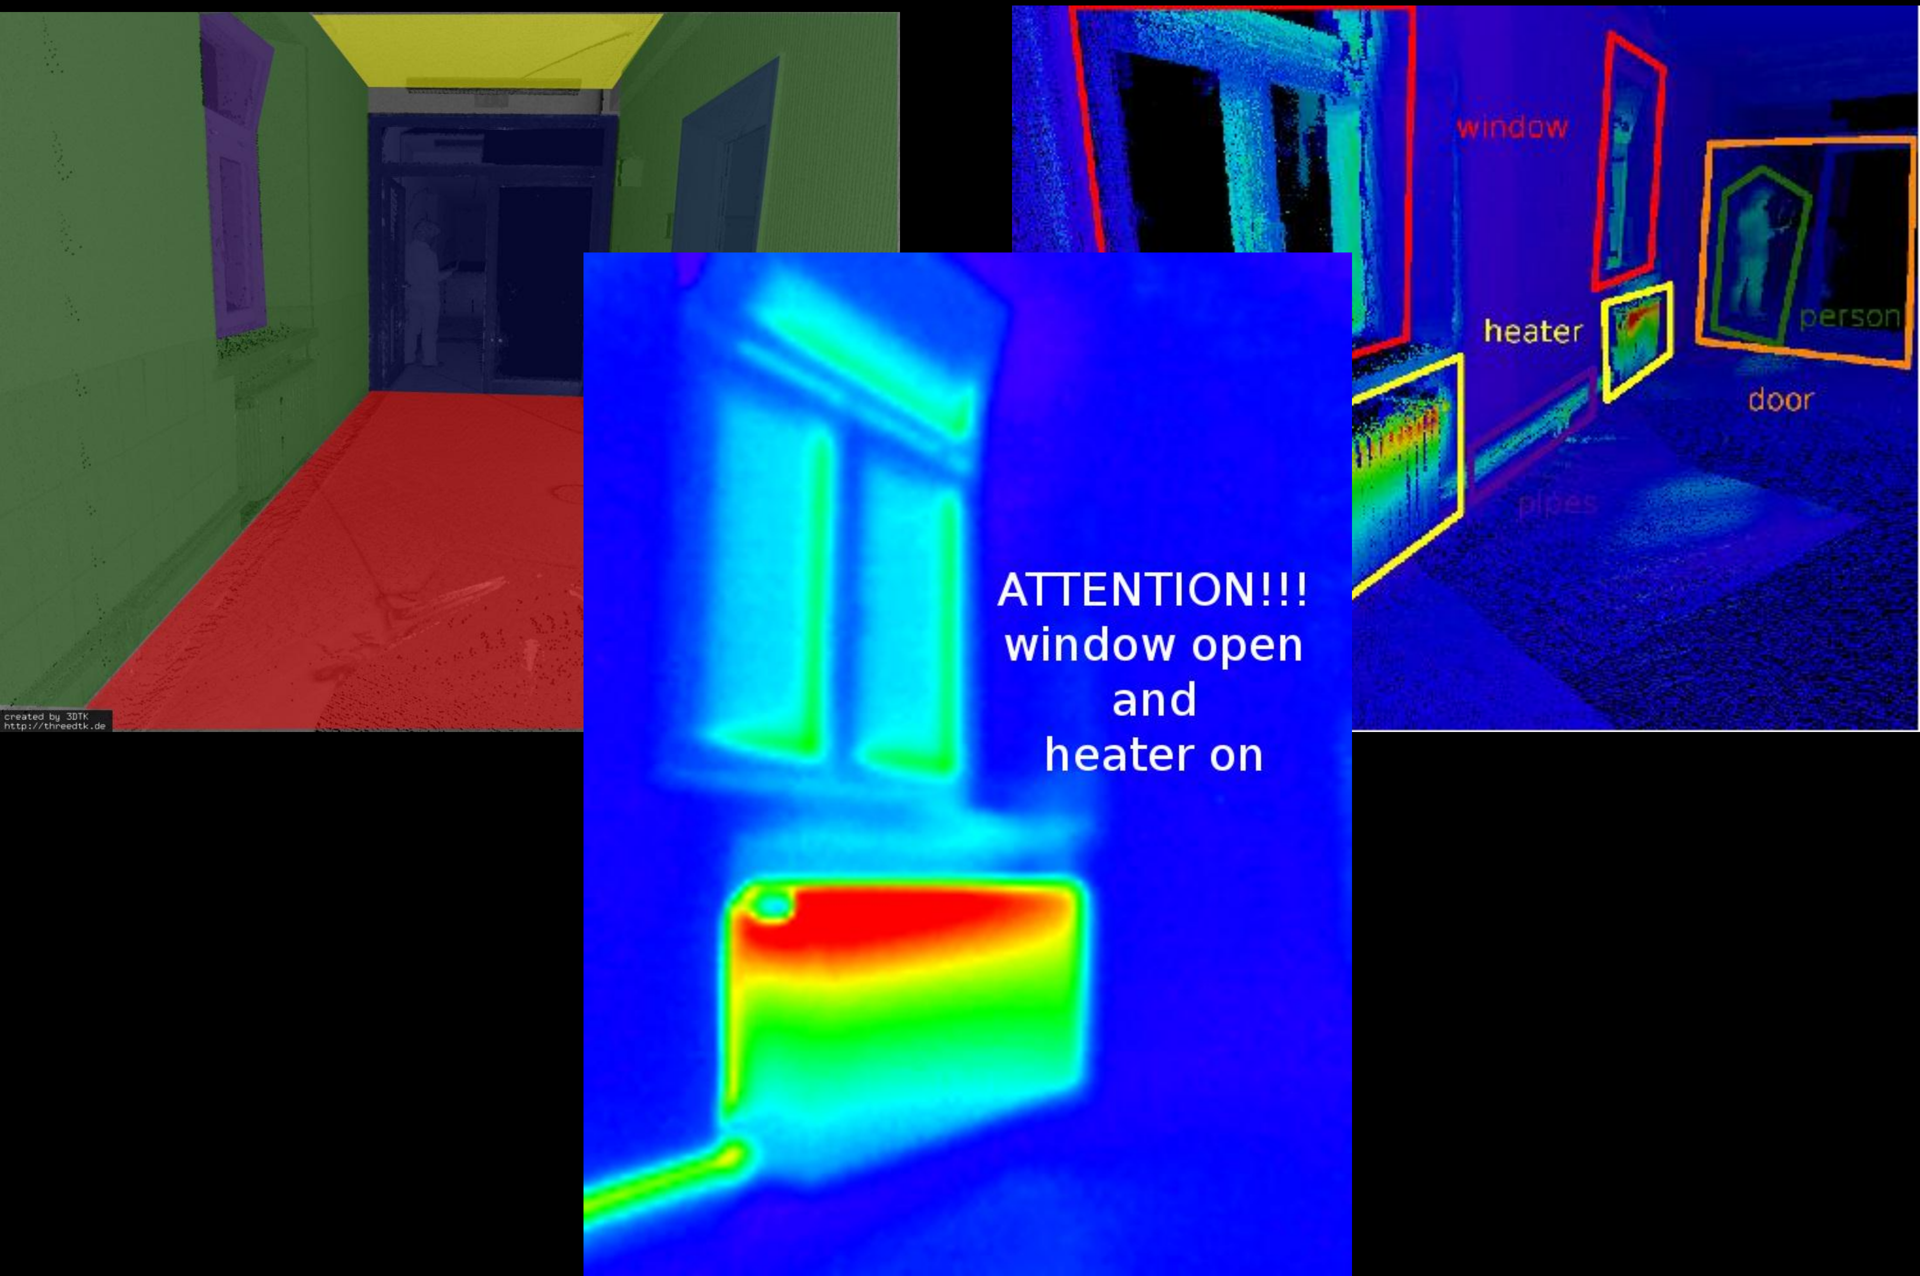

created by 3DTK
http://threedtk.de
window
heater
person
door
pipes
ATTENTION!!!
window open
and
heater on

References

- Please visit my homepage @

  www.nuechti.de

- Please visit our youtube page @

  www.youtube.com/user/AutomationAtJacobs

- Please check out our Open Source project

  3DTK – The 3D Toolkit

  http://slam6d.sourceforge.net