ICRA 2015

Conference Report

Seattle, WA
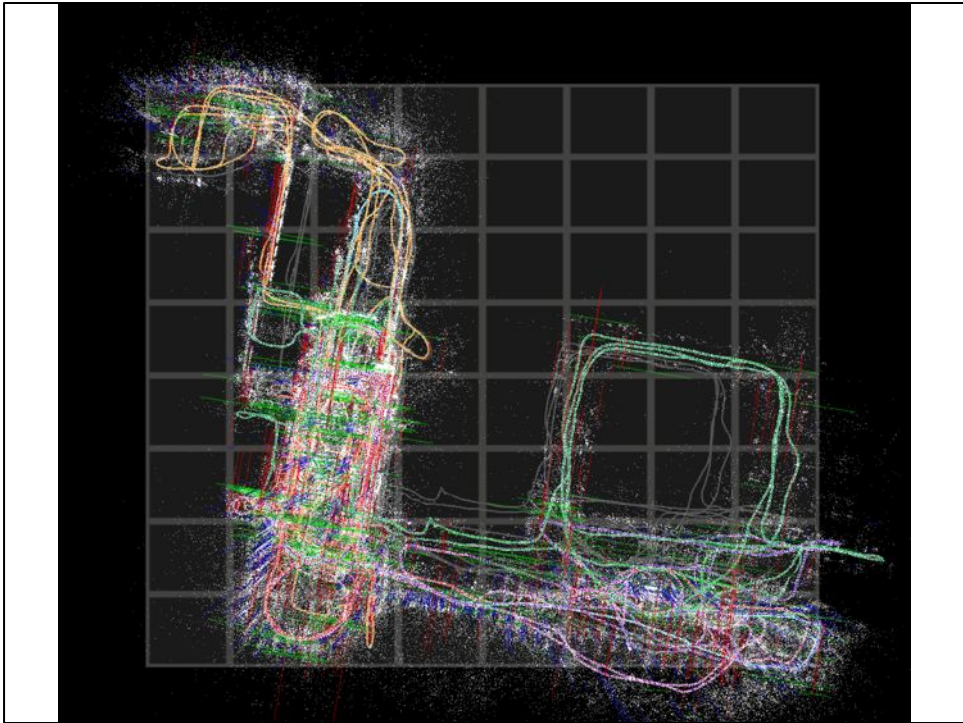
Conference Venue, courtesy of http://onionmaps.info/

Lessons

- Going to your first conference can be hard!

- Get introduced by people you know

- Have something to present

- Become a volunteer?

- Knowing about other people's work helps

Ruzena Bajcsy

Ruh-jee-na Bye-chee

I also learned…
How to pronounce Ruzena Bajcsy

Lessons

- Wolfram Burgard (still) wins ICRA

- Pieter Abbeel a close second

- The classical AI people are disgruntled with the kids these days
  - [(See here)](#)

- Burgard had 14 papers, Abbeel 13
- Darwin Caldwell 10, Howie Choset 11 (19 if you count workshops!), Vijay Kumar 12, Paul Newman 11, Daniela Rus 11, Roland Siegwart 11
- AI folks want more research in symbolic logic/reasoning and symbol grounding, and they also seem to think there's a lack of interdisciplinary / systems-integration projects like Shakey

Conference Format

- 12 tracks!
- 3-minute lightning talk + 1 minute questions
- 30-minute interactive session
- 3 plenaries
- 12 keynotes
- Late Breaking Results poster session (new)
- 940 papers
- 41% accept rate

Conference Format

- 3 Banquet Receptions
- Networking Lunches
- 41 Workshops!
- 66 Exhibitors
- Career Fair (new)
- 4 Robot Challenges
- Tours: Boeing, UW, Blue Origin, Amazon, VA Hospital

Conference Strategy

- Stick to what you know
- See the best papers
- Branch out occasionally
- Make time for the exhibit hall?
- Prioritize networking whenever given the choice

Favorite Videos

- [Innovation: Past, Present, and Future](#) — Dean Kamen
    - "They asked me to talk about innovation. Well, that's like talking about love: Everyone wants it, and no one knows what it is."
- [Taming the Swarm](#) — Radhika Nagpal
- [Reel of Best Paper Finalists](#)

Learning Contact-Rich Manipulation Skills With Guided Policy Search

- Sergey Levine, Nolan Wagener, Pieter Abbeel (UC Berkeley)
- Best Manipulation Paper
- Modified version of LQG to optimize trajectory for a task.
- Instead of learning parameters of a hand-engineered control function, learn parameters of a neural network that *directly* maps states -> motor torques.
- Allows learning of arbitrary nonlinear policy, even allows for some generalization to previously unseen target positions.
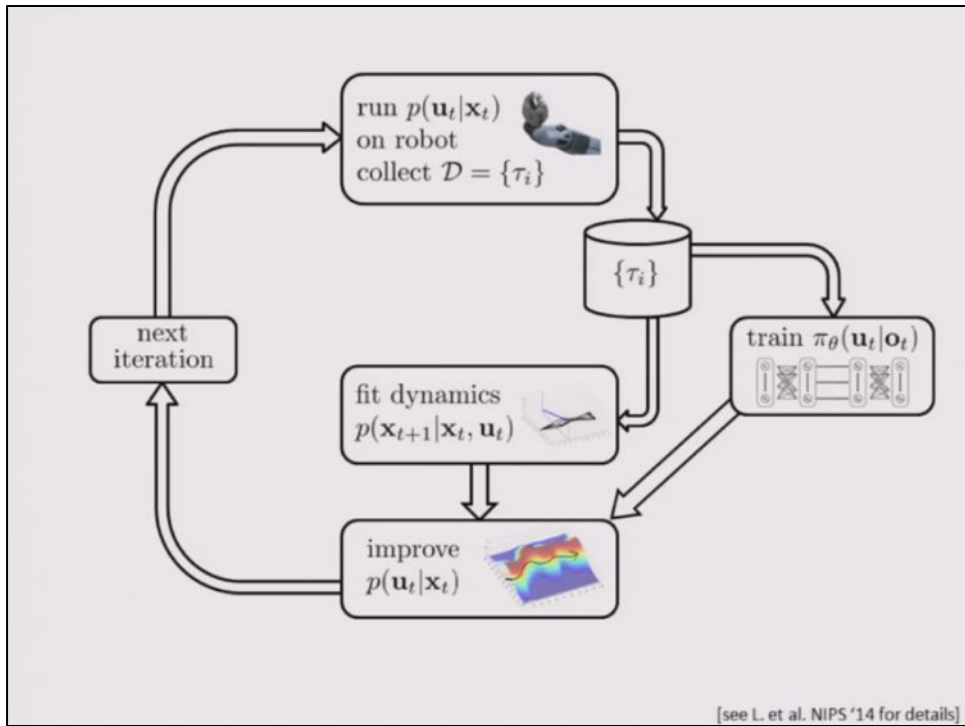
Loop for simultaneously learning good policy & good approximation of the dynamics

End-to-End Training of Deep Visuomotor Policies

- Sergey Levine, Chelsea Finn, Trevor Darrell, Pieter Abbeel (UC Berkeley)
- Target position is no longer provided to the system—instead, visual features are given.
- Deep network learns mapping of features to joint torques.

Pedestrian Detection with a Large-Field-Of-View Deep Network

Pedestrian Detection with a Large-Field-Of-View Deep Network

- Anelia Angelova, Alex Krizhevsky, and Vincent Vanhoucke (Google)
- 3 stages: LFOV, fast ped detector, slow ped detector (deep convnet)
- 3.6 FPS
- LFOV generates proposals, avoids sliding window or other proposal methods
- More generally, can be used to speed up any classifier by generating good proposals.

Real-Time Grasp Detection Using Convolutional Neural Networks
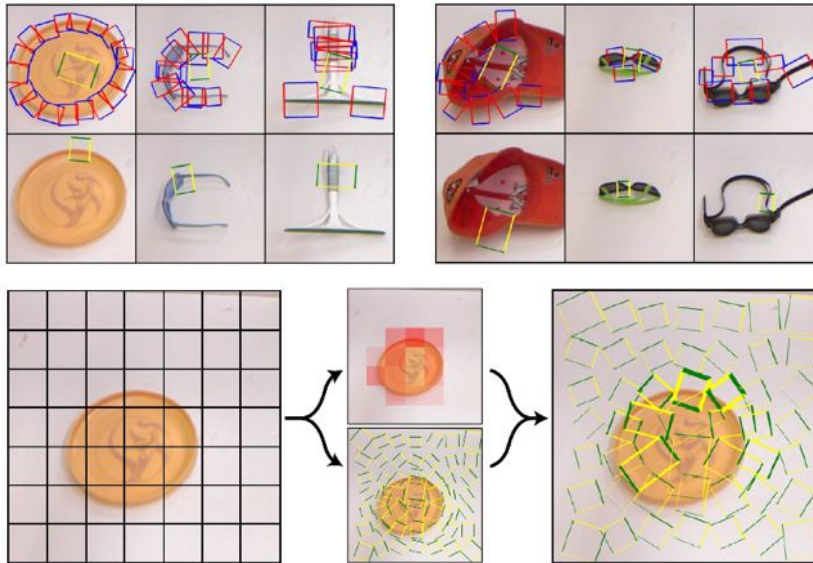
Real-Time Grasp Detection Using Convolutional Neural Networks

- Joseph Redmon, Anelia Angelova (UW, Google)
- Similar concept: avoid sliding window using deep convnet.
- Basically the same as object detection, except there's one extra parameter for orientation.
- 13 FPS
- Significantly better (87% vs 76%) while also being way faster (~3 orders of magnitude!)

Learning Legged Swimming Gaits from Experience

- David Meger, Juan Camilo Gamboa Higuera, Anqi Xu, Philippe Giguere, and Gregory Dudek (McGill)
- RL initialized in the simulator and then further optimized on the bot
- The dynamics is also refined during real-world trials, not just the policy
- Uses PILCO: efficient model-based policy search
- GP is used to approximate forward dynamics

Grasping without Squeezing: Shear Adhesion Gripper with Fibrillar Thin Film

- Elliot W. Hawkes, David L. Christensen, Amy Kyungwon Han, Hao Jiang, and Mark R. Cutkosky (Stanford)

Learning by Observation for Surgical Subtasks: Multilateral Cutting of 3D Viscoelastic and 2D Orthotropic Tissue Phantoms

- Adithyavairavan Murali, Siddarth Sen, Ben Kehoe, Animesh Garg, Seth McFarland, Sachin Patil, W. Douglas Boyd, Susan Lim, Pieter Abbeel, Ken Goldberg (UC Berkeley)
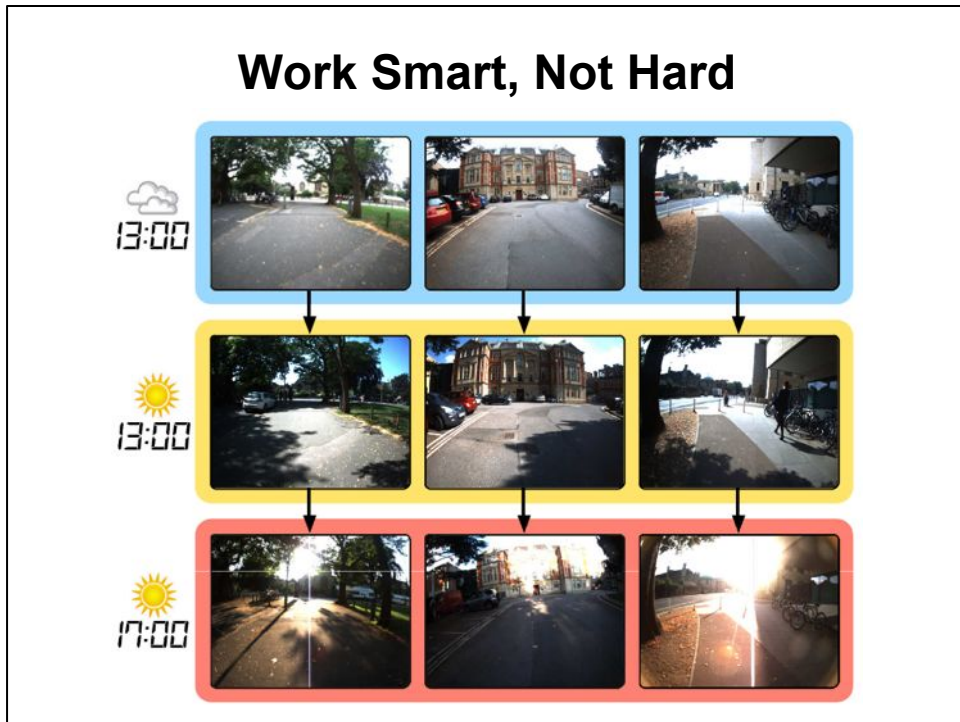
Work Smart, Not Hard: Recalling Relevant Experiences for Vast-Scale but Time-Constrained Localisation

- Chris Linegar, Winston Churchill, and Paul Newman (Oxford)
- Experience graphs: Appearance-based maps where you can have multiple nodes in the same geographic place because their appearance differs.
- Normally, to localise you need to test against all these experiences.
  - Search all neighbors to the previous experience and find a stereo match for the current image.
  - This is an anytime search that might be aborted before a good match is found.
- Instead, they learn a simple ranking policy to predict which experience to look at to produce a good localisation.
- Need to connect together experiences that occur at the same time of day, and use this memory to predict which experiences to use based on the current conditions.
- Enables long-term (perhaps life-long) autonomy.
- There's lots of other interesting work on experience graphs if you're interested.
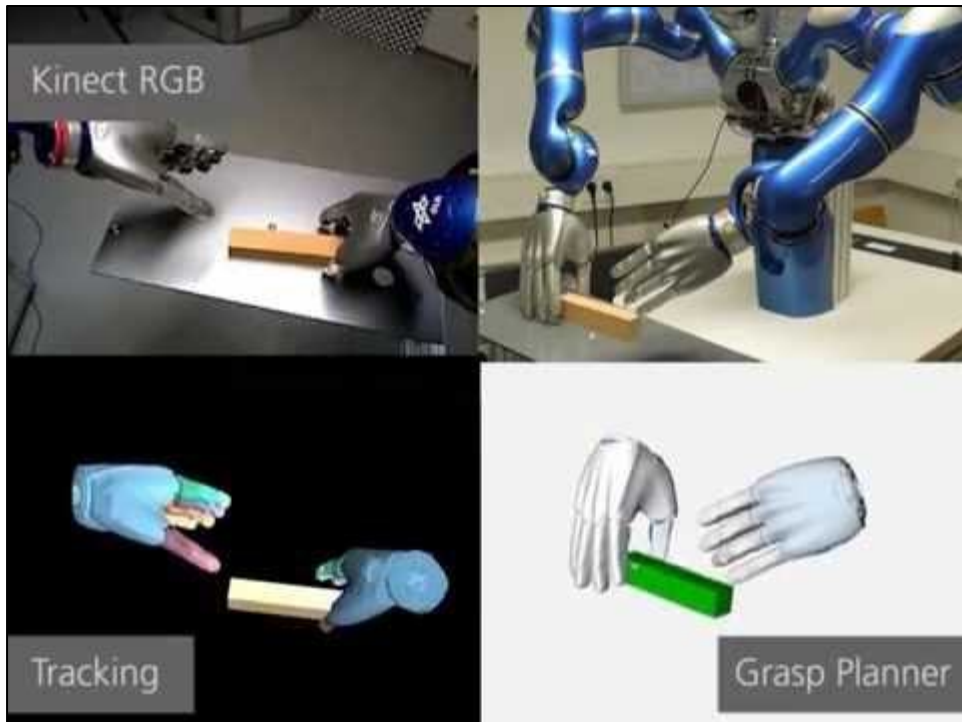
Depth-Based Tracking with Physical Constraints for Robot Manipulation

- Tanner Schmidt, Katharina Hertkorn, Richard Newcombe, Zoltan Marton, Michael Suppa, Dieter Fox (University of Washington)
- Adding physical constraints improves the performance of tracking.

Sequence-Level Object Candidates Based on Saliency
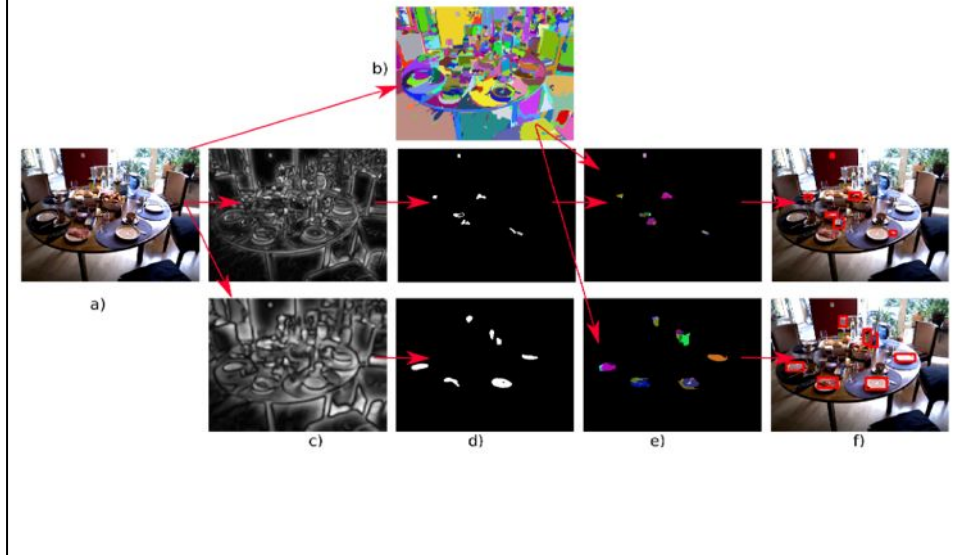
Sequence-Level Object Candidates Based on Saliency for Generic Object Recognition on Mobile Systems

- Esther Horbert, German M. Garcıa, Simone Frintrop, Bastian Leibe (RWTH Aachen)
- "Contrary to previous approaches, we address the candidate generation problem at the level of entire video sequences instead of at the single image level. This enables us to group candidates for similar objects and to automatically filter out inconsistent regions."
- They also have a novel per-frame proposal algorithm.
- Taken together with the video processing, they achieve higher recall and fewer false positives.
- Really nice to see someone incorporating video & motion priors—these are the advantages that robotics has that it should be incorporating

Reconstructed and segmented scene
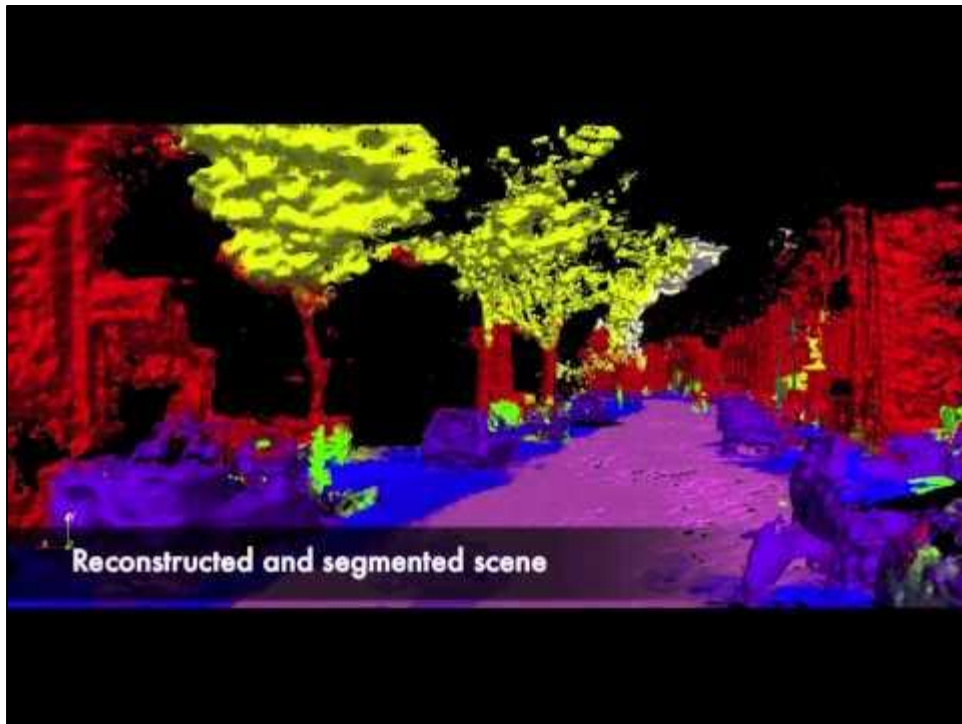
Incremental Dense Semantic Stereo Fusion for Large-Scale Semantic Scene Reconstruction

- Vibhav Vineet, Ondrej Miksik, et al. (Stanford, Oxford)
- "[We] present what to our knowledge is the first system that can perform dense, large-scale, outdoor semantic reconstruction of a scene in (near) real time."

Dynamic Multi-Heuristic A*

- Fahad Islam, Venkatraman Narayanan, and Maxim Likhachev (CMU)
- Bears similarity to simulated annealing
- Separate searches are done in parallel with different heuristics.
- But the searches can share states, so duplicate states are never expanded. Linear search time is achieved.
- Whenever some heuristic gets stuck, the other heuristics can help it escape that local minimum.
- Only one heuristic needs to be admissible, the others can be any arbitrary heuristic.

Towards a Data-Driven Approach to Human Preferences in Motion Planning
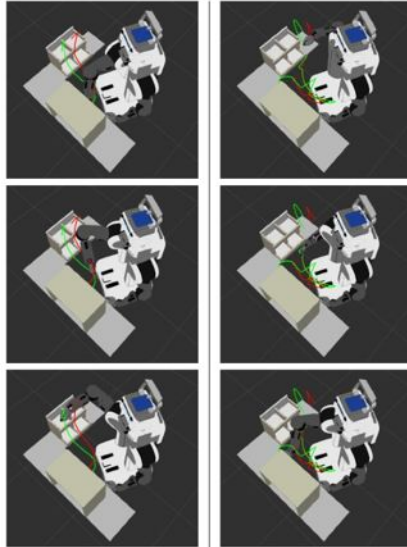
Towards a Data-Driven Approach to Human Preferences in Motion Planning

- Arjun Menon, Pooja Kacker, Sachin Chitta (SRI International)
- Had Mechanical Turkers rate trajectories by Efficiency, Elegance, Smoothness, and Overall
- Tried to match people's preferred trajectories using random forests and boosting
- Showed features which are useful for classifying good trajectories
- Would have been nice to see a comparison to optimal trajectories. In what situations do humans prefer non-optimal?
- I think this is what will need to be done in the future to plan legible, predictable, and human-like movements.
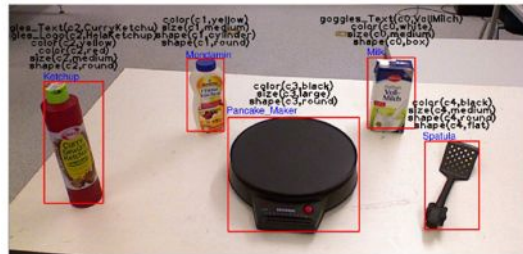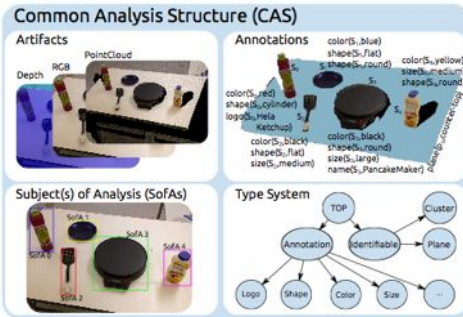
RoboSherlock: Unstructured Information Processing for Robot Perception

- Michael Beetz, Ferenc Balint-Benczedi, Nico Blodow, Daniel Nyga, Thiemo Wiedemeyer, Zoltan-Csaba Marton (University of Bremen, TUM)
- Best Service Robotics Paper
- Similar to RoboBrain, but makes a bit more sense to me
- Can be equipped w/ ensembles of perception algorithms
- Can request objects that have specific properties/attributes—acts like a query
- Example of "Unstructured Information Management", like Watson
- Multiple algorithms take unstructured data and annotate it—like multiple experts all adding their expert knowledge

Analysis Engine: Pose estimation from oriented 3D bounding boxes

$pose(c_i, T)$
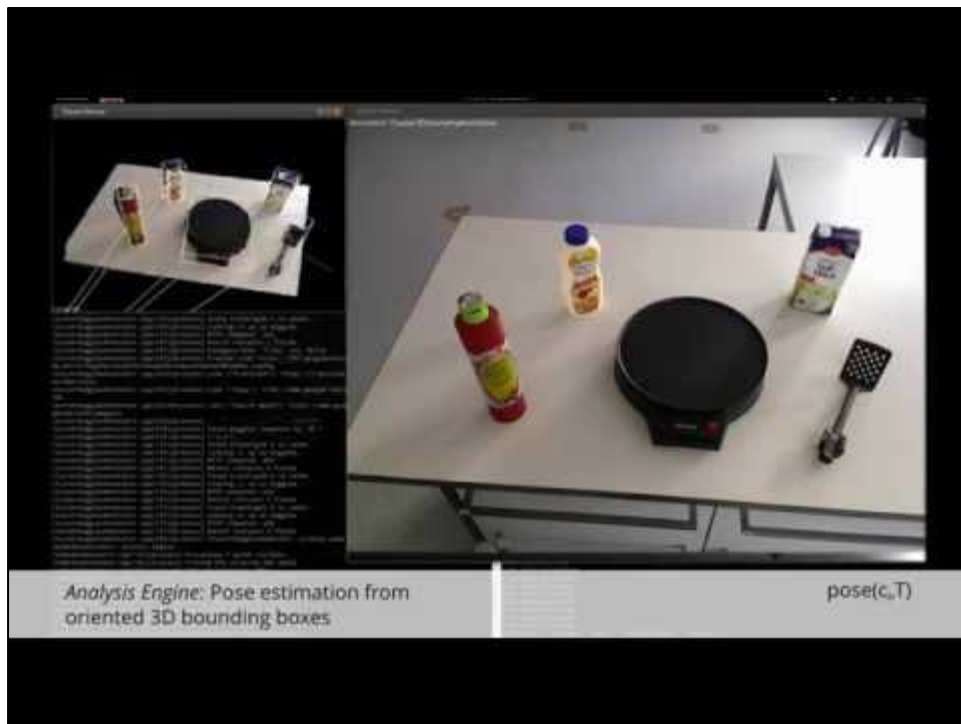
More RoboSherlock

Open-EASE -- a Knowledge Processing Service for Robots and Robotics/AI Researchers

- Michael Beetz, Moritz Tenorth and Jan Winkler (University of Bremen)
- Used by RoboSherlock
- Knowledge base that can be queried with Prolog
- Experiment trials are recorded in the knowledge base and lots of information about what happened during the trials & what the robot was perceiving can be recalled.

amazon Picking Challenge
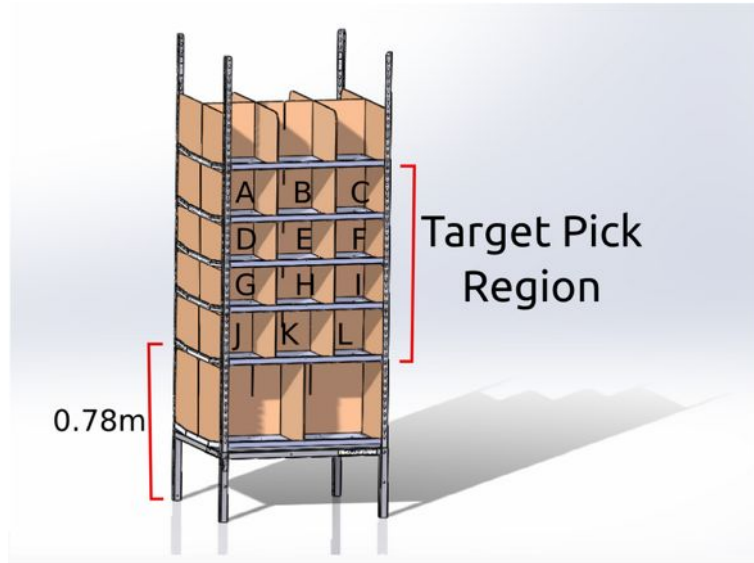
ICRA 2015, Seattle WA

Videos
- [Team RBO Original speed](#)
- [Team RBO Sped up x4](#)
- [Team MIT](#)

The challenge is over! Congrats to our winners: Team RBO (#1), Team MIT (#2), and Team Grizzly (#3)

amazon Picking Challenge
ICRA 2015, Seattle WA

oreo_mega_stuf	champion_copper_plus_spark_plug	expo_dry_erase_board_eraser	genuine_joe_plastic_stir_sticks
crayola_64_ct	mommys_helper_outlet_plugs	sharpie_accent_tank_style_highlighters	stanley_66_052 +3
cheezit_big_original	paper_mate_12_count_mirado_black_warrior	feline_greenies_dental_treats	elmers_washable_no_run_school_glue
first_years_take_and_toss_straw_cup +2	highland_6539_self_stick_notes	mark_twain_huckleberry_finn +3	kyjen_squeakin_eggs_plush_puppies +1
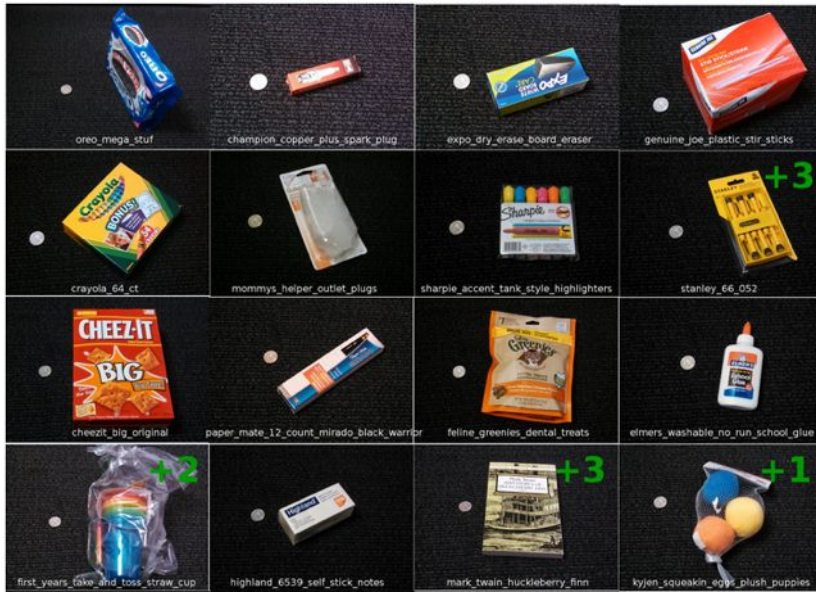
amazon Picking Challenge

ICRA 2015, Seattle WA

Moving a target item from a multi-item shelf bin into the order bin	+20 points
Moving a target item from a double-item shelf bin into the order bin	+15 points
Moving a target item from a single-item shelf bin into the order bin	+10 points
Target Object Bonus	+(0 to 3) points
Moving a non-target item out of a shelf bin (and not replacing it in the same bin)	−12 points
Damaging any item or packaging	−5 points
Dropping a target item from a height above 0.3 meters	−3 points

Robotic Vision Workshops

- Dieter Fox: The 100-100 Tracking Challenge
- Tim Barfoot: We shouldn't treat geometry and appearance as separate things.
- Michael Milford: Argues for event-based cameras + neuromorphic sensing.
- John Leonard: We're still a far way off from autonomous driving (where on the ROC curve do we need to land for it to be worth it?)

Other observations
- Vision is incredibly popular!
- Lots of papers doing object proposals / saliency segmentation
- Neuromorphic sensors are an interesting new paradigm (there was a half-day workshop specifically on this topic)

Motion Planning Workshops

- It's not all about sampling-based planning.
  - Some (like Marco Pavone) are looking at deterministic sampling schemes (basically equivalent to grid-based search).
  - Some (like Maxim Likhachev) use heuristic search-based planning, not roadmaps or sampling.
- Dan Halperin: Please work on non-Euclidean spaces (non-holonomic systems).
- Kris Hauser: Please work on high-D cases. Even with sampling-based methods, this still requires clever tricks to improve performance.