Perception and Planning for Autonomous Mobile Robots in Complex Environments

Sven Behnke

University of Bonn, Germany

Computer Science Institute VI

Autonomous Intelligent Systems

Some of Our Cognitive Robots

- Equipped with many sensors and DoFs
- Demonstration in complex scenarios

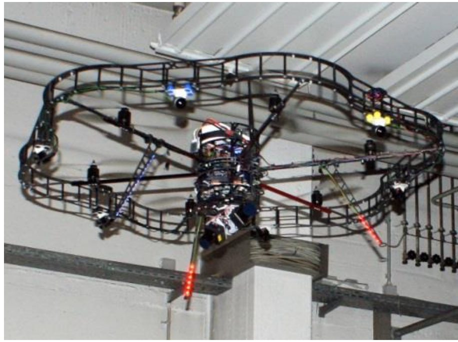

MAV

Soccer robot

Service robot

Exploration robot

Picking robot

Cognitive Service Robot Cosero

Table-top Analysis and Grasp Planning

- Detection of clusters above horizontal plane
- Two grasps (top, side)

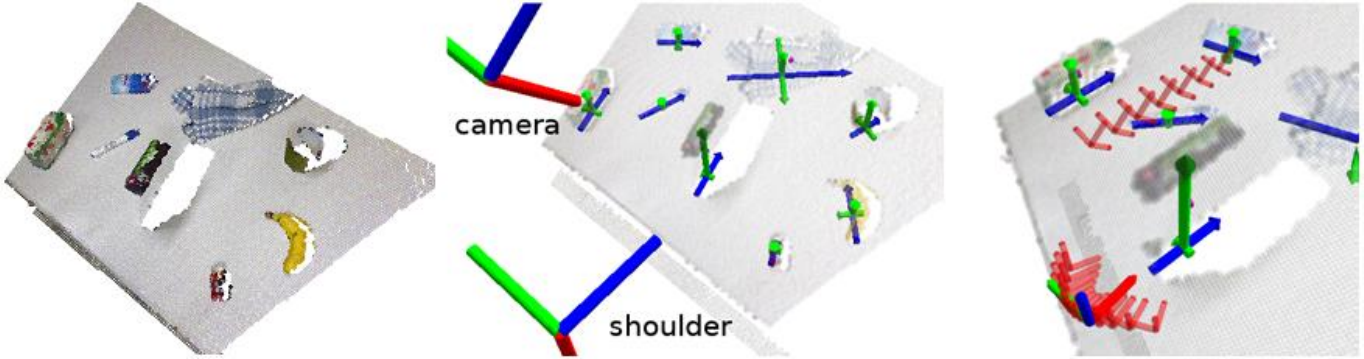

- Flexible grasping of many unknown objects

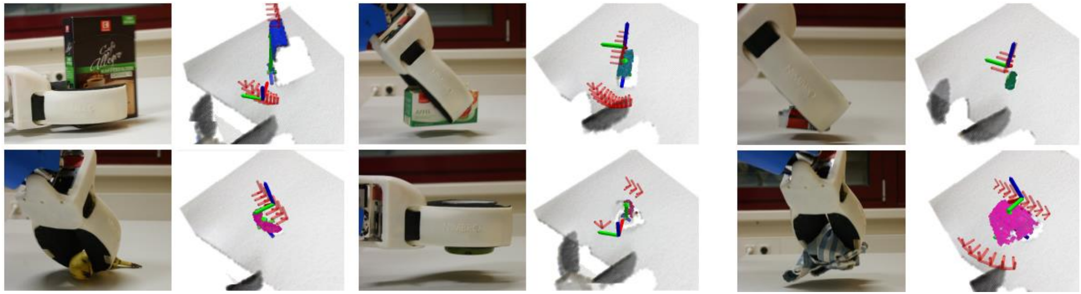

[Stückler et al, Robotics and Autonomous Systems, 2013]

3D Mapping by RGB-D SLAM

[Stückler, Behnke: Journal of Visual Communication and Image Representation 2013]

- Modelling of shape and color distributions in voxels
- Local multiresolution
- Efficient registration of views on CPU
- Global optimization

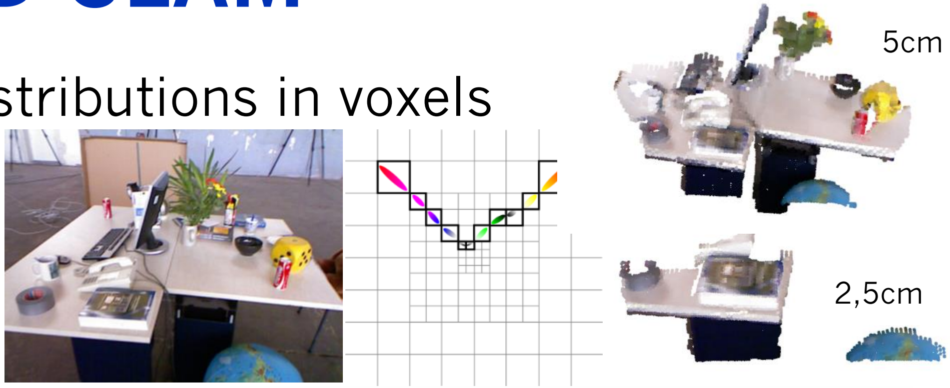

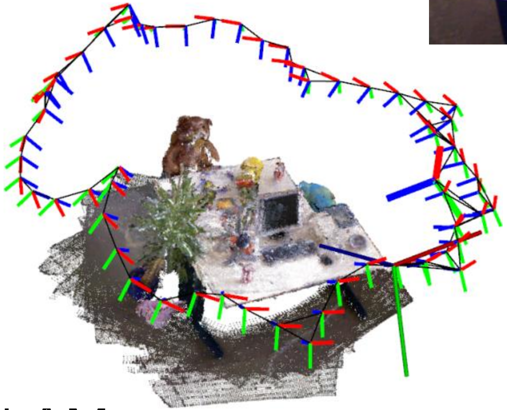

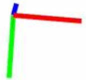

- Multi-camera SLAM

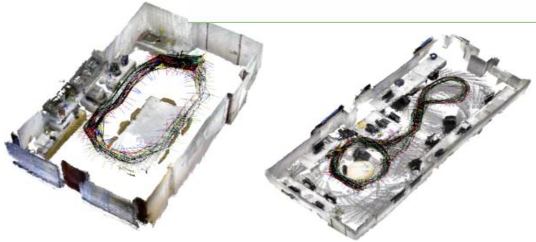

[Stoucken]

Learning and Tracking Object Models

- Modeling of objects by RGB-D-SLAM

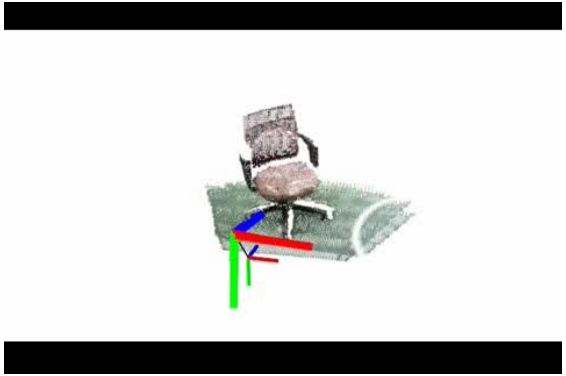

- Real-time registration with current RGB-D frame

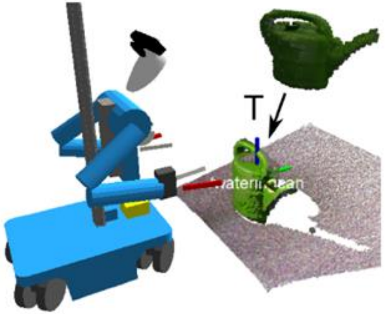

Deformable RGB-D-Registration

- Based on Coherent Point Drift method [Myronenko & Song, PAMI 2010]
- Multiresolution Surfel Map allows real-time registration

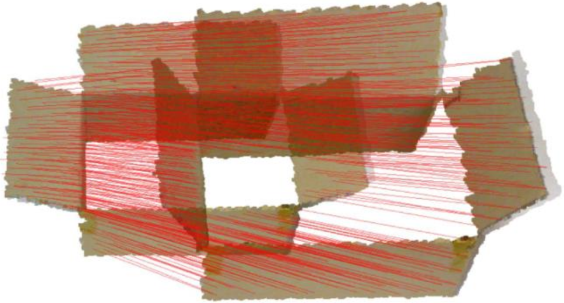

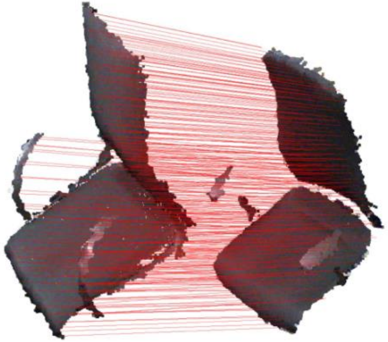

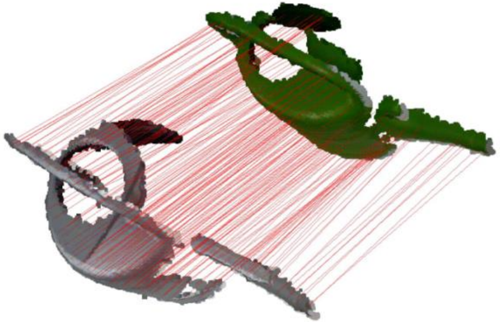

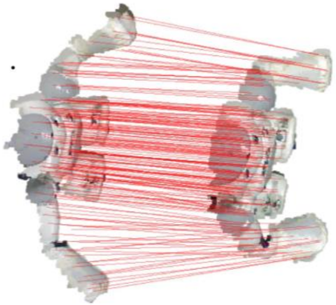

Transformation of Poses on Object

- Derived from the deformation field

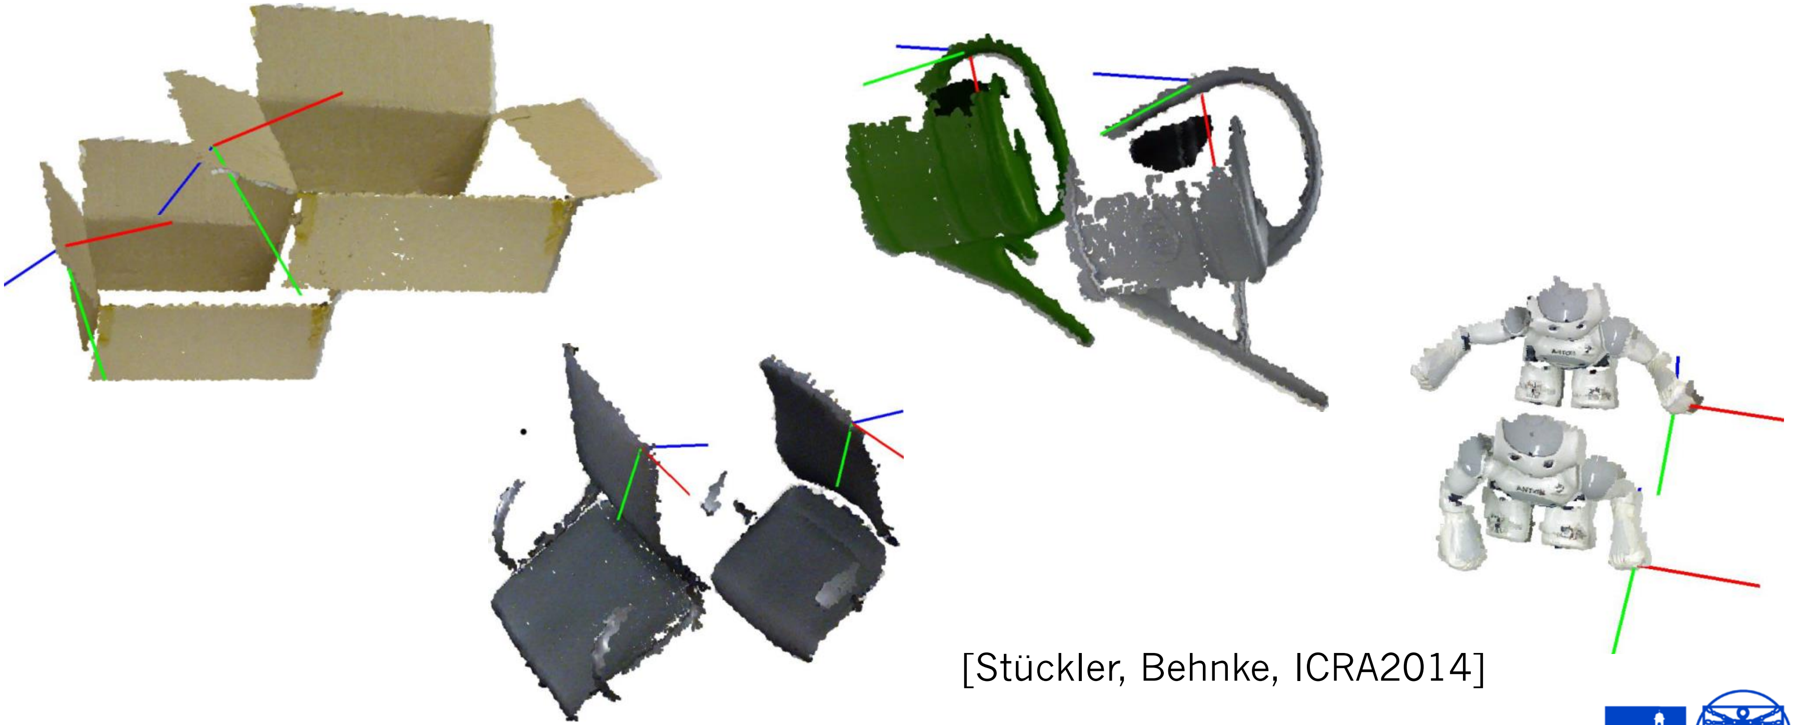

[Stückler, Behnke, ICRA2014]

Grasp & Motion Skill Transfer

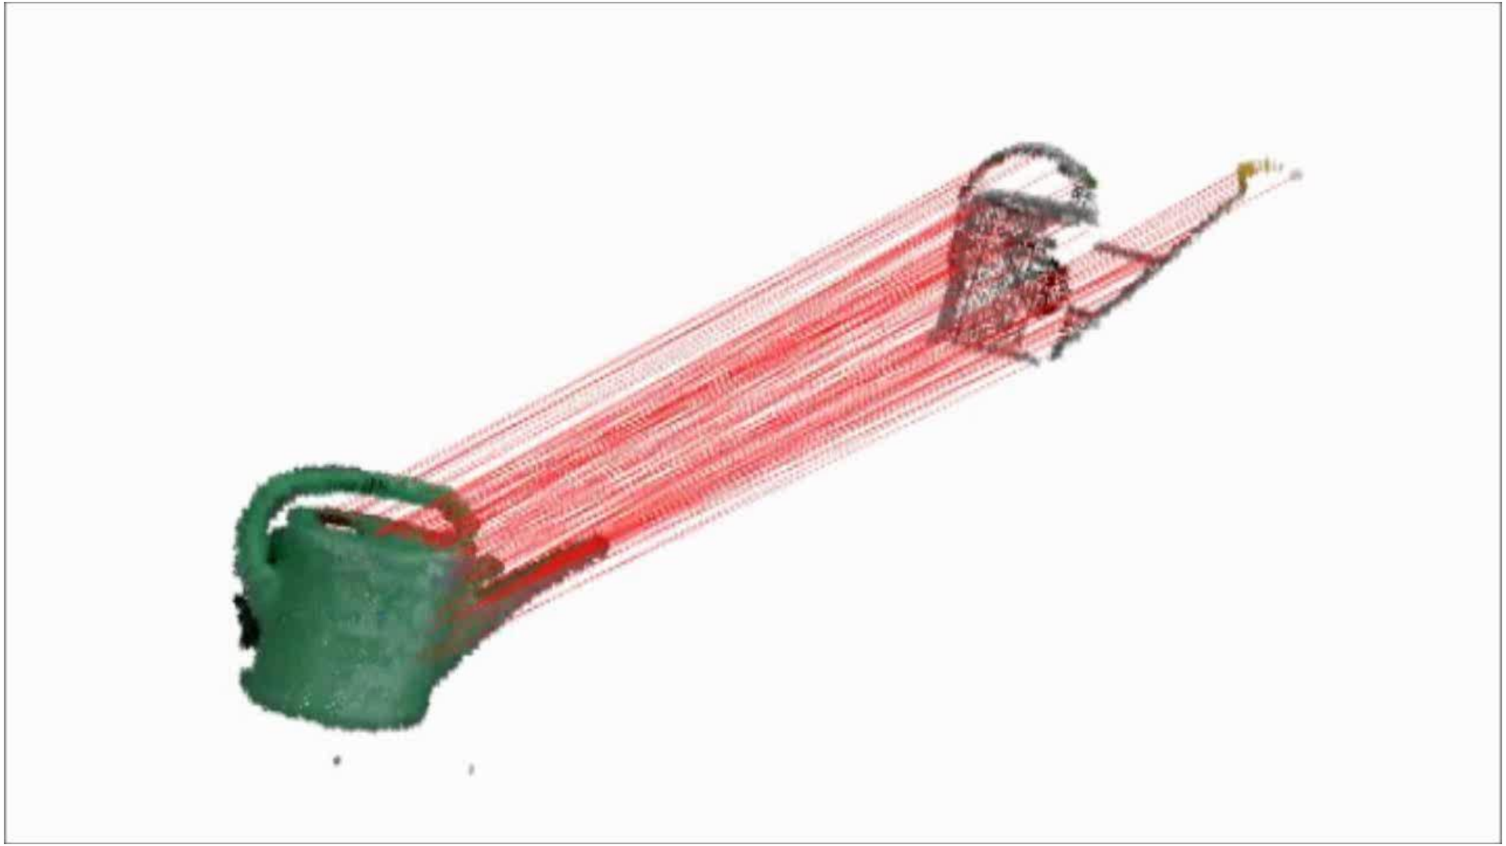

[Stückler, Behnke, ICRA2014]

Tool use: Bottle Opener

- Tool tip perception

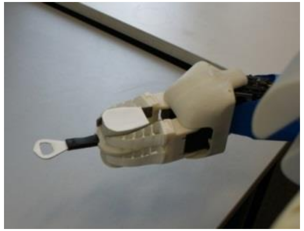

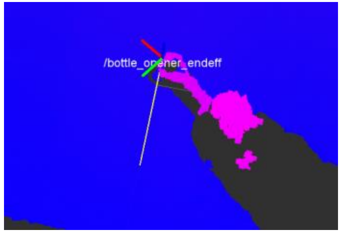

- Extension of arm kinematics
- Perception of crown cap
- Motion adaptation

[Stückler, Behnke, Humanoids 2014]

Hierarchical Object Discovery trough Motion Segmentation

- Simultaneous object modeling and motion segmentation

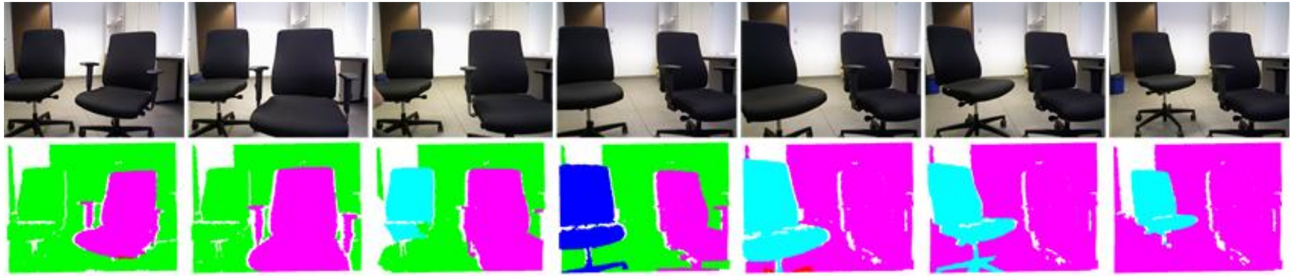

- Inference of a segment hierarchy

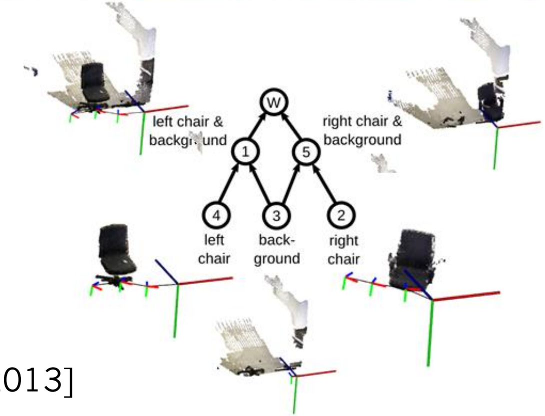

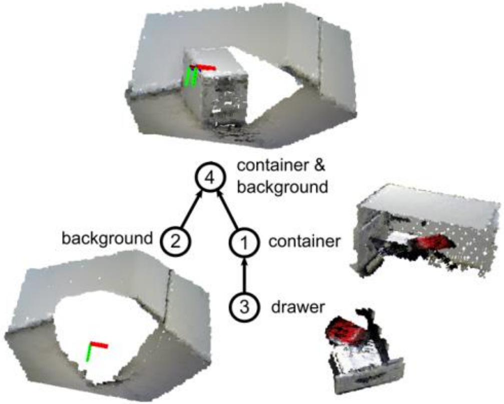

[Stückler, Behnke: IJCAI 2013]

Semantic Mapping

- Pixel-wise classification of RGB-D images by random forests
- Compare color / depth of regions
- Size normalization
- 3D fusion through RGB-D SLAM
- Evaluation on NYU depth v2

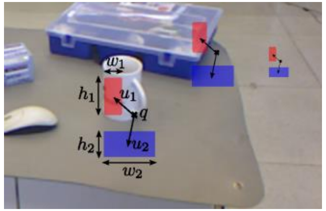

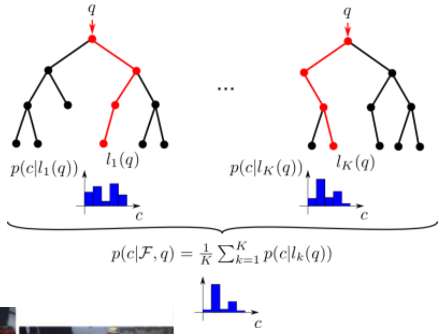

[Stückler, Biresev, Behnke: IROS 2012]

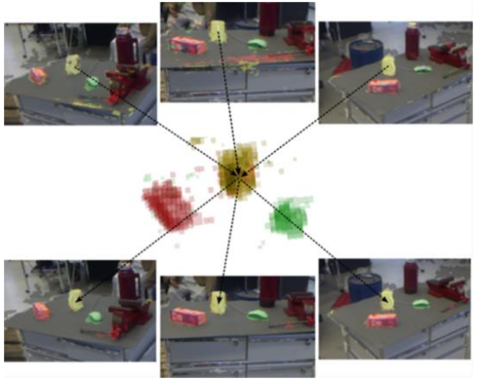

Ground truth

Segmentation

Accuracy in %	Ø Classes	Ø Pixels
Silberman et al. 2012	59,6	58,6
Couprie et al. 2013	63,5	64,5
Random forest	65,0	68,1
3D-Fusion	**66,8**	

Learning Depth-sensitive CRFs

- SLIC+depth super pixels
- Unary features: random forest
- Height feature

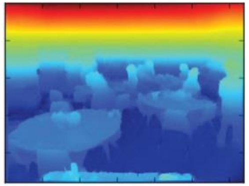

- Pairwise features
  - Color contrast
  - Vertical alignment
  - Depth difference
  - Normal differences

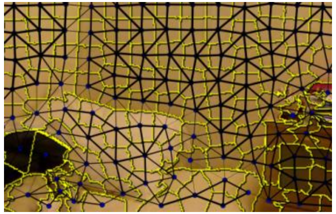

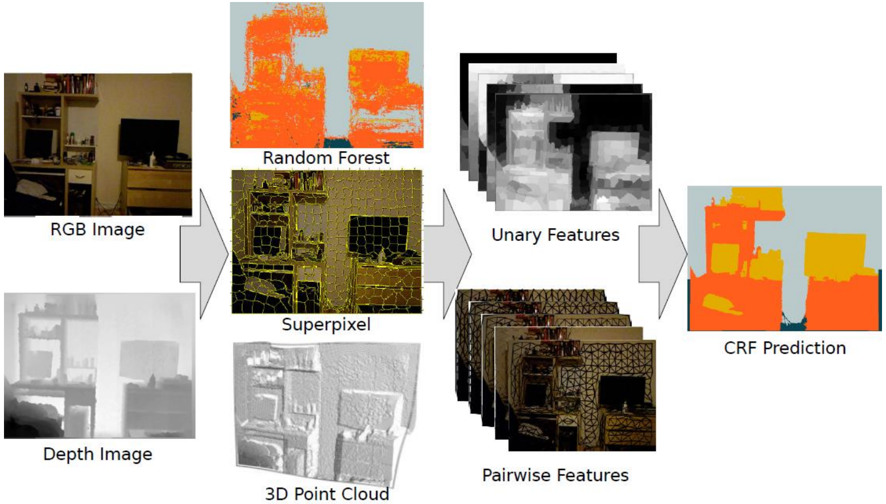

- Results:

	class average	pixel average
RF	65.0	68.3
RF + SP	65.7	70.1
RF + SP + SVM	70.4	70.3
RF + SP + CRF	**71.9**	**72.3**
Silberman *et al.*	59.6	58.6
Couprie *et al.*	63.5	64.5

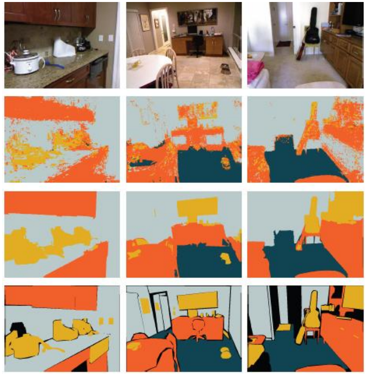

Random forest

CRF prediction

Ground truth

[Müller and Behnke, ICRA 2014]

Deep Learning

- Learning layered represen-tations

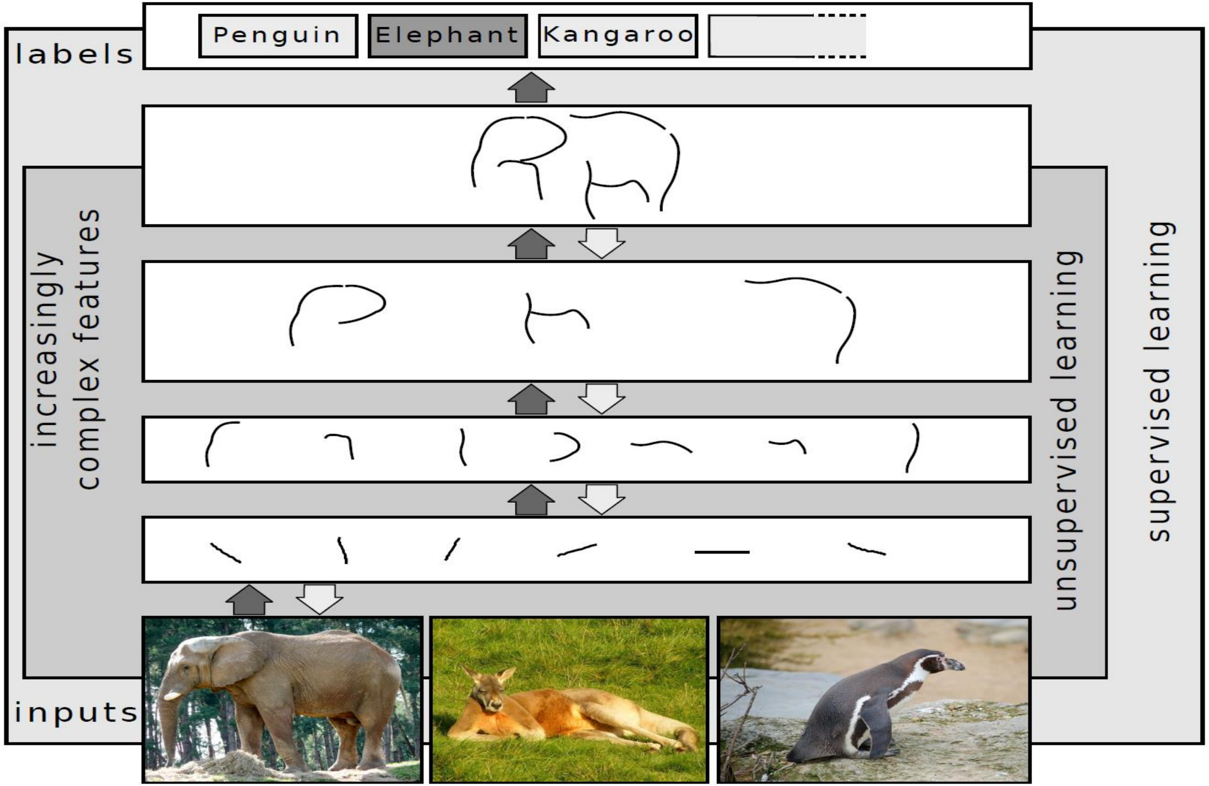

[Schulz; Behnke, KI 2012]

Object-class Segmentation

[Schulz, Behnke, ESANN 2012]

- Class annotation per pixel

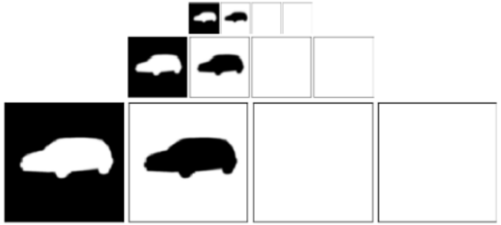

- Multi-scale input channels

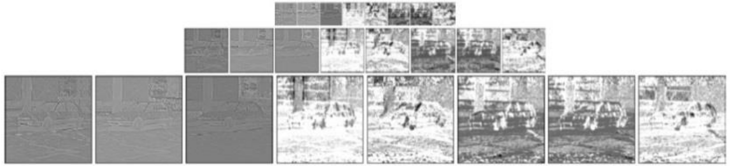

- Evaluated on MSRC-9/21 and INRIA Graz-02 data sets

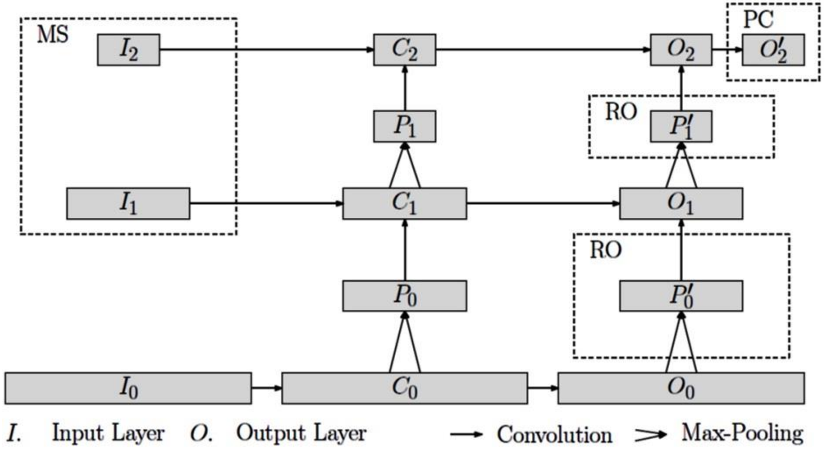

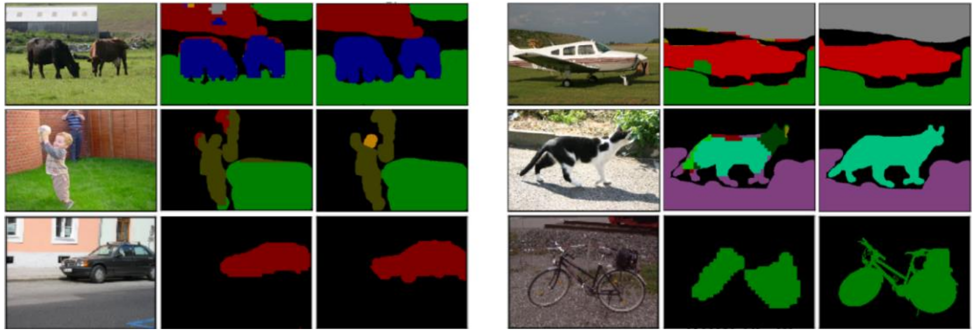

Input Output Truth

Object Detection in Natural Images

- Bounding box annotation
- Structured loss that directly maximizes overlap of the prediction with ground truth bounding boxes
- Evaluated on two of the Pascal VOC 2007 classes

[Schulz, Behnke, ICANN 2014]

RGB-D Object-Class Segmentation

- Covering windows segmented with CNN
- Scale input according to depth, compute pixel hight

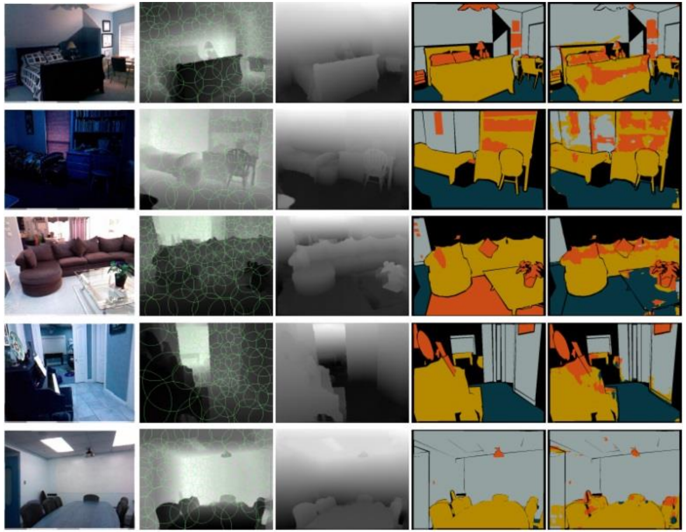

Method	floor	struct	furnit	prop	Class Avg.	Pixel Acc.
CW	84.6	70.3	58.7	52.9	66.6	65.4
CW+DN	87.7	70.8	57.0	53.6	67.3	65.5
CW+H	78.4	74.5	55.6	62.7	67.8	66.5
CW+DN+H	93.7	72.5	61.7	55.5	70.9	70.5
CW+DN+H+SP	91.8	74.1	59.4	63.4	72.2	71.9
CW+DN+H+CRF	93.5	80.2	66.4	54.9	**73.7**	**73.4**
Müller et al.[8]	94.9	78.9	71.1	42.7	71.9	72.3
Random Forest [8]	90.8	81.6	67.9	19.9	65.1	68.3
Couprie et al.[9]	87.3	86.1	45.3	35.5	63.6	64.5
Höft et al.[10]	77.9	65.4	55.9	49.9	62.3	62.0
Silberman [12]	68	59	70	42	59.7	58.6

CW is covering windows, H is height above ground, DN is depth normalized patch sizes. SP is averaged within superpixels and SVM-reweighted. CRF is a conditional random field over superpixels [8]. Structure class numbers are optimized for class accuracy.

[Schulz, Höft, Behnke, ESANN 2015]

Neural Abstraction Pyramid

[Behnke, Rojas, IJCNN 1998]
[Behnke, LNCS 2766, 2003]

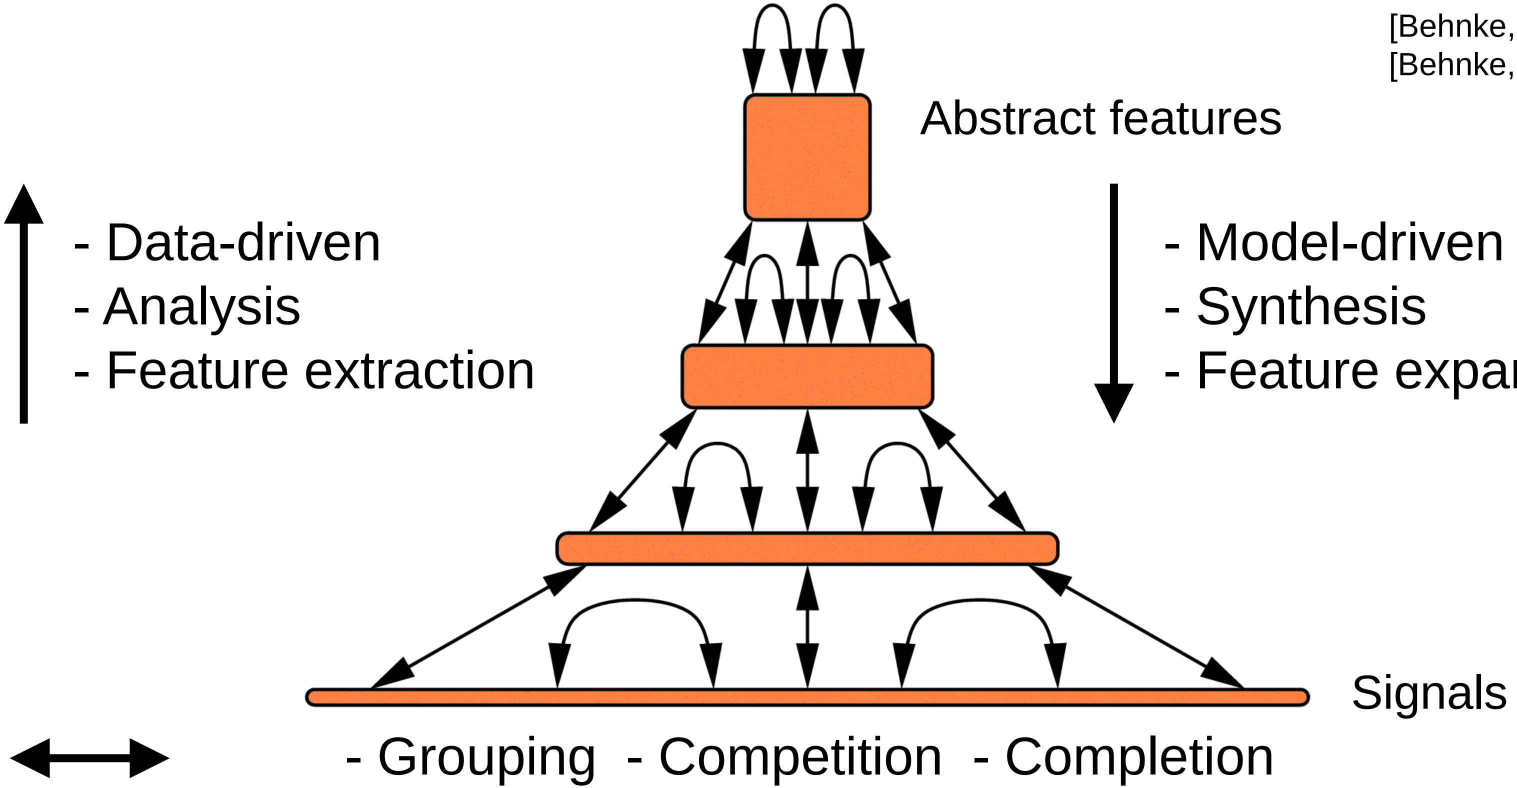

Iterative Image Interpretation

- Interpret most obvious parts first
- Use partial interpretation as context to resolve local ambiguities

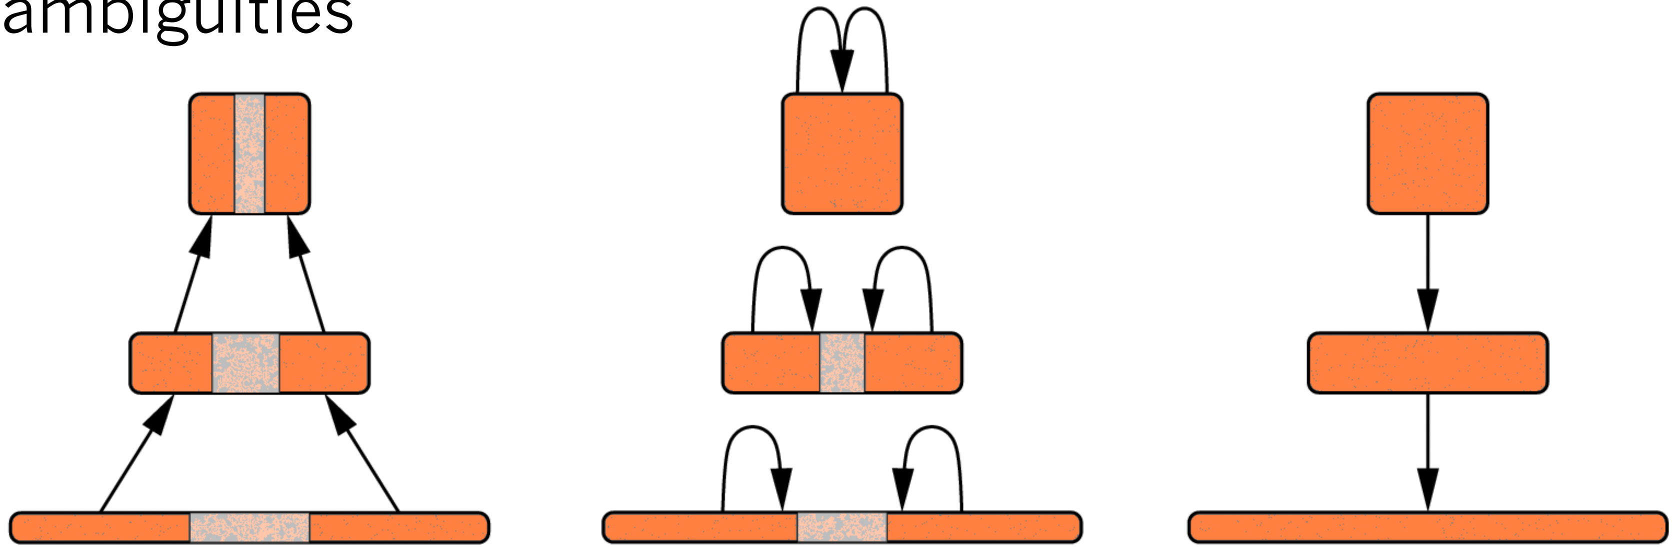

Neural Abstraction Pyramid for RGB-D Video Object-class Segmentation

- Recursive computation is efficient for temporal integration

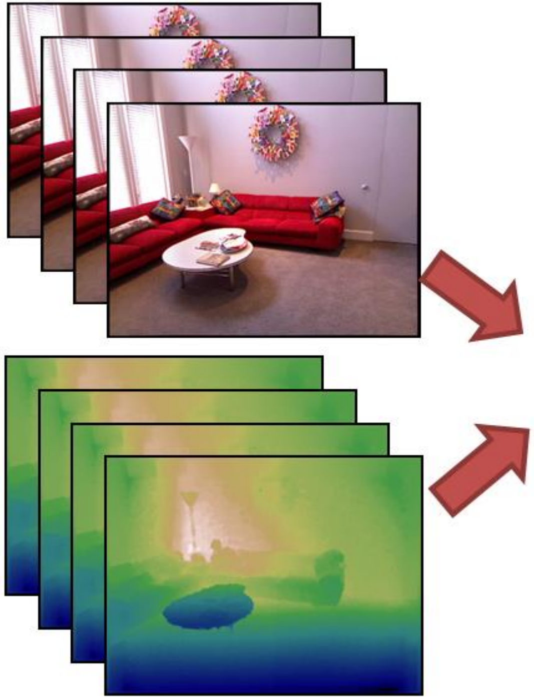

Neural Abstraction Pyramid

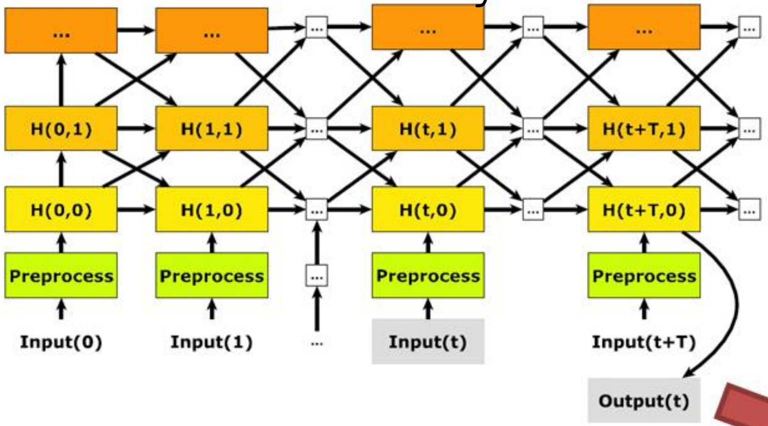

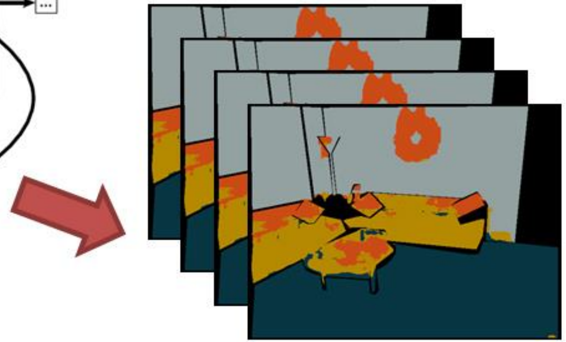

[Pavel, Schulz, Behnke, Neural Networks 2017]

Geometric and Semantic Features for RGB-D Object-class Segmentation

- New **geometric** feature: distance from wall
- **Semantic** features pretrained from ImageNet
- Both help significantly

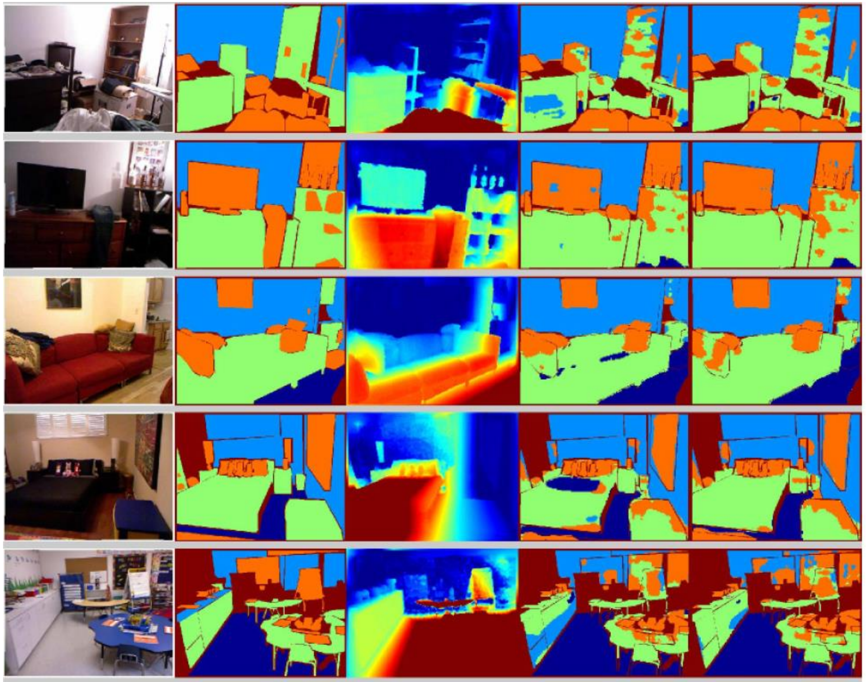

RGB Truth DistWall OutWO OutWithDistWall

[Husain et al. RA-L 2016]

Semantic Segmentation Priors for Object Discovery

- Combine bottom-up object discovery and semantic priors
- Semantic segmentation used to classify color and depth superpixels
- Higher recall, more precise object borders

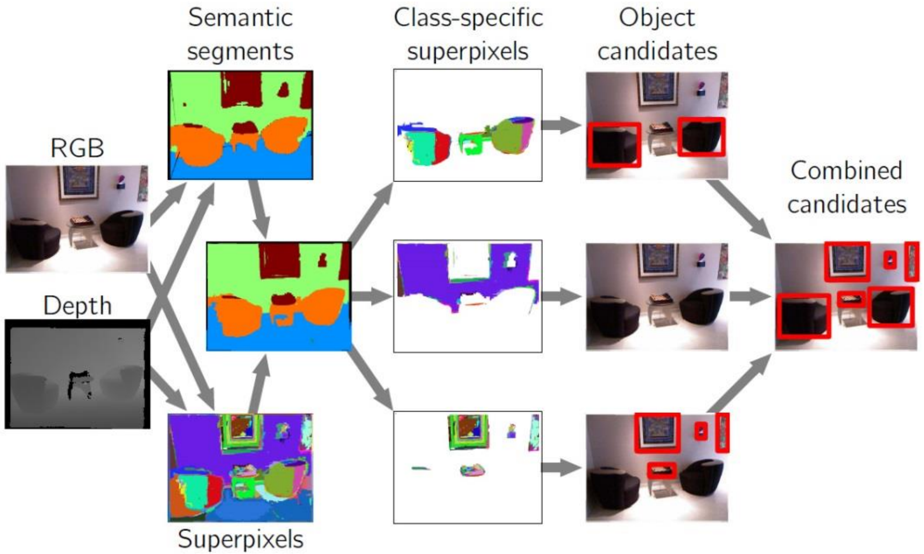

[Garcia et al. ICPR 2016]

RGB-D Object Recognition and Pose Estimation

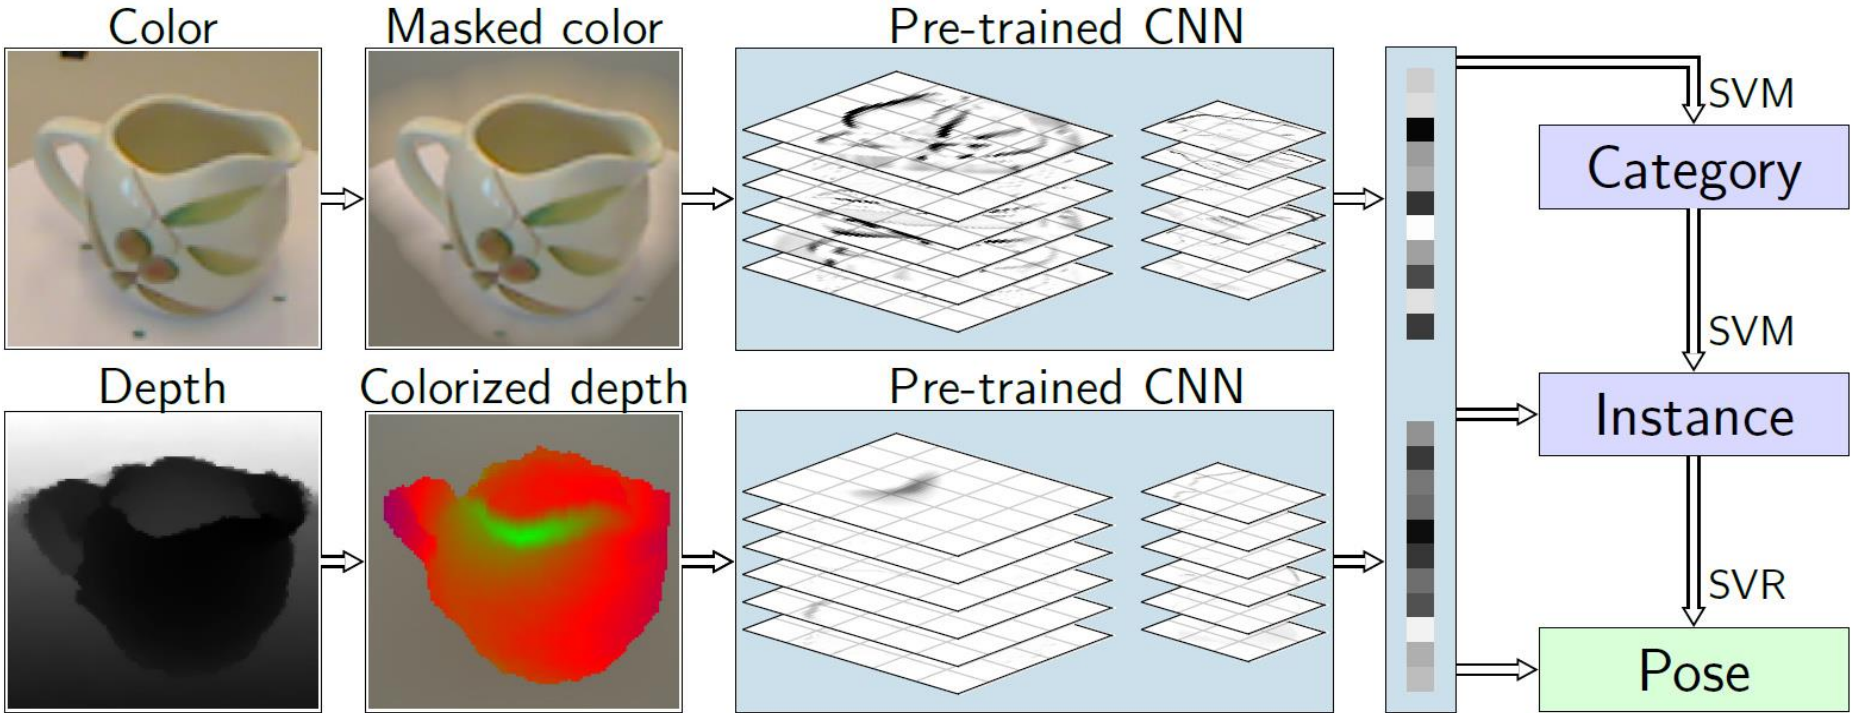

[Schwarz, Schulz, Behnke, ICRA2015]

Canonical View, Colorization

- Objects viewed from different elevation
- Render canonical view

- Colorization based on distance from center vertical

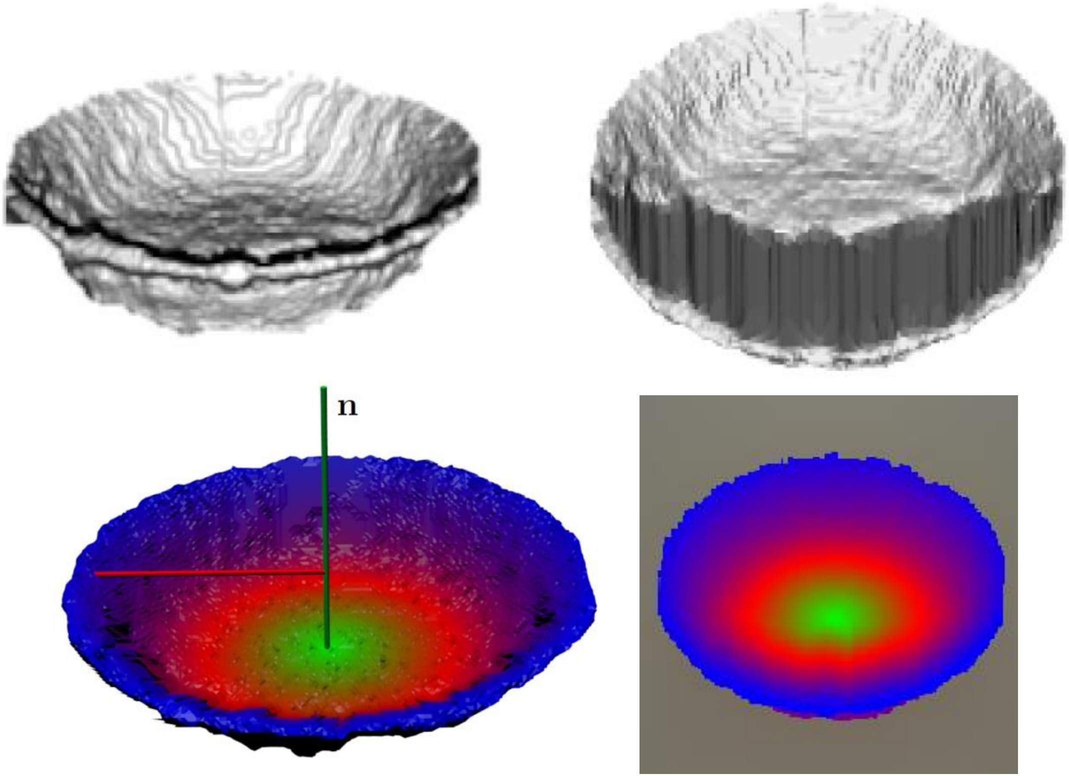

[Schwarz, Schulz, Behnke, ICRA2015]

Pretrained Features Disentangle Data

- t-SNE embedding

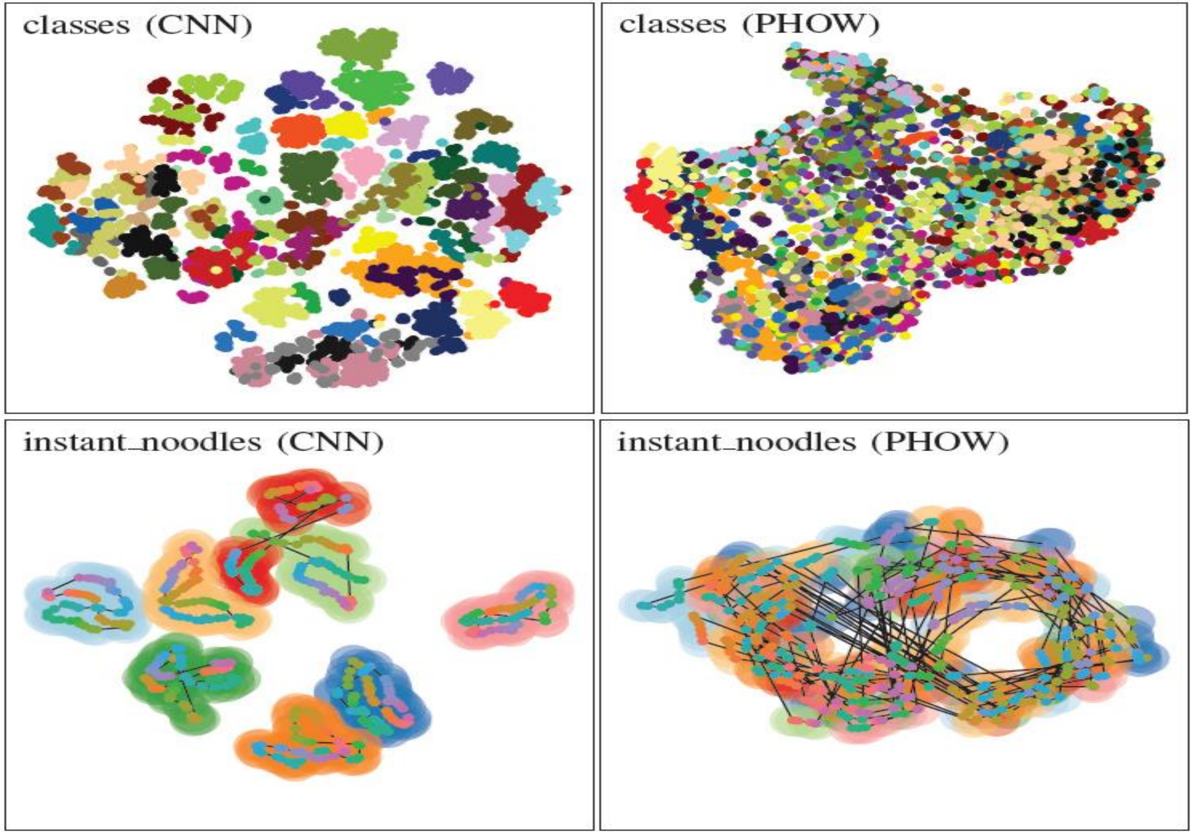

[Schwarz, Schulz, Behnke ICRA2015]

Recognition Accuracy

- Improved both category and instance recognition

Method	Category Accuracy (%) RGB	Category Accuracy (%) RGB-D	Instance Accuracy (%) RGB	Instance Accuracy (%) RGB-D
Lai *et al.* [1]	74.3 ± 3.3	81.9 ± 2.8	59.3	73.9
Bo *et al.* [2]	82.4 ± 3.1	87.5 ± 2.9	**92.1**	92.8
PHOW[3]	80.2 ± 1.8	—	62.8	—
Ours	$\mathbf{83.1 \pm 2.0}$	88.3 ± 1.5	92.0	**94.1**
Ours	$\mathbf{83.1 \pm 2.0}$	$\mathbf{89.4 \pm 1.3}$	92.0	**94.1**

- Confusion:

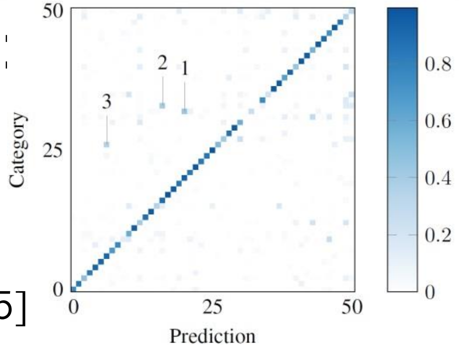

1: pitcher / coffe mug

2: peach / sponge

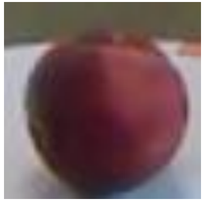

[Schwarz, Schulz, Behnke, ICRA2015]

Autonomous Flight Near Obstacles

- Multimodal obstacle detection
  - 3D laser scanner

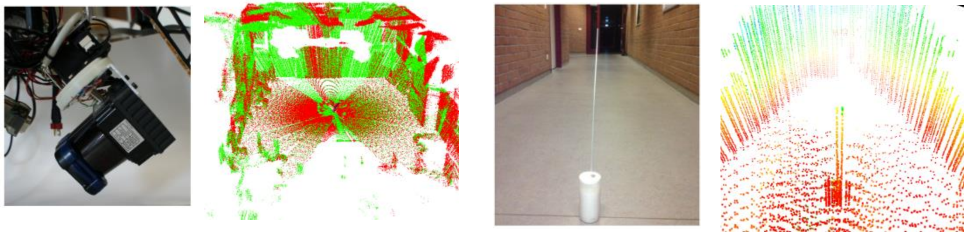

  - Stereo cameras

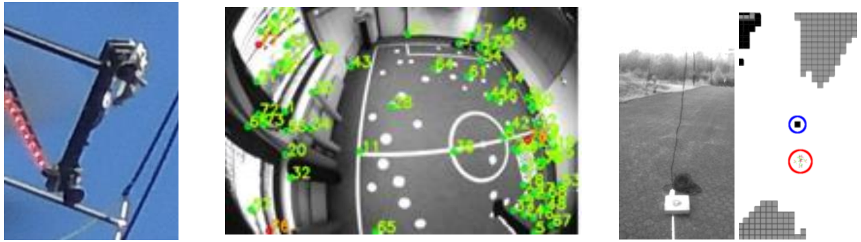

  - Ultrasound

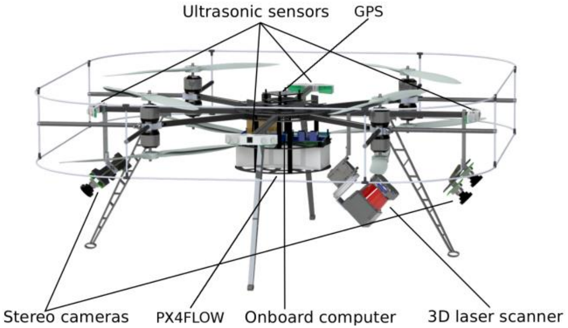

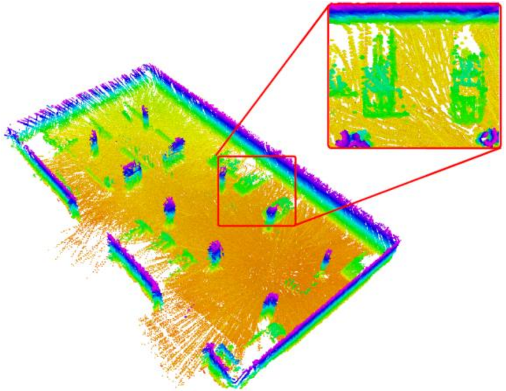

[Droeschel et al.: Journal of Field Robotics, 2015]

universitätbonn ais

Egocentric Laser-based 3D Mapping

- Motion compensation

Distorted

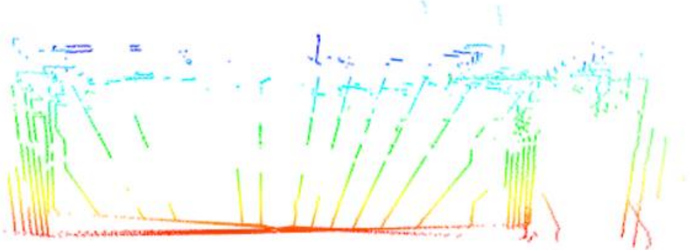

Undistrorted

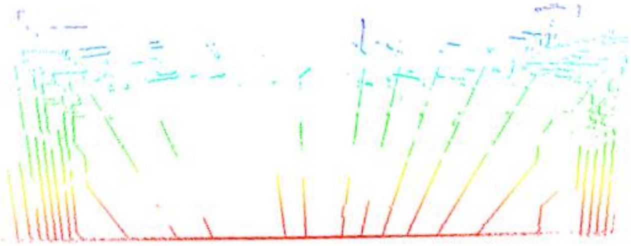

- Local multiresolution surfel maps

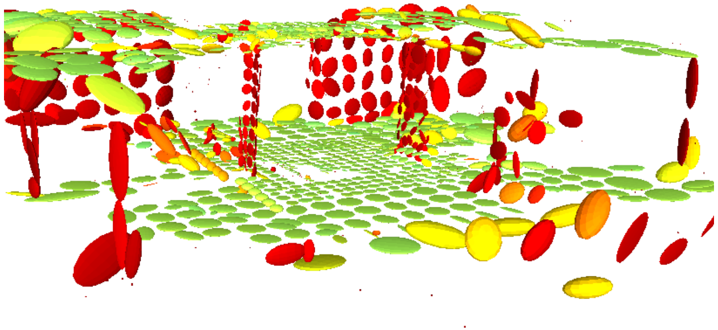

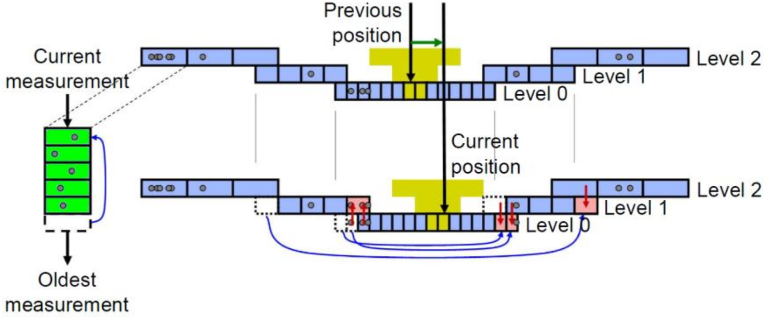

universitätbonn ais

Allocentric 3D Map

- Registration of egocentric maps
- Global optimization of registration error by GraphSLAM

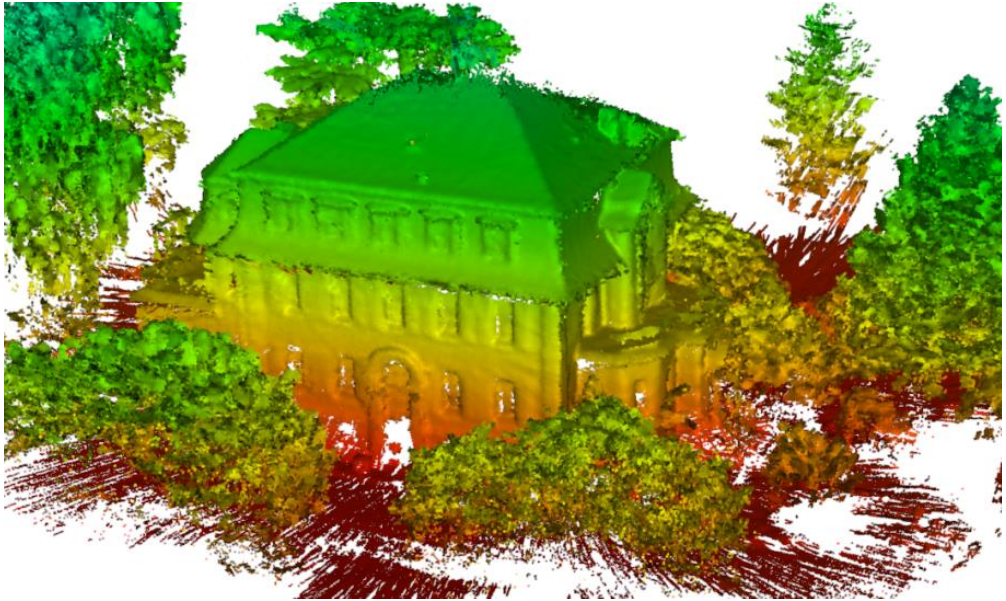

[Droeschel et al. JFR 2016]

Hierarchical Navigation

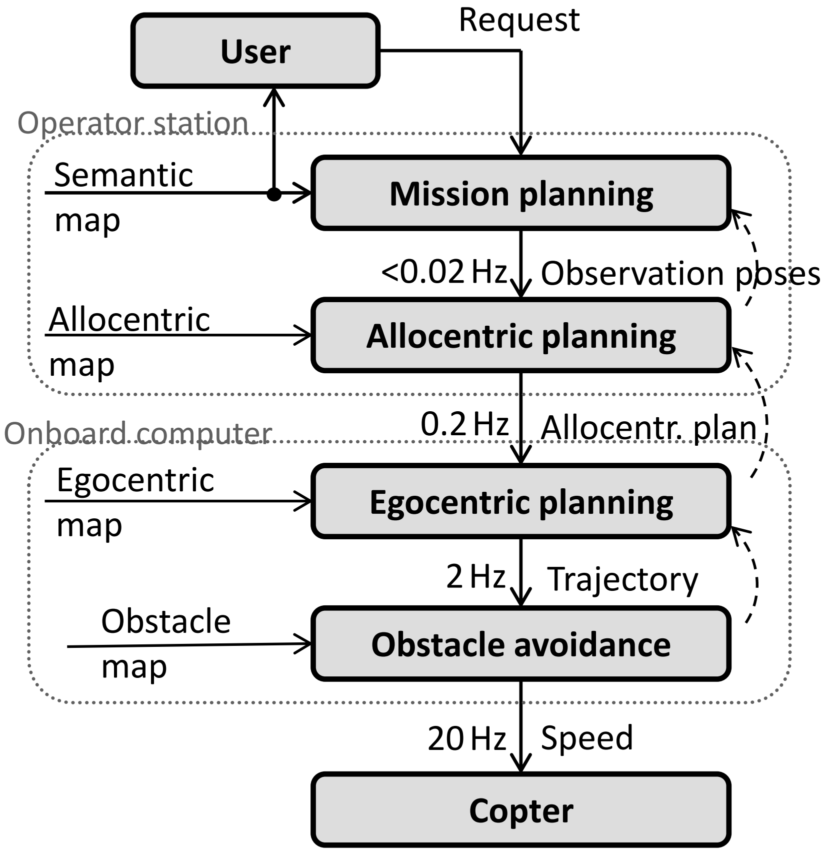

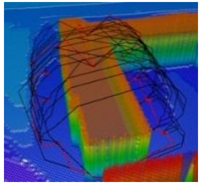

Mission plan

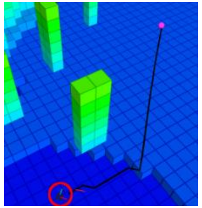

Allocentric planning

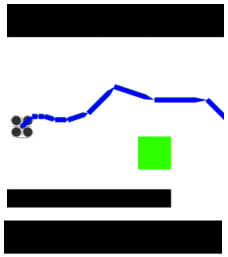

Egocentric planning

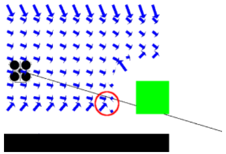

Obstacle avoidance

Mapping on Demand

Autonomous Flight to Planned View Poses

3D Simultaneous Localization and Mapping

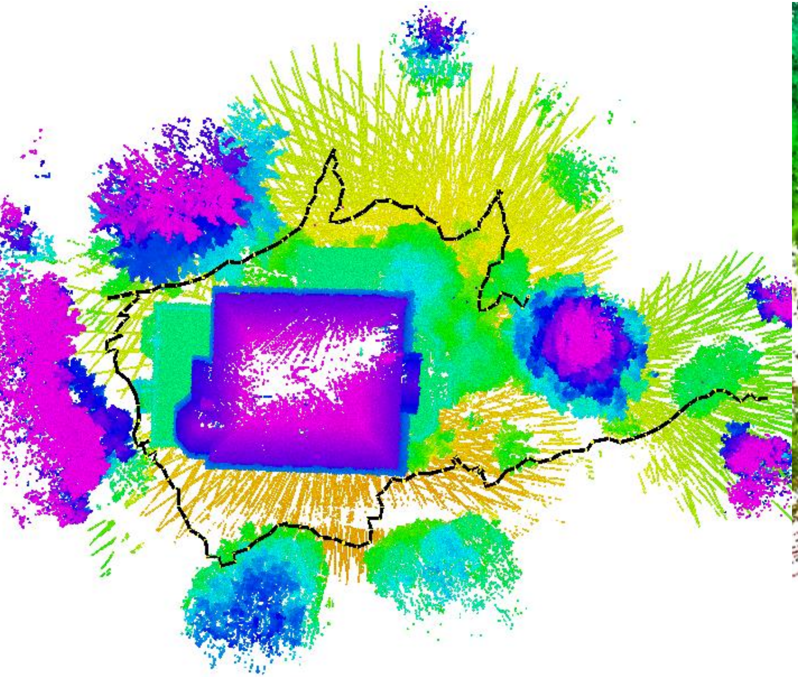

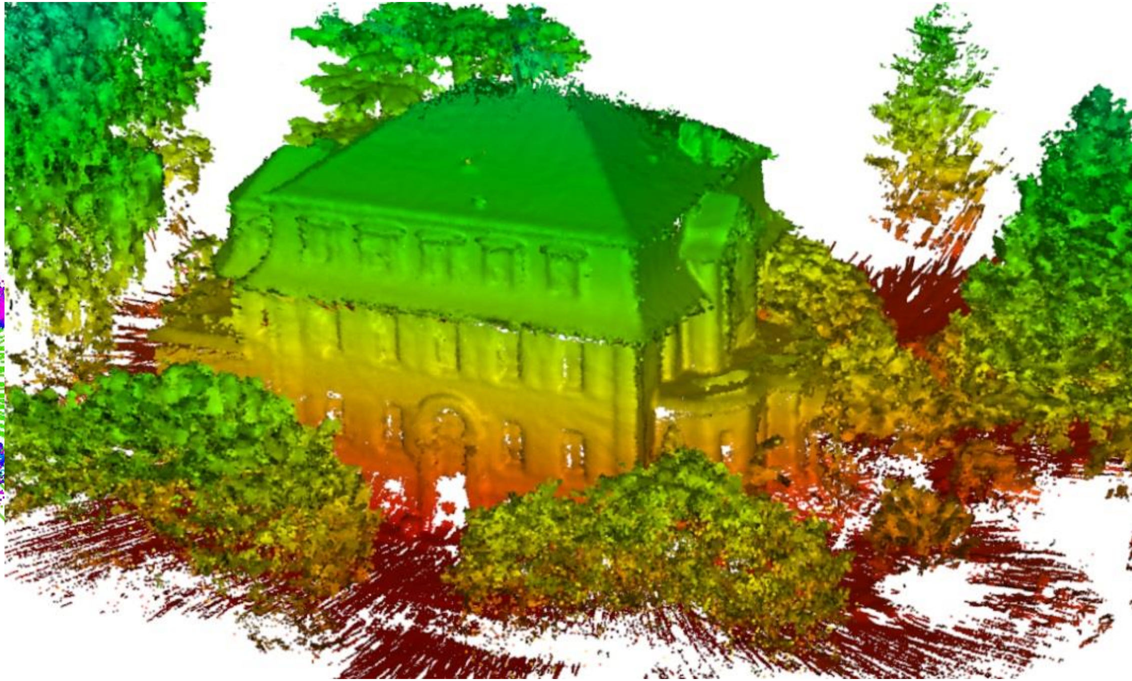

Autonomous Flight in Warehouses

- Dual 3D laser scanner

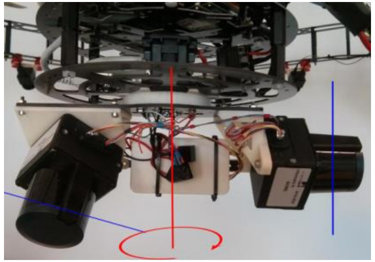
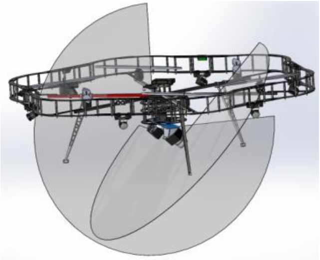
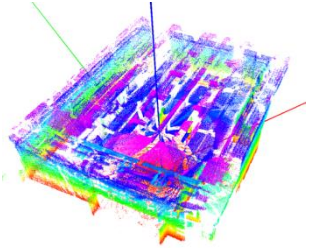
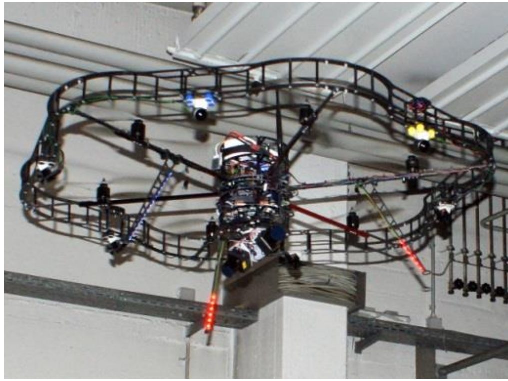

- Omnidirectional cameras

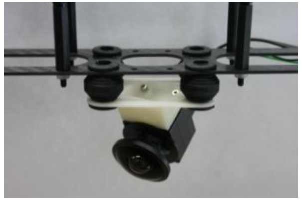
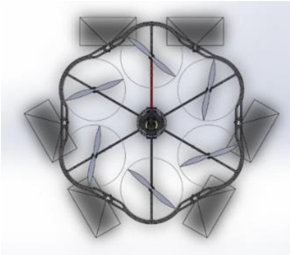
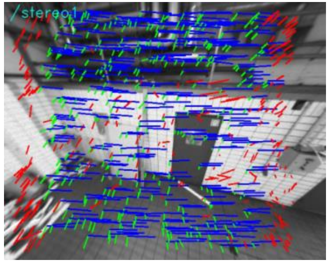

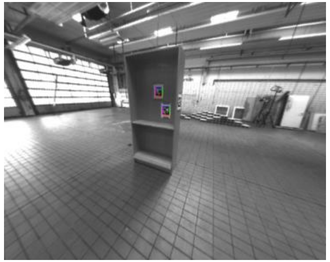
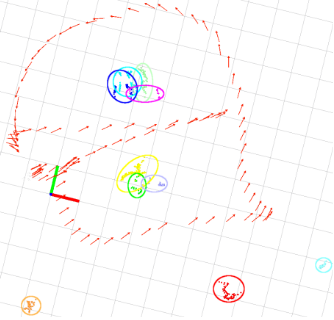

- RFID reader

3D Map

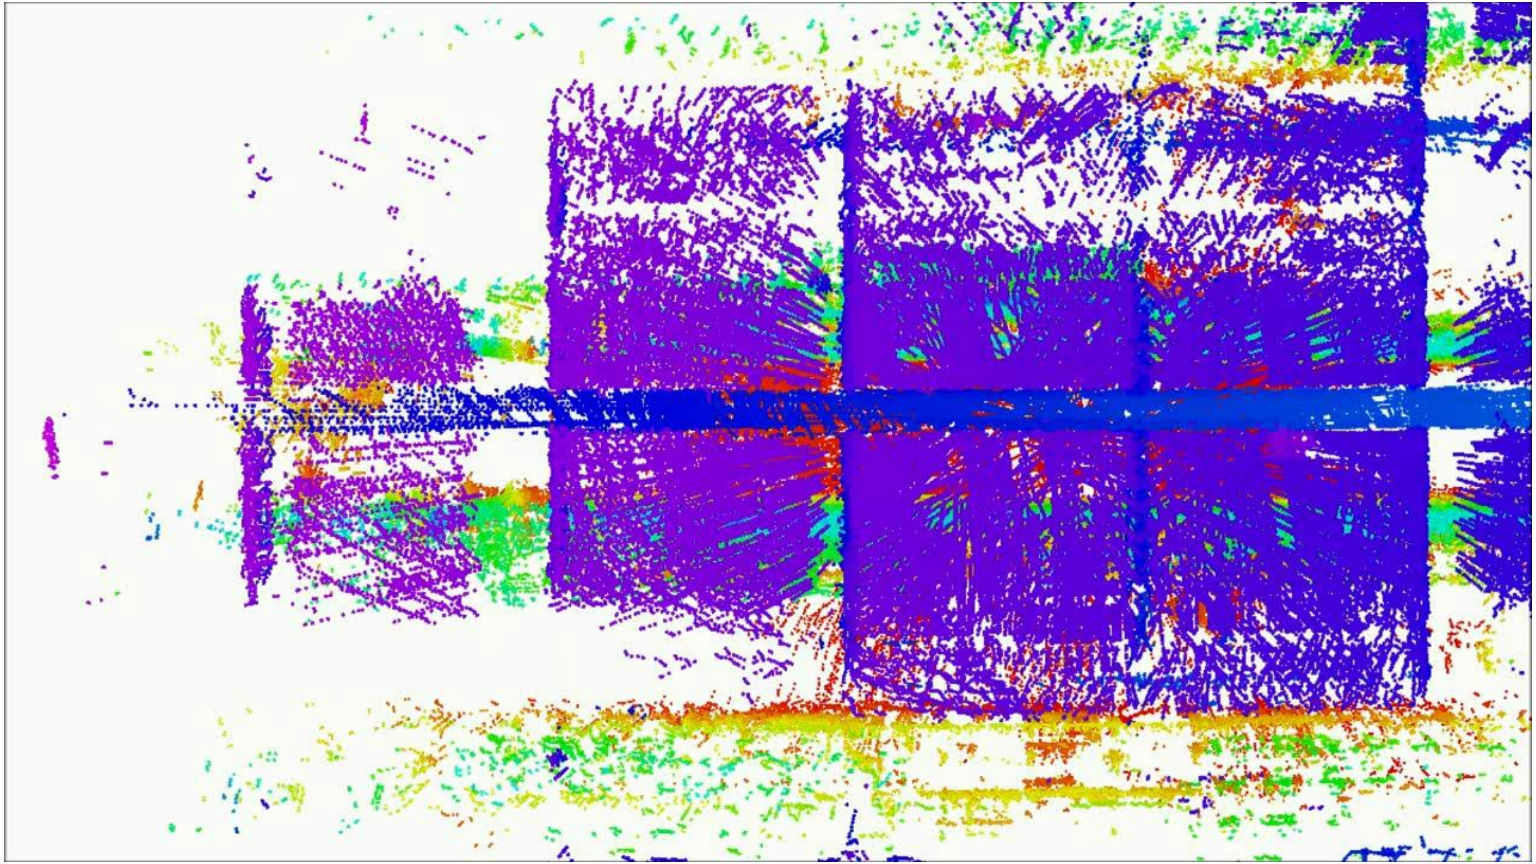

Localization

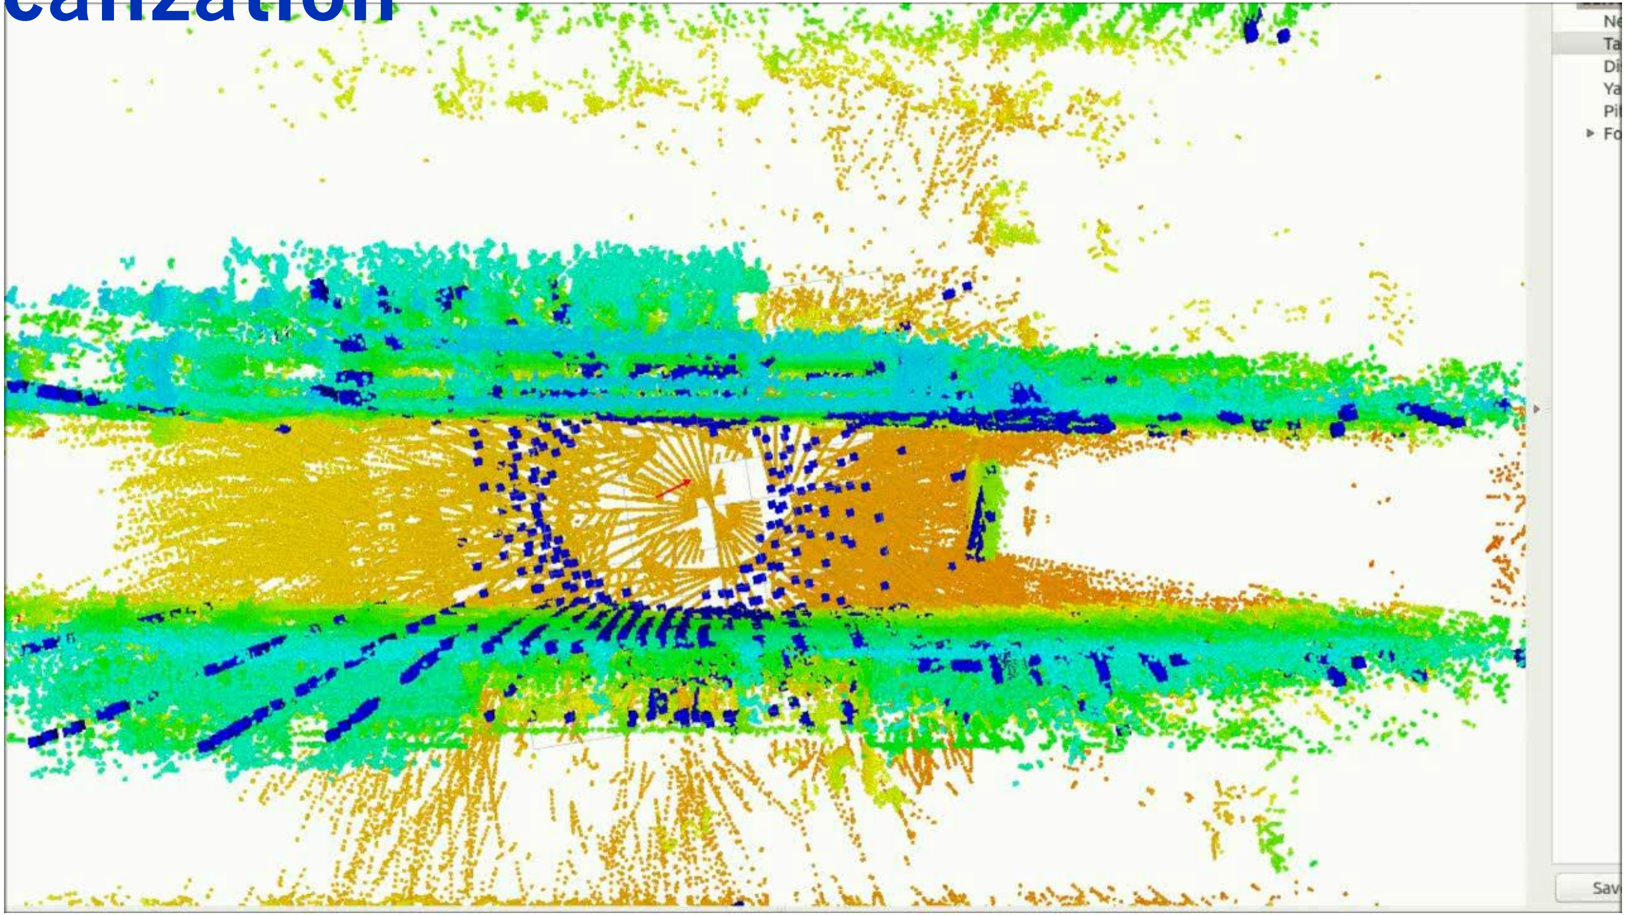


Autonomous Mission in Warehouse

DJI Matrice 600 with Velodyne Puck

InventAIRy Final Demonstration

EuRoC Challenge 3: ChimneySpector

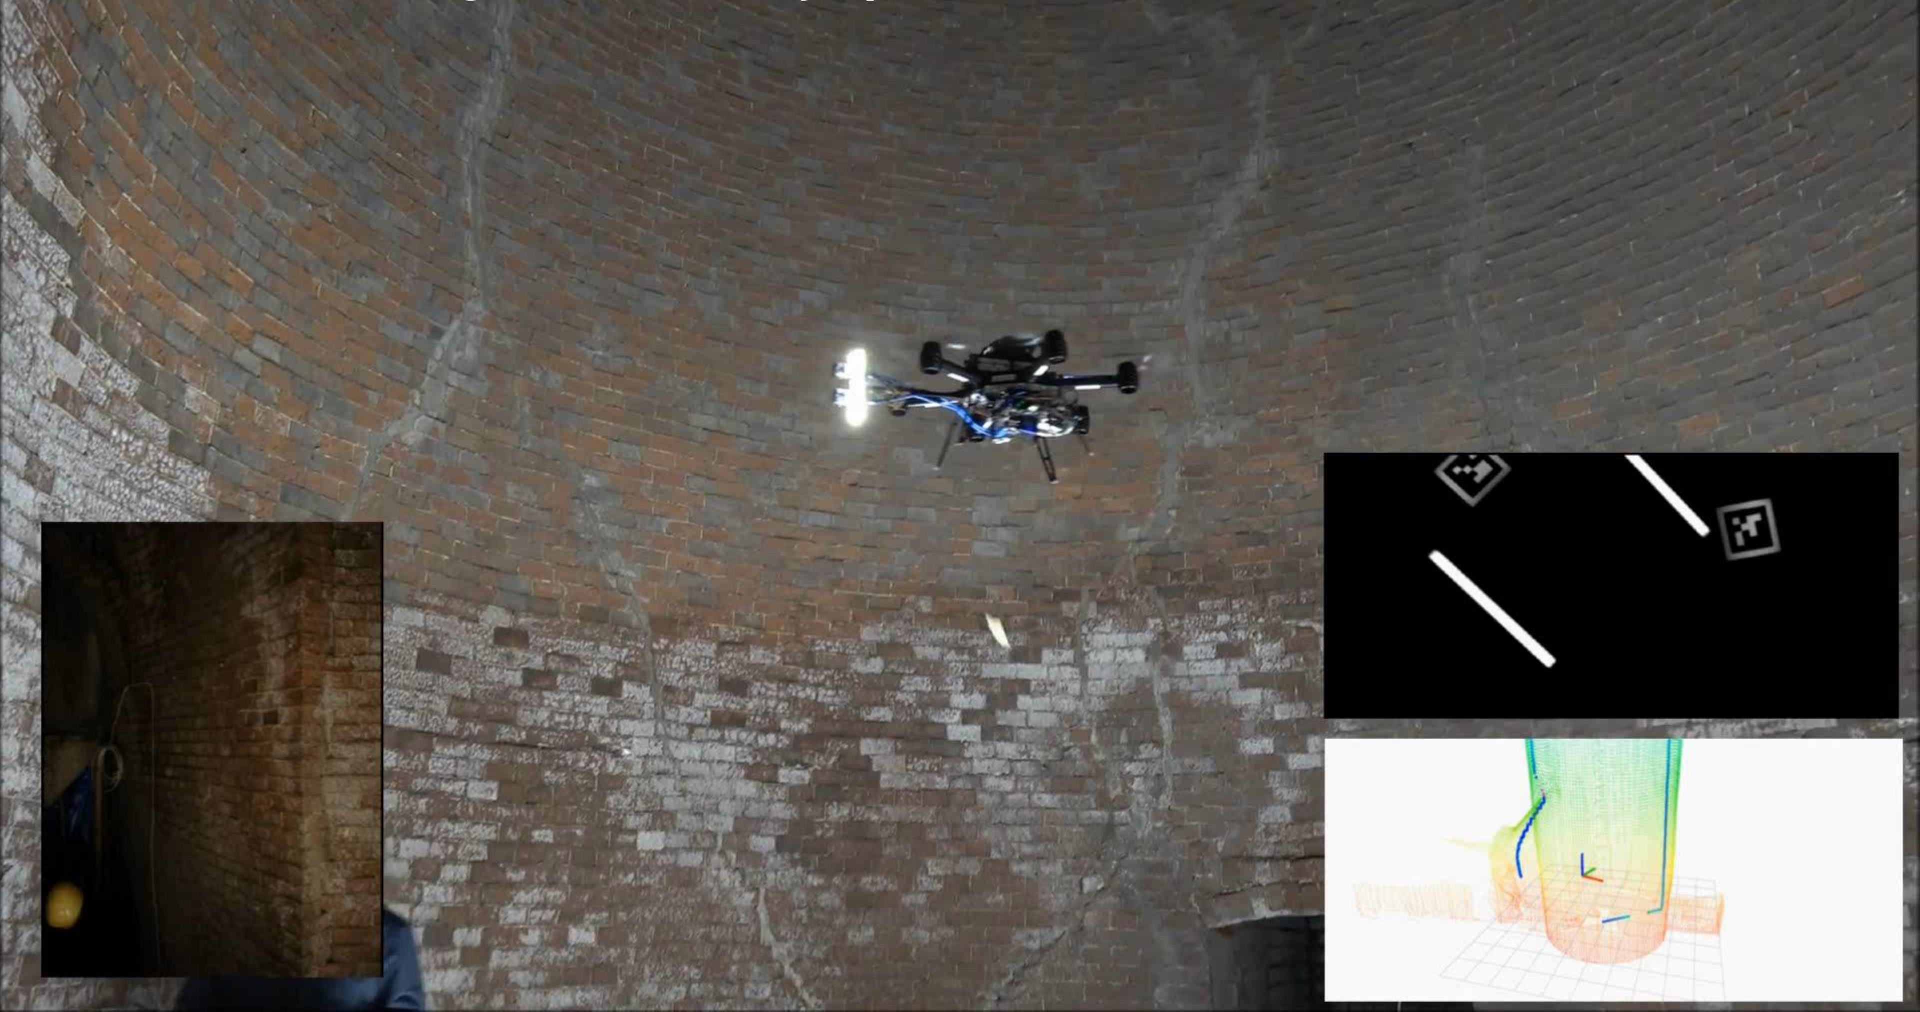

Mobile Manipulation Robot Momaro

- Four compliant legs ending in pairs of steerable wheels
- Anthropomorphic upper body
- Sensor head
  - 3D laser scanner
  - IMU, cameras

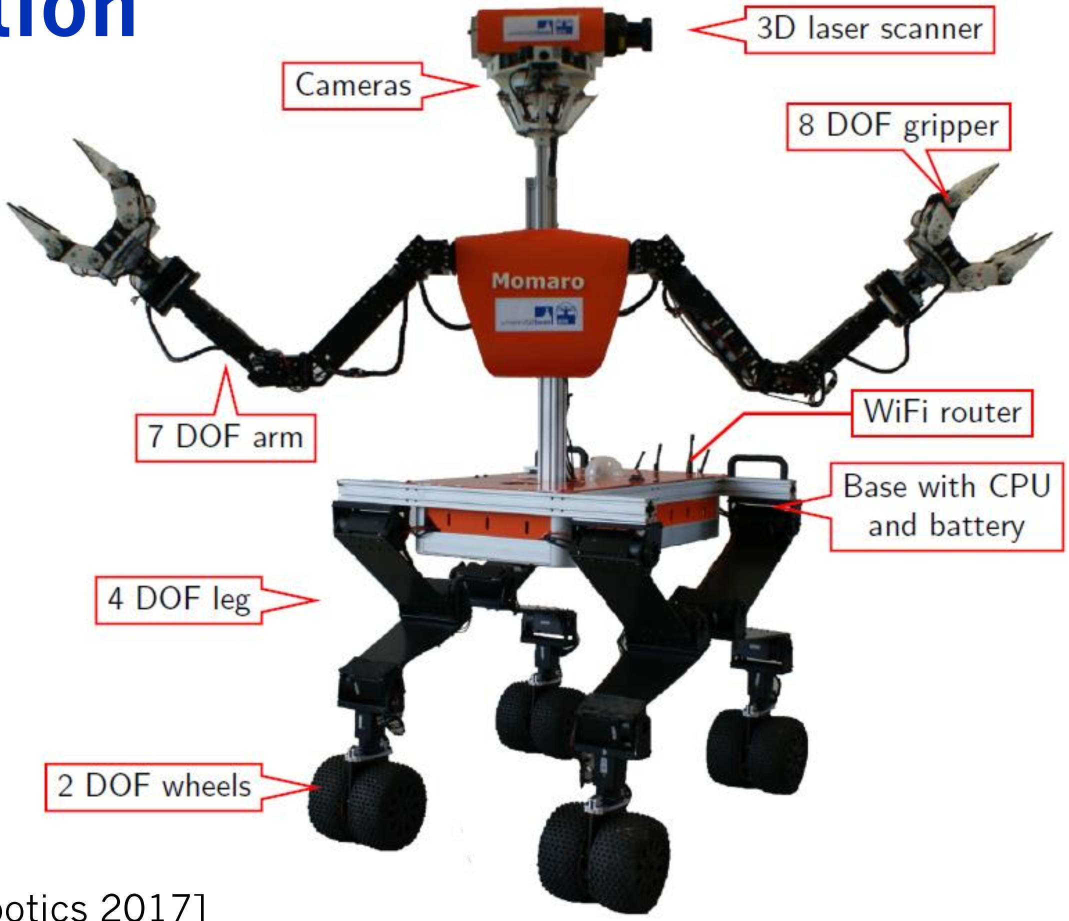

[Schwarz et al. Journal of Field Robotics 2017]

Driving a Vehicle

Egress

Manipulation Operator Interface

- 3D head-mounted display

- 3D environ-ment model + images

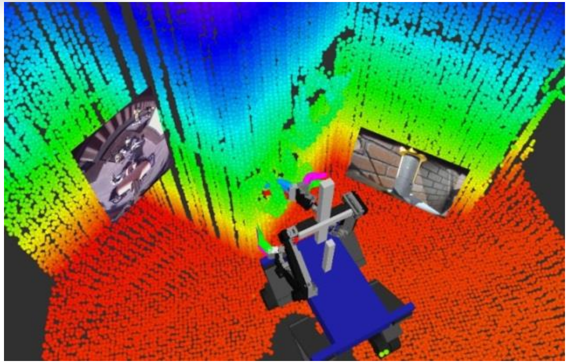

- 6D magnetic tracker

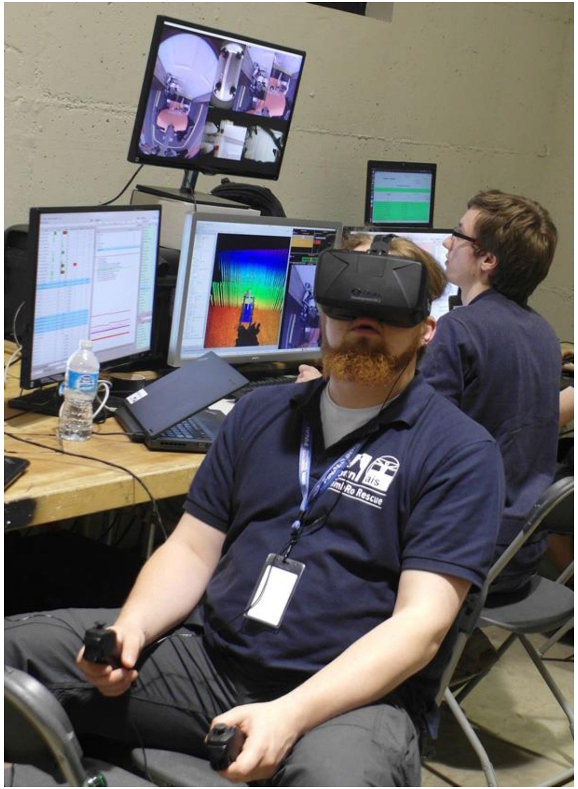

[Rodehutskors et al., Humanoids 2015]

Opening a Door

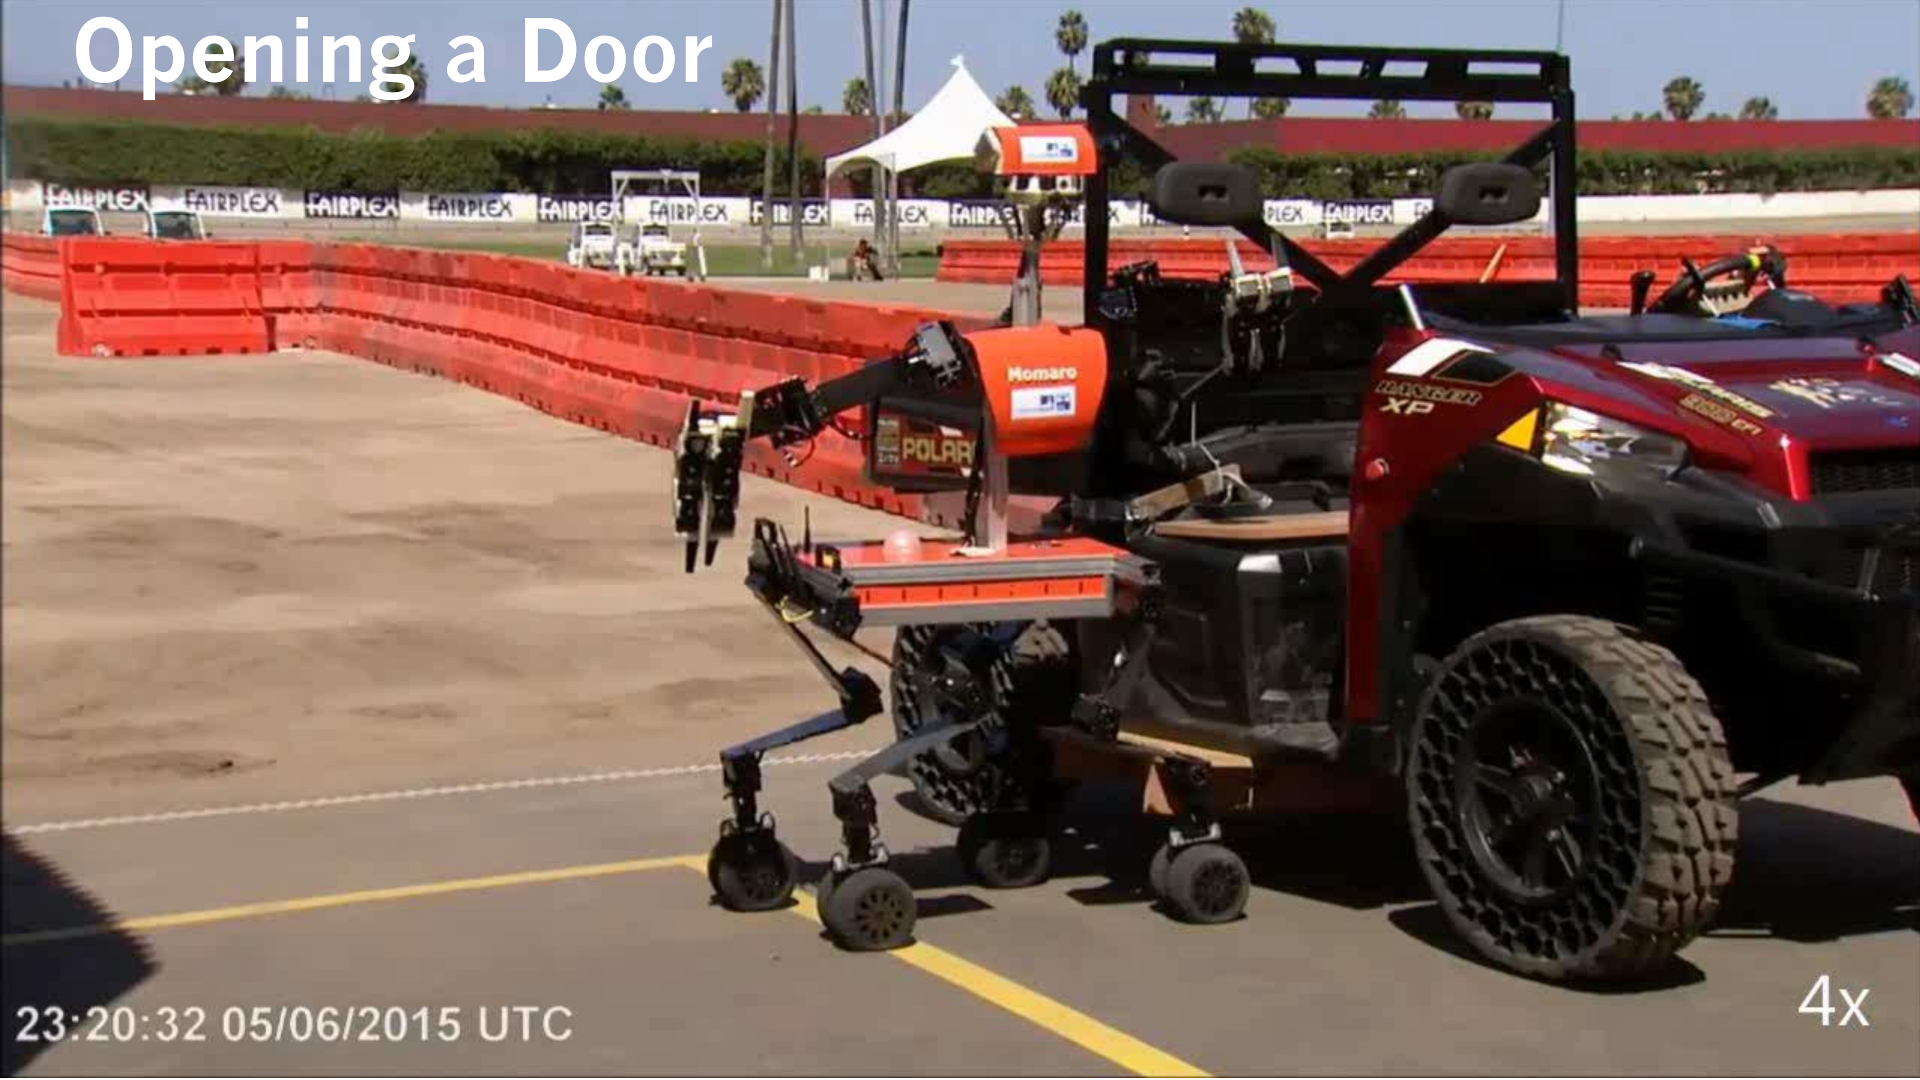

Local Multiresolution Surfel Map

- Registration and aggregation of 3D laser scans
- Local multi-resolution grid
- Surfel in grid cells

[Droeschel et al., Robotics and Autonomous Systems 2017]

3D scan

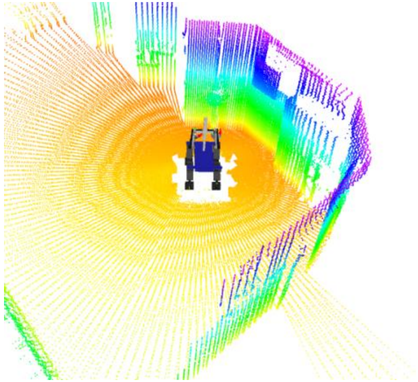

Aggregated scans

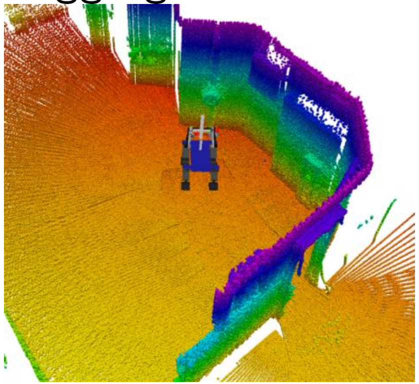

Multiresolution grid

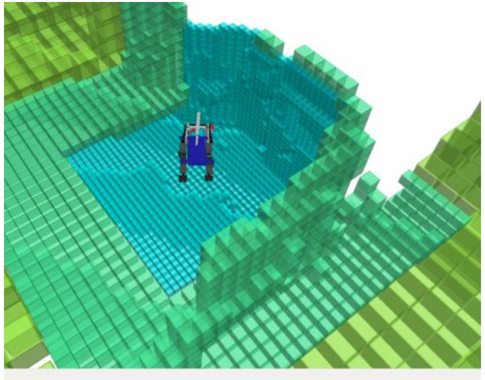

Surfels

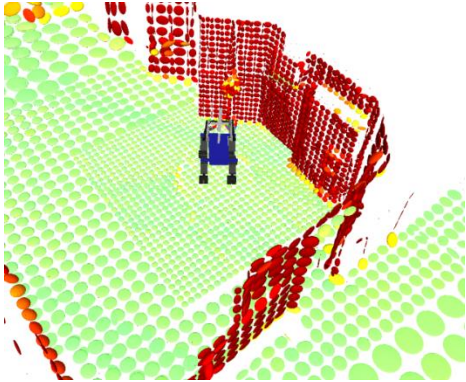

Filtering Dynamic Objects

- Maintain occupancy in each cell

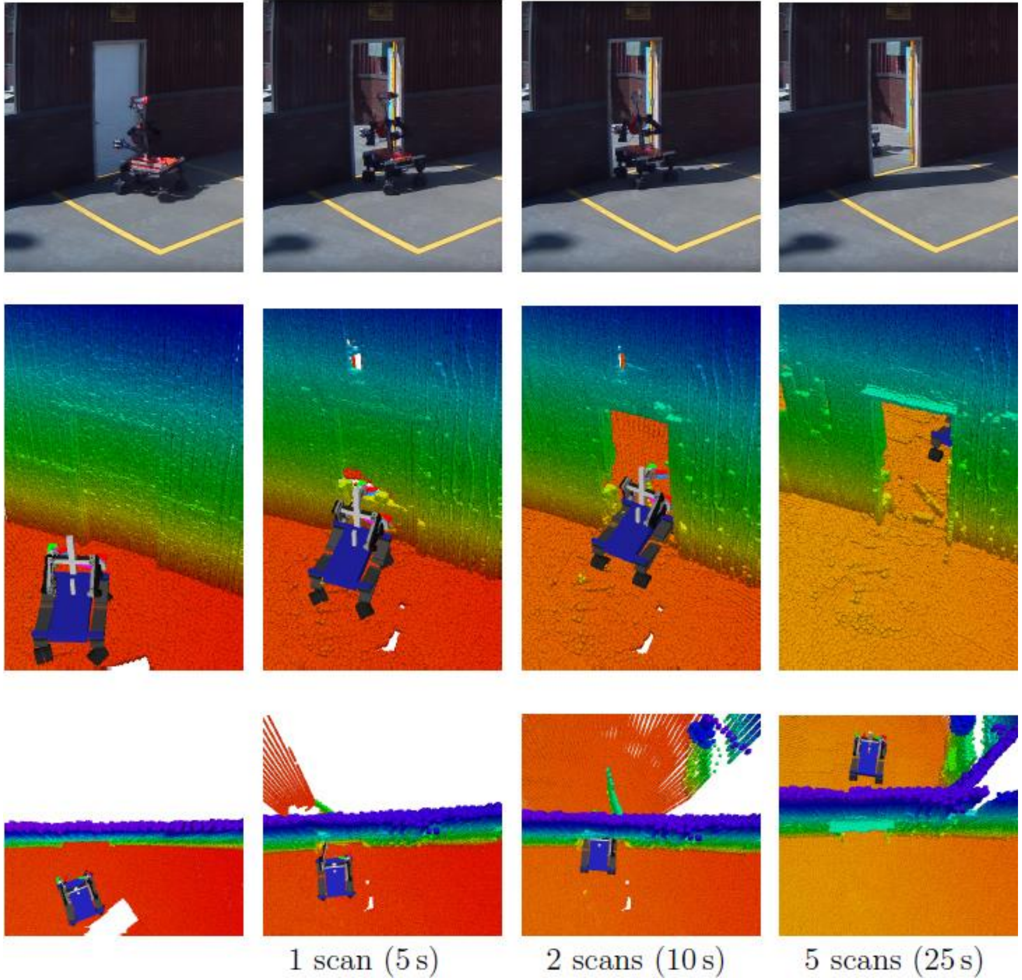

universitätbonn ais

Allocentric 3D Mapping

- Registration of egocentric maps by graph optimization

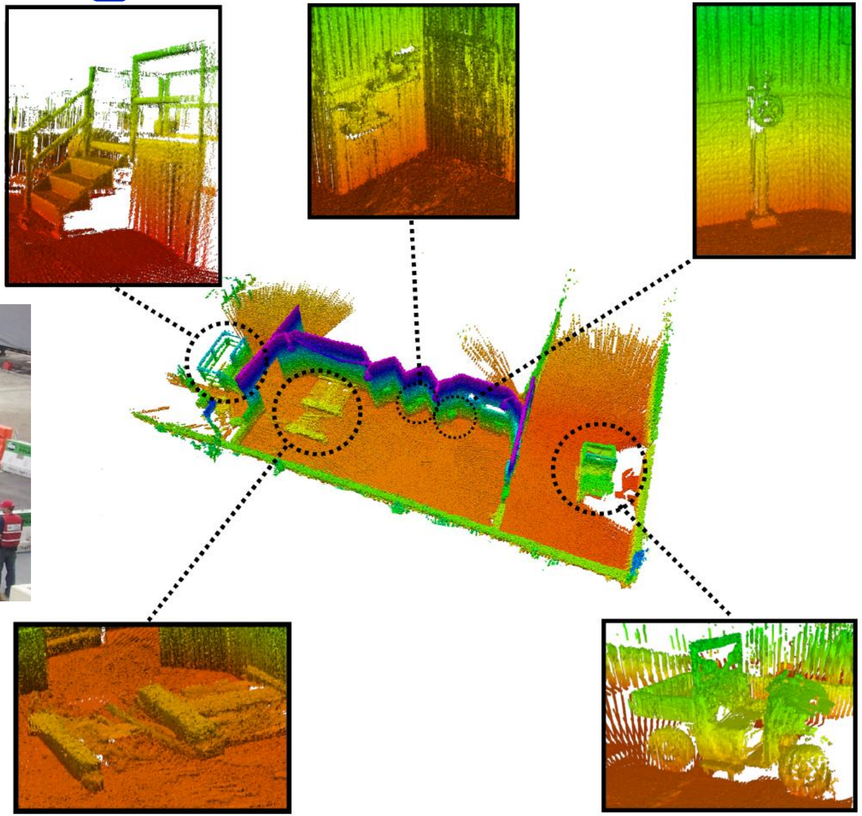

[Droeschel et al., Robotics and Autonomous Systems 2017]

Valve Turning Interface

- Align wheel model with 3D points using interactive marker

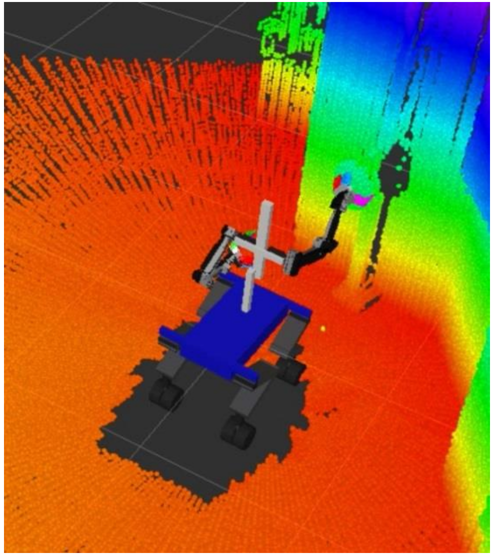

[Schwarz et al. Journal of Field Robotics 2017]

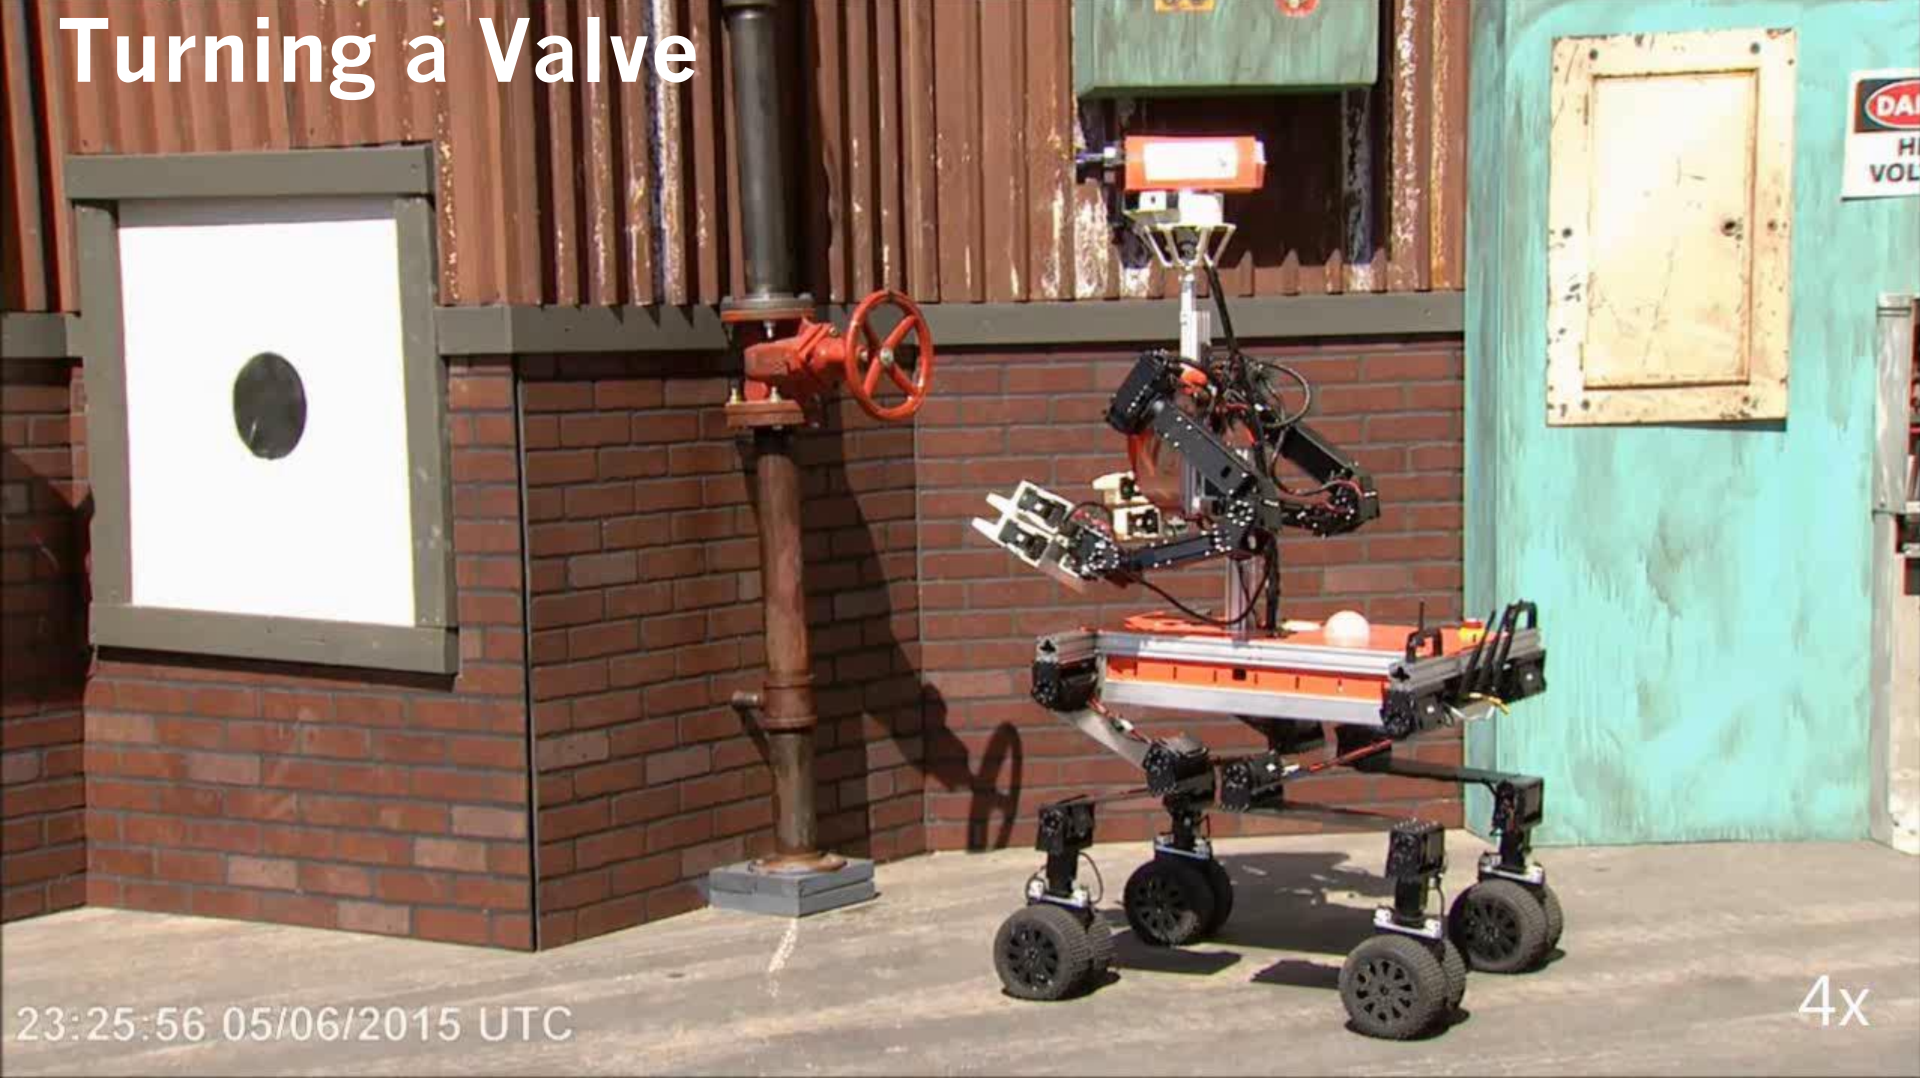
Turning a Valve
23:25:56 05/06/2015 UTC
4x

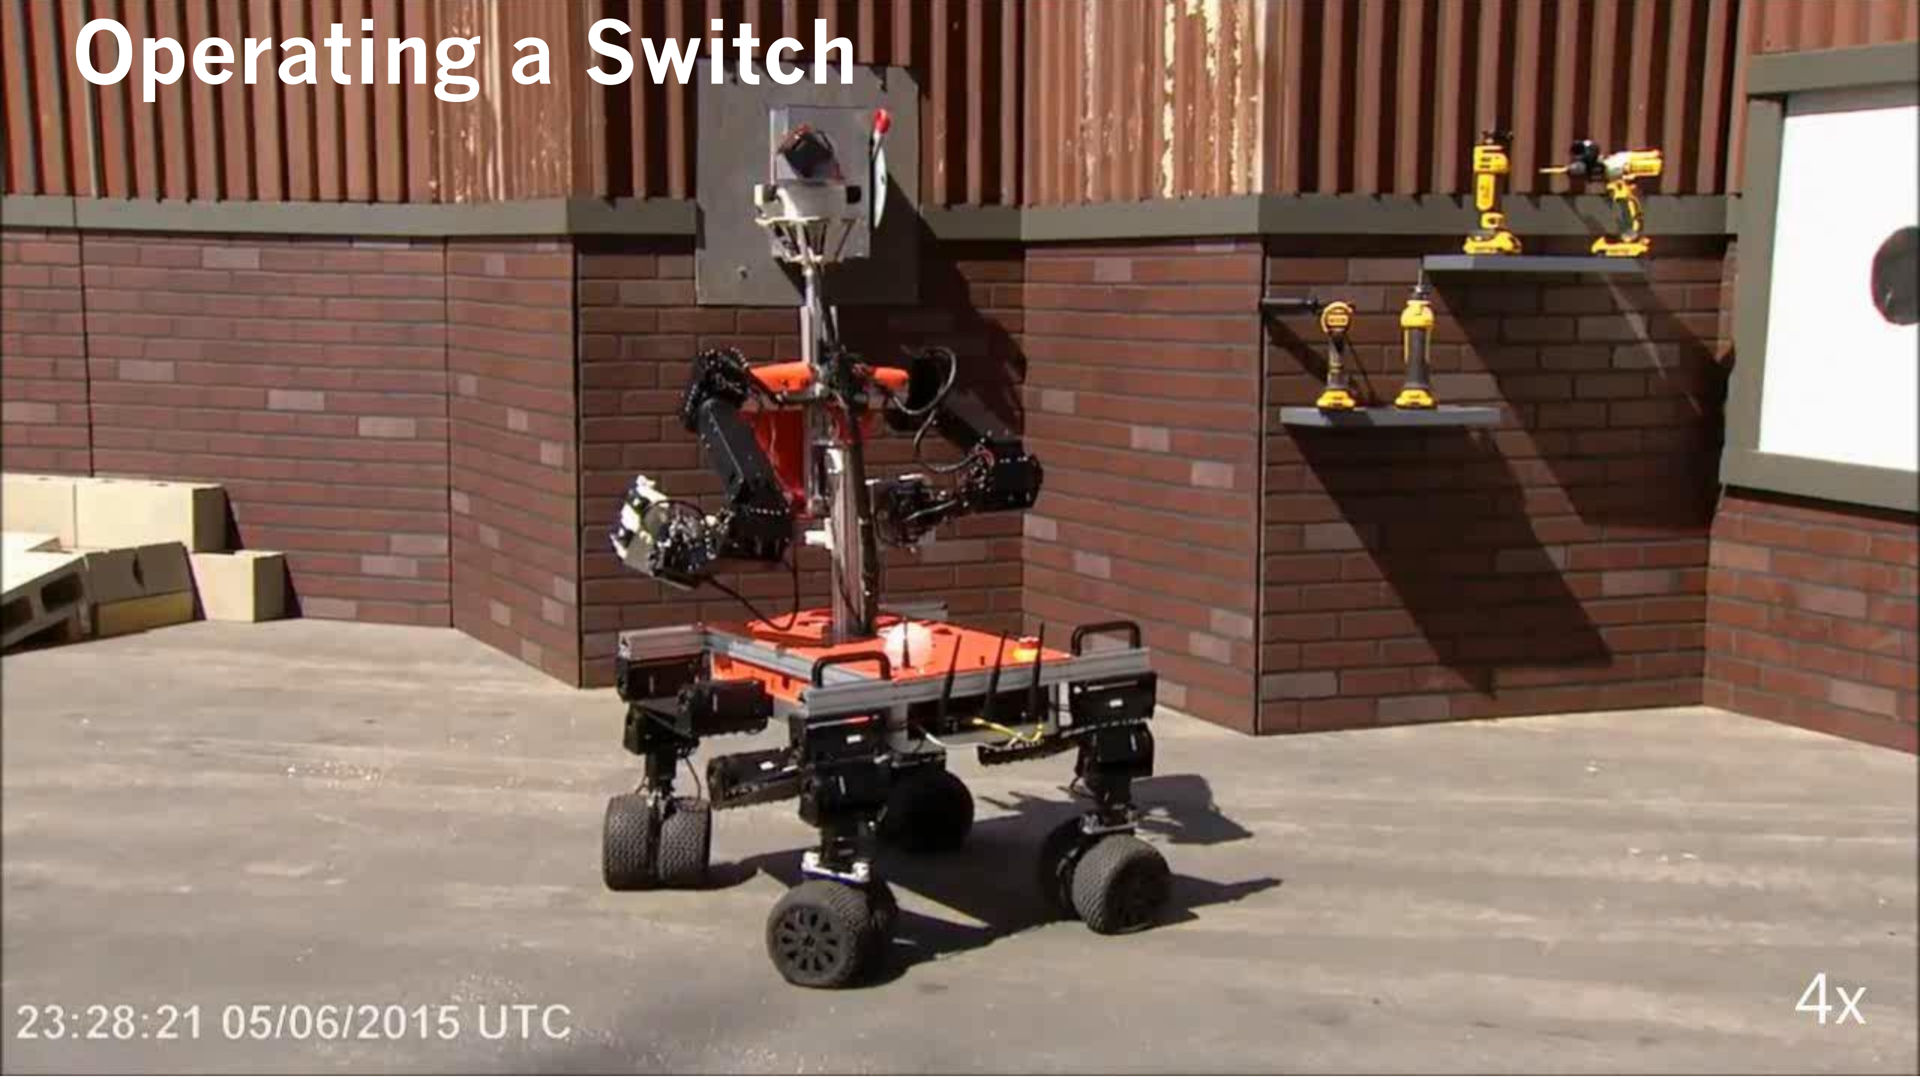
Operating a Switch
23:28:21 05/06/2015 UTC
4x


Plug Task
02:23:20 07/06/2015 UTC
4X

Debris Tasks

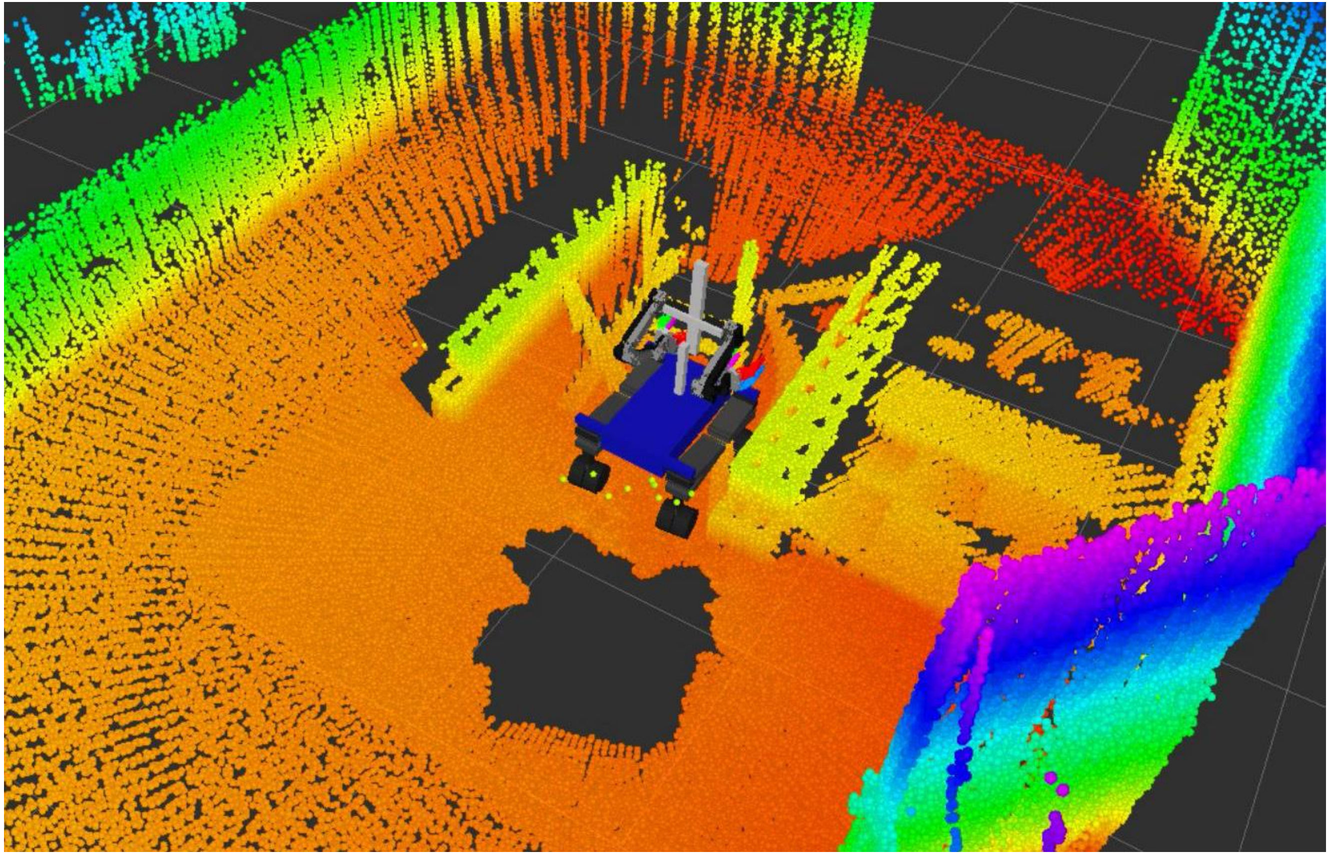

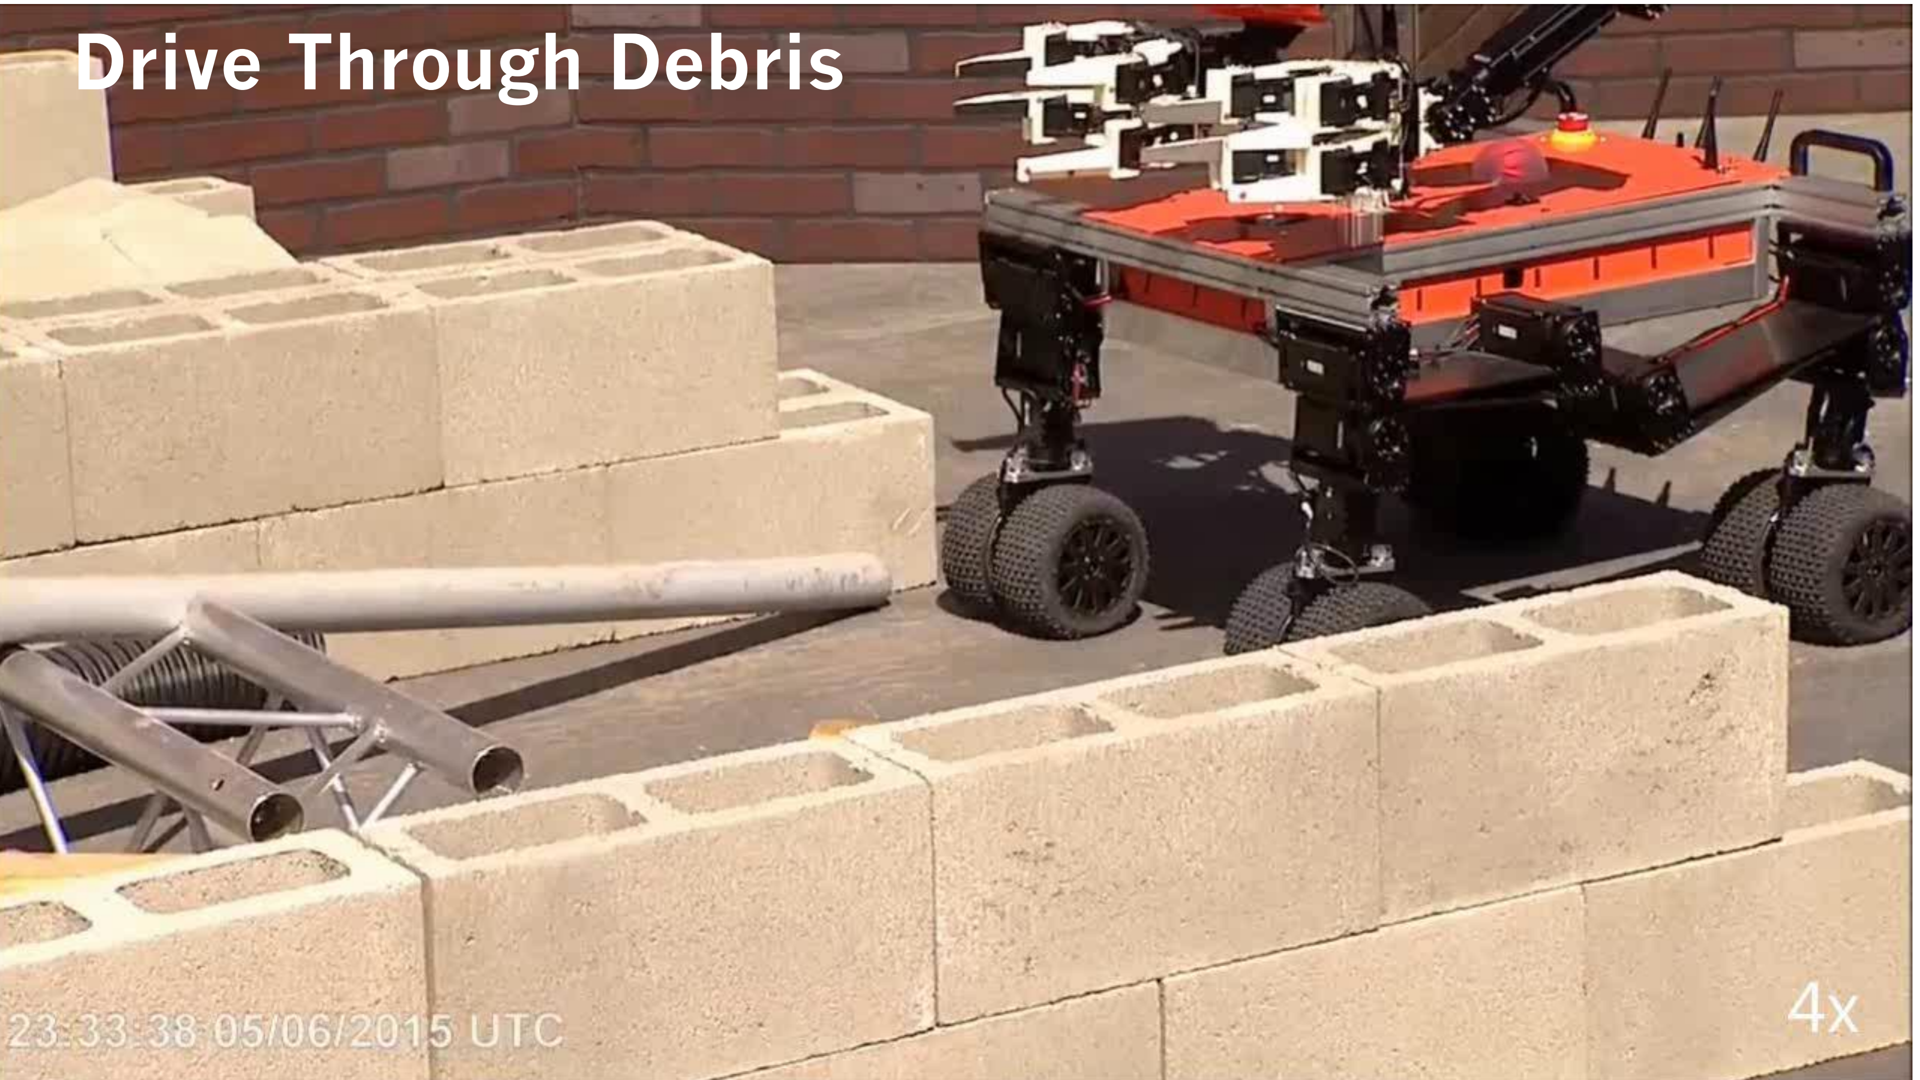
Drive Through Debris
23:33:38 05/06/2015 UTC
4x

Cutting Drywall

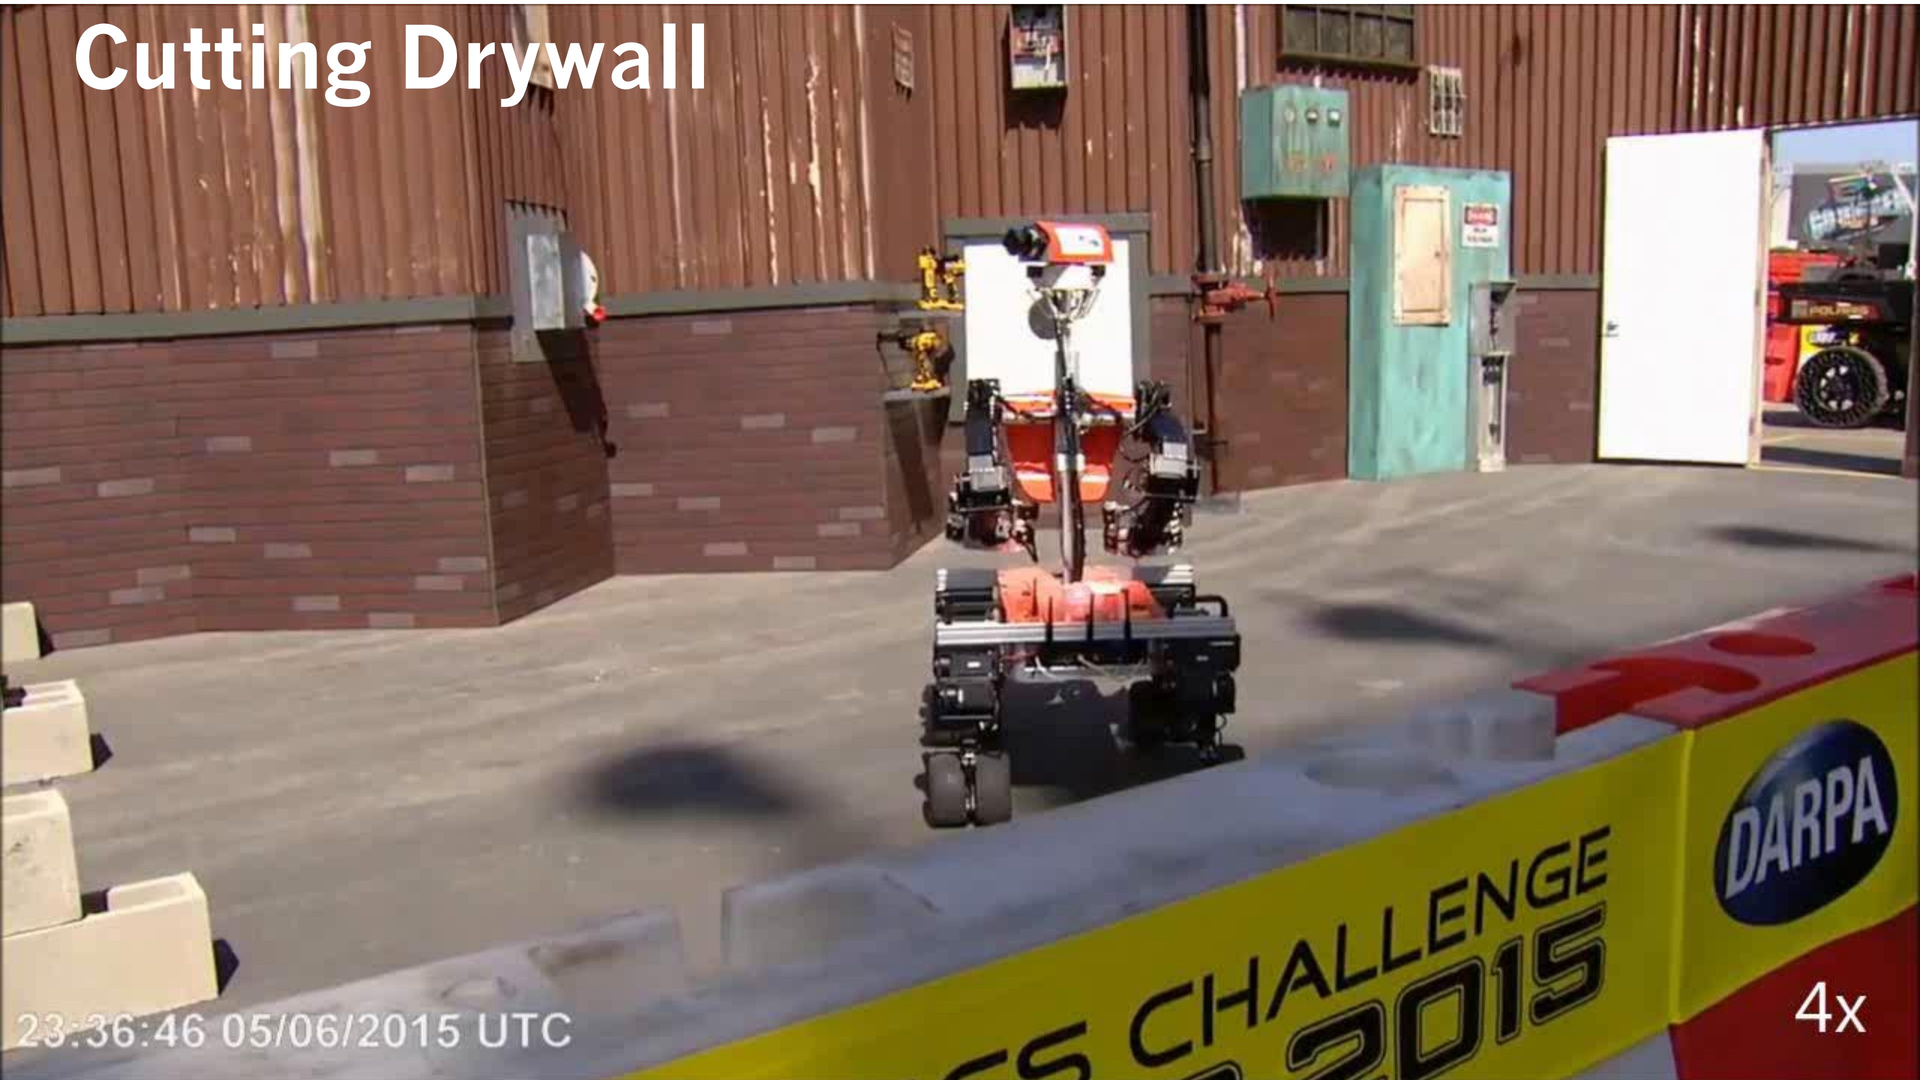

Team NimbRo Rescue

Best European Team (4th place overall), solved seven of eight tasks in 34 minutes

Stair Climbing

- Determine leg that most urgently needs to step
- Weight shift: sagittal, lateral, driving changes support
- Step to first possible foot hold after height change

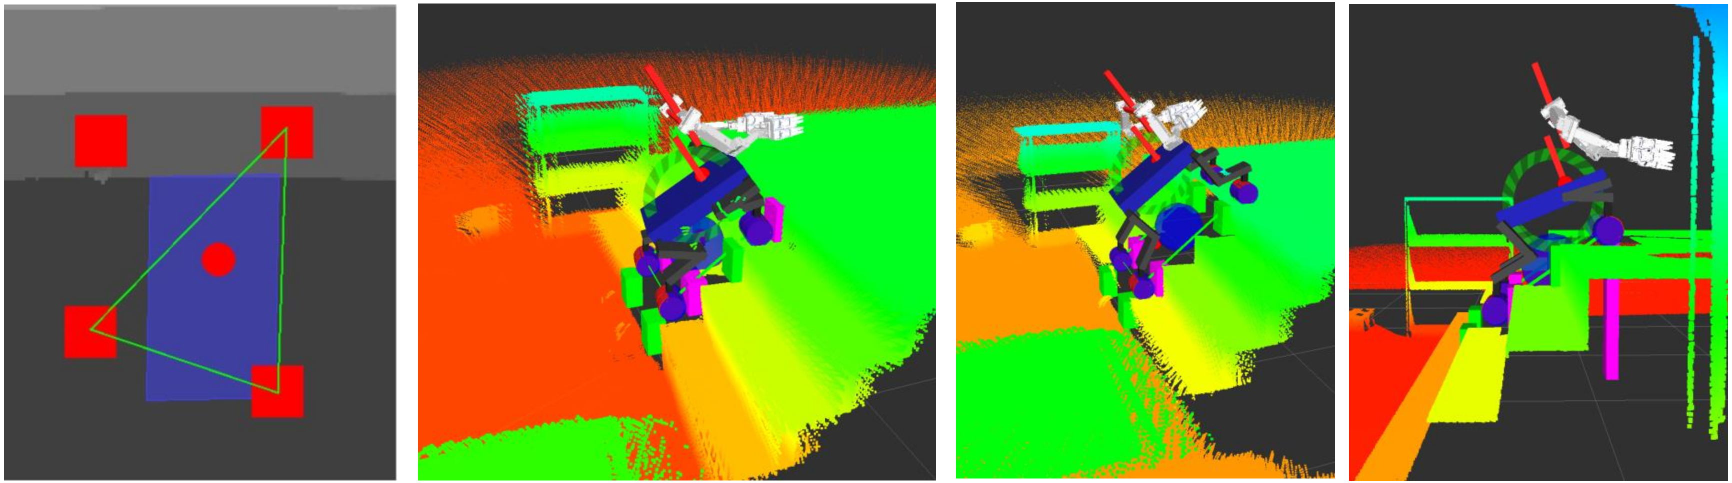

[Schwarz et al., ICRA 2016]


Stair Crawling

Hose Connecting Task

- Bimanual task
  - Grab the left hose with the left gripper,
  - Grab the right hose with the right gripper, and
  - Connect both hoses
- 10/11 trials successful
- Execution time

Task	Time [min:s]				
	Avg.	Median	Min.	Max.	Std. Dev.
Left grasp	0:44	0:38	0:27	1:20	0:16
Right grasp	0:45	0:40	0:34	1:04	0:10
Connect	1:36	1:32	1:07	2:04	0:21
Total	3:04	2:57	2:21	3:51	0:28

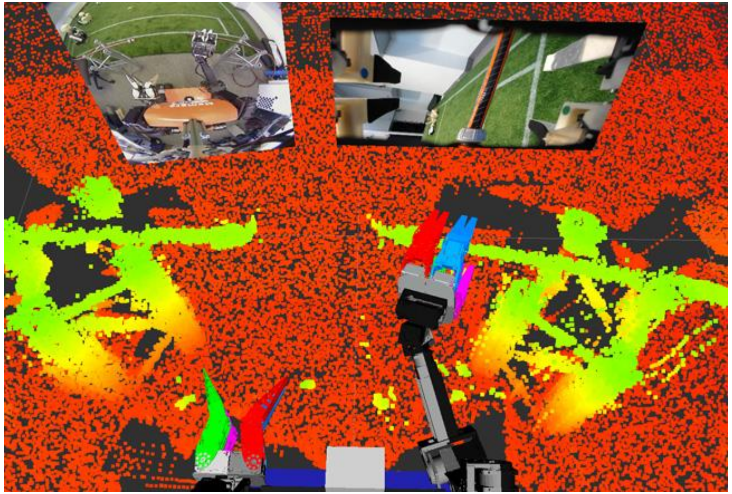

[Rodehutskors et al., Humanoids 2015]

DLR SpaceBot Cup 2015

- Mobile manipulation in rough terrain

[Schwarz et al., Frontiers on Robotics and AI 2016]

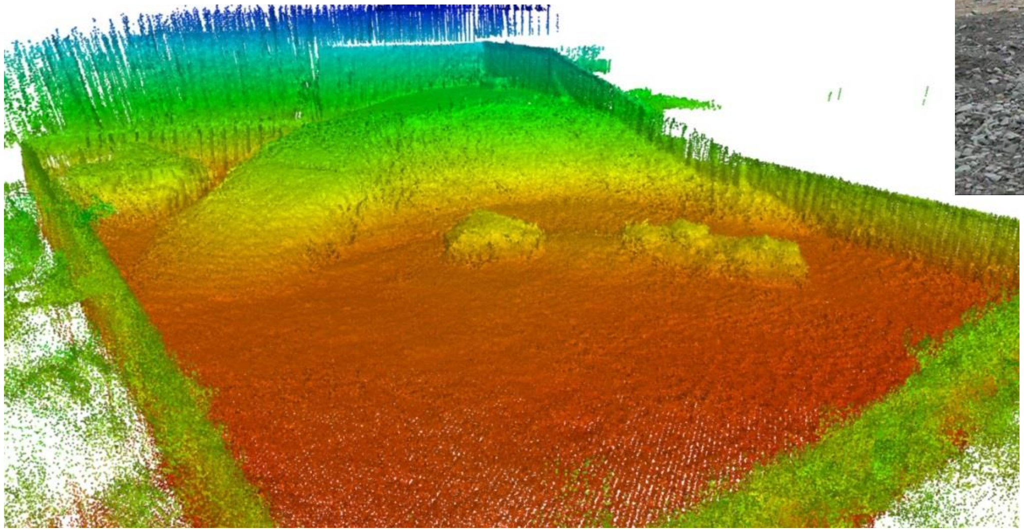

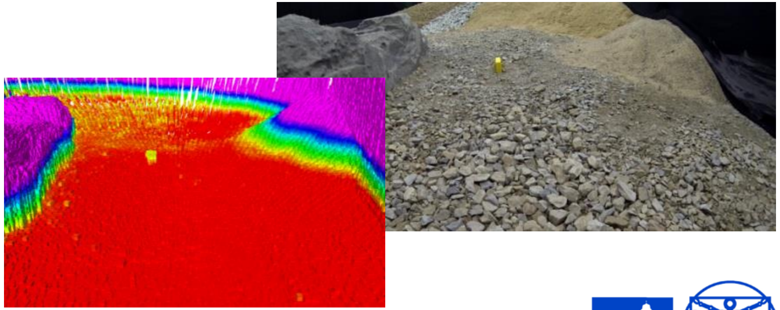

universitätbonn ais


DLR SpaceBot Camp 2015
8X

Autonomous Mission Execution

- 3D mapping, localization, mission and navigation planning

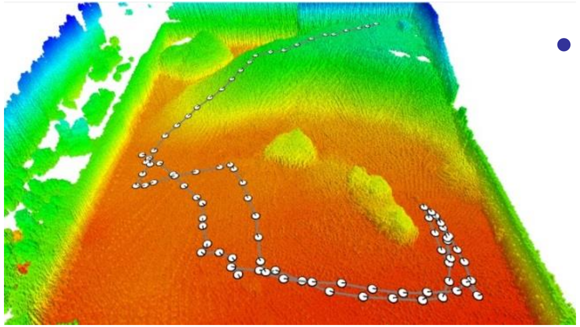

- 3D object perception and grasping

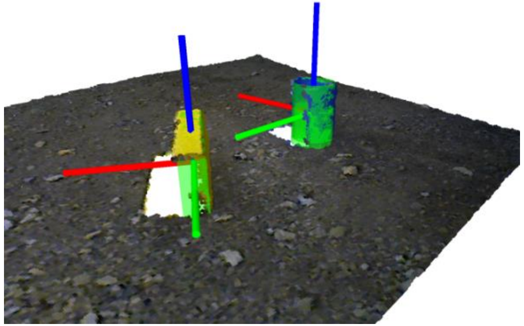

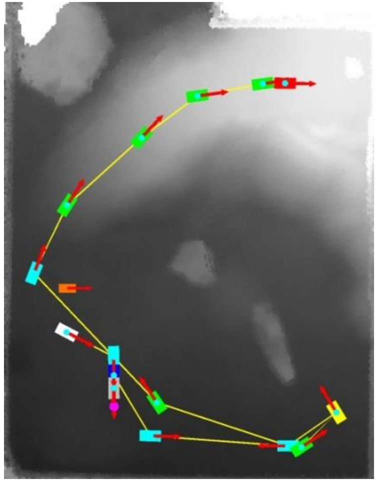

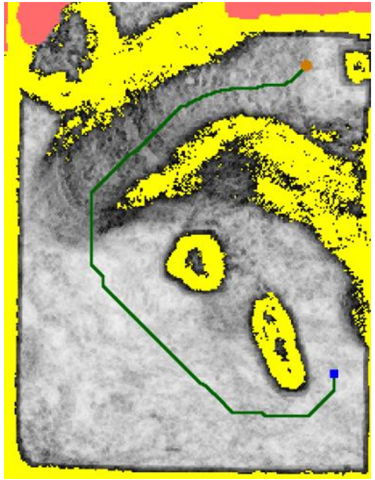

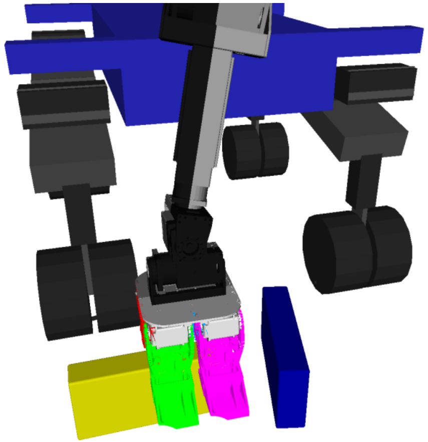

[Schwarz et al. Frontiers 2016]

Navigation Planning

- Costs from local height differences
- A* path planning

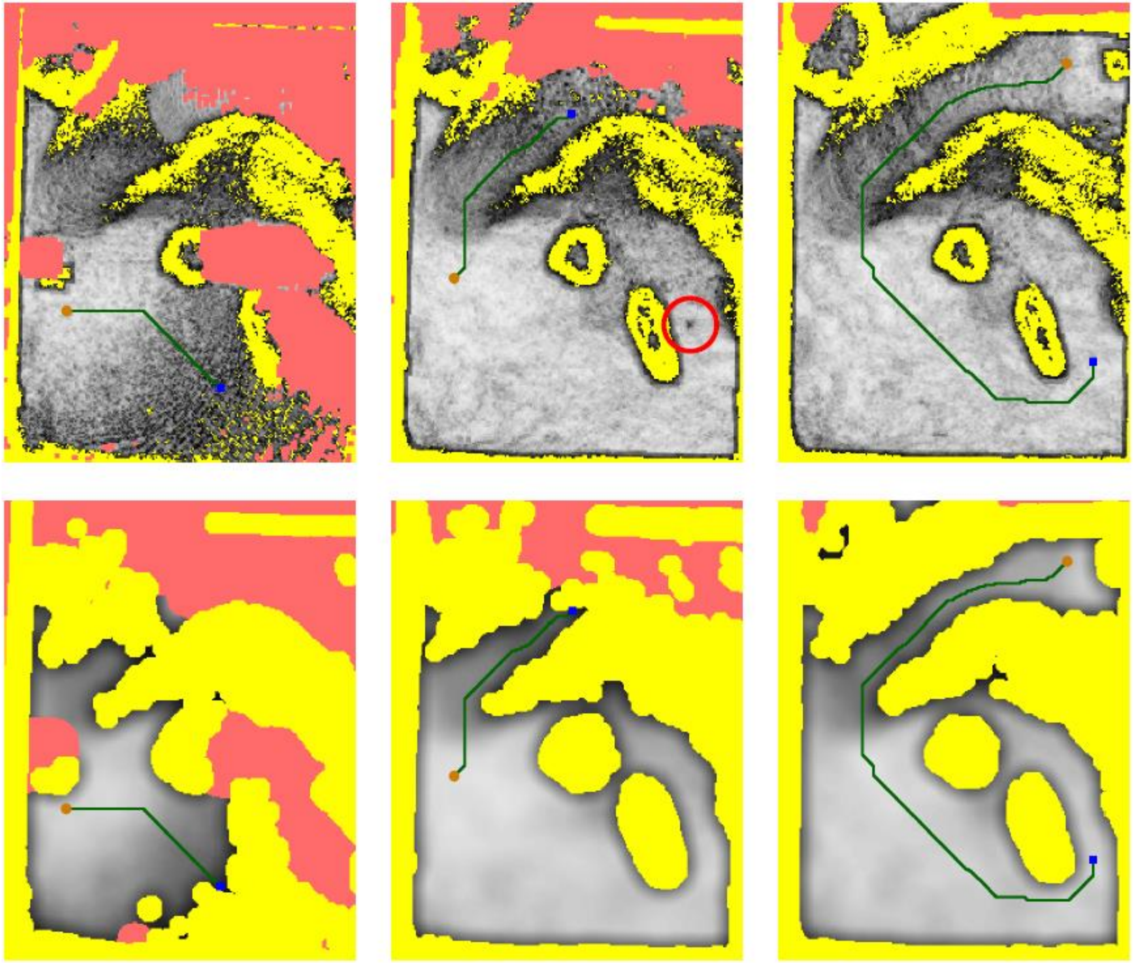

[Schwarz et al., Frontiers in Robotics and AI 2016]

3D Map

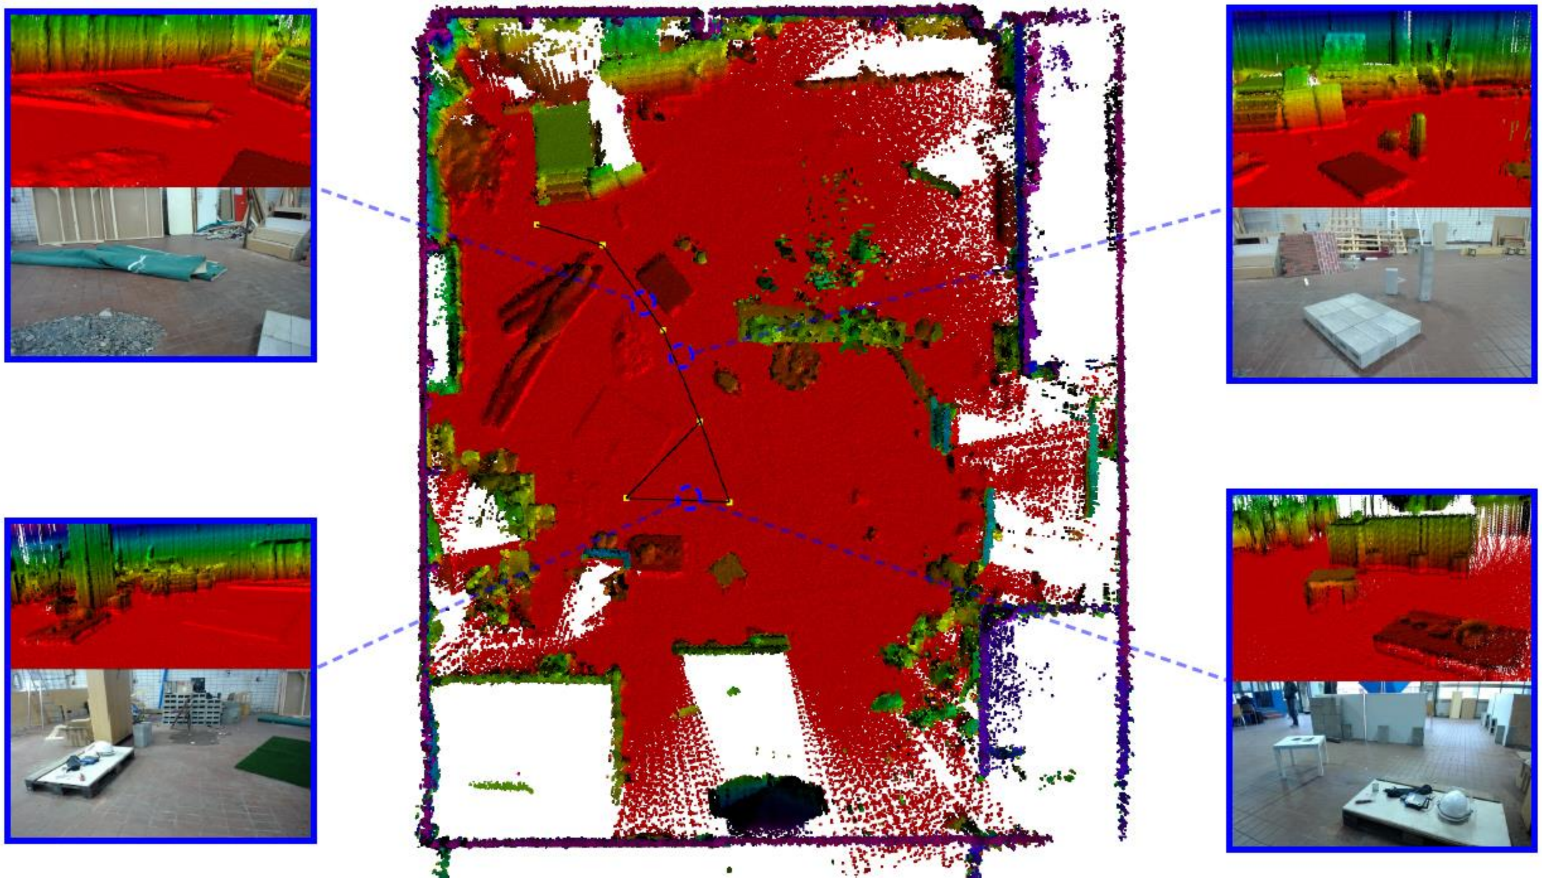

New Sensor Head

- Continuously rotating Velodyne Puck VLP-16
  - 300,000 3D points/s
  - 100 m range
  - Spherical field of view
- Three wide-angle color cameras (total FoV 210×103°)
- Kinect V2 RGB-D camera on pan-tilt unit

universitätbonn ais

3D Map of Indoor+Outdoor Scene

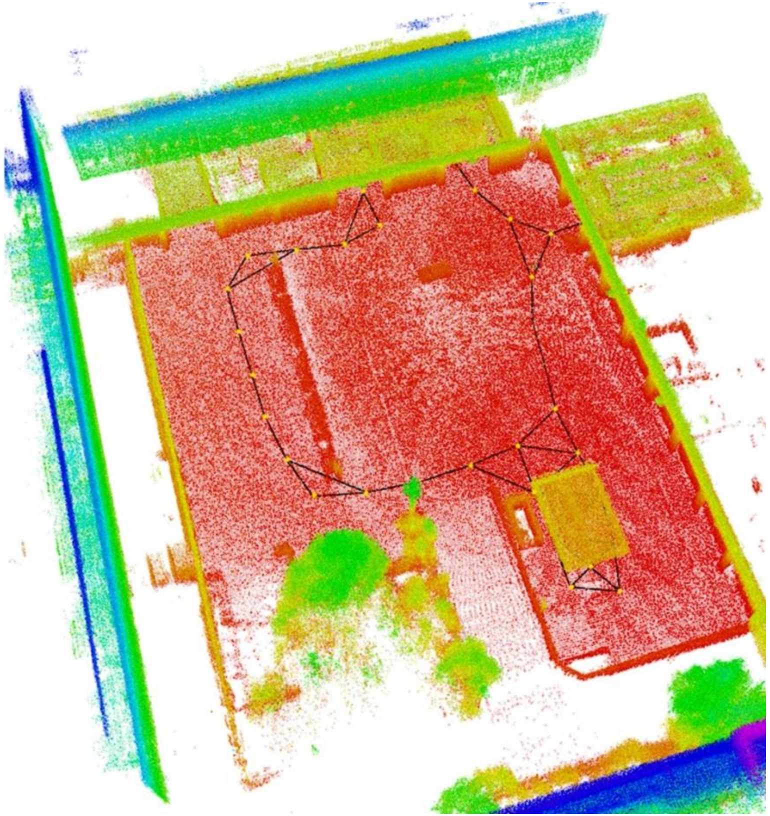

[Droeschel et al., Robotics and Autonomous Systems 2017]

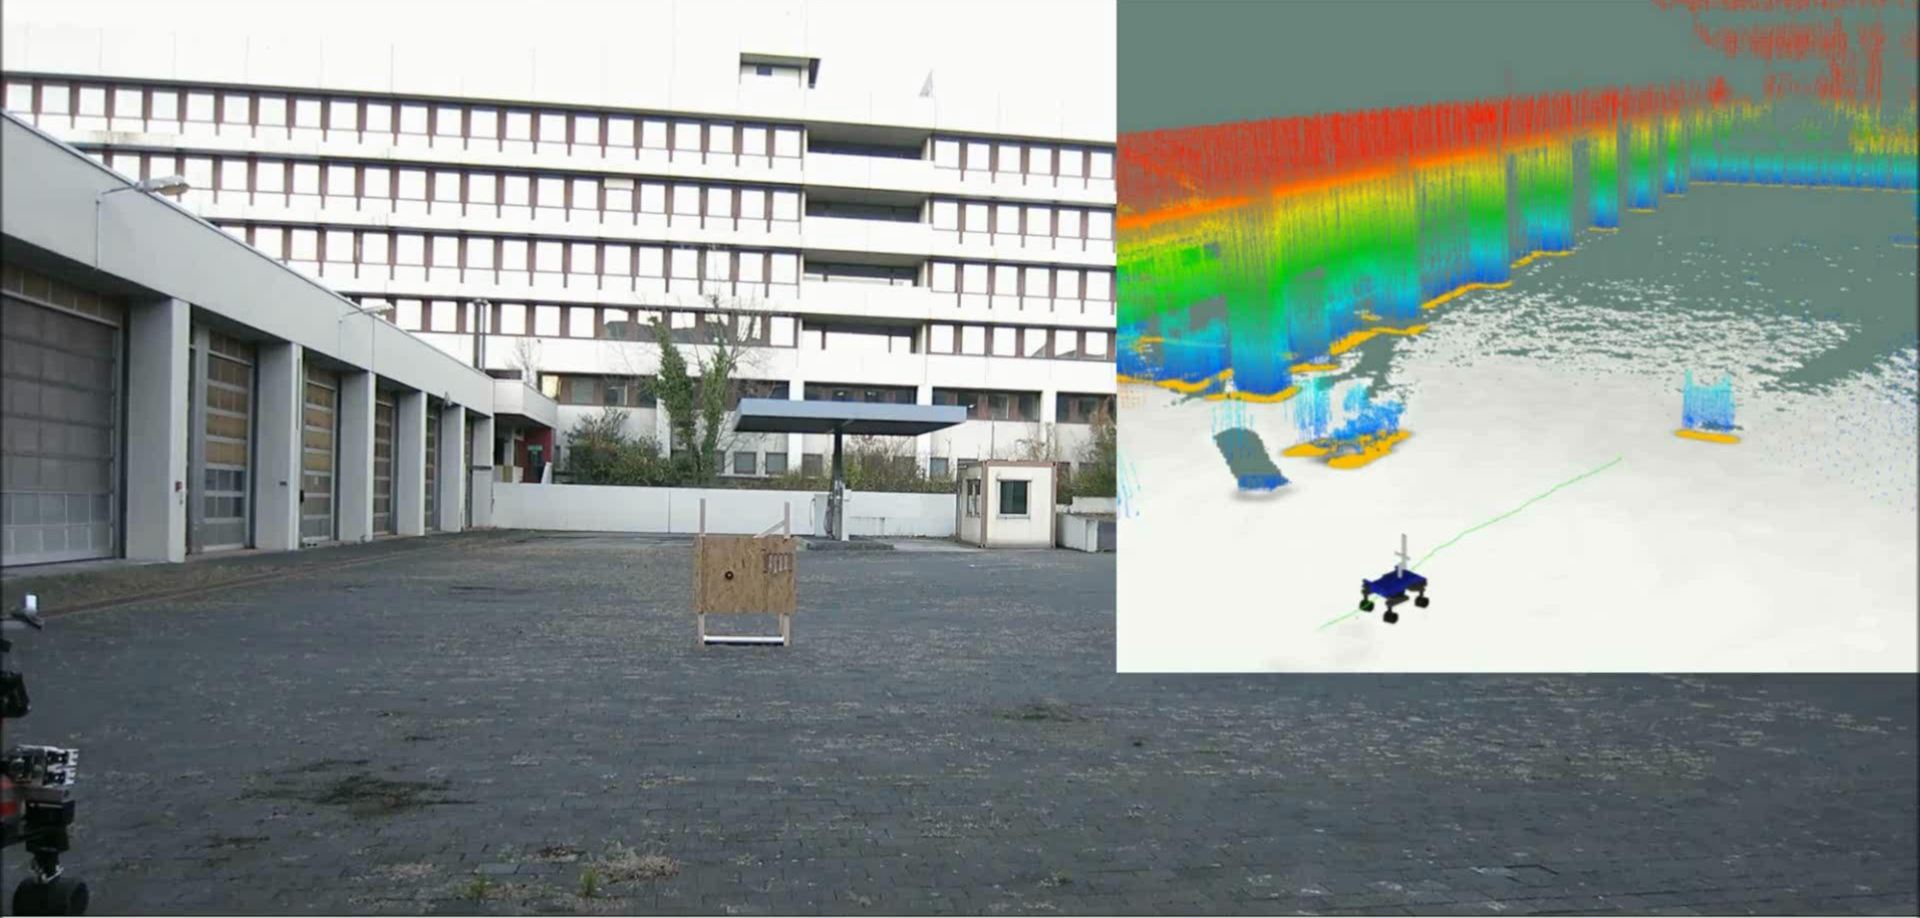

Navigation in allocentric laser map (colored points)

Using a Wrench to Turn a Valve

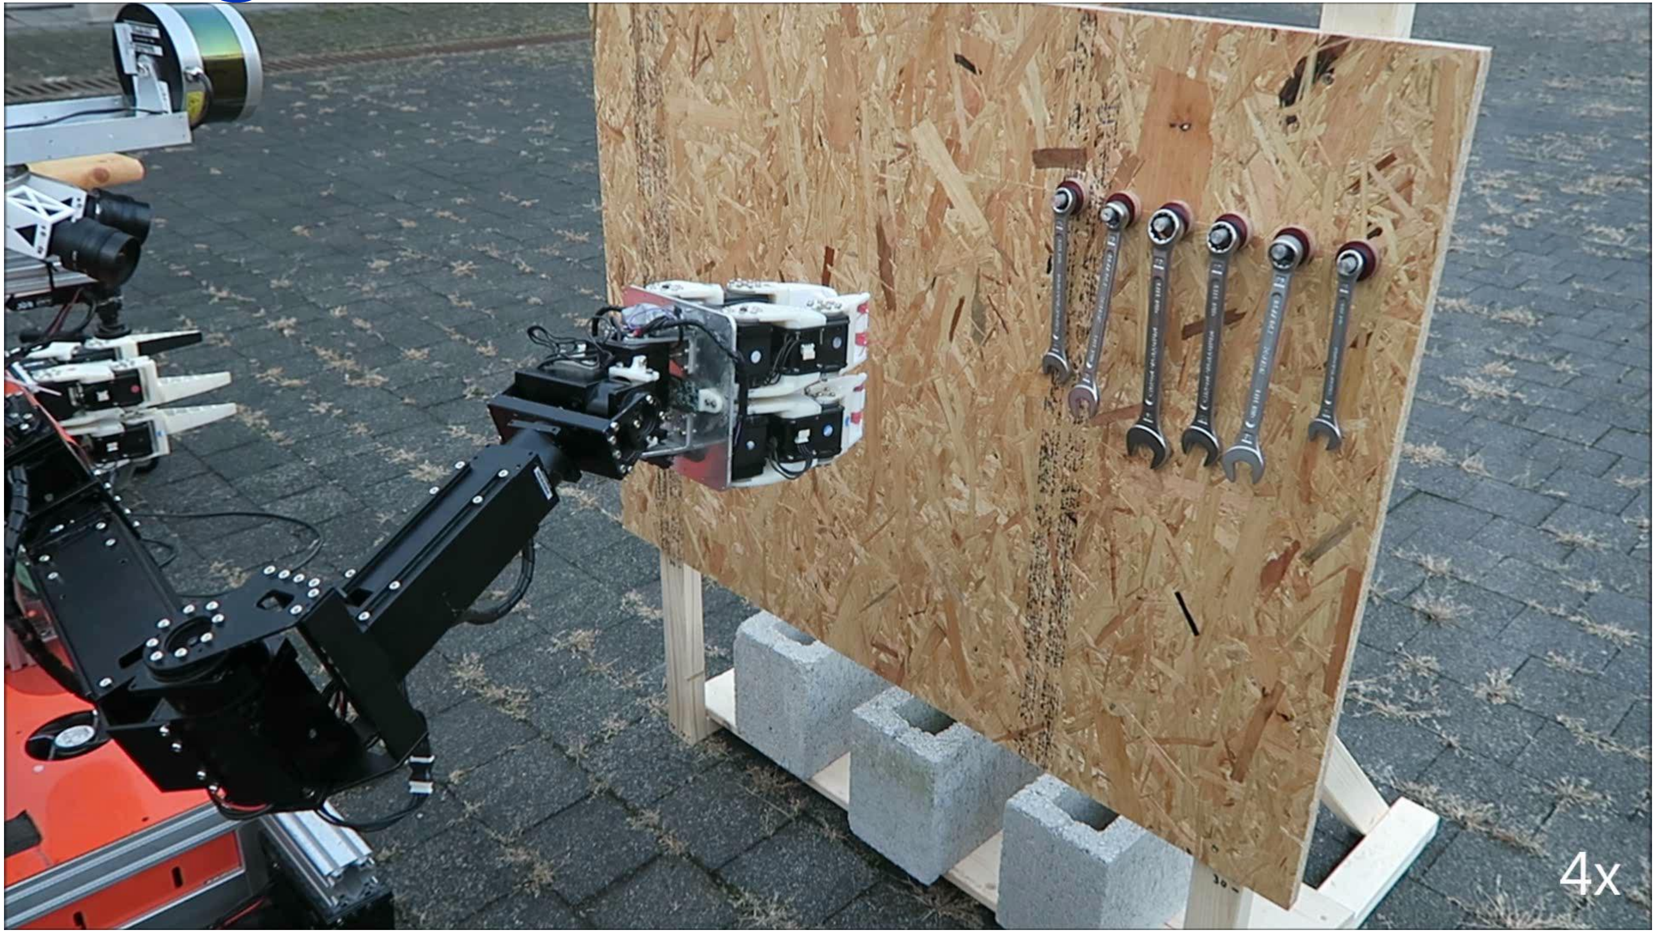

MBZIRC Challenge 2

Mario Robot Manipulator

- 6DoF arm (UR5)
- Stereo cameras (Pointgray)
- ToF camera (PMD picoflexx)
- Two-finger gripper

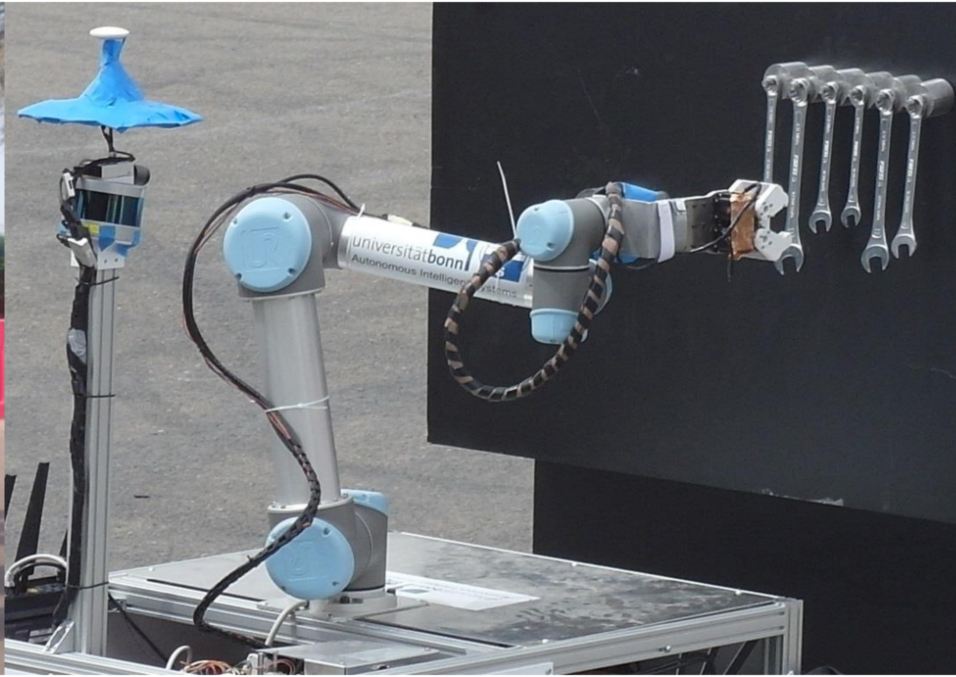

Wrench Selection: Detection of Tool Ends

Valve Stem Registration

- Picoflexx depth
- Euclidean clustering
- Rotating calipers for estimating valve stem angle and size

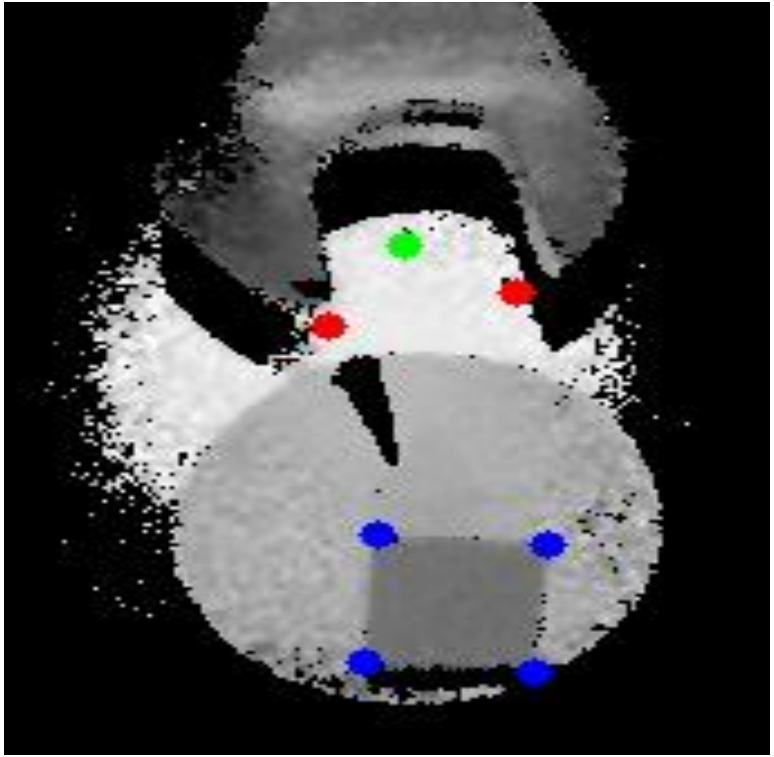

MBZIRC Challenge 1

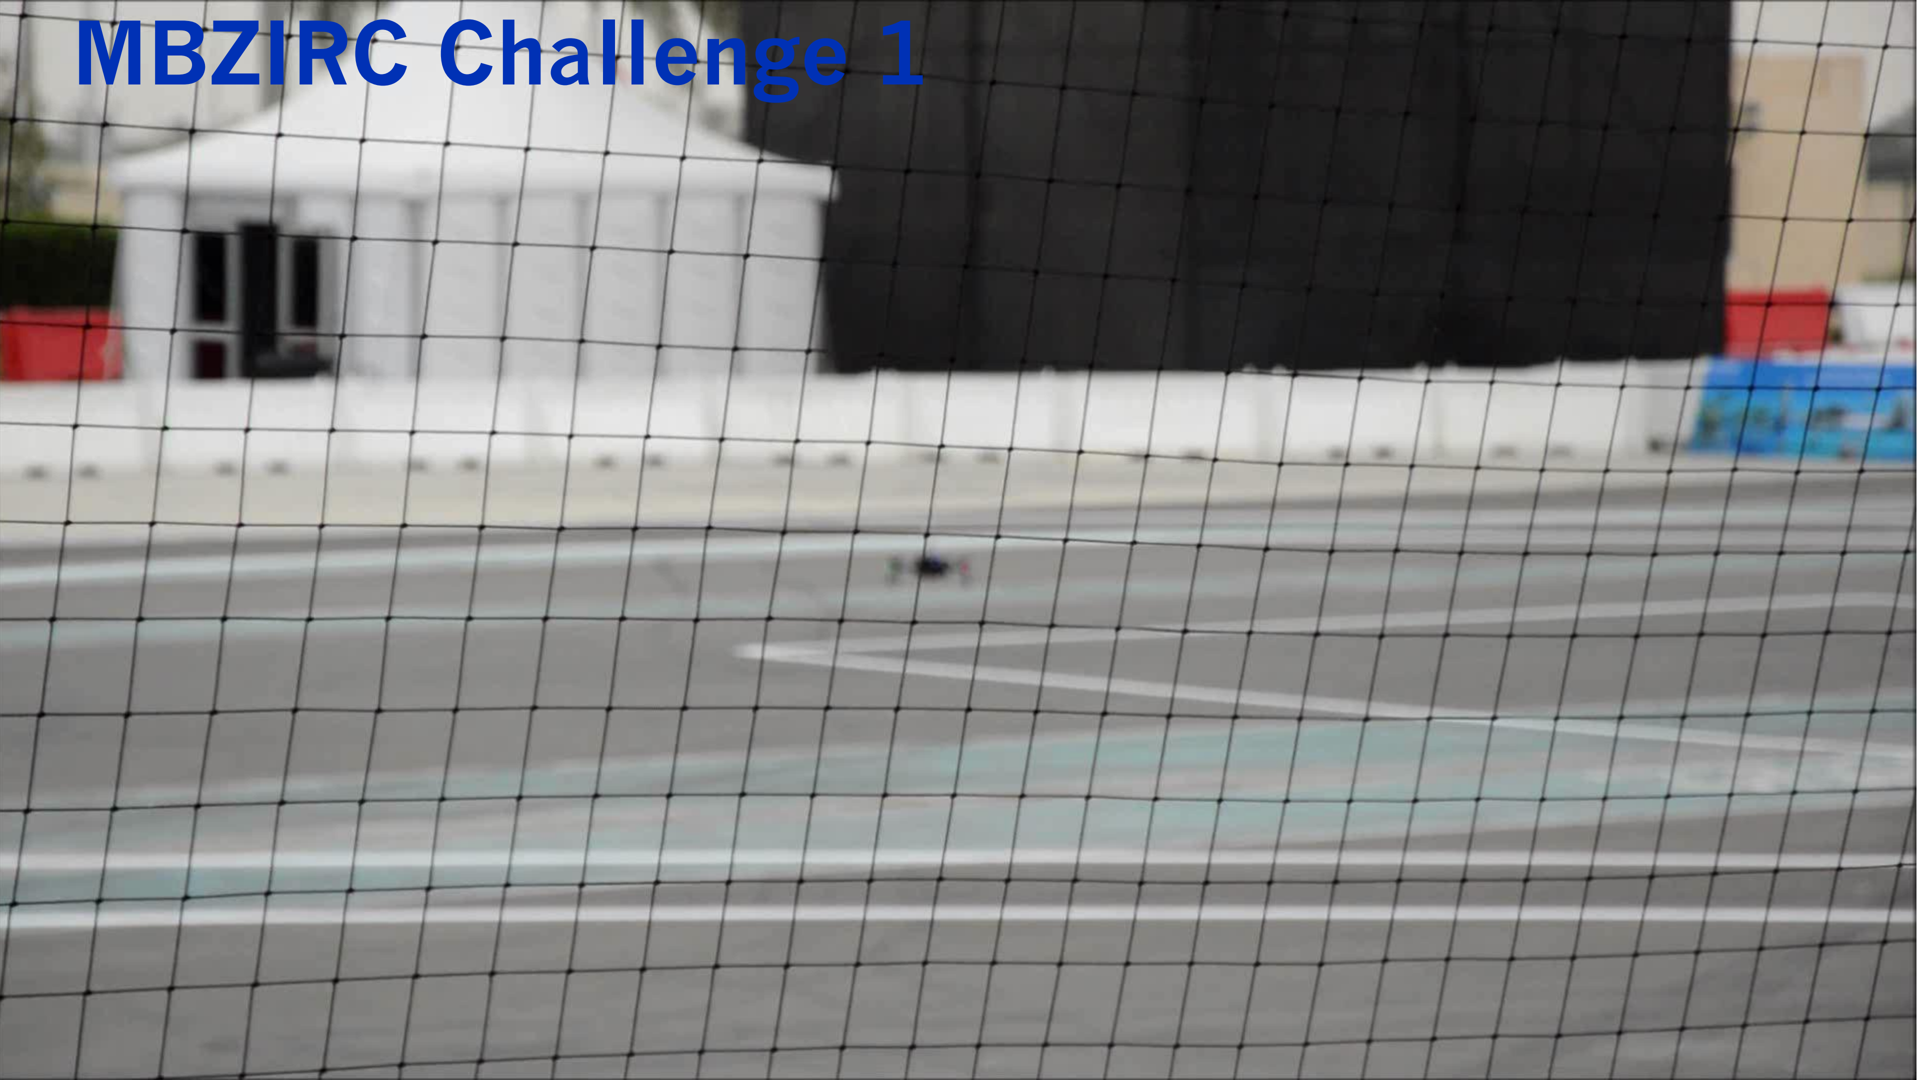

Landing Pattern Detection

- Line-based detection in two cameras

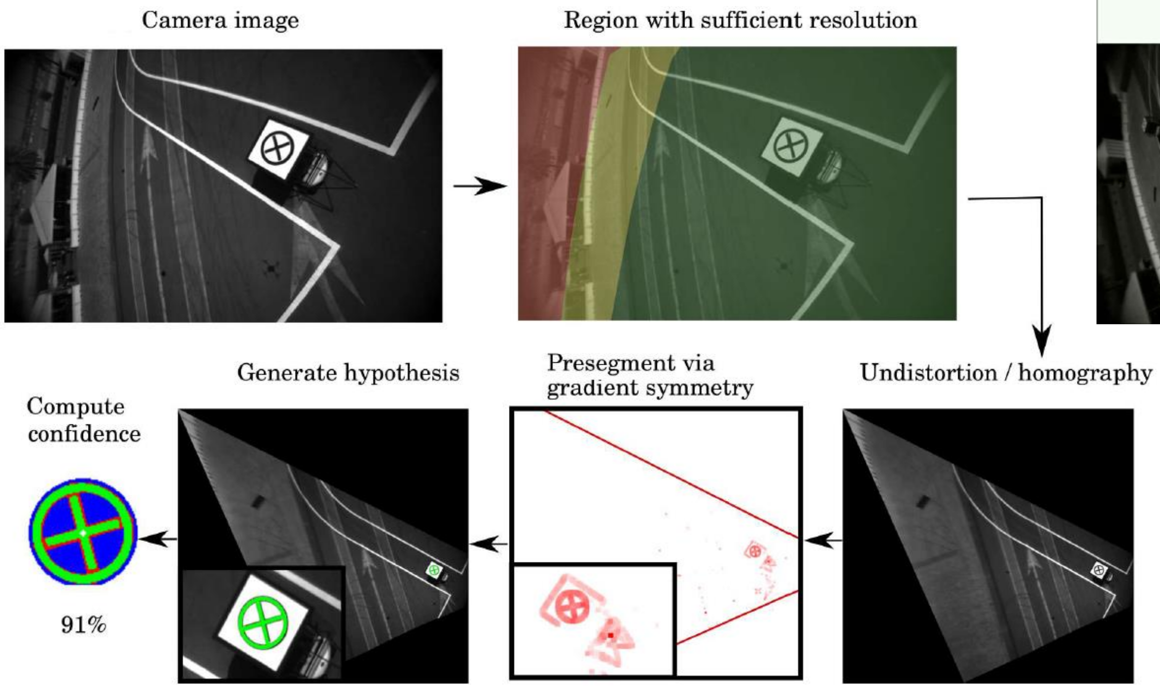

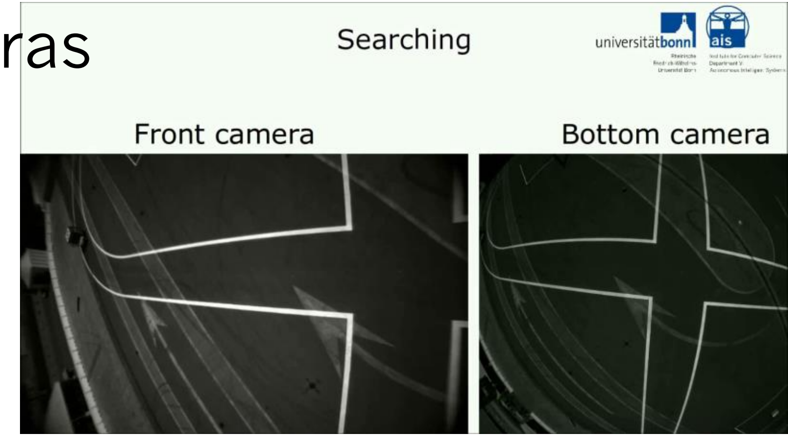

[Beul et al. ECMR 2017]

Time-optimal Interception Trajectory

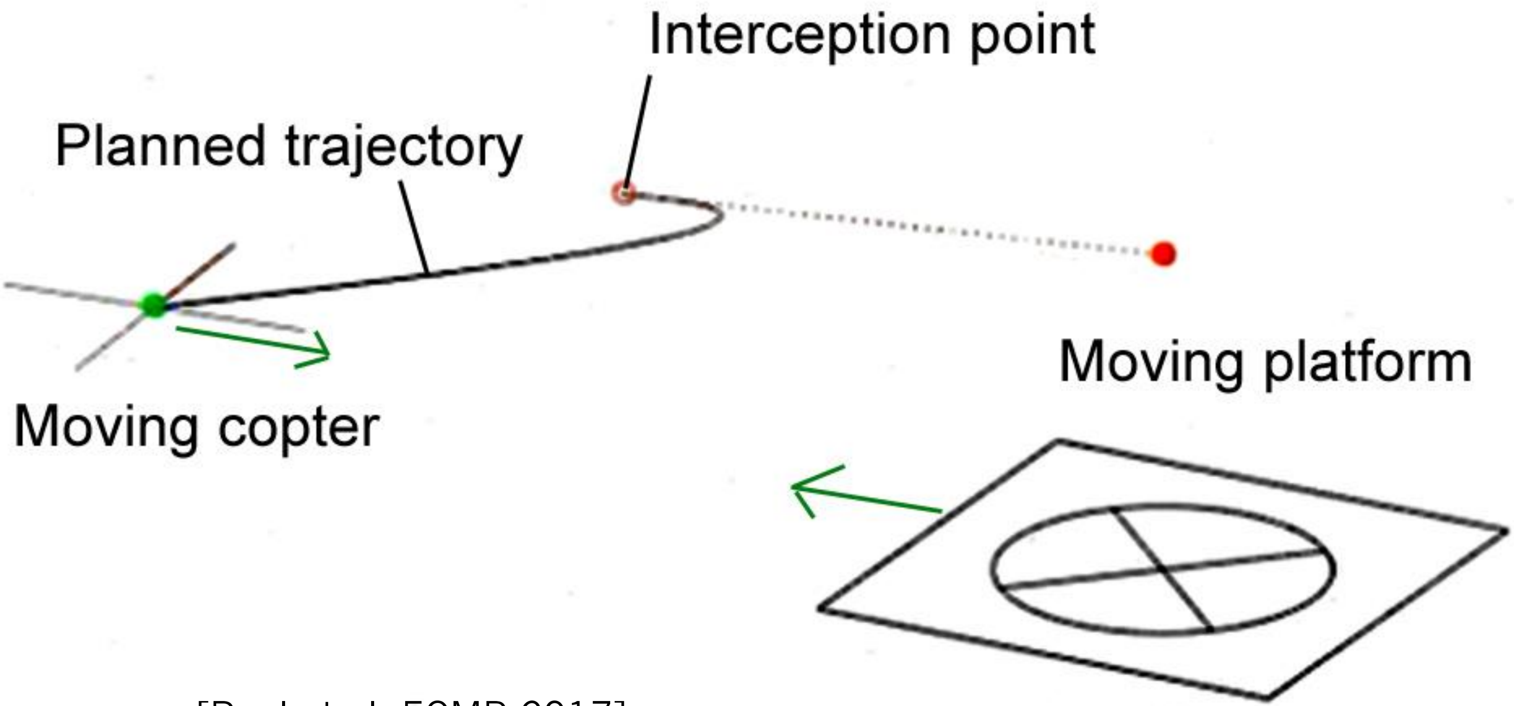

[Beul et al. ECMR 2017]

Picking Copter DJI Matrice 100

- Wide-angle downward looking color camera
- Electromagnetic gripper
- Laser-distance sensor to ground
- Dual-core PC

[Nieuwenhuisen et al. ECMR 2017]

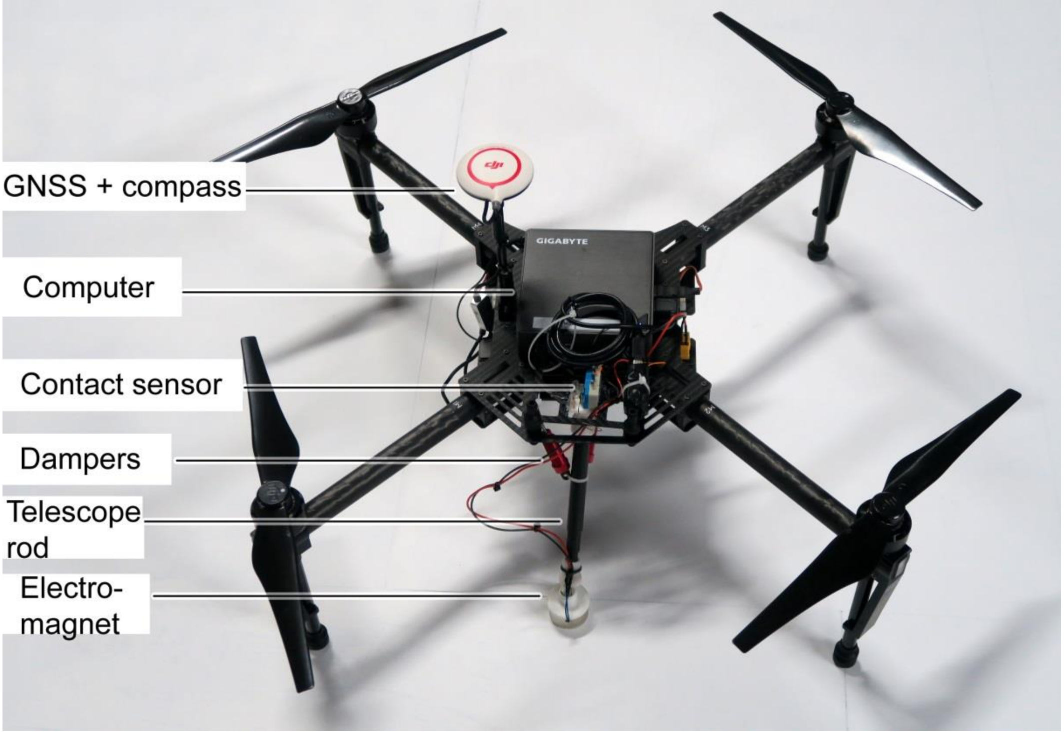

MBZIRC Challenge 3

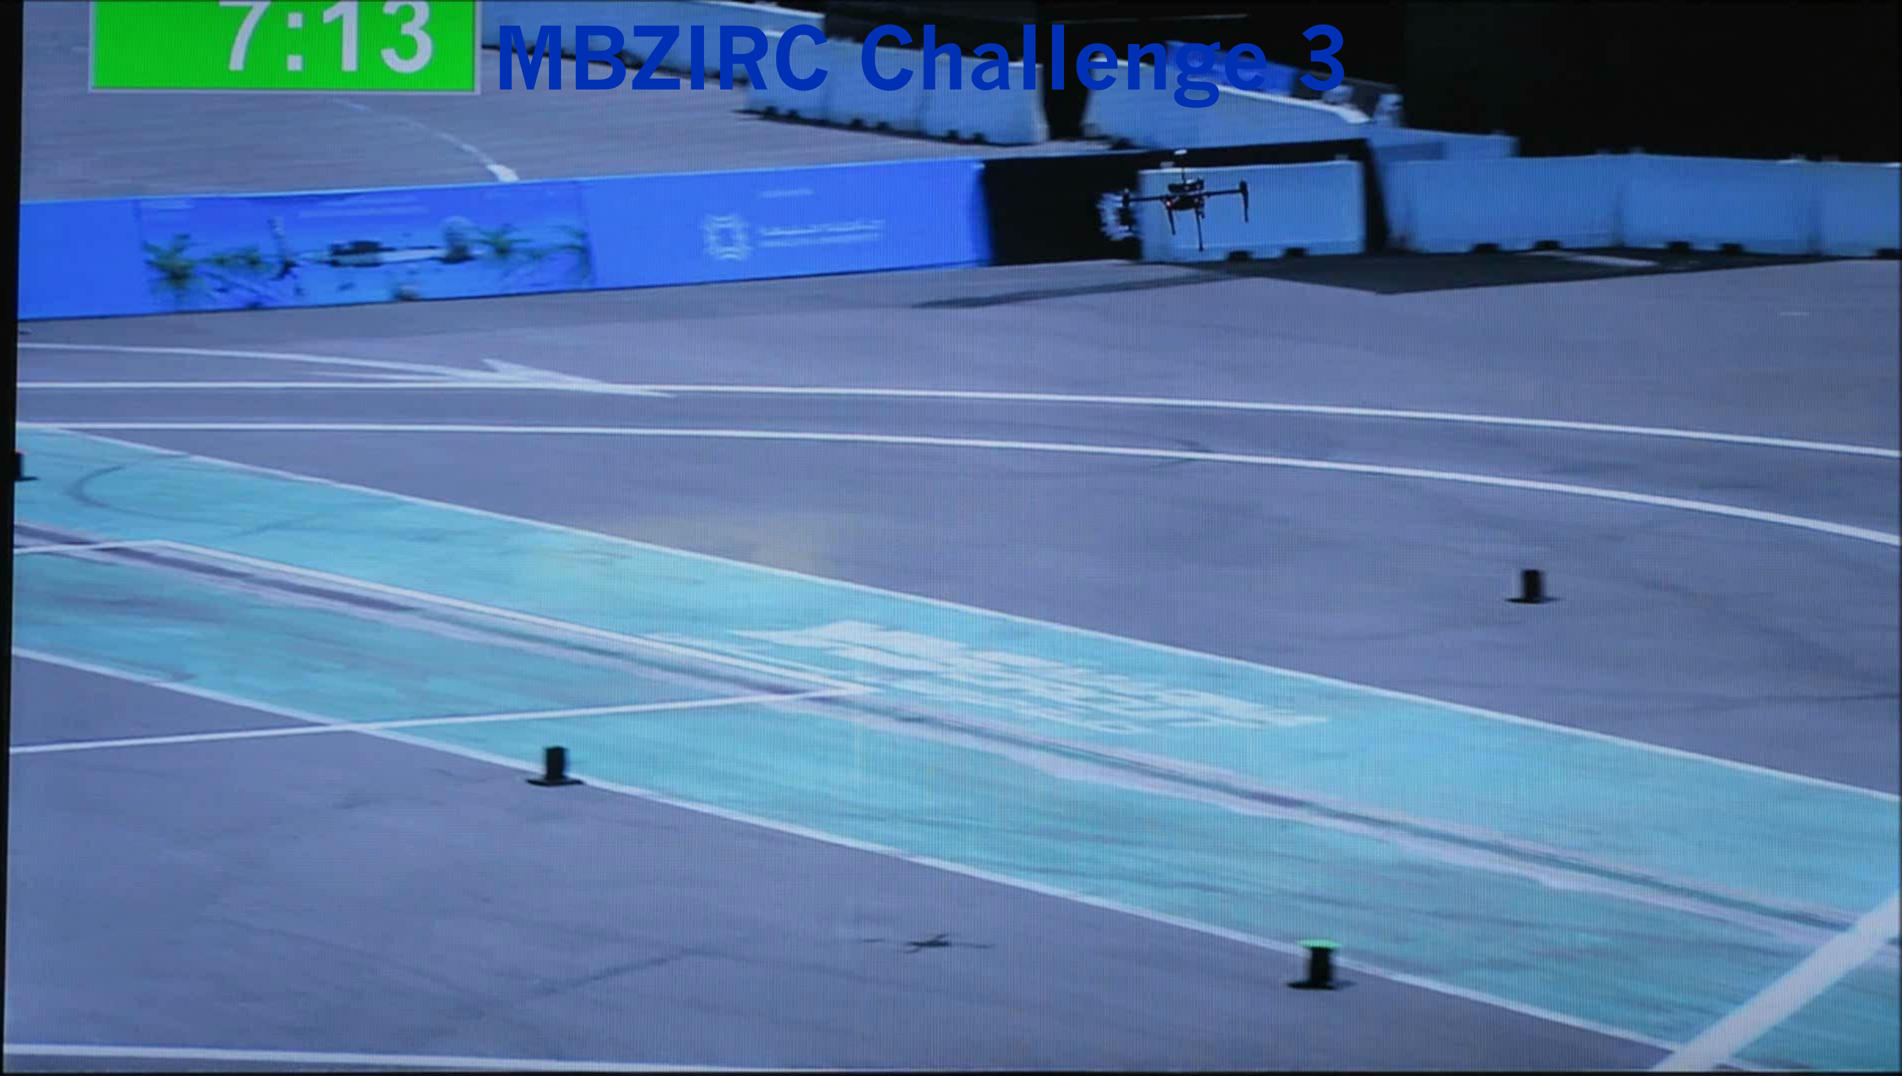

Pickable Object and Drop-box Detection

- Probabilistic color segmentation
- RANSAC-like drop-box detection

[Nieuwenhuisen et al. ECMR 2017]

Drop box

Color segmentation

Raw image

MBZIRC Team NimbRo

H2020 Project CENTAURO

Robust Mobility and Dexterous Manipulation in Disaster Response by Fullbody Telepresence in a Centaur-like Robot

universität bonn · iit ISTITUTO ITALIANO DI TECNOLOGIA · Scuola Superiore Sant'Anna di Studi Universitari e di Perfezionamento · KTH VETENSKAP OCH KONST ROYAL INSTITUTE OF TECHNOLOGY · LINKÖPINGS UNIVERSITET · RWTH AACHEN UNIVERSITY · progenox · KHG

CENTAURO Objective

Development of a Human-robot system where a human operator is telepresent with its whole body in a Centaur-like robot, which is capable of robust locomotion and dexterous manipulation in the rough terrain and austere conditions characteristic of disasters

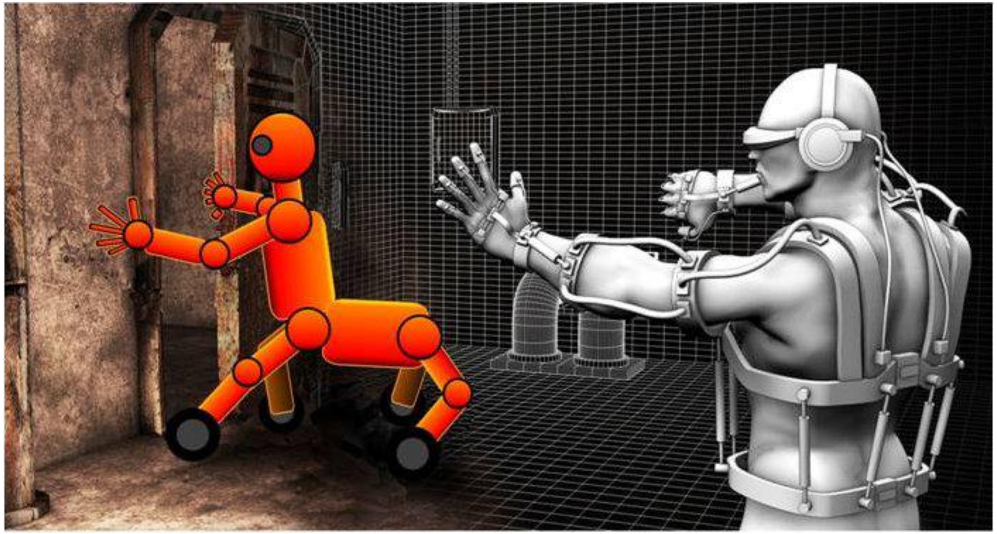

CENTAURO Approach

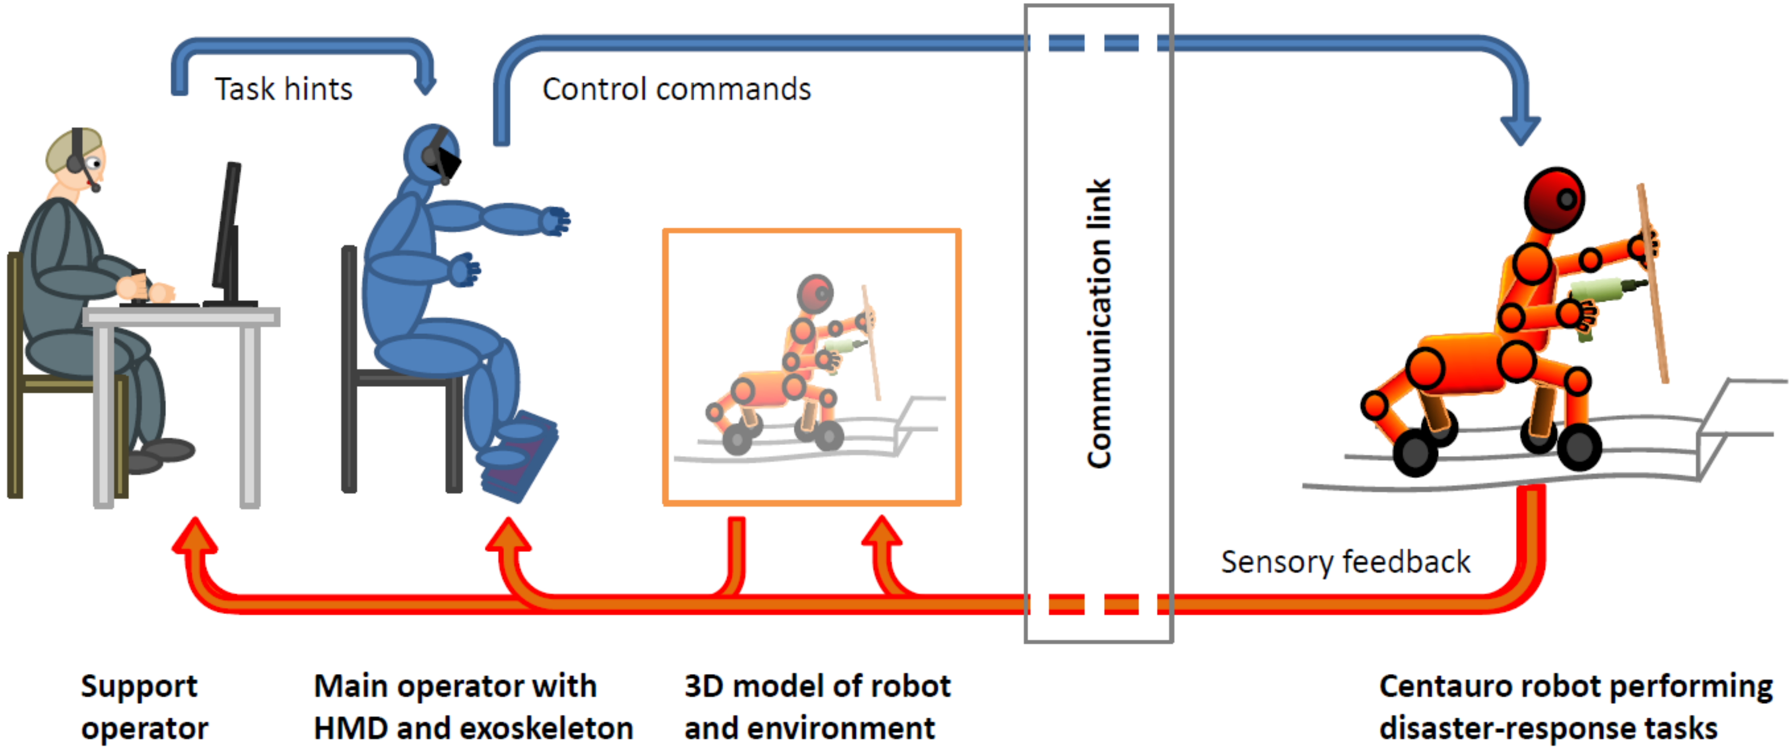

Centauro Robot Upper Body

- Strong

- Fast

[Giusti et al. ICRA 2017]

Centauro Robot Upper Body

- Serial-elastic actuators (SEA) => Compliant & adaptive

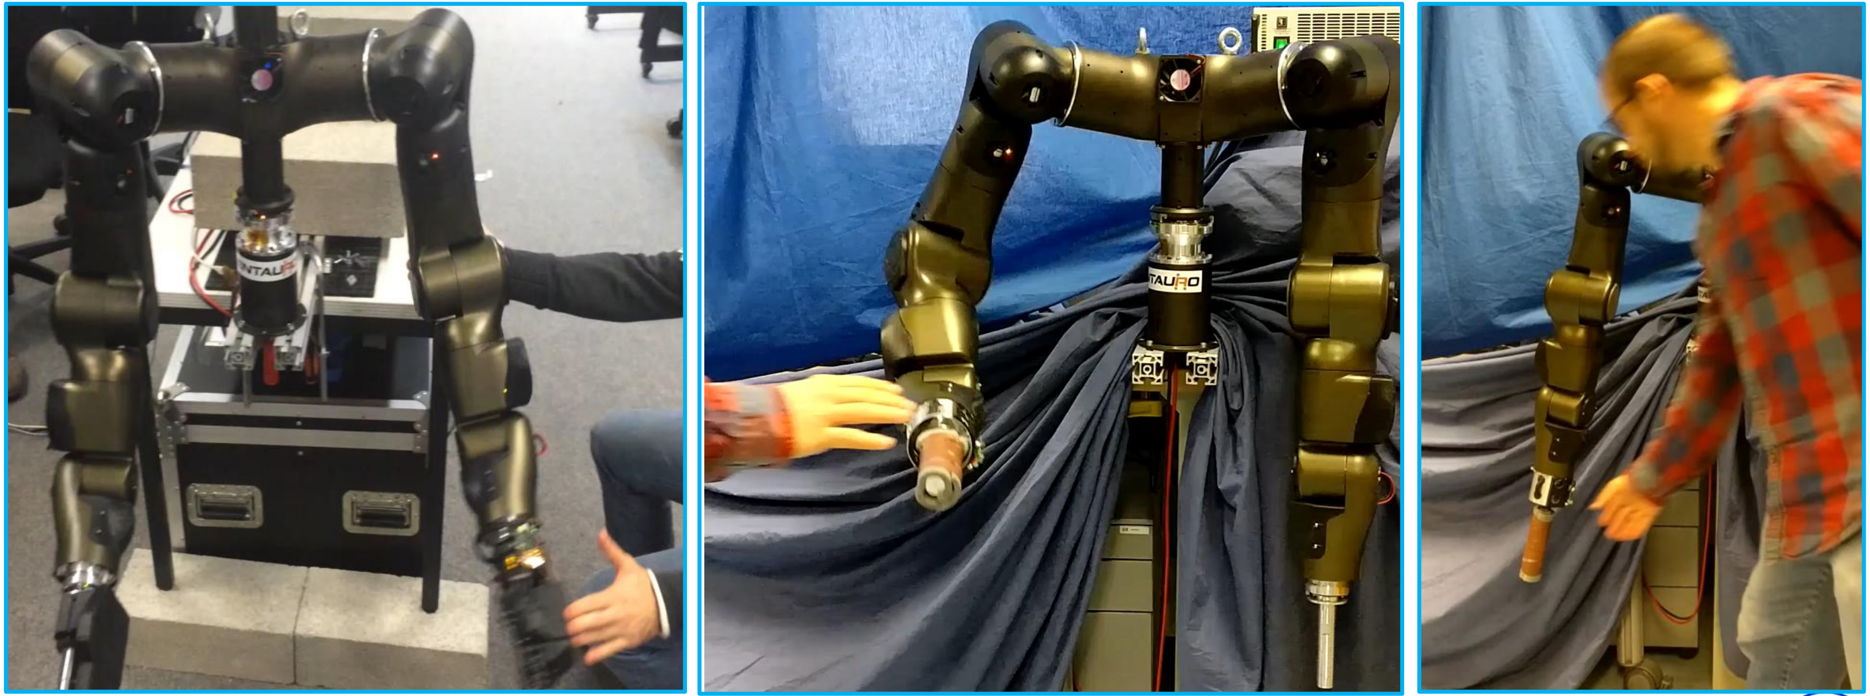

[Giusti et al. ICRA 2017]

universitätbonn ais

Centauro Upper Body: Resilient

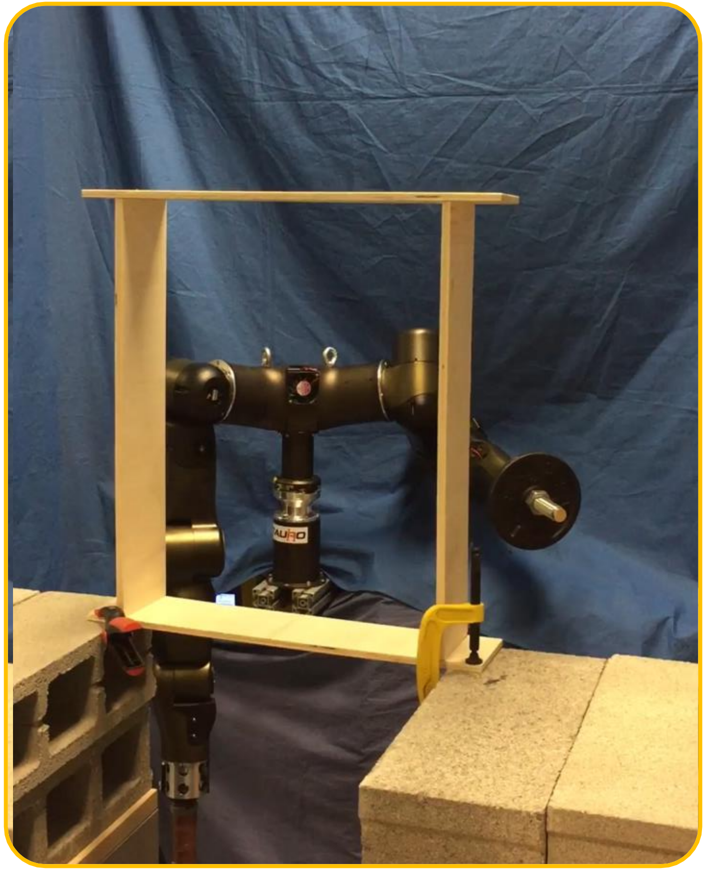

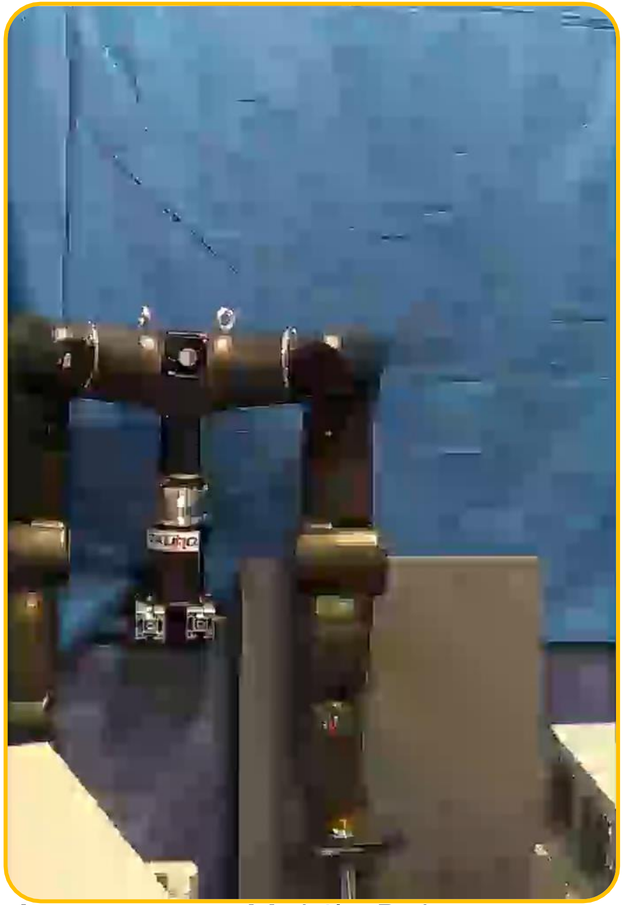

[Giusti et al. ICRA 2017]

Centauro Robot

- Some unpublished material on the Centauro robot has been removed from this version of the slides.
- Please check https://www.centauro-project.eu for updates.

Main Operator Telepresence Interface

- Tendon-driven dual-arm exoskeleton
- Active wrist with differential tendon transmission
- Underactuated hand exoskeleton
- Head-mounted display
- Foot pedals

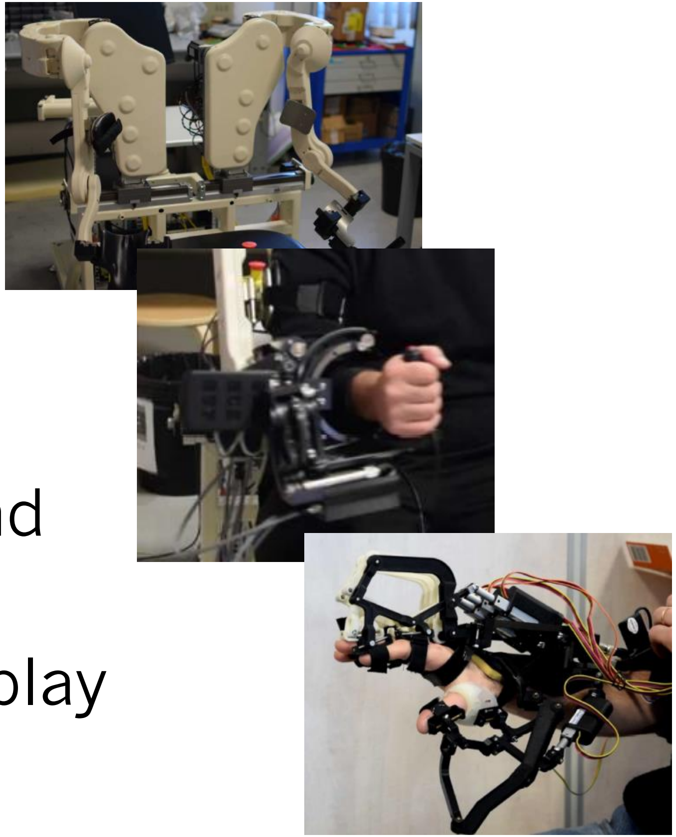

[Frisoli et al., SSSA 2017]

Centauro Upper Body Bi-manual Teleoperation

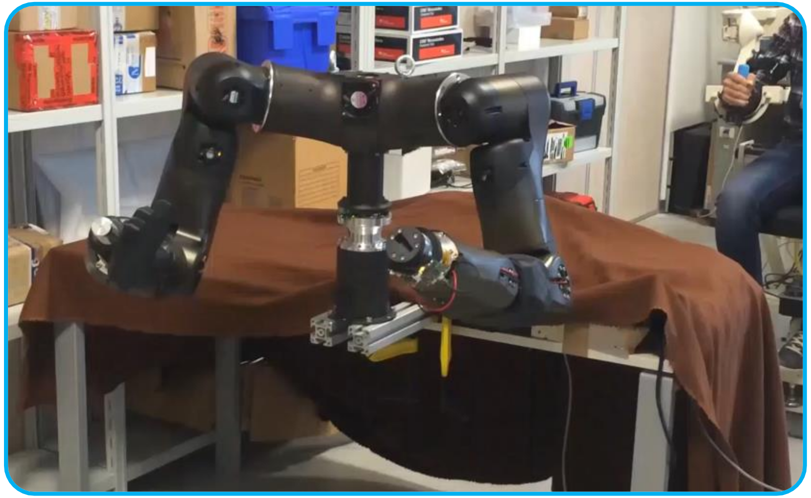

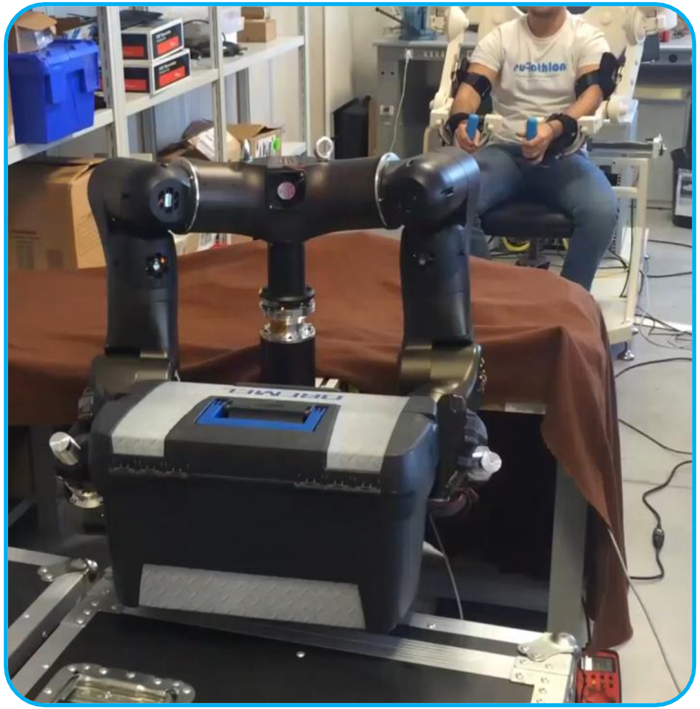

[Tsagarakis et al. (IIT) + Frisoli et al. (SSSA), 2016]

Locomotion Planning Considering Robot Footprint

- Costs for individual wheel pairs from height differences
- Base costs
- Non-linear combination yields 3D (x, y, θ) cost map

[Klamt and Behnke, IROS 2017]

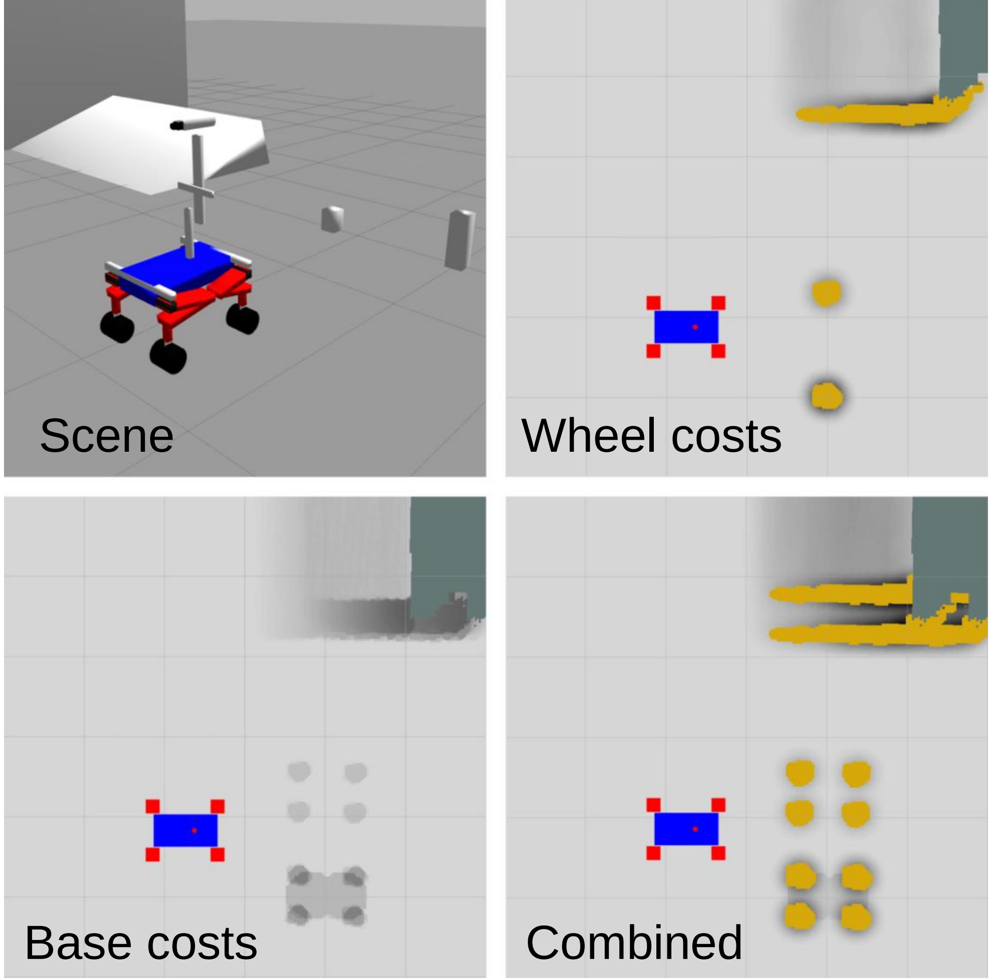

3D Driving Planning (x, y, θ): A*

- 16 driving directions

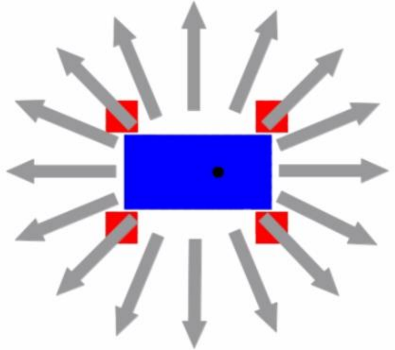

- Orientation changes

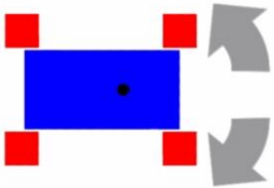

=> Obstacle between wheels

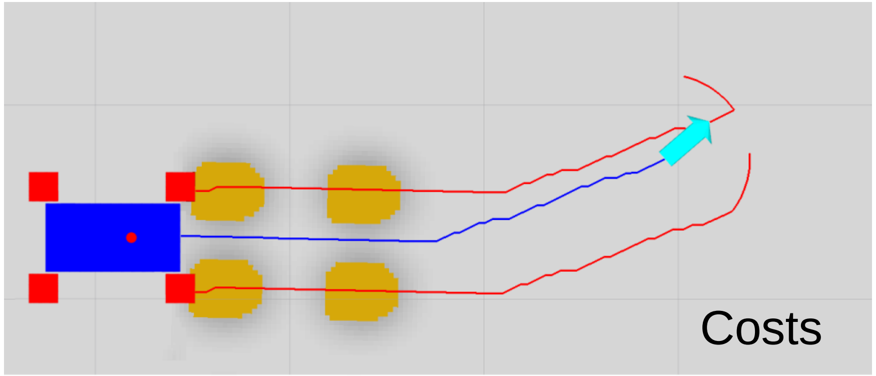

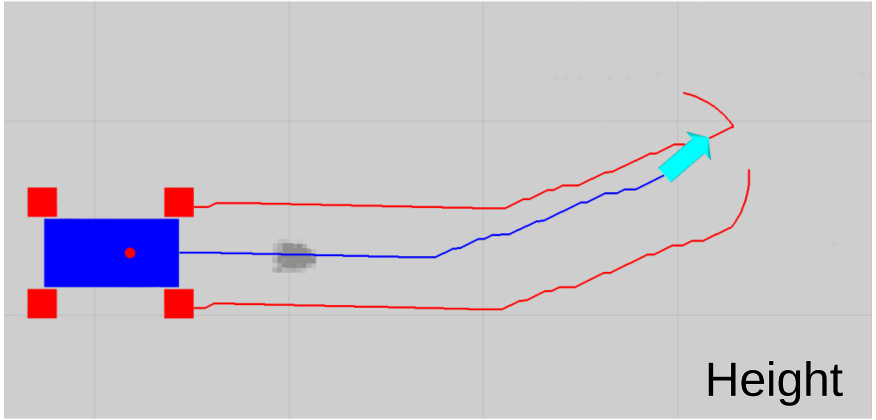

[Klamt and Behnke, IROS 2017]

Making Steps

- If not drivable obstacle in front of a wheel
- Step landing must be drivable
- Support leg positions must be drivable

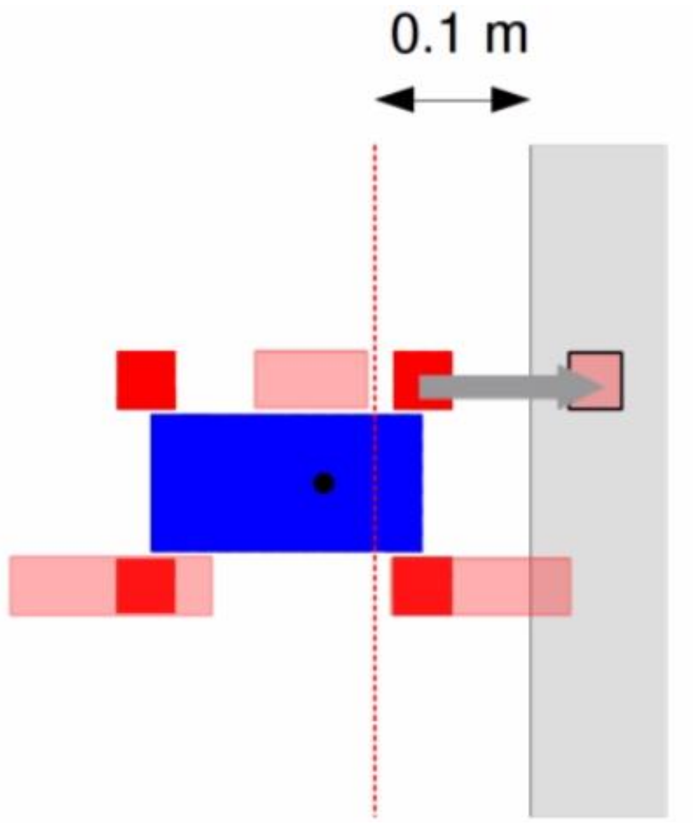

[Klamt and Behnke: IROS 2017]

Expanding Abstract Steps to Detailed Motion Sequences

[Klamt and Behnke: IROS 2017]

Planning for Challenging Scenarios

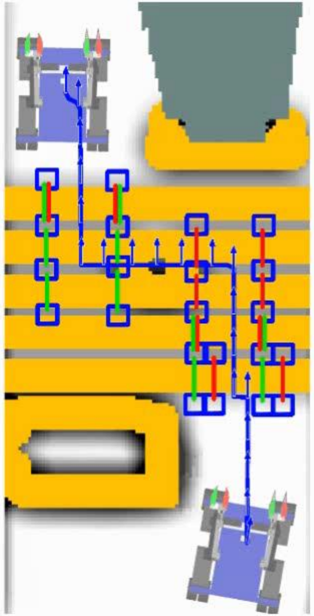

[Klamt and Behnke: IROS 2017]

CENTAURO Workspace Perception Data Set

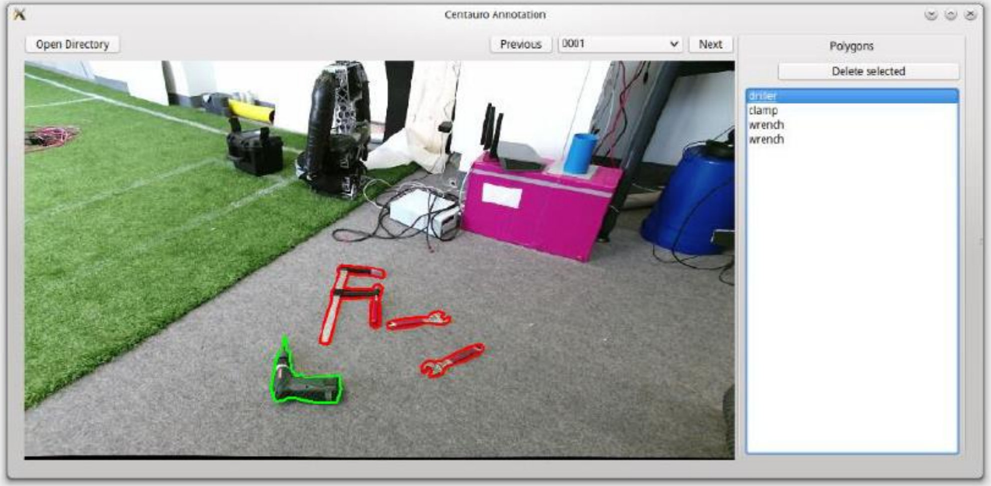

129 frames, 6 object classes

https://www.centauro-project.eu/data_multimedia/tools_data

Deep Learning Object Detection

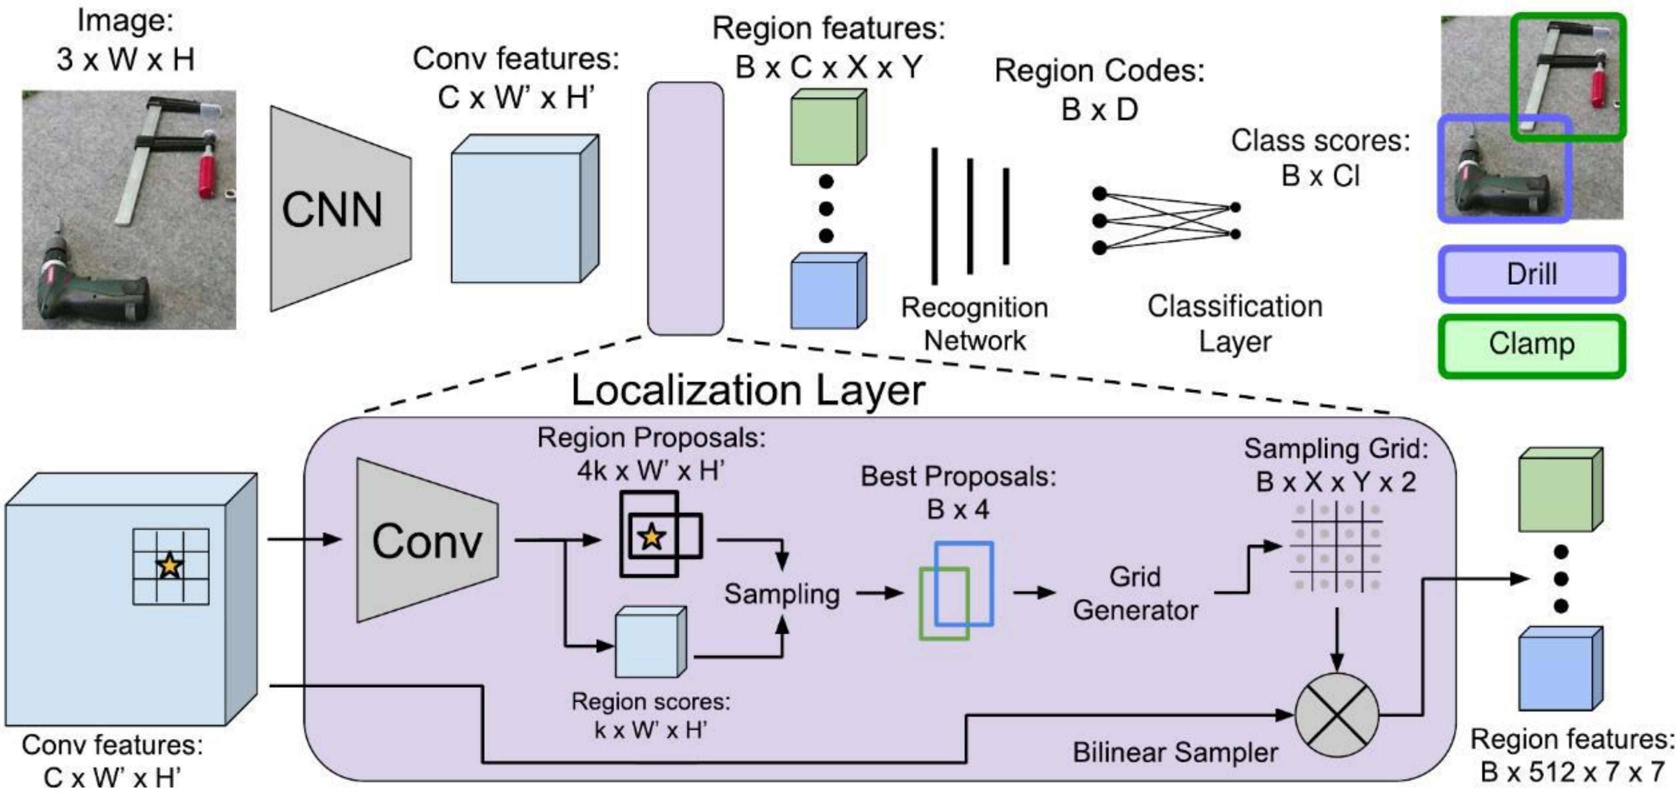

[Johnson et al. 2015]

Tool Detection Results

[Schwarz et al. IJRR 2017]

Resolution	Clamp	Door handle	Driller	Extension	Stapler	Wrench	Mean
	AP / F1	AP / F1	AP / F1	AP / F1	AP / F1	AP / F1	AP / F1
720×507	0.881/0.783	0.522/0.554	0.986/0.875	1.000/0.938	0.960/0.814	0.656/0.661	0.834/0.771
1080×760	0.926/0.829	0.867/0.632	0.972/0.893	1.000/0.950	0.992/0.892	0.927/0.848	0.947/0.841
1470×1035	0.913/0.814	0.974/0.745	1.000/0.915	1.000/0.952	0.999/0.909	0.949/0.860	0.973/0.866

Tools Detection Examples

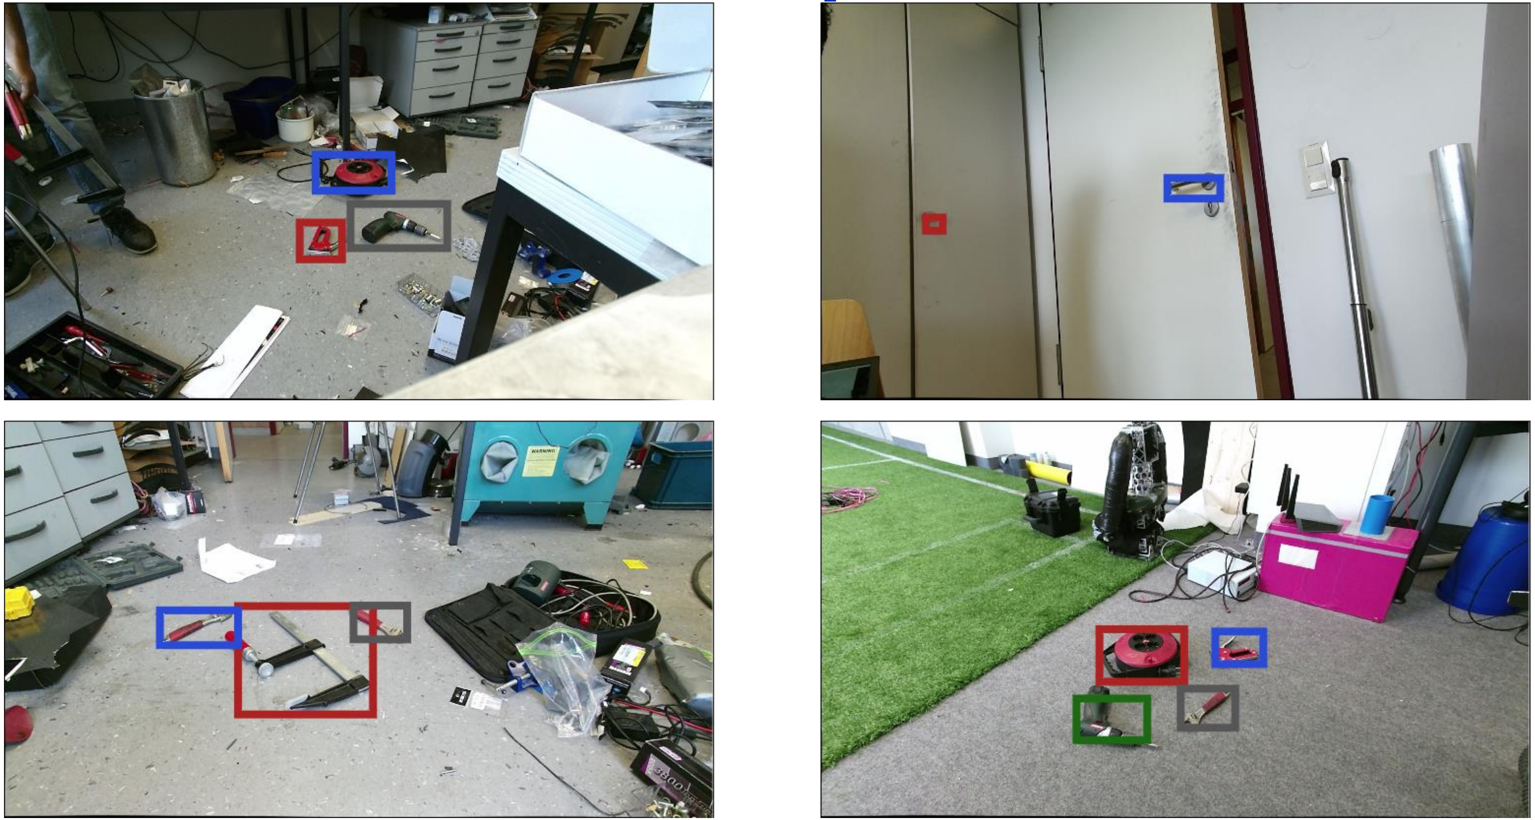

[Schwarz et al. IJRR 2017]

Semantic Segmentation

- Deep CNN

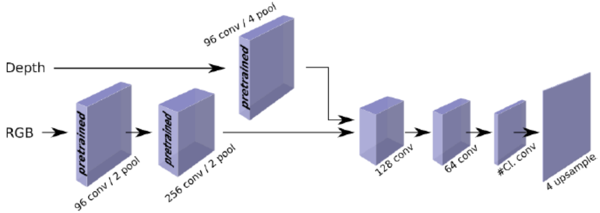

[Husain et al. RA-L 2016]

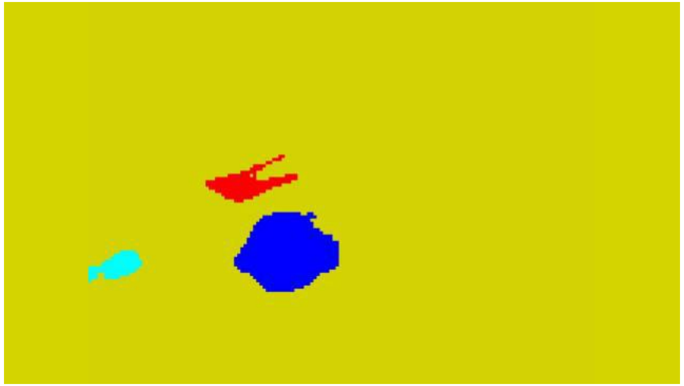

Pixel-wise accuracy:

Clamp	Door handle	Driller	Extension	Stapler	Wrench	Background	Mean
0.727	0.751	0.769	0.889	0.775	0.734	0.992	0.805

3D Object Modeling and 6D Pose Estimation

- Build 3D model on turn table
- Generate proposals
- Register to test image

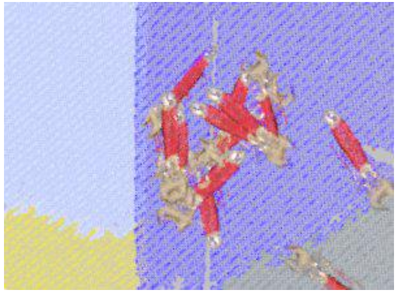

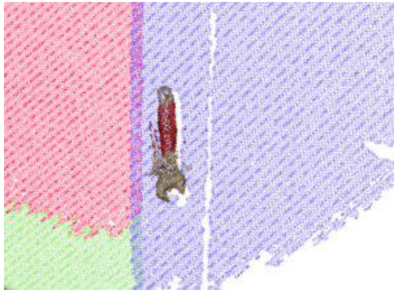

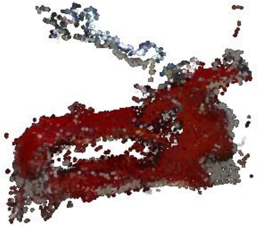

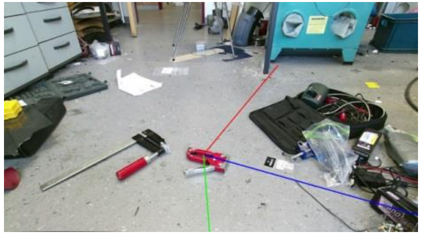

[Aldoma et al., ICRA 2013]

Schunk Five-finger Hand SVH

- Anthropomorphic hand
- 9 DoF

Rviz Interface with Interactive Markers

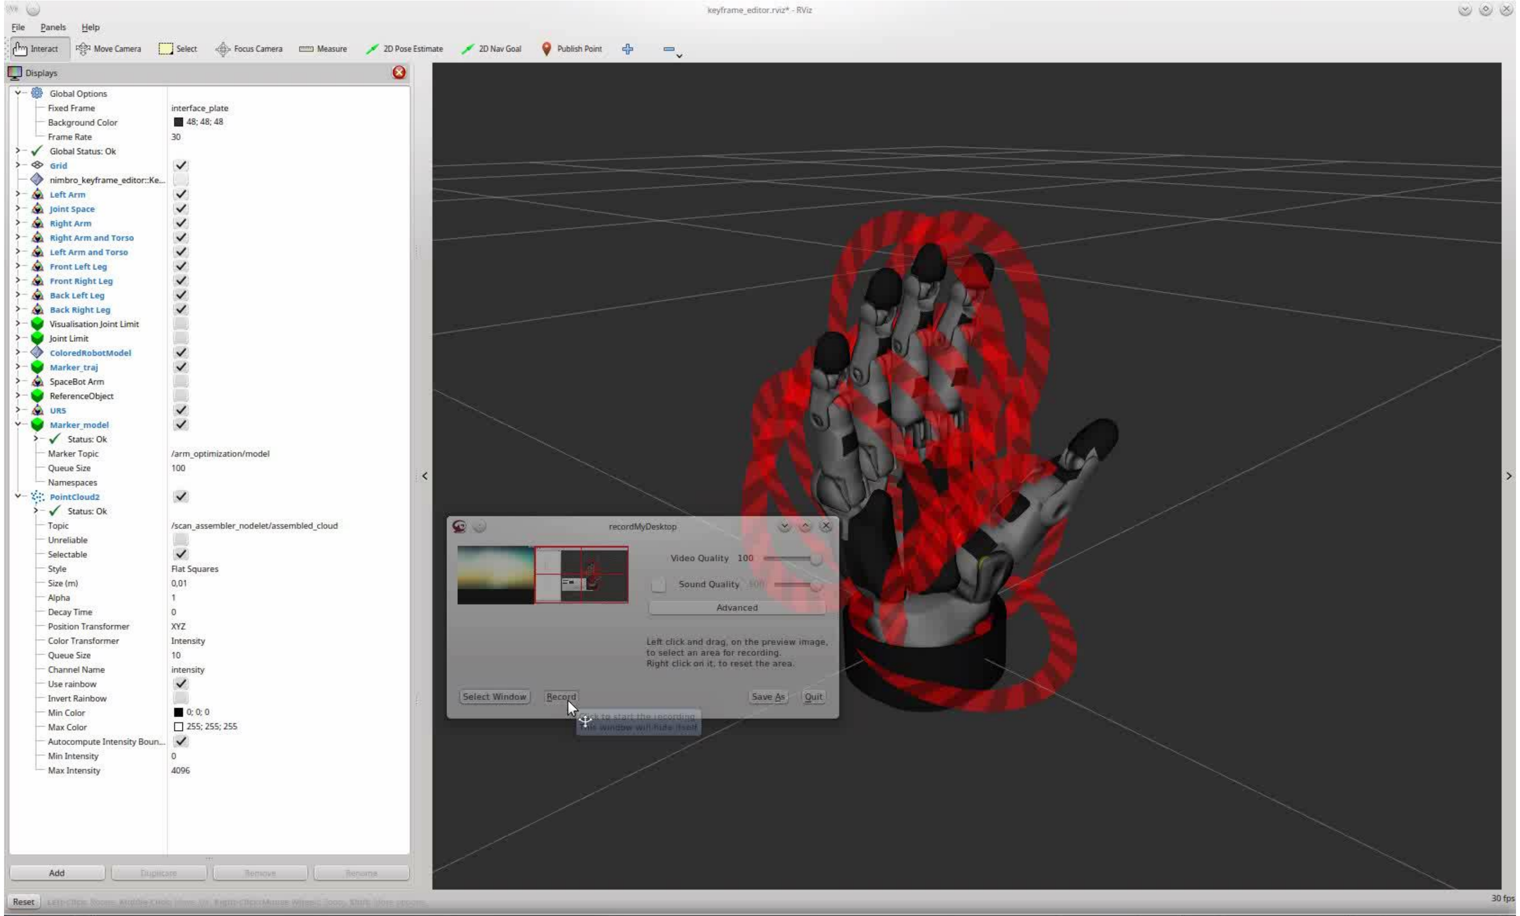

universitätbonn ais

Grasping the Drill and Switching it On

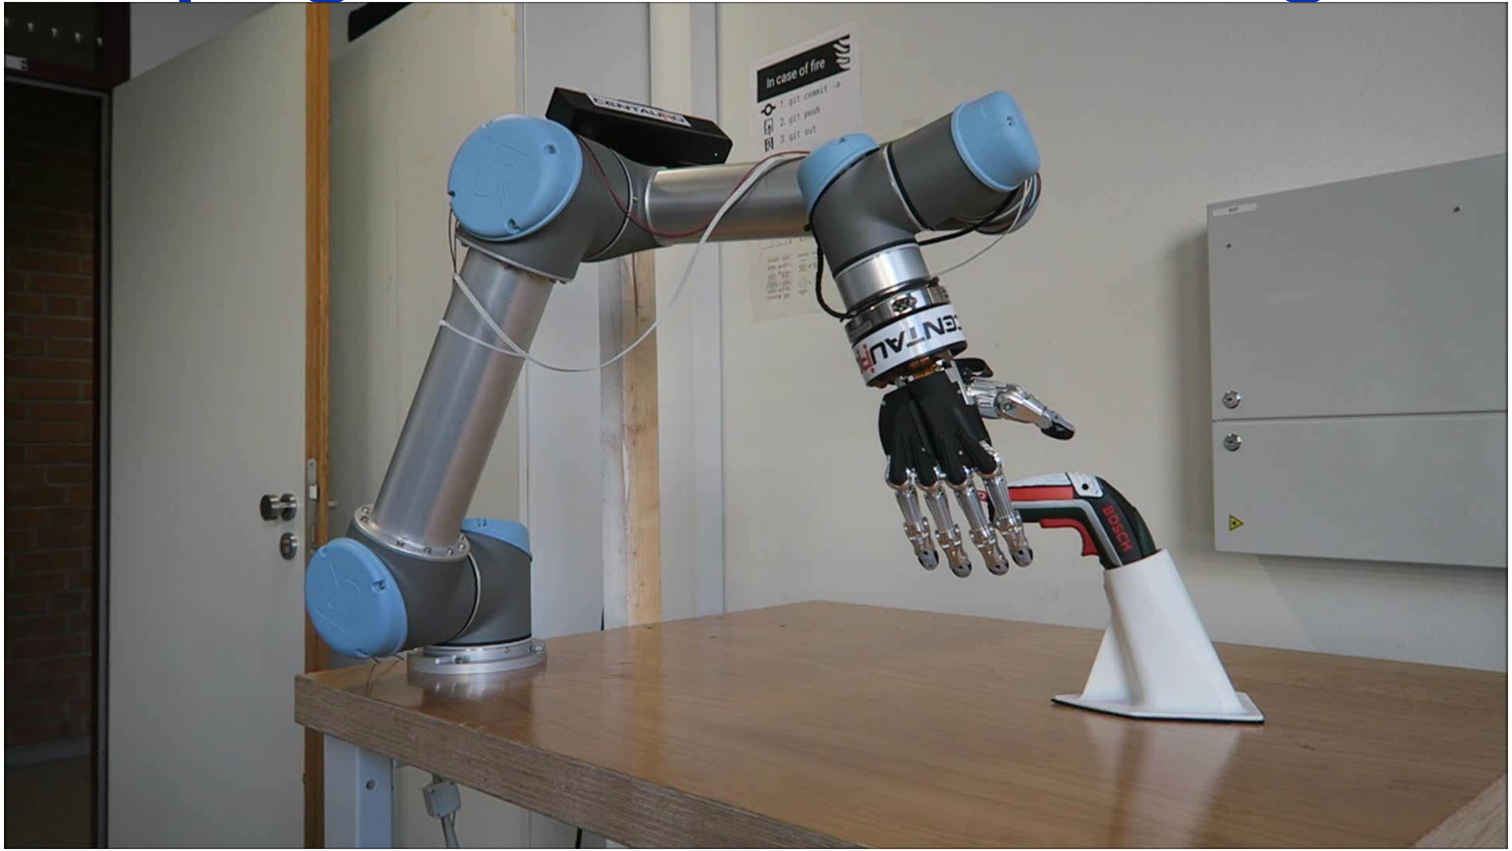

Transfer of Manipulation Skills

- Objects belonging to the same **category** can be handled in a very similar manner.

Transfer of Manipulation Skills

- **Knowledge Transfer**

Learning a Latent Shape Space

- Non-rigid registration of instances and canonical model
- Principal component analysis of deformations

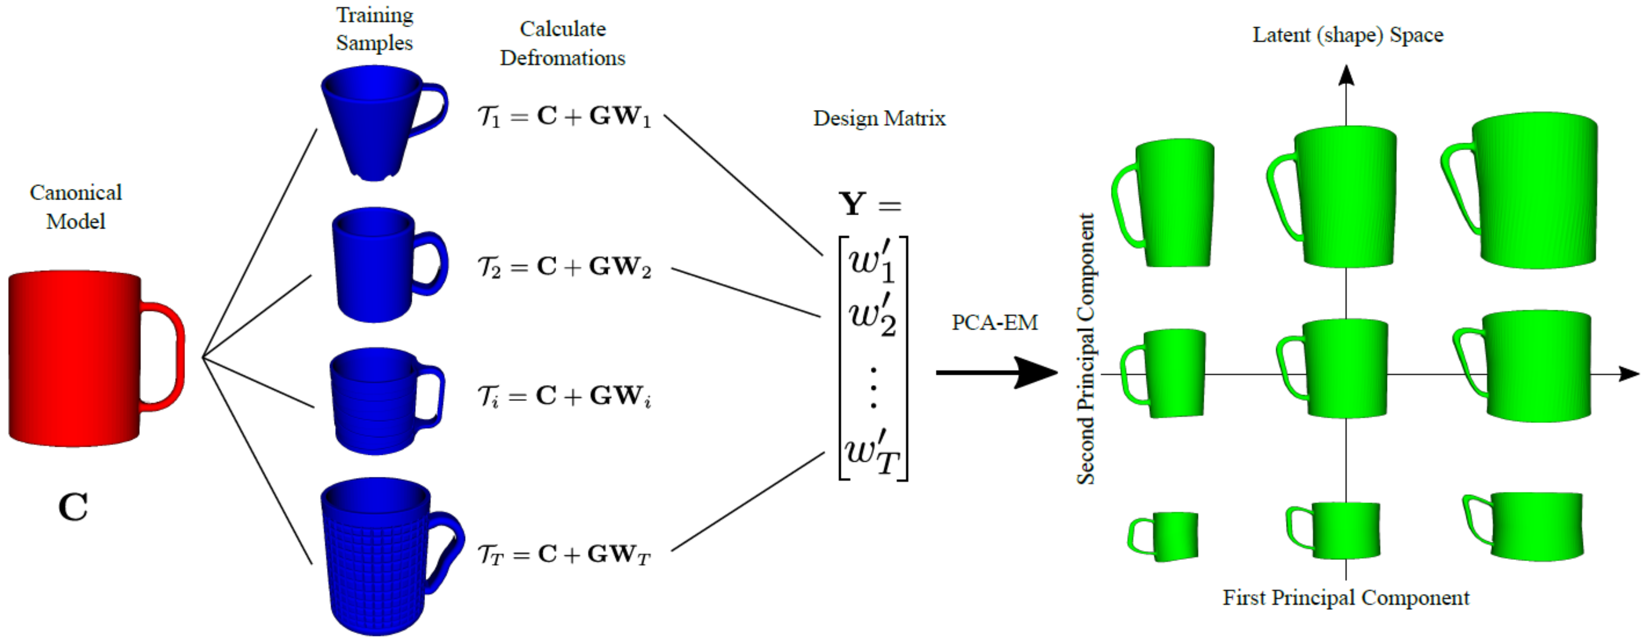

Interpolation in Shape Space

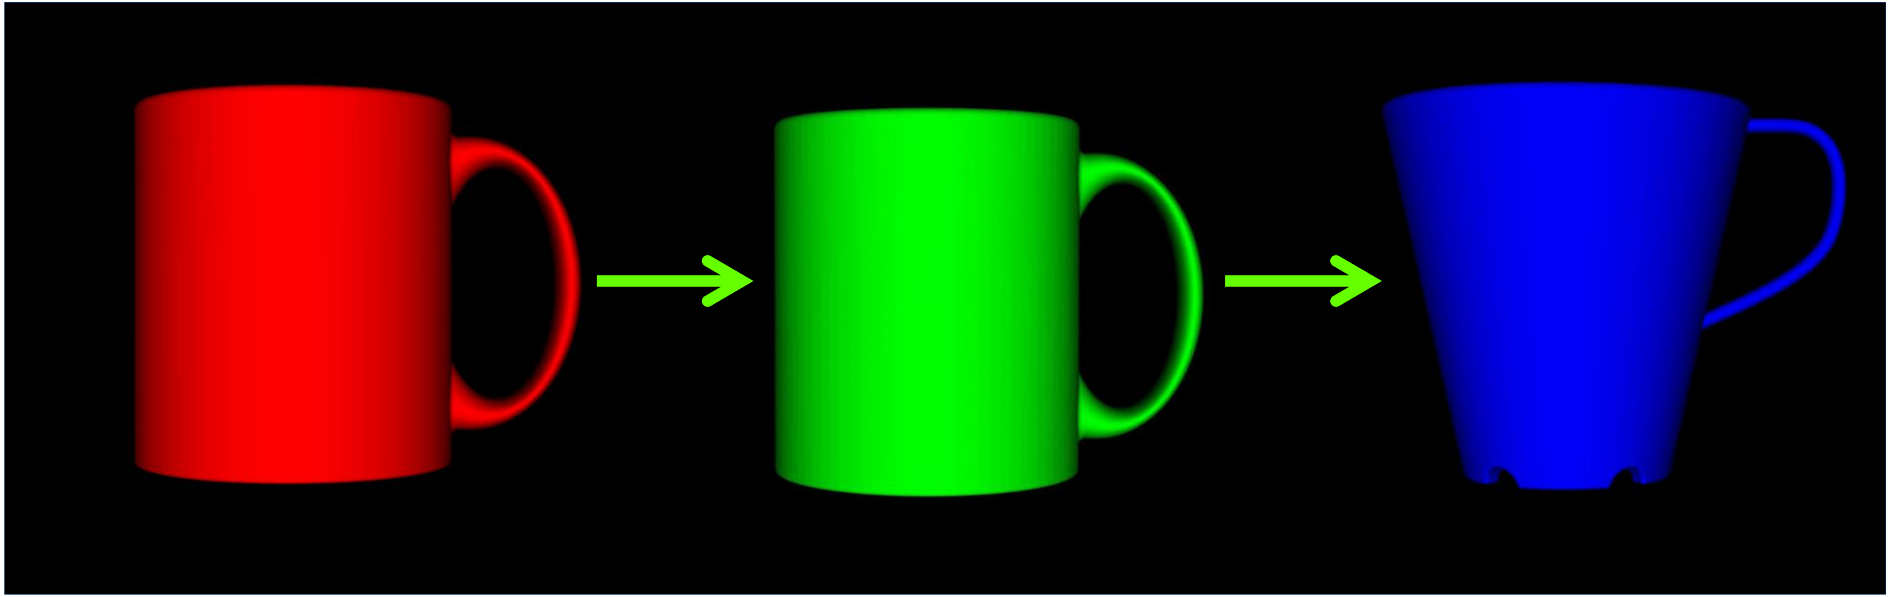

Interpolation in Shape Space

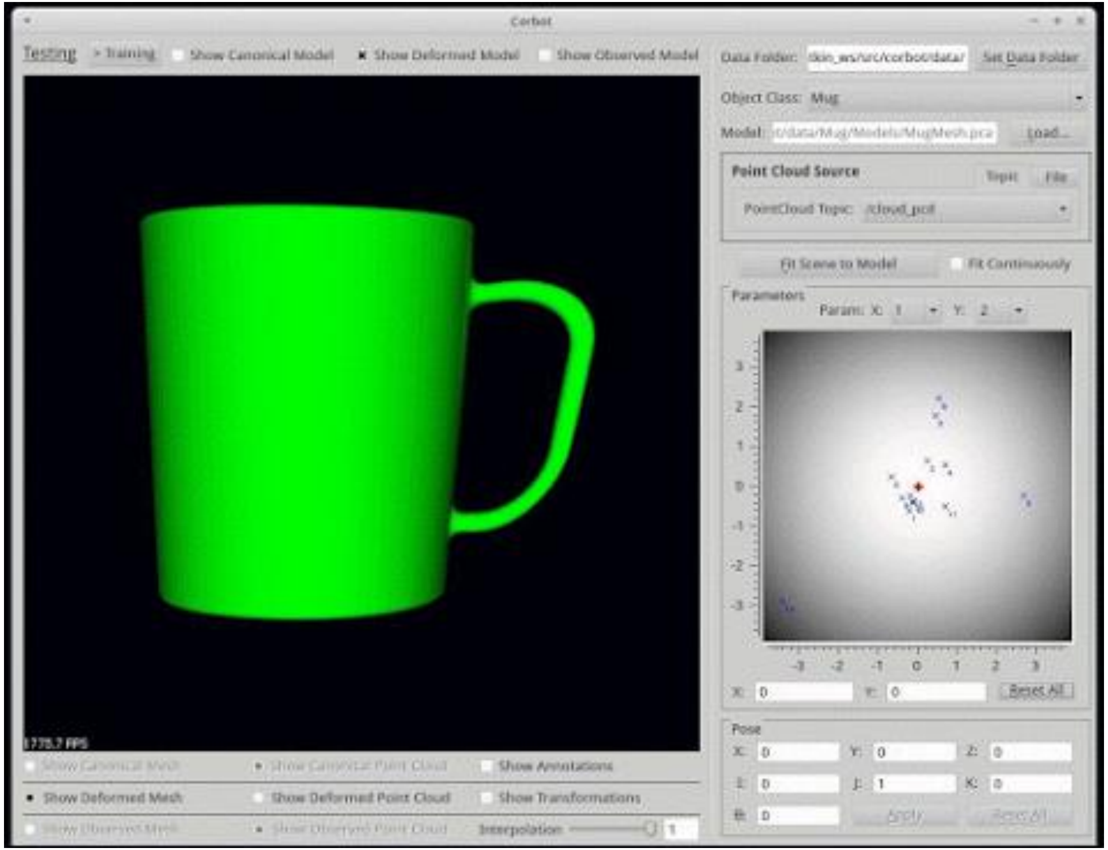

Shape-aware Non-rigid Registration

- Partial view of novel instance
- Deformed canonical model

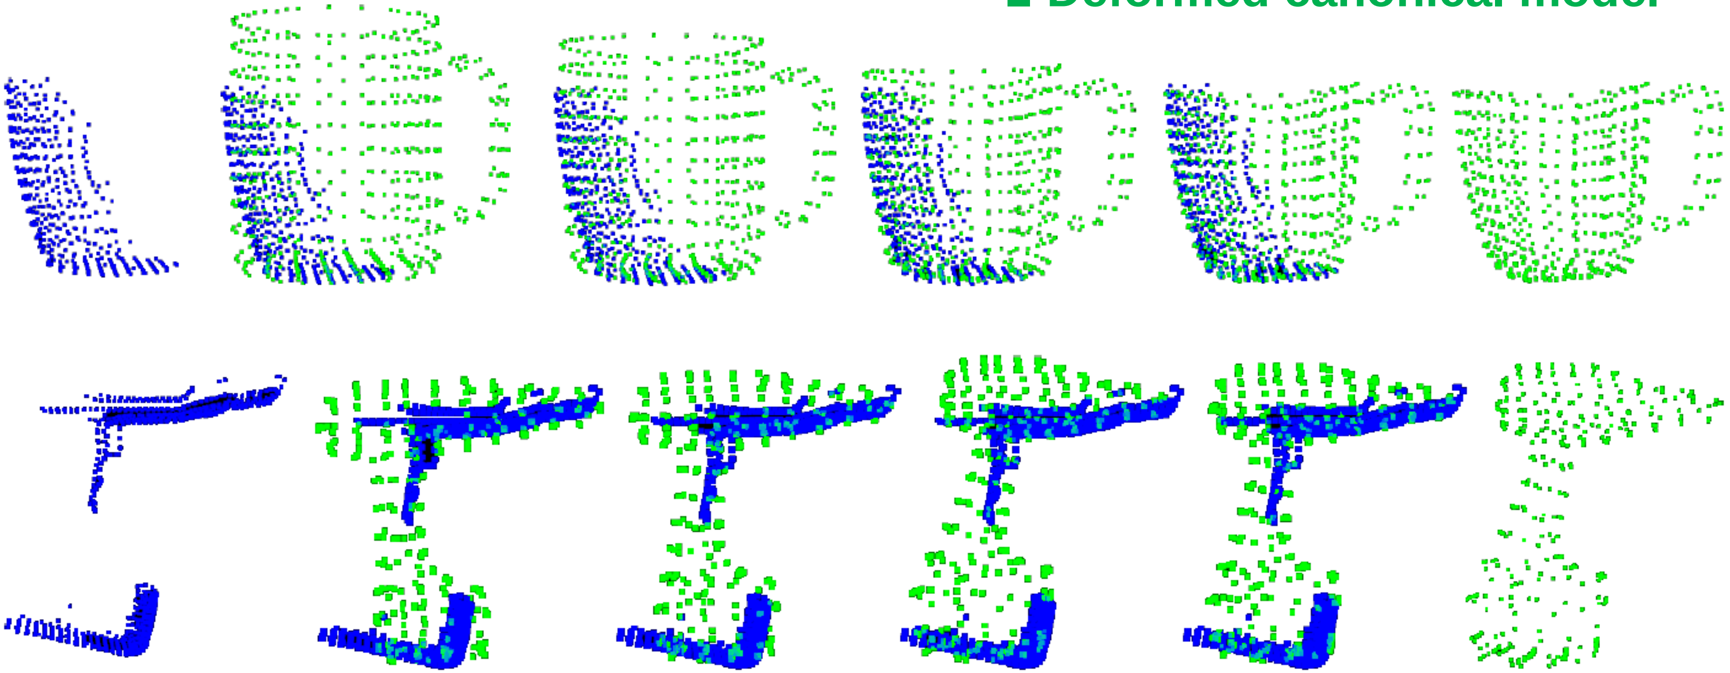

Transference of Grasping Skills

Warp grasping information

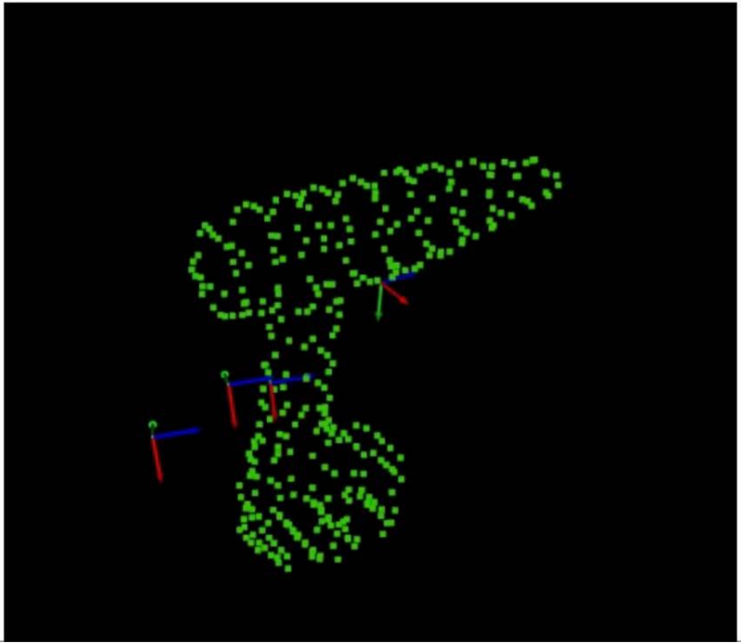

Manipulation Trajectory Optimization

- Extended stochastic trajectory optimization (STOMP), 8 DoF
- Weighting multiple objectives, e.g. speed, obstacles, torque, ...

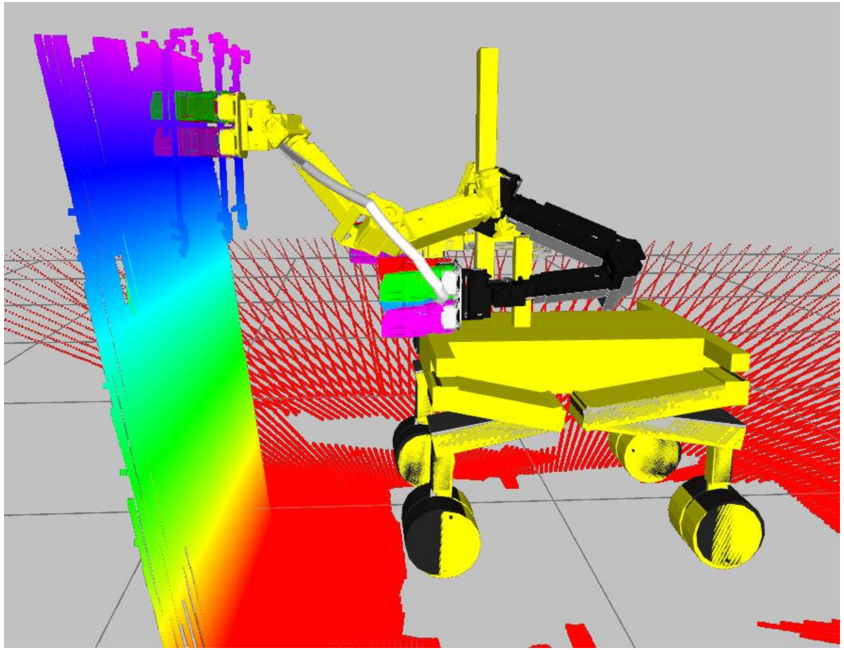

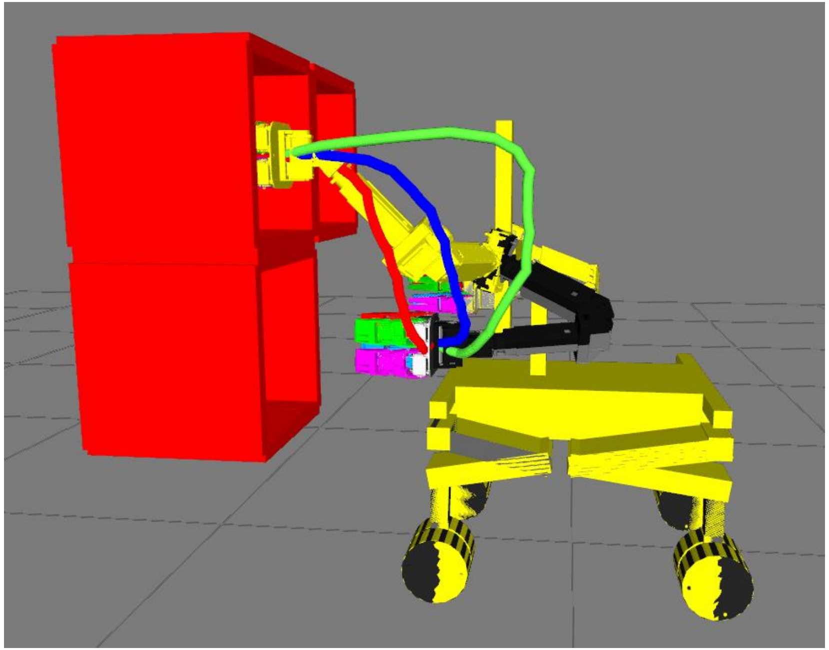

[Pavlichenko and Behnke, IROS 2017]

Momaro Reaching for an Object

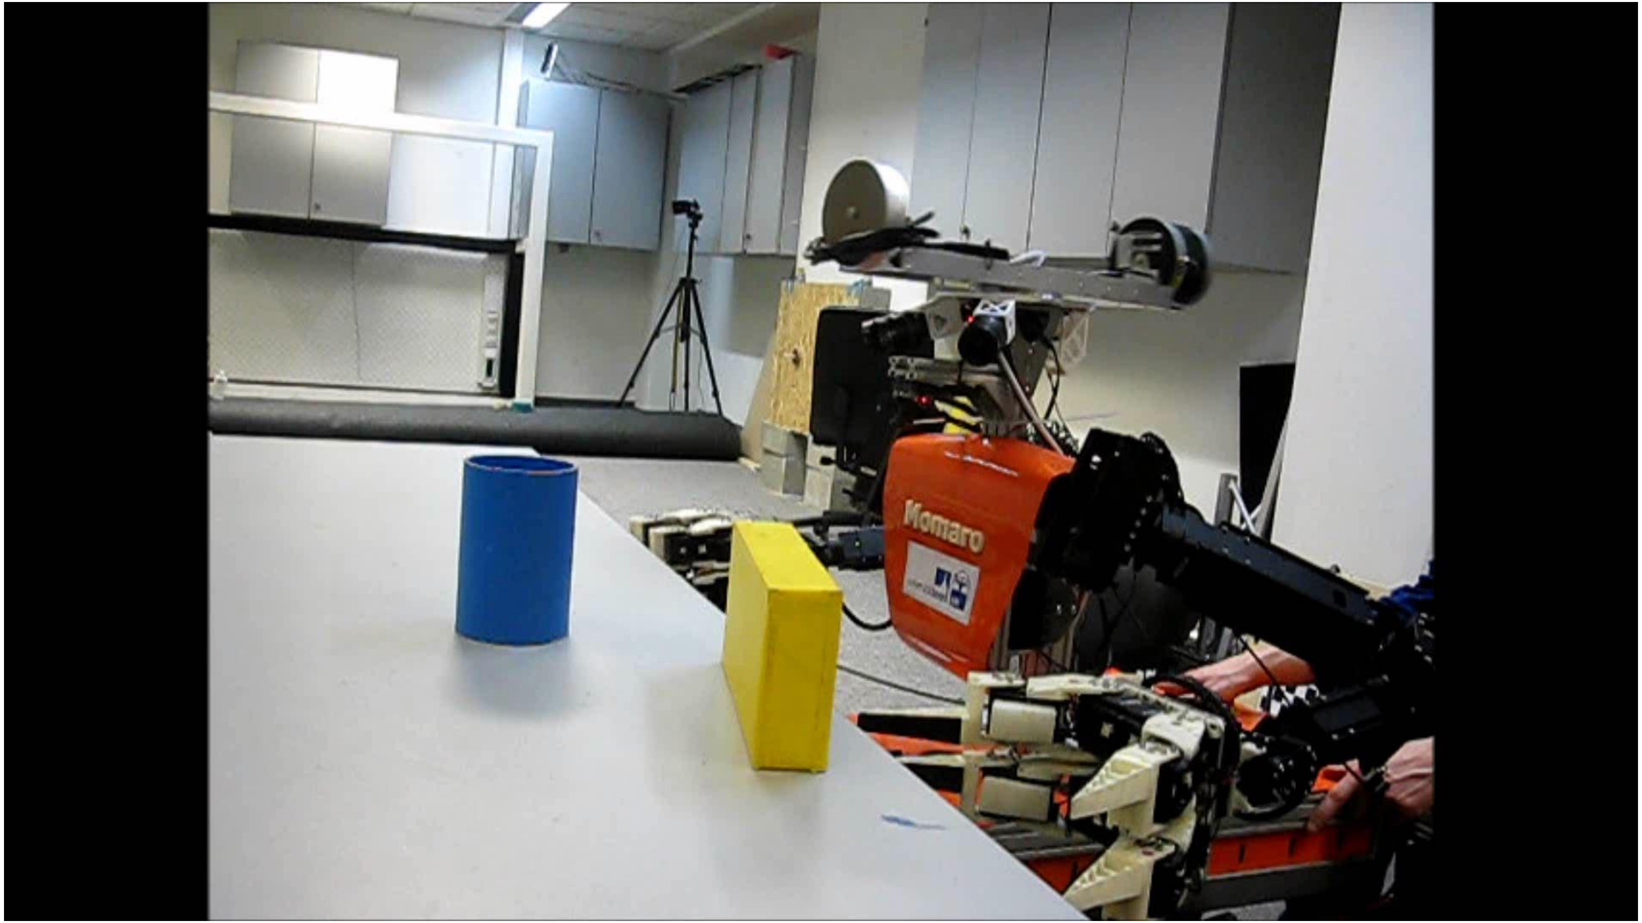

Shelf Experiment

[Pavlichenko and Behnke, IROS 2017]

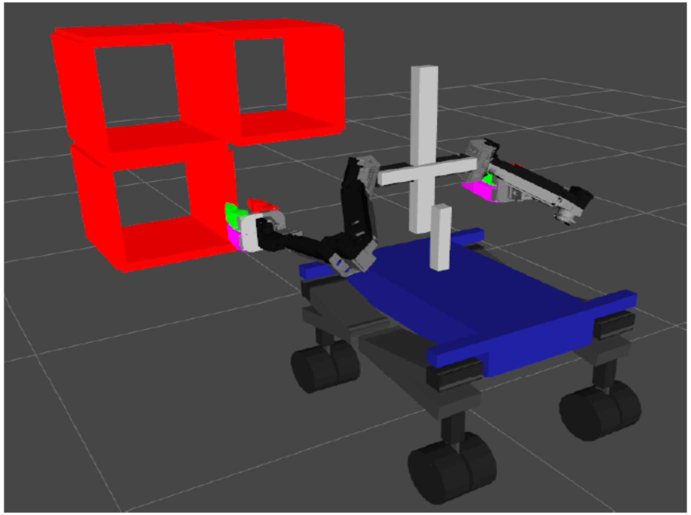

- Four configurations
  - 12 planning tasks
  - 100 executions for each task
- 3 difficulty levels:
  - Easy
  - Hard (gripper deeper)
  - Hard constrained (endeffector orient.)

	Difficulty level					
	Easy		Hard		Hard constrained	
Algorithm	success rate	runtime [s]	success rate	runtime [s]	success rate	runtime [s]
LBKPIECE	0.94	2.47 ± 1.08	0.93	2.46 ± 0.85	-	-
STOMP-Industrial	0.87	0.87 ± 0.86	0.76	$1.47 \pm 1,01$	-	-
RRTConnect	0.97	0.29 ± 0.18	0.96	0.85 ± 0.58	0.97	1.22 ± 1.04
STOMP-New	1.0	0.09 ± 0.02	1.0	0.18 ± 0.11	0.99	0.28 ± 0.21

Corridor Experiment

- Two difficulty levels:
  - Easy
  - Hard
- 100 trials each

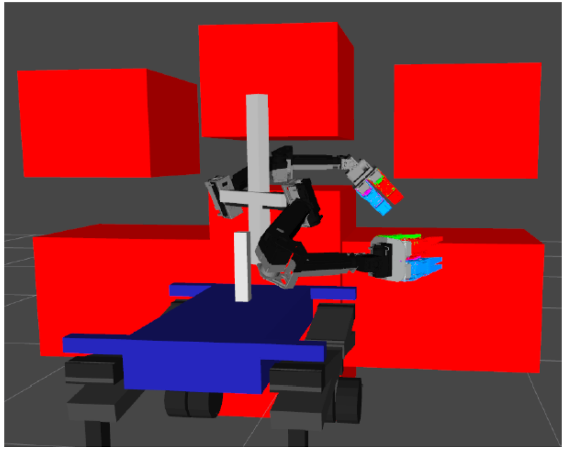

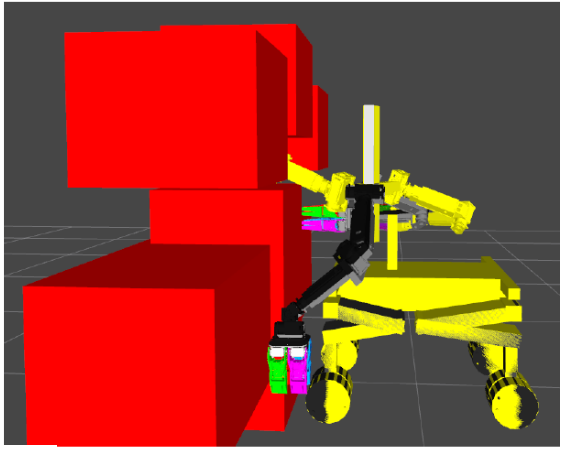

	Difficulty level			
	Easy		Hard	
Algorithm	success rate	runtime [s]	success rate	runtime [s]
LBKPIECE	0.65	6.97 ± 2.58	0.50	7.82 ± 2.58
RRTConnect	0.08	9.64 ± 1.27	0.06	9.71 ± 1.56
STOMP-Industrial	0.00	2.82 ± 0.07	0.00	2.85 ± 0.08
STOMP-New	0.78	1.89 ± 1.44	0.18	3.64 ± 1.29

Velocity Profiles

- Start pose can have velocity
- Continuous replanning possible

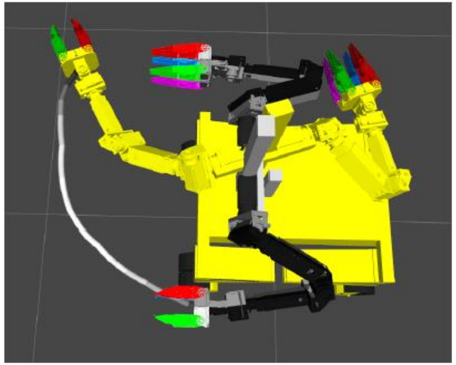

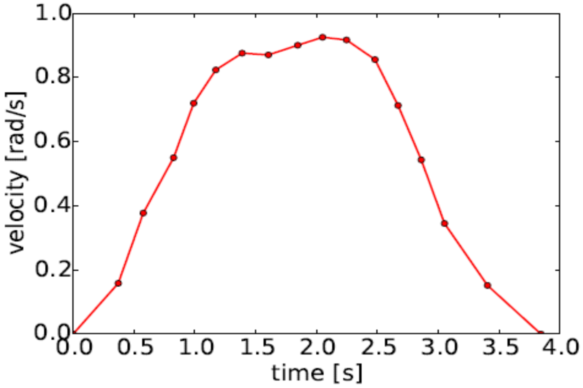

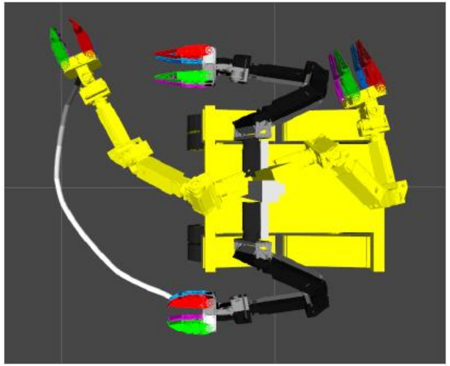

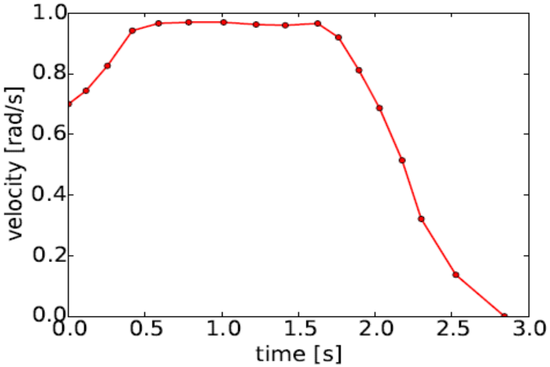

RefineNet

[Lin et al. CVPR 2017]

- Increase resolution by using features from the higher resolution
- Corse-to-fine semantic segmentation

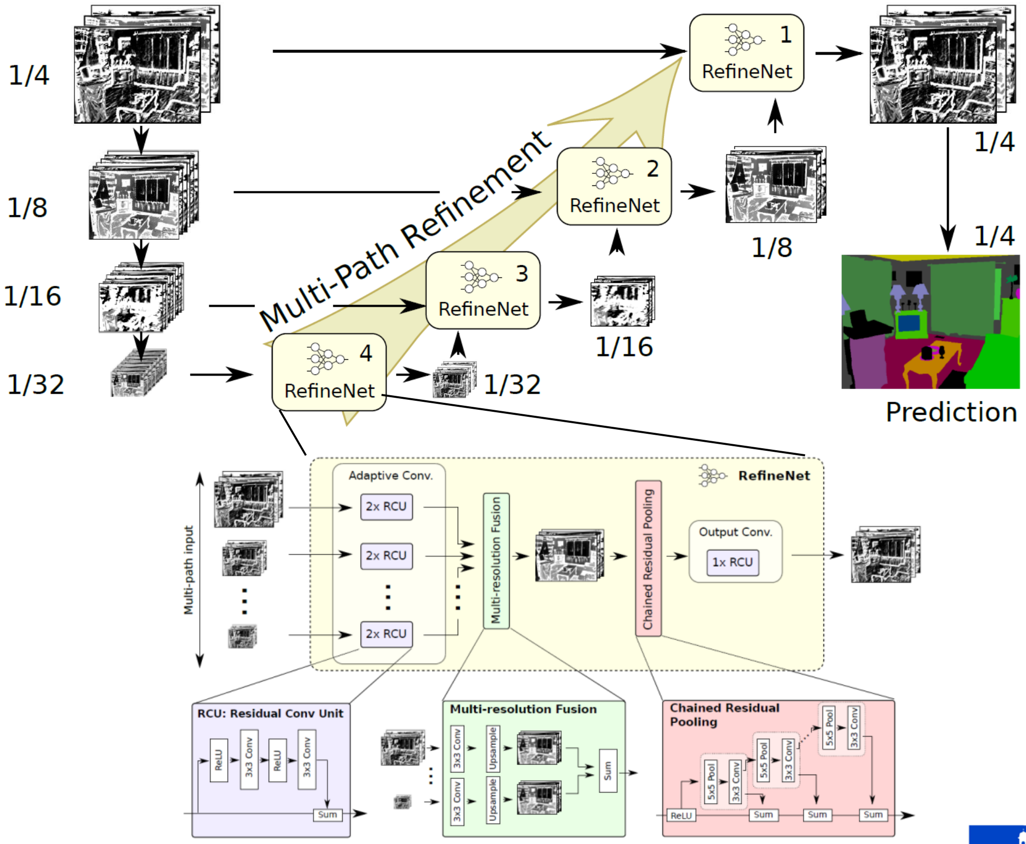

universitätbonn ais

ARC 2017 Perception Example

Object Capture and Scene Rendering

- Turn table + DLSR

Rendered scenes

Amazon Robotics Challenge 2017

- Quick learning of novel objects
- Training with rendered scenes

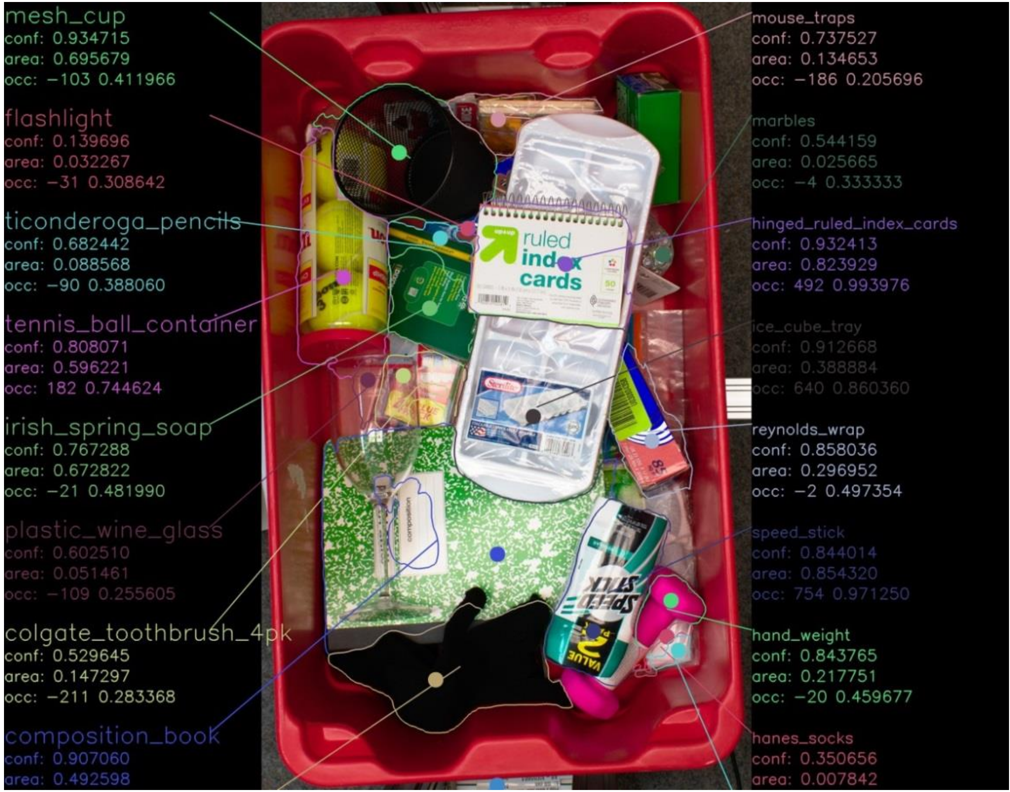

Object Pose Estimation

- Use upper layer of RefineNet as input
- Predict pose coordinates for one segment

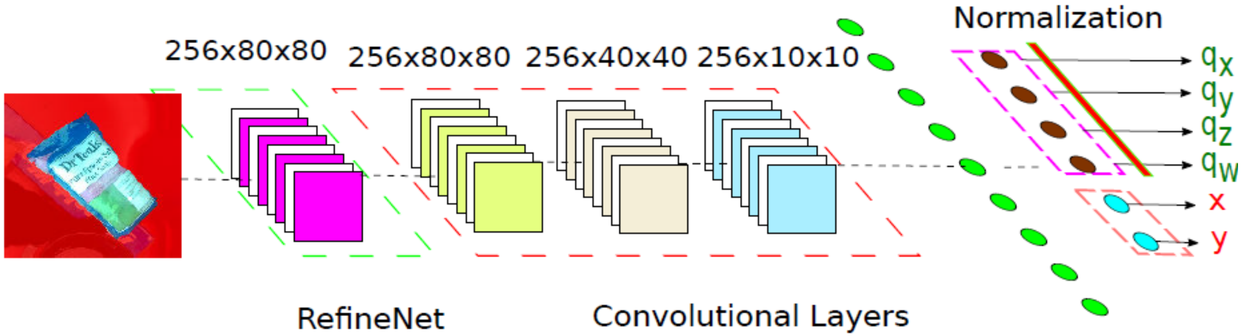

Input

Predicted pose

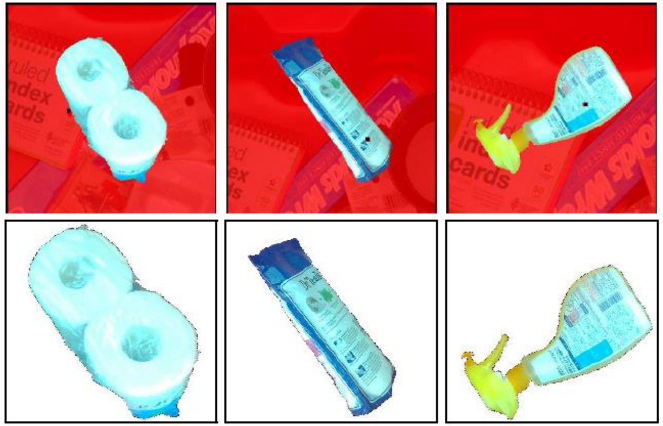

CENTAURO

Conclusions

- Developed high-performance platforms for challenging scenarios
  - Mobile manipulation robots
  - Micro aerial vehicles
- Teleoperation is flexible, but demanding and error-prone
- Autonomy for common navigation and manipulation tasks needed, challenges include
  - 3D mapping
  - Semantic perception
  - High-dimensional motion planning

We are Hiring!

- PhD students and postdocs
- **ais.uni-bonn.de/jobs.html**

universitätbonn ais

Questions?