Real-time Perception meets Reactive Motion Generation

Presenter: Omeed Tehrani

Publication Date: 01/22/2018

Date: Thursday, November 3rd

Authors:





Daniel Kappler
Moonshot Factory

Franziska Meier
Facebook AI

Jan Issac
Lula / NVIDIA

Jim Mainprice
Univ. of Stuttgart

Cristina Cifuentes
Max Plank Institute





Manuel Wüthrich
Max Plank Institute

Vincent Berenz
Max Plank Institute

Stefan Schaal
Max Plank / UCLA

Nathan Ratliff
Lula Robotics /
NVIDIA / Max Plank

Jeannette Bohg
Stanford Univeristy

Motivation and Main Problem

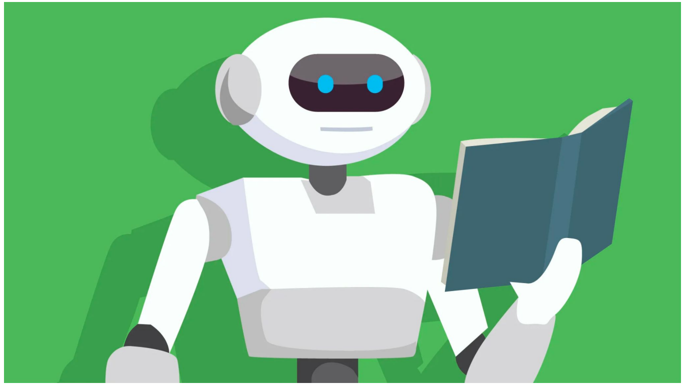

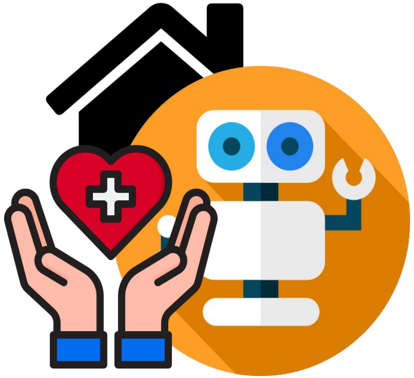

General robot autonomy has been envisioned to be applied to things such as healthcare, in-home and workplace assistance, disaster relief, etc.

For this to be successful, robots must be able to manipulate and grasp objects in unstructured environments when **uncertainty** is present.

One key is perceptual awareness.

An essential capability as humans is to understand and perceive objects around us in real-time, and make decisions quickly. We have a *sense* of our surroundings.

Creating a perceptually aware system can lead to **natural reactions** in **uncertain** scenarios. A reaction could encompass split-second decisions such as a change of motion on a task.

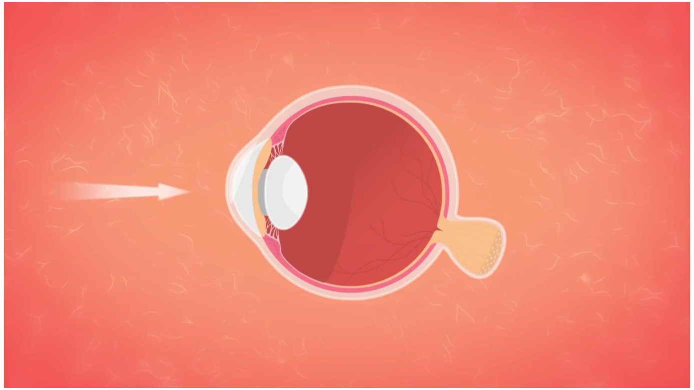

Autonomous Driving

1. For some systems such as autonomous vehicles, integration of perception with motion has worked quite well. This is because it is a **low-dimensional control** problem.
2. At the time of this paper, was at the brink of becoming a consumer end product.

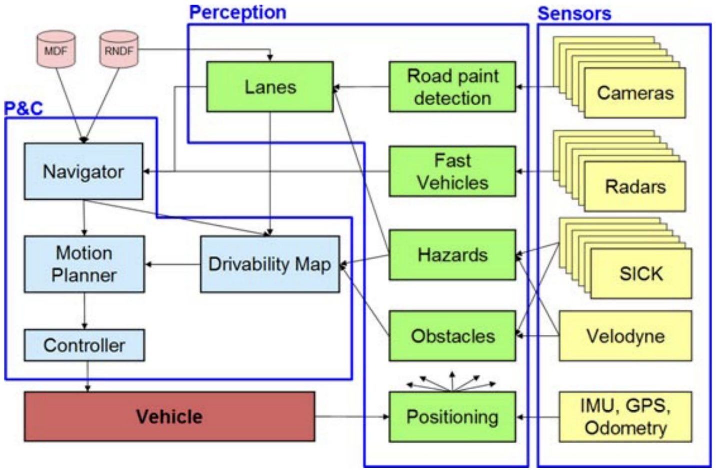

So why have we not integrated perceptual feedback and motion generation to general robotic systems to enable reactive behavior in the presence of uncertainty?

The Key Problem

With tasks that require controlling high degrees of freedom (DoF) + a physical interaction with the environment….

The question is **NOT** *why should we*, it is how can we effectively integrate perception and motion generation to these systems?

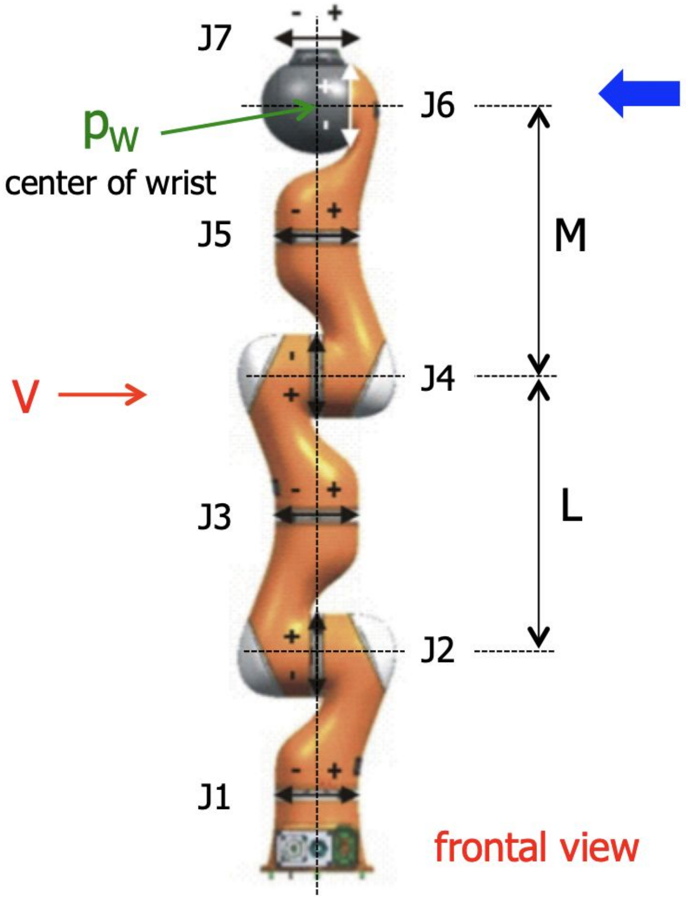

Key Inspirations and Related Work

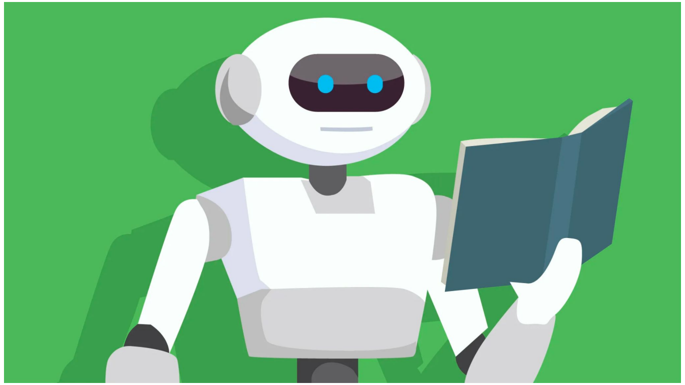

2015 Amazon Picking Challenge (APC):

Autonomously picking 12/25 objects and placing into storage container within 20 minutes. Robot has knowledge of which objects are contained in each of the shelves 12 bins, but not exact arrangement in the bin.

Fig. 2 The 25 objects from the challenge

2015 DARPA Robotics Challenge (DRC):

3 stage competition to develop robots for assisting humans in natural and man-made disasters. In final stage:

1. Robot drives through an obstacle course to dest and exits car.
2. Enters a building through door
3. Turns a valve
4. Cuts a hole in a wall using a power tool, navigates over debris or rubble
5. Finishes by climbing stairs

Lessons Learned

Trade-off between Planning vs. Feedback

Planning: searches within a world model to find verifiable
Feedback: can help reduce *uncertainty*, help find *local solutions,* and is less computationally expensive.
Lesson: If the manipulation task does not require global path planning, it is recommended to explore perceptual feedback.

Trade-off between Modularization vs. Integration

Incorrect modularization of a problem can lead to unnecessary complexity, therefore, until we can formulate clear modularity for unstructured environments, building **tightly integrated systems** is essential.

Relevant/Related Work

[12]: R. A. Brooks, "Elephants don’t play chess", Robot. Auton. Syst., vol. 6, no. 1/2, pp. 3-15, 1990.

[17]: S. Levine, C. Finn, T. Darrell and P. Abbeel, "End-to-end training of deep visuomotor policies", J. Mach. Learn. Res., vol. 17, no. 1, pp. 1334-1373, 2016: learning motion policies directly from perceptual feedback in form of raw camera images and the system joint state, e.g.,

[5]: C. Eppner et al., "Lessons from the amazon picking challenge: Four aspects of robotic systems building", Proc. Robot. Sci. Syst., pp. 4831-4835, 2016: **Amazon Picking Challenge**

[7]: M. Wüthrich, P. Pastor, M. Kalakrishnan, J. Bohg and S. Schaal, "Probabilistic object tracking using a range camera", Proc. 2013 IEEE/RSJ Int. Conf. Intell. Robots Syst., pp. 3195-3202, 2013: **previous work used in this work.**

[8] C. Garcia Cifuentes, J. Issac, M. Wüthrich, S. Schaal and J. Bohg, "Probabilistic articulated real-time tracking for robot manipulation", IEEE Robot. Autom. Lett., vol. 2, no. 2, pp. 577-584, Apr. 2017:

[9] N. Ratliff, M. Toussaint and S. Schaal, "Understanding the geometry of workspace obstacles in motion optimization", Proc. 2015 IEEE Int. Conf. Robot. Autom., pp. 4202-4209, 2015.

[10] J. Mainprice, N. Ratliff and S. Schaal, "Warping the workspace geometry with electric potentials for motion optimization of manipulation tasks", Proc. 2016 IEEE/RSJ Int. Conf. Intell. Robots Syst., pp. 3156-3163, 2016.

[31] N. D. Ratliff, J. Issac and D. Kappler, "Riemannian motion policies", Jan. 2018: **previous work used in this work.**

[28] R. B. Rusu and S. Cousins, "3d is here: Point cloud library (PCL)", Proc. 2011 IEEE Int. Conf. Robot. Autom., pp. 1-4, 2011: **used in this work**

Key Ideas from this body of work:

1. Postulations since early 90’s that integration of real-time perception and reactive motion is beneficial.
2. **This paper is unique because:**
    a. experimental scenarios are complex due to environment dynamics linked to general robot autonomy.
    b. performs quantitative evaluation on this level of integration.
3. They leverage a range of past work on visual tracking and motion planning to make this system work.
4. No teleoperation (DRC winner used it).

The Problem Setting

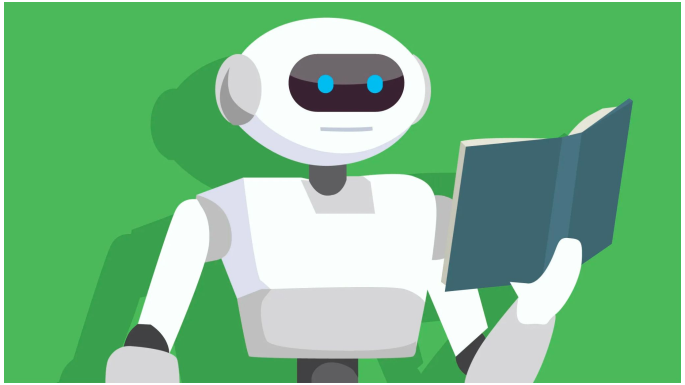

Approach and Goal of Research

Goal: This is a systems paper. The goal is to use empirical evidence to quantify the benefit of integrating real-time perceptual feedback and reactive motion generation in **dynamic** manipulation scenarios for high DoF systems.

Approach: Inspired by lessons from 2015 competitions, *three architectures* are chosen. Their instantiations and *feedback components* are defined, compared with experimentation, and then evaluated with discussion.

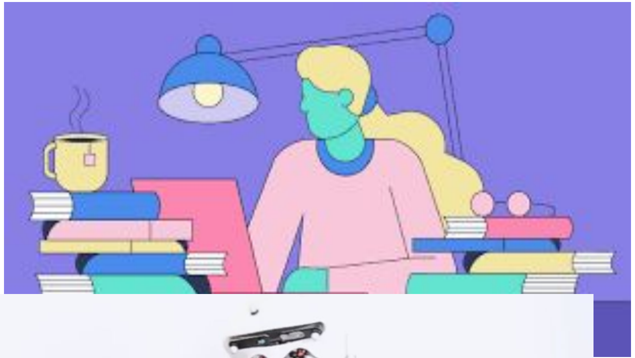

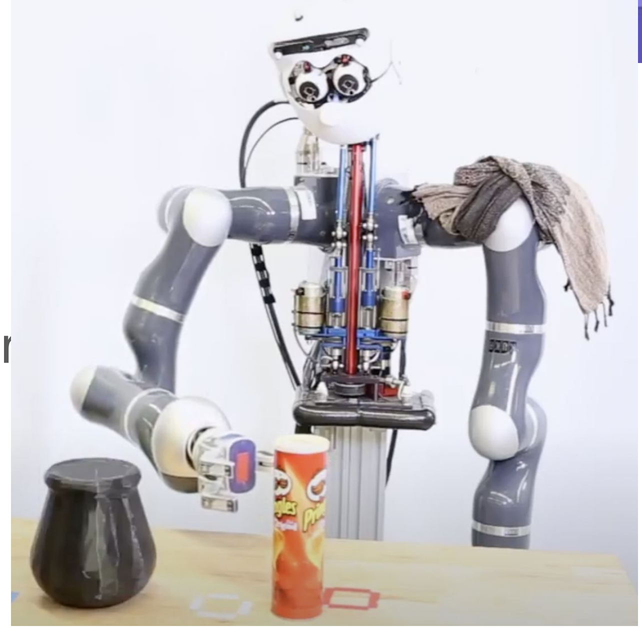

Architectures

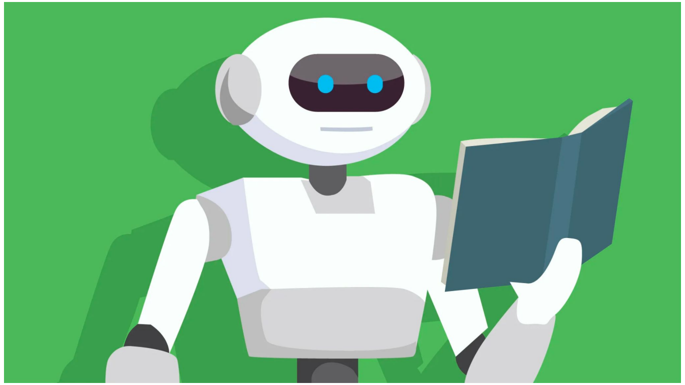

The Diagram

This diagram represents **3 time steps**. Shows how information flows between the perception and the motion generation modules.

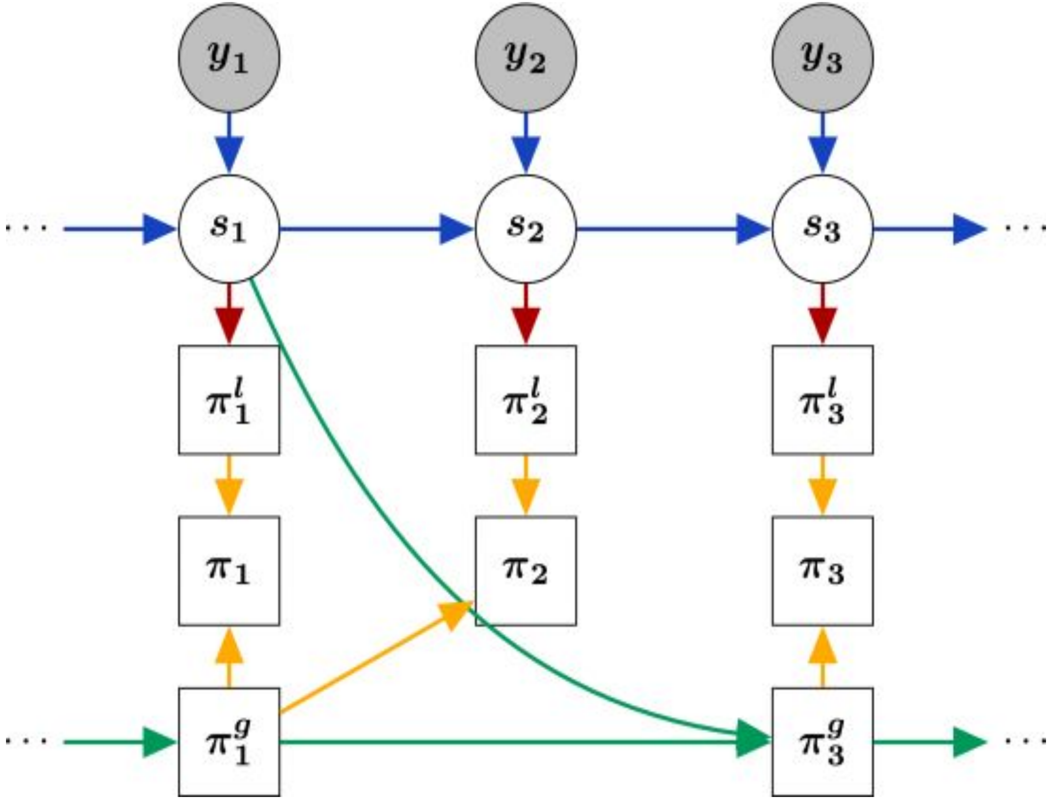

Visualization of Diagram

one time step

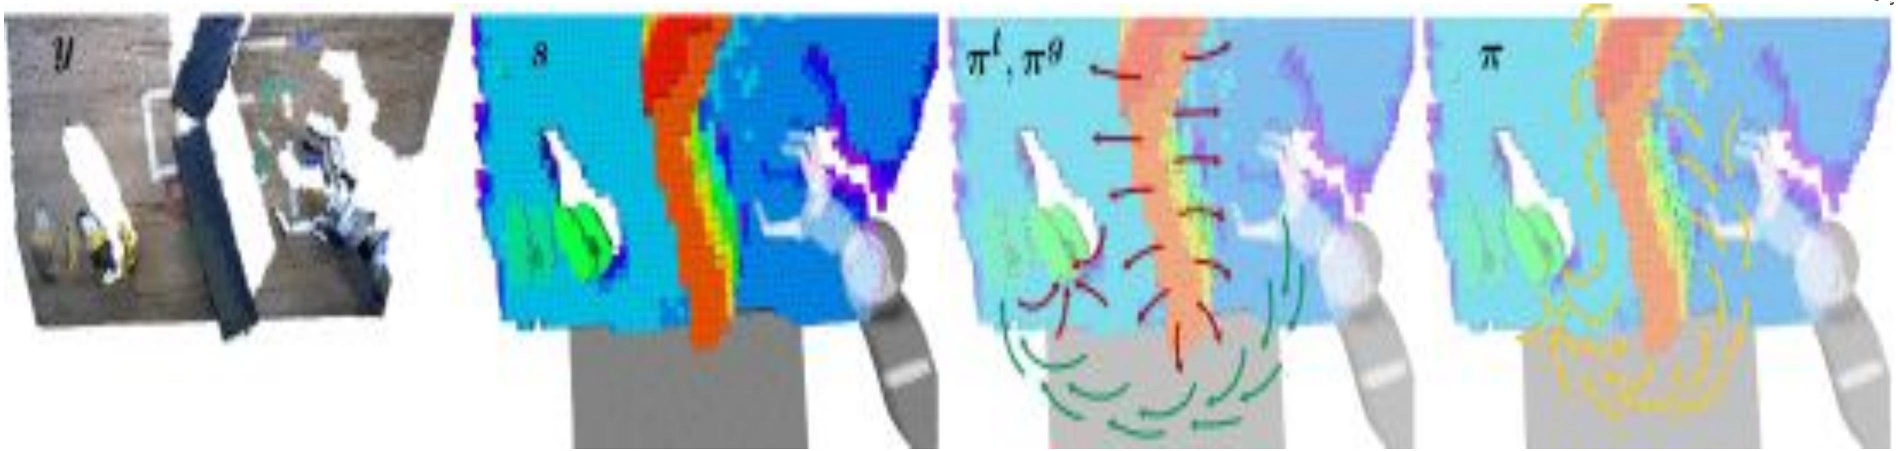

y: observed sensory input that helps us continuously infer s.

sensor information is integrated into some world state, represented as **s**.

π (l) → locally reactive
π (g) → globally reactive
π → reactive planning

Sense-Plan-Act (Architecture #1)

1. Perception models the environment (**depicted by blue arrow**)
2. Motion planner finds an optimal, collision-free path (**green arrow**) that is tracked by a stiff/accurate controller.

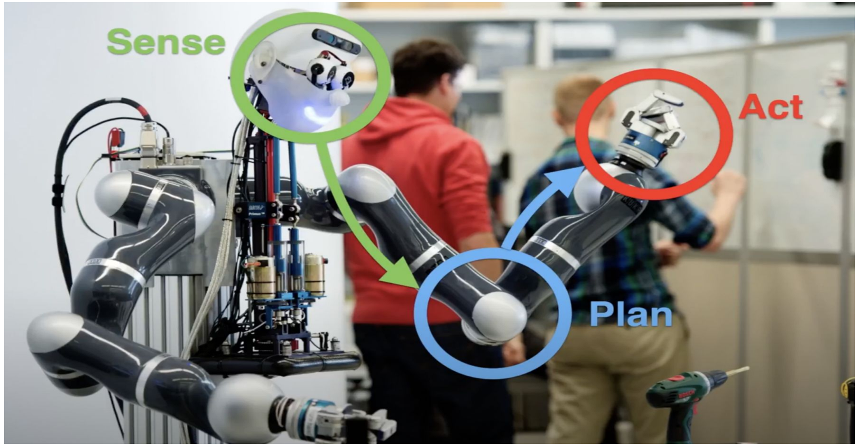

- Visual feedback is only is considered at the beginning of the task.
- This is a **strongly modularized** architecture. Has advantage of simple and solvable subproblems.
- Popular for high DoF systems.

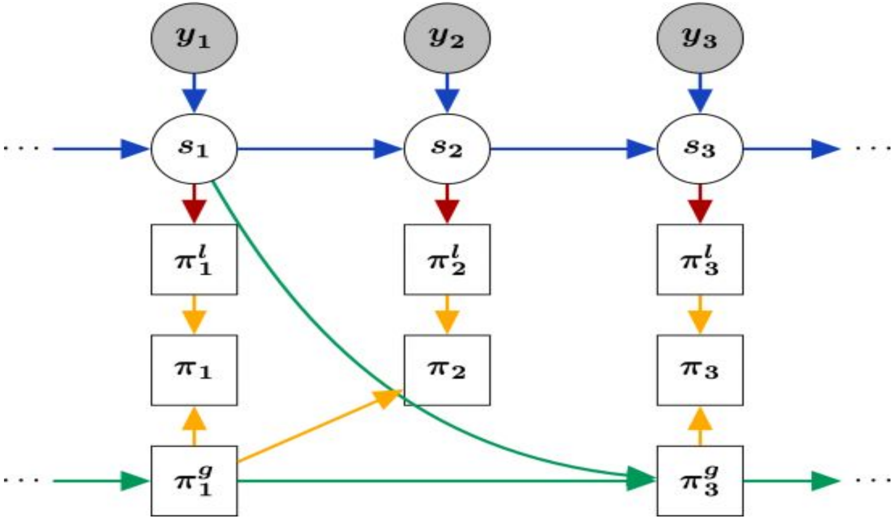

Locally Reactive Control (Architecture #2)

1. Perception models the local geometry around the current manipulator pose.
2. Computes a **local policy** to get next control command (**red arrow**).

- Relies on pure visual feedback.
- Local policies enable reactive motion behavior and robustness to uncertainty.
- Unfortunately, **susceptible to getting stuck in a local minima.** You will see this in the upcoming video.

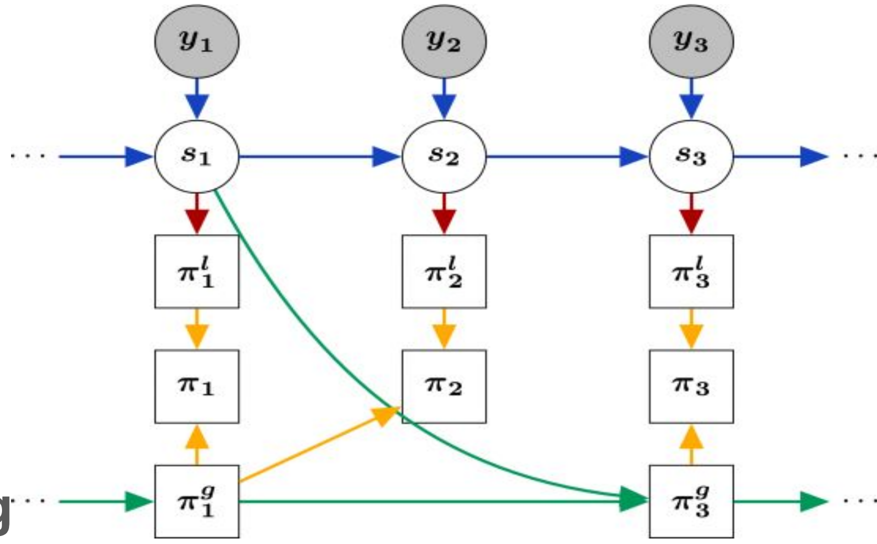

Now, envision a hybrid system, that combines these two to have both **local control** mixed with **motion planning**.

Reactive Planning (Architecture #3)

- In this architecture, the idea is to mix reactive **motion planning** and **locally reactive control**.

1. Motion representations of global (**green**) and local (**red**) policies are merged.
2. Motion generation policy is created (**yellow**).

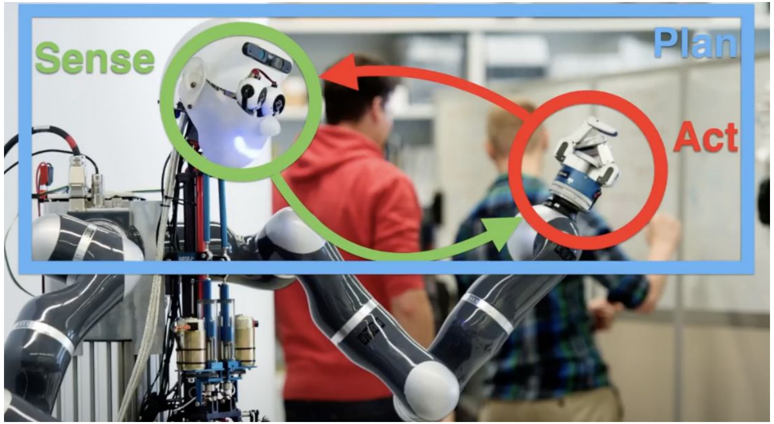

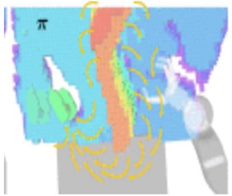

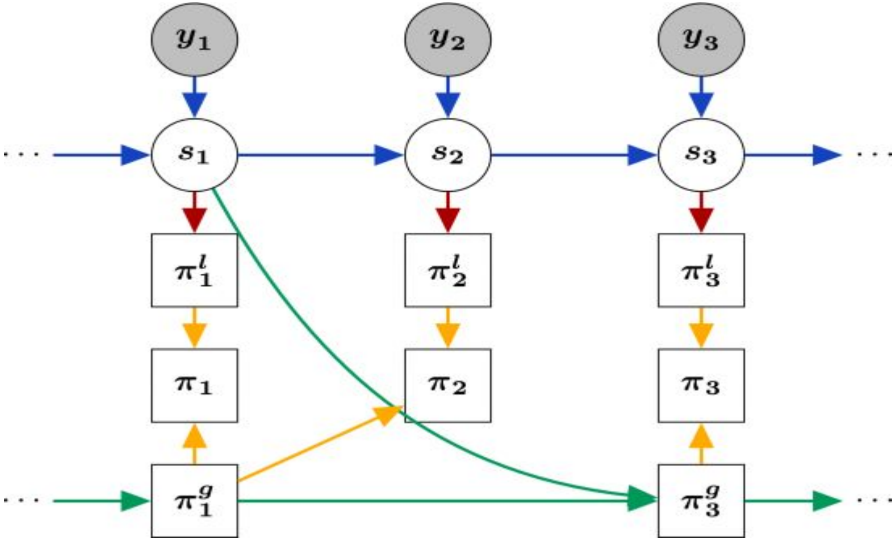

TEXAS
The University of Texas at Austin

Let's recap. We've defined the architectures. But to perform a fair quantitative evaluation in this research, each proposed architecture relies on the exact **same feedback components** for real-time perception and reactive motion generation. Let's discuss how those components are defined, then watch everything in action with videos.

Feedback Components

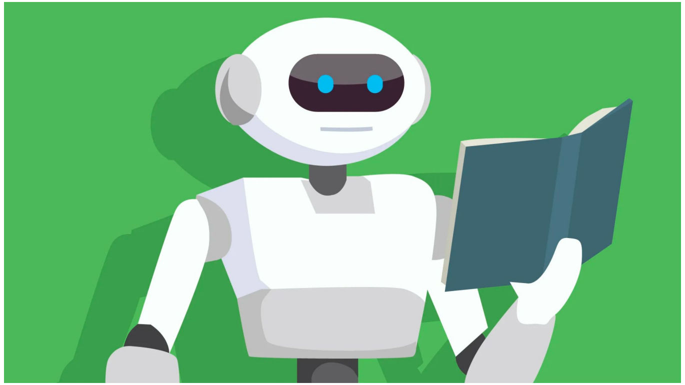

Important Note: We will discuss all feedback integration methods used in this system. But, the goal is to serve as a requirement outline for any alternative methods created and used in the future.

Preface to Visual Tracking of Target Objects (Feedback Component #1)

Manuel Wüthrich

Stefan Schaal

arXiv > cs > arXiv:1505.00241

Computer Science > Robotics

[Submitted on 1 May 2015]

Probabilistic Object Tracking using a Range Camera

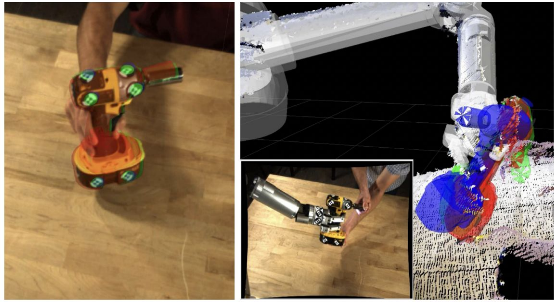

Fig. 6. The fiducial tracker uses camera images to track the round fiducial markers (left). The obtained pose estimate (green) is used as a baseline comparison. The forward kinematics model (blue) is misplaced because of deliberate disturbances (small image). The proposed approach achieves accurate tracking (red) even in the presence of occlusions.

Big Ideas:

1. Probabilistic approach (Bayesian Networks + Rao-Blackwellised particle filtering) for object tracking using range cameras
2. **robust to occlusion**
3. **fast for real-time tracking** on a singular core. [7]

The Algorithm (extra information)

- From the previous time step we have a set of particles $\{r_{1:t-1}^{(l)}\}$ distributed according to $p(r_{1:t-1}|z_{1:t-1}, u_{1:t-1})$ and for each of these particles we know $p(o_{t-1}^{i}|r_{1:t-1}^{(l)}, z_{1:t-1}, u_{1:t-1})$. Furthermore we know the control u_t which is applied during the current time step and we observe a depth image z_t.
- For each particle in $\{r_{1:t-1}^{(l)}\}$
  - We draw a sample $r_t^{(l)}$ from the pose process model $p(r_t|r_{t-1}^{(l)}, u_t)$.
  - We compute the likelihood $p(z_t|r_{1:t}^{(l)}, z_{1:t-1})$ according to Eq. 10.
  - We update the occlusion probabilities $p(o_t^{i}|r_{1:t}^{(l)}, z_{1:t}, u_{1:t})$ for each pixel according to Eq. 7.
- We resample the particles according to the likelihoods. We thus now have a set of particles $\{r_{1:t}^{(l)}\}$ distributed according to $p(r_{1:t}|z_{1:t}, u_{1:t})$, and the corresponding occlusion probabilities $p(o_t^{i}|r_{1:t}^{(l)}, z_{1:t}, u_{1:t})$. We can now go to the next time-step and repeat the procedure above.

Visual Tracking of Target Objects (Feedback Component #1)

1. **Goal:** to effectively estimate pose of objects the robot seeks to manipulate.
2. **Past Issues:** some approaches **are not fast** for real-time tracking and direct contact occlusion.
3. **Mitigation Approach:**
    a. Leverage **probabilistic methodology [7]** to visually track the target object.
    b. **Some Facts:** Assumes knowledge of the 3D target object, represented as triangle meshes, depth images are taken as input and compressed into 6 DoF object poses.

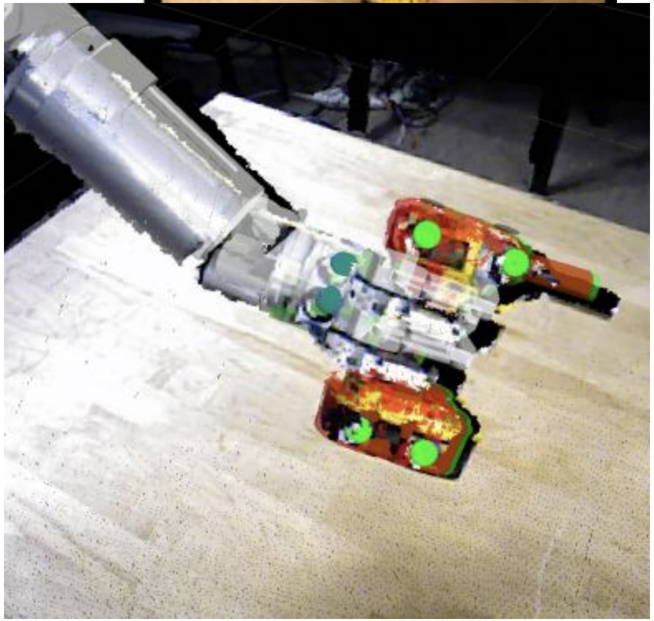

Preface to Visual Robot Tracking (Feedback Component #2)

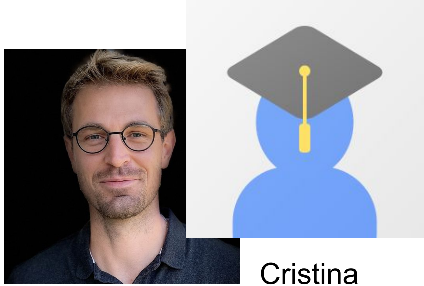

Manuel Wüthrich

Cristina Cifuentes

Stefan Schaal

Jeannette Bohg


arXiv > cs > arXiv:1610.04871

Computer Science > Robotics

[Submitted on 16 Oct 2016 (v1), last revised 25 Nov 2016 (this version, v2)]

Probabilistic Articulated Real-Time Tracking for Robot Manipulation

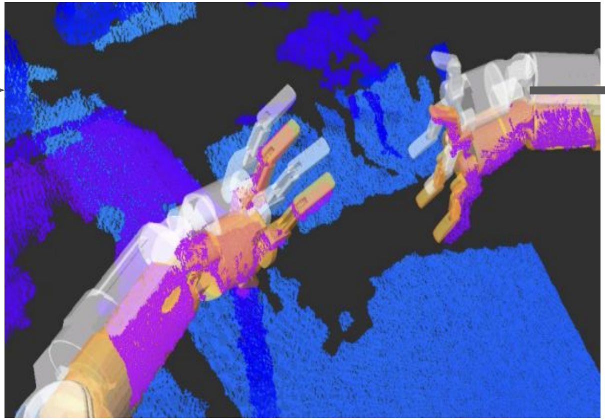

Big Ideas:

1. Probabilistic filtering approach (recursive bayesian estimation) that fuses joint measurements with depth images to track robotic arm. [8]
2. Robust, precise and keeps tractability in mind.
3. Computationally efficient and good for real-time implementation due to Coordinate Particle Filtering and parallelization of depth image likelihoods. [16] [20].

Visual Robot Tracking (Feedback Component #2)

1. **Goal:** precise positioning of robot arm with respect to its sensing of environment & target object.
2. **Past Issues:** real robotic platforms have inaccuracy with kinematic modeling and joint parameters, leading to incorrect predictions of **end-effector pose** relative to camera.
3. **Mitigation Approach:**
   a. Continuously estimate *robot configuration, object pose and workspace geometry at 1 KHz* relative to camera on head of robot.
   b. Estimations produced through fusion of *depth images* and *joint angles* using **probabilistic methodology [8]**.

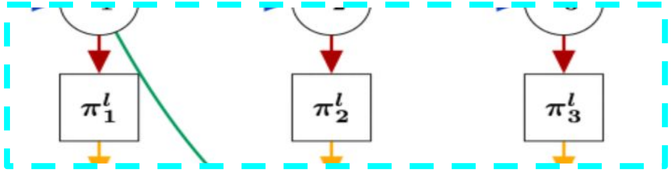

Modeling Unstructured Workspace Obstacles (Feedback Component #3)

1. **Goal:** for collision free motion, robot needs to effectively be aware of the geometry of environment.
2. **Past Issues:**
   a. Octomaps (open-source framework to generate volumetric 3D environment models) uses Octotrees for mapping.
   b. This is quite **computationally expensive**. Does not enable reactiveness [27].

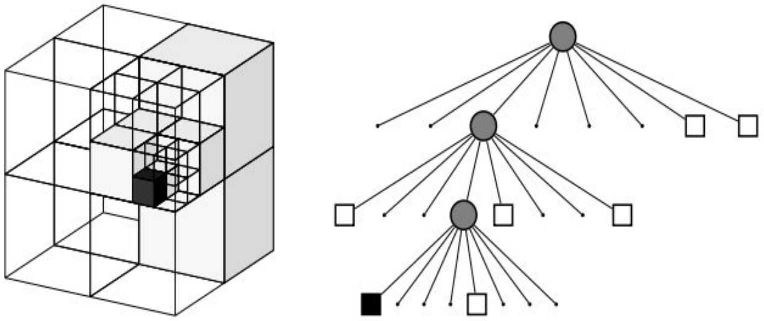

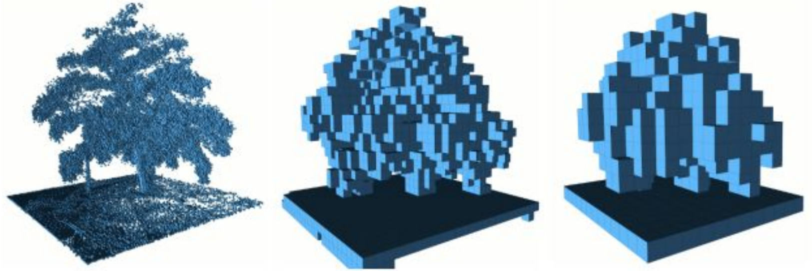

Modeling Unstructured Workspace Obstacles Cont. (Feedback Component #3)

Mitigation Approach:

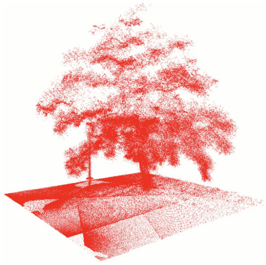

Transform point clouds into world coordinate points.

Remove all points outside of the robots workspace.

Filter for statistical outliers using point cloud library [28]

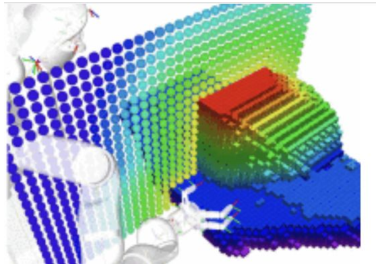

Take this transformed point cloud and convert to binary occupancy grid

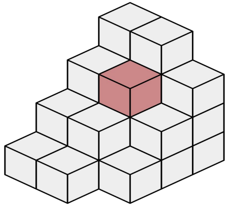

Voxels [27]

1. Point clouds processed into voxels, representing a voxelgrid.
2. Set to *empty* for robot arm and tracked object. Empty voxels will not be considered in the octotree and will not be part of the collision mappings.
3. Occluded regions are set **occupied** using ray casting

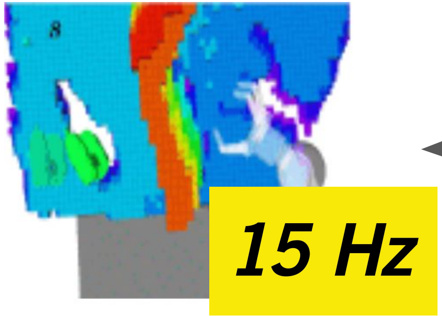

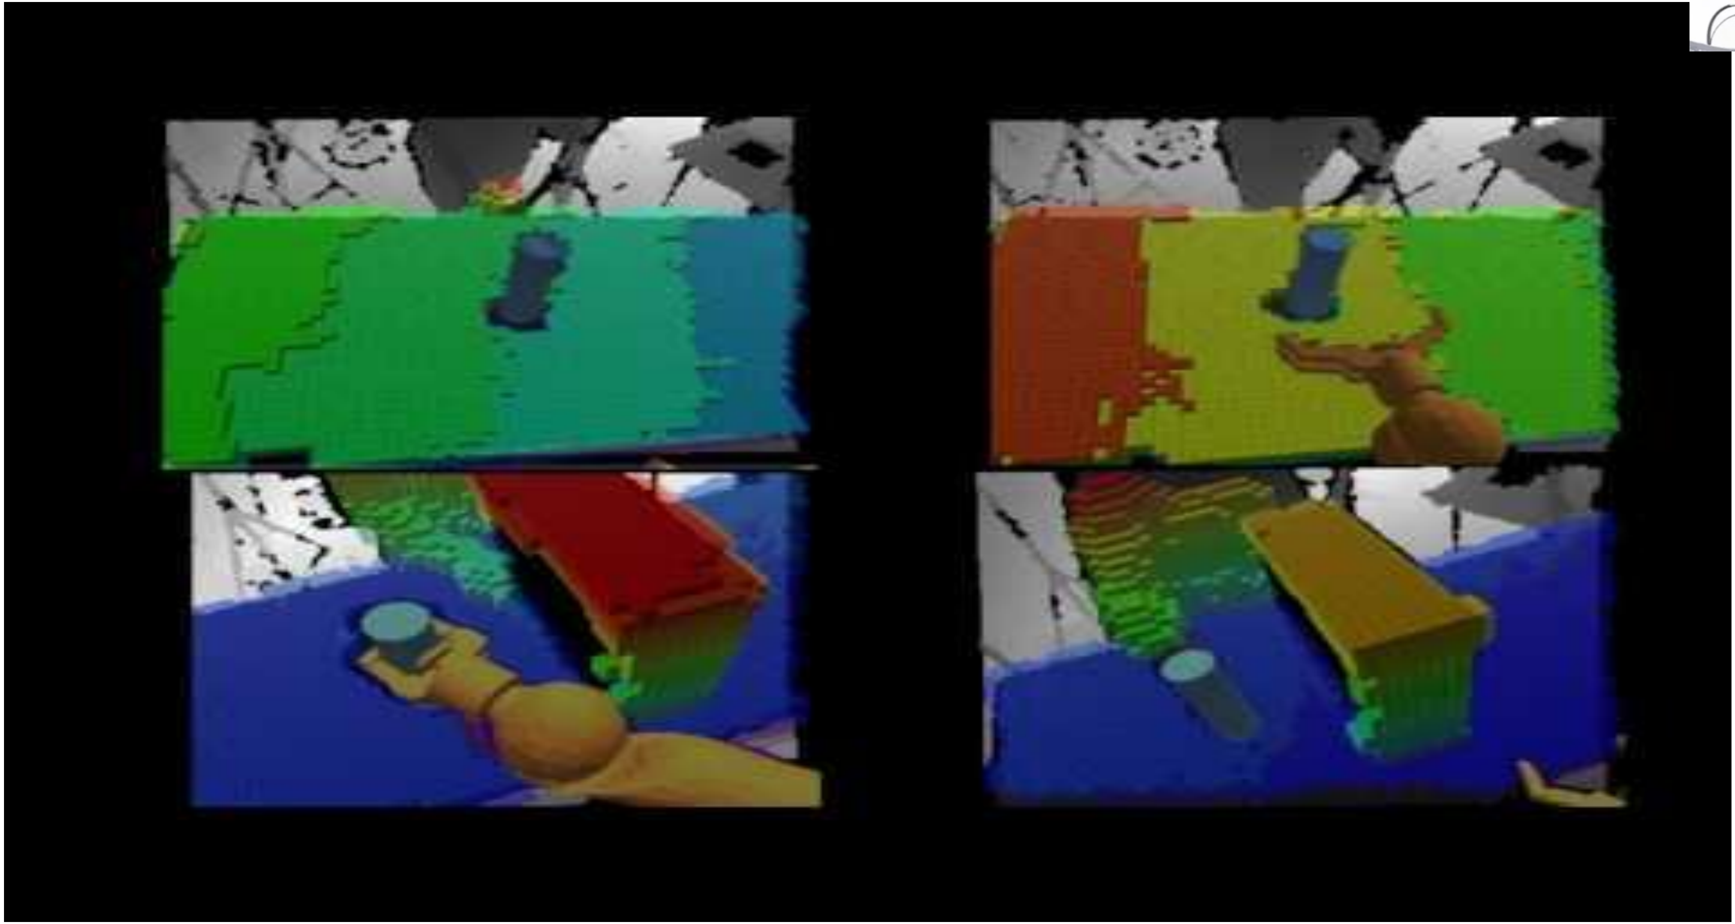

Alright, so we've discussed the architectures and how the environment is perceived by the robot, but how do we perform collision checking and generate motion (as seen in the previous video)?

Signed Distance Fields (SDFs)

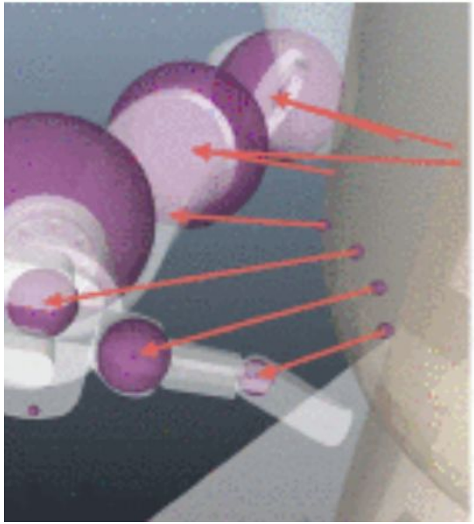

SDFs can be used to describe target object, table, obstacles, robot, etc. SDFs also allow to define proper **Riemannian metrics** for measuring path length in workspaces populated by obstacles (used by RP)

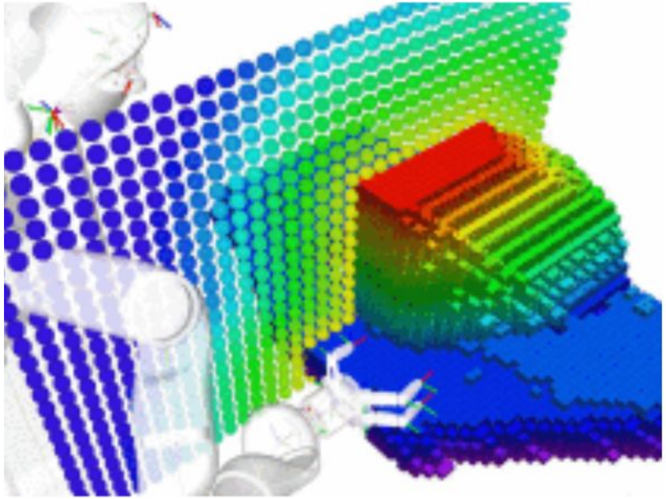

They are computed using analytical formations and distance propagation. Simple shapes enable this.

SDFs are positioned based on:

1. 6 DoF pose estimations from probabilistic methodology [7].
2. forward kinematics and joint angle estimations from probabilistic methodology [8].
3. SDFs are combined to produce collision free movement.

Theory of Motion Generation

RMP Framework: New approach to representing and transforming motion policies that preserves their geometry, leading to an optimal control system.

1. Individual controllers can be myopically designed to control only a small portion of the problem where the geometry is well understood.

RMP: Second-order dynamical system (acceleration field or motion policy) is coupled with a corresponding **Riemannian metric** [31].

Riemannian metric: defines which directions in the space the policy cares about most.

Theory of Motion Generation (cont.)

We can **combine RMPs** to solve the problem of **motion generation.**

Jacobian matrix of differentiable mappings

$$\ddot{\mathbf{q}}^{\text{desired}} = \arg\min_{\ddot{\mathbf{q}}} \frac{1}{2} \sum_{i=1}^{n} ||\ddot{\mathbf{x}}_i^{\text{desired}} - \mathbf{J}_i \ddot{\mathbf{q}}||^2_{\mathbf{A}_i}$$

desired motion policy in our configuration space

BIG IDEA: motion policy maps positions and velocities to accelerations, while the metric captures the directions in the space important to the policy [31]

Riemannian metric

Theory of Motion Generation (cont.)

We can **combine RMPs** to solve the problem of **motion generation.**

$$\text{RMP}\ (\mathbf{f}_i, \mathbf{A}_i)$$

$$\ddot{\mathbf{q}}^{\text{desired}} = \arg\min_{\ddot{\mathbf{q}}} \frac{1}{2} \sum_{i=1}^{n} ||\ddot{\mathbf{x}}_i^{\text{desired}} - \mathbf{J}_i \ddot{\mathbf{q}}||_{\mathbf{A}_i}^2$$

Jacobian matrix of differentiable mappings

desired motion policy in our configuration space

combine all our desired accelerations while taking **associated metrics** into account.

desired acceleration vector fields

Riemannian metric

Motion Generation

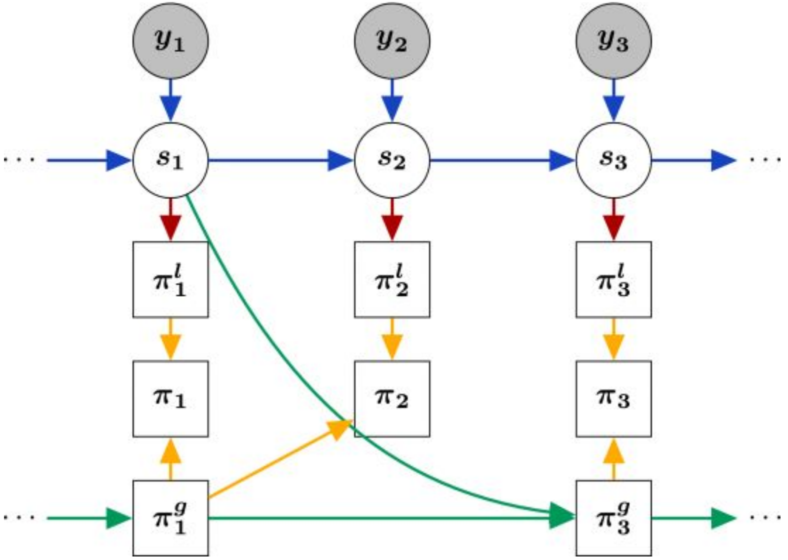

Locally Reactive Control →

Basis: Collision and obstacle avoidance controllers are combined using RMP equation to create reactive behavior in local geometry.

Reactive Planning →

Basis:

1. Uses a RieMO [9] that continuously runs, tracking local minimum based on feedback → integrates info over 3 second time horizon.
2. Optimizer summarizes its policies as Linear Quadratic Regulators (LQRs), kinematically.
3. Integrated with other controllers using **motion generation module.**
4. Continuous optimization is slower than locally reactive control, so to mitigate delays, LQR is sent to **an optimal region** to find either global or local policy, whichever is good enough.

Experimental Setup

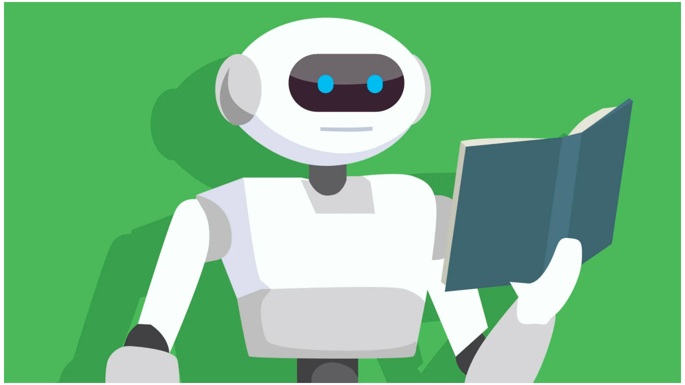

Compare 3 architectures: locally reactive control and sense-plan-act (baselines), reactive planning

Planning time: overall planning time of sense-plan-act is limited to 2 seconds for all experiments

Hardware Setup: fixed-base, manipulation platform equipped w

1. Two 7-DoF Kuka LWR IV arr
2. Three fingered Barrett Hand
3. RGB-D camera (Asus Xtion) an active humanoid head.

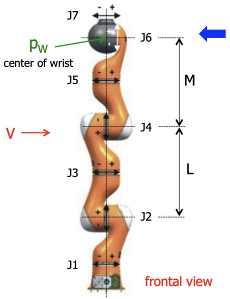

executed on a PC that runs Xenomai, which is a real-time framework for Linux.

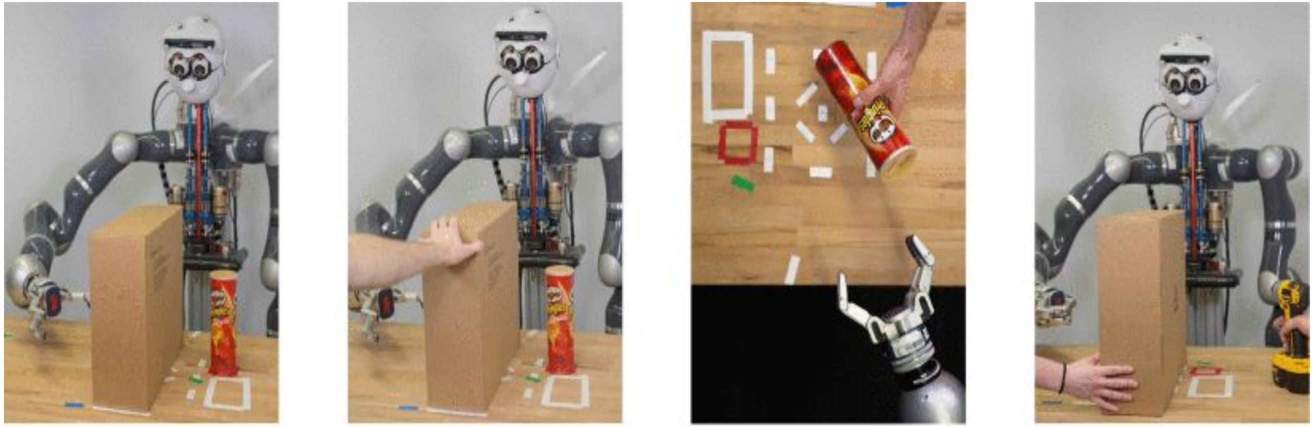

(a) (b) (c) (d)

4 experimental scenarios:

(a) Static pick and place. (b) Dynamic pick and place. (c) Dynamic grasping. (d) Dynamic pointing.

Pick and Place in Static Environment

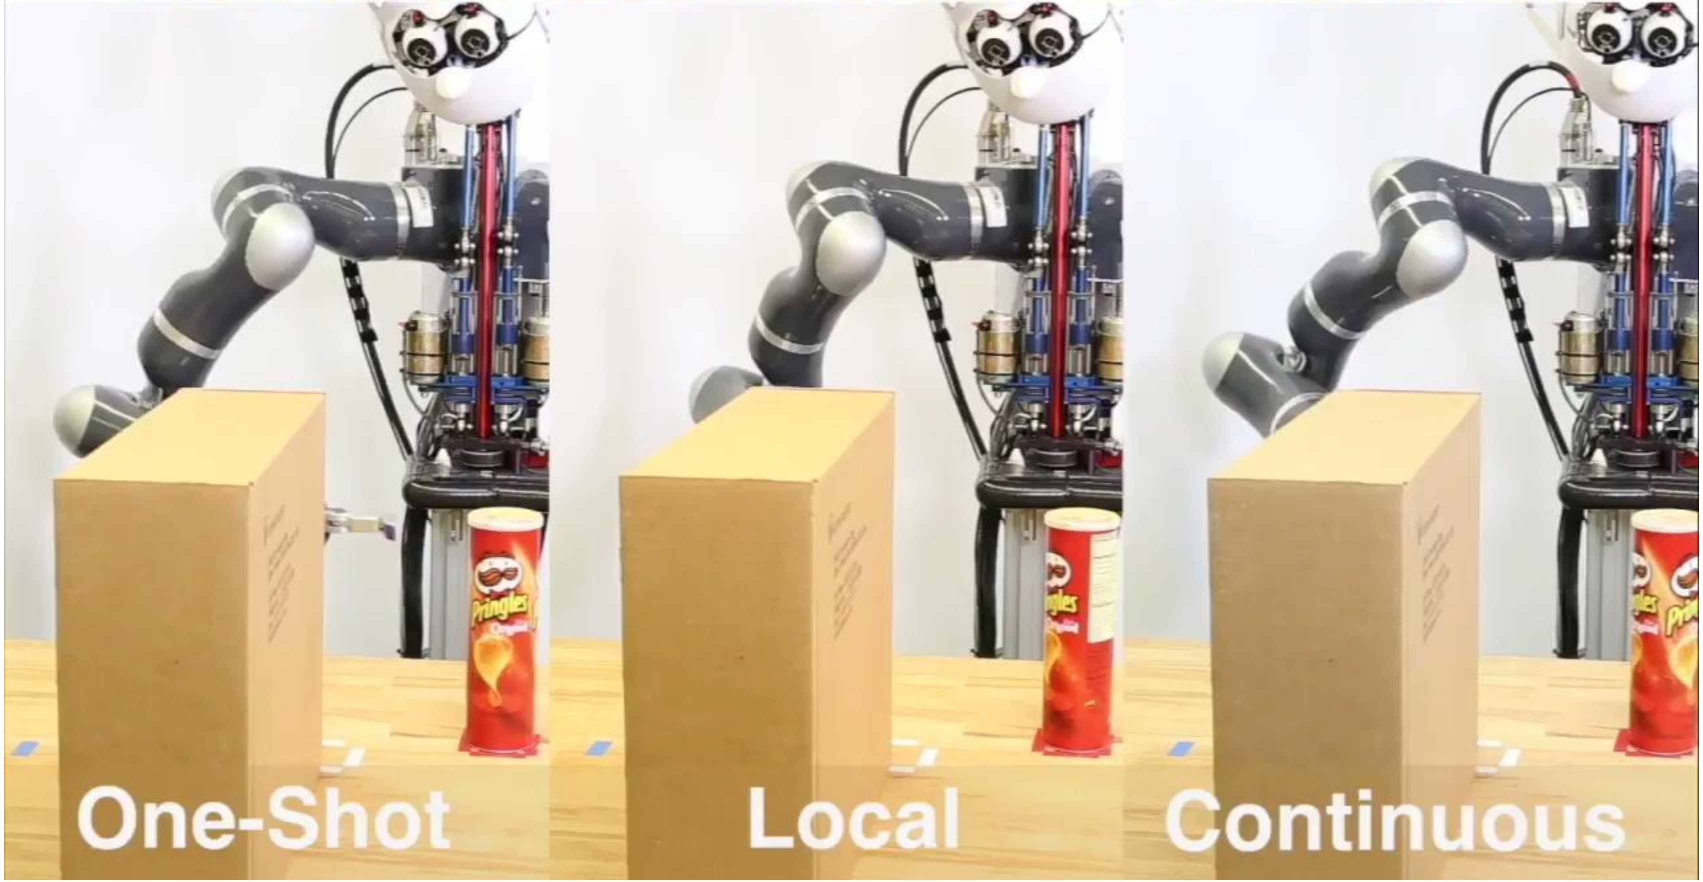

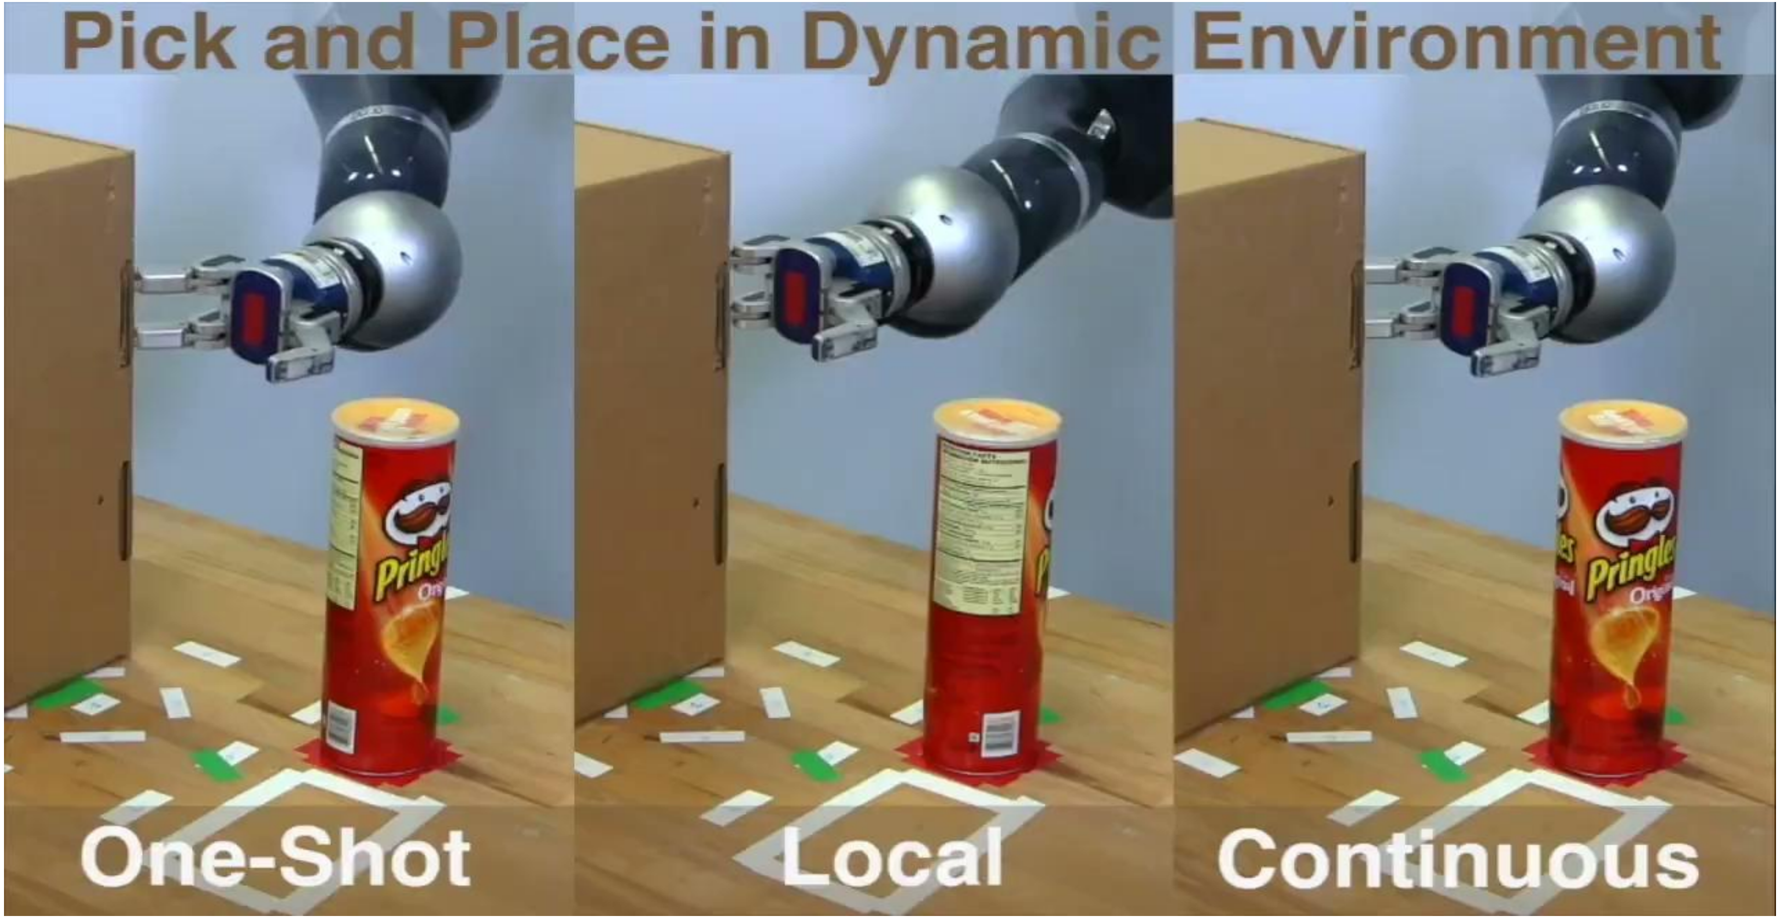
Pick and Place in Dynamic Environment
One-Shot
Local
Continuous

Static and Dynamic Results / Discussion

Difficulty	l. react. c.	react. pl.	s-p-a
static -10	100% (3)	100% (3)	100% (3)
static 0	100% (3)	100% (3)	100% (3)
static 5	67% (3)	100% (3)	100% (3)
static 10	33% (3)	100% (3)	67% (3)
static 15	17% (6)	50% (6)	17% (6)
dynamic	100% (3)	100% (3)	0% (3)

Difficulty	l. react. c.	react. pl.	s-p-a
static -10	13.64 s	13.41 s	21.05 s
static 0	14.29 s	13.44 s	21.76 s
static 5	15.83 s	14.26 s	20.16 s
static 10	20.95 s	18.32 s	21.51 s
static 15	28.75 s	21.06 s	18.27 s
dynamic	17.95 s	15.44 s	-

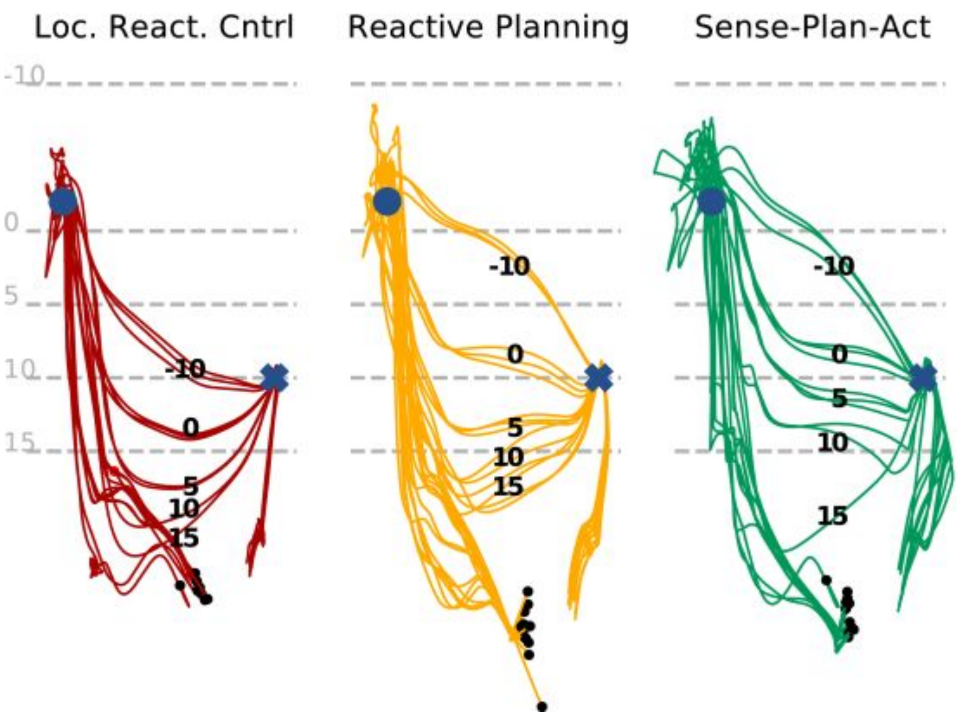

Dynamic Graphs

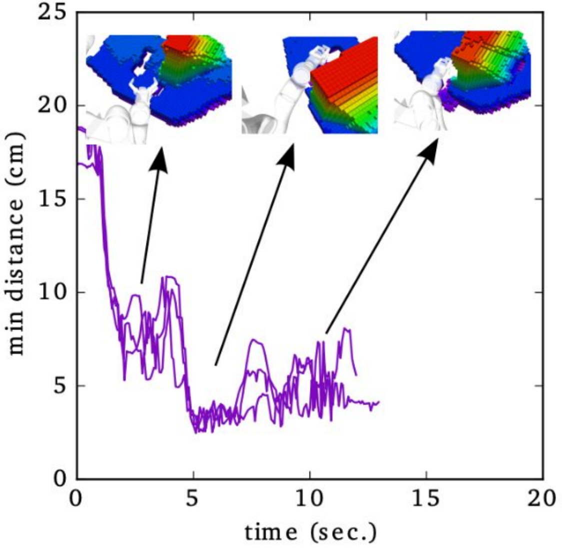

(a) Locally Reactive Control

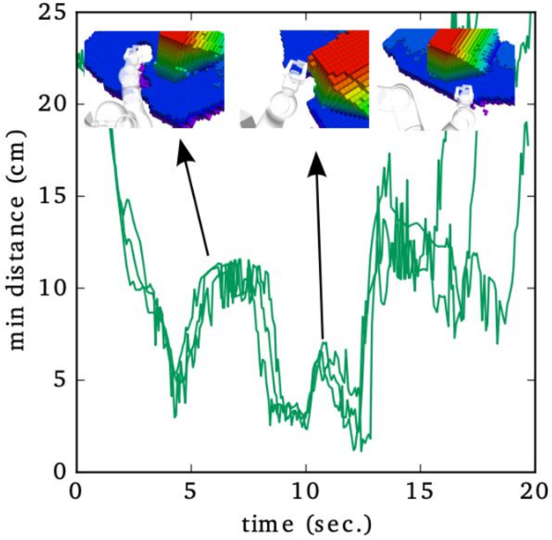

(b) Reactive Planning

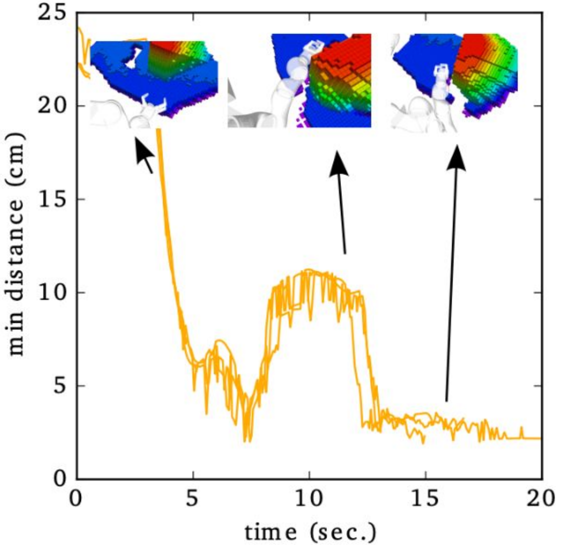

(c) Sense-Plan-Act

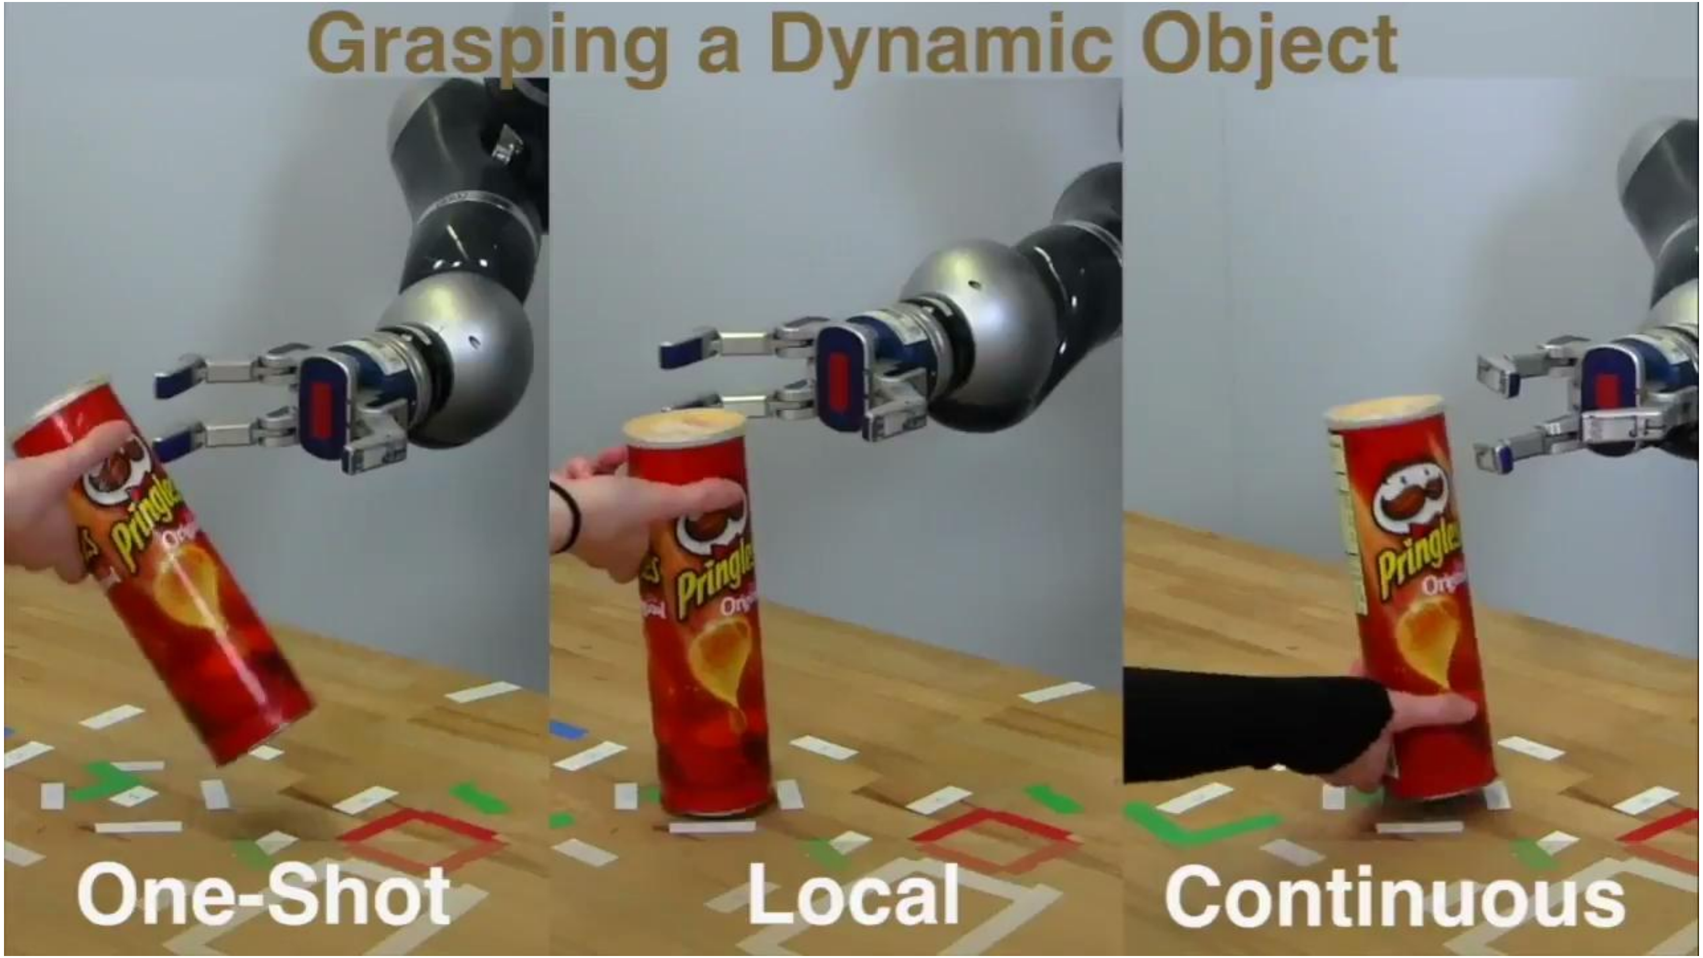
Grasping a Dynamic Object
One-Shot
Local
Continuous

Grasping With Dynamic Targets Results / Discussion

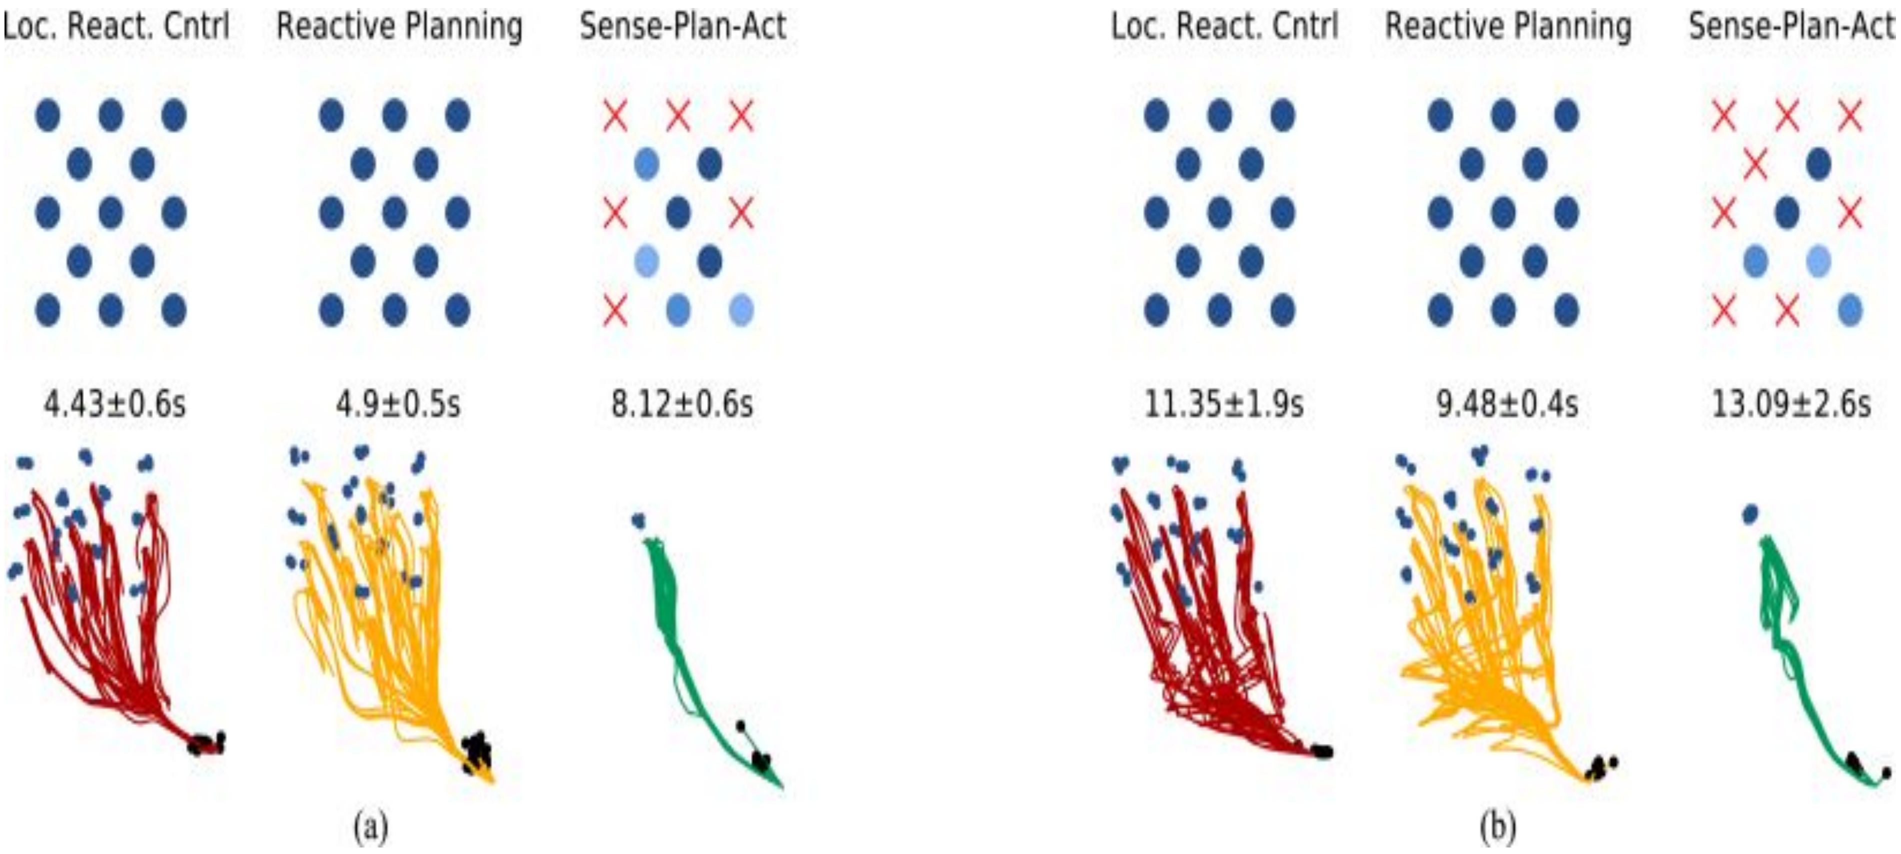

Pointing in Dynamic Environments With Dynamic Target

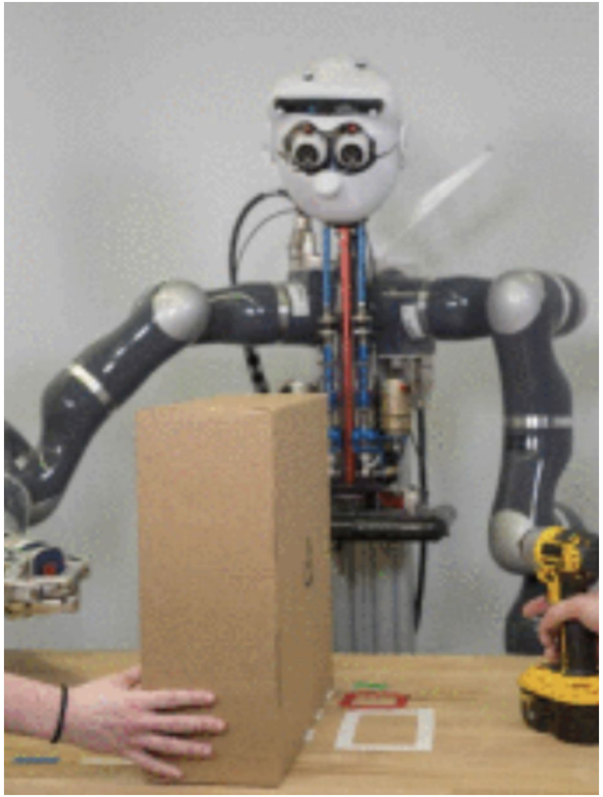

Task has 4 levels of complexity:

1. Static environment, no obstacles, drill is not moving.
2. Blocking box introduced during execution.
3. Obstacle is moved into the way so that arm must take a huge detour. More exaggerated blocking.
4. Start with blocking obstacle. In the middle of movement, obstacle is removed while also changing the orientation of the drill by 90 degrees, for pointing approach adaption.

Difficulty	l. react. c.	react. pl.	s-p-a
static	100% (3)	100% (3)	100% (3)
straight	100% (3)	100% (3)	0% (3)
diagonal	100% (3)	100% (3)	0% (3)
turning	100% (3)	100% (3)	0% (3)

Brief Overview / Discussion of Results

1. Reactive planning achieves better performance in **complex** environments, but is even good in static environments.
2. Locally reactive controller at a high perception rate does good too.
3. More beneficial to have fast feedback than accurate world representations.
4. Limiting information transfer between components helps!
5. Data association is important.

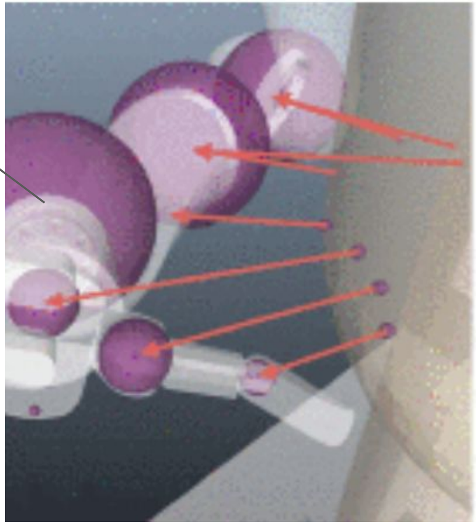

Critiques, Limitations and Open Issues

1. **Probabilistic Object Tracking:** assumes knowledge of the 3-dimensional shape of the objects of interest for visual tracking. Is this really "general robot autonomy"?
2. **SDFs:** rely on distinct shape primitives (simple geometrics) to allow for analytic SDF formations. In a real world environment, is this really the case?
3. **Experimentation:** probabilistic methodology is robust to specific occlusion, but did not experiment with objects touching the end effector during movement of target.
   a. <u>Not a lot of dynamic experiments.</u>

Critiques, Limitations and Open Issues (Cont).

1. **Visuals:** diagrams for three architectures could have been more clear, could have proposed a final diagram with how all the feedback components are integrated.
2. **Planning Time:** limited SPA to 2 seconds, and claim it was chosen empirically. I believe that there needs to be more explanation on why this was the case.
3. ***Grasping:*** Lack of information on the grasping integration. They did a big experiment on it but did not explain it well.

Future Work for Paper / Reading

Since this paper, it has been cited 71 times! **<u>Some selections from 2021 and 2022:</u>**

1. An efficient locally reactive controller for safe navigation in visual teach and repeat missions (2022)
2. Leveraging Contact Forces for Learning to Grasp
3. RMPflow: A Computational Graph for Automatic Motion Policy Generation

Future Proposed Work:

1. Look into direct contact interaction with haptics.
2. Learn representations of perceptual data for improved performance.
3. Use contact constraints over motion optimization.

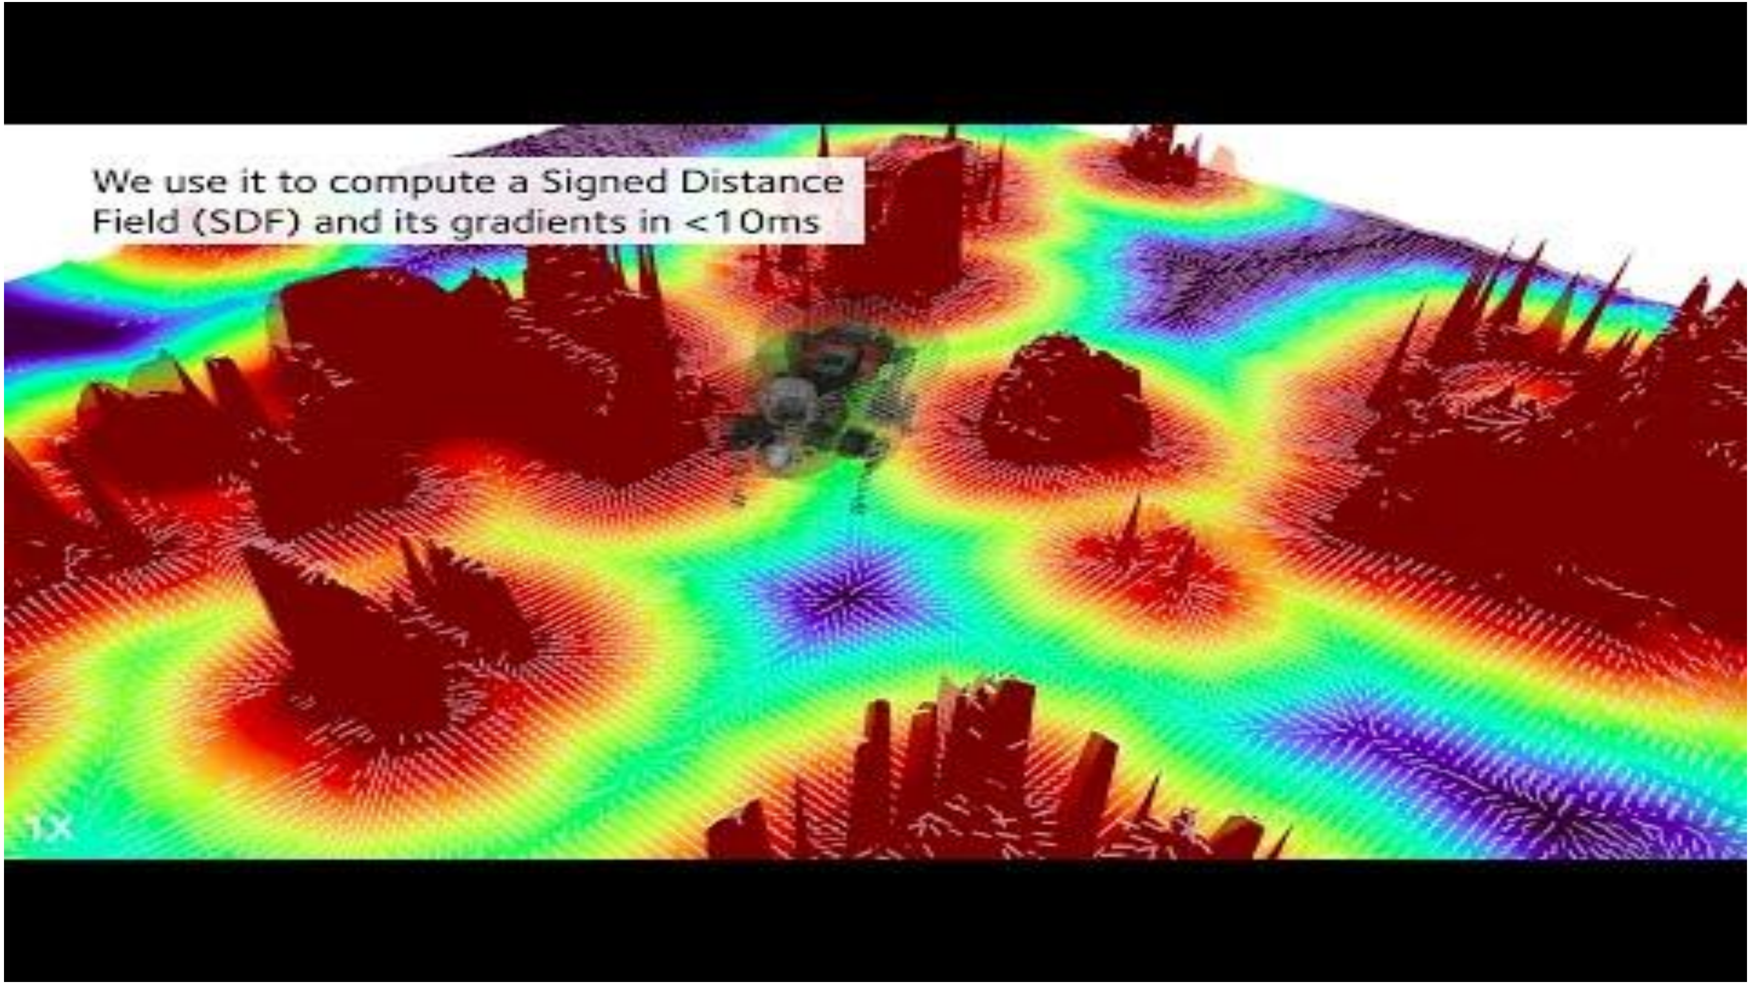
We use it to compute a Signed Distance Field (SDF) and its gradients in <10ms
1X

Summary / Key Insights

THE BIG IDEA : integrating real-time feedback on different time scales is essential to have **safe and successful** manipulation in unstructured and dynamic environments filled with uncertainty [ICRA, 2018].

ITS ROLE IN ROBOTICS: reactive motion generation and its success shows that trying to maximize “one-shot” system accuracy is not necessary, focus should shift to reactive systems. [ICRA, 2018]

Problem the reading is discussing: integrating perception and motion generation to general robotic systems for safe and efficient manipulation

Why is it important and hard: increased degrees of freedom paired with unstructured and dynamic environments can make tasks challenging and unsafe, but general robot autonomy is dependent on its success.

What is the key limitation of prior work: lessons learned from robotics challenges found there are tradeoffs between modularization, tight integration, feedback and planning. Did not quantitatively evaluate and experiment the way this paper did.

Extended Readings / Links / Credits:

REAL TIME PERCEPTION MEETS REACTIVE (THROUGH MAXPLANK): **https://am.is.mpg.de/publications/2017_rss_system**

Perceptually Driven Autonomous Vehicle: http://acl.mit.edu/papers/LeonardJFR08.pdf

Riemannian Motion Policies: https://arxiv.org/pdf/1801.02854.pdf

ICRA Spotlight Video: https://www.youtube.com/watch?v=tSe7RxIr9I8&t=7s

Alternate that have cited this one (with year filtering functionality):

https://scholar.google.com/scholar?oi=bibs&hl=en&cites=9104856553503528663&as_sdt=5

2015 DARPA Robotics Challenge (VIDEO): https://www.youtube.com/watch?v=8P9geWwi9e0

Lessons from DARPA Challenge (PAPER): https://doi.org/10.1002/rob.21674

Lessons from Amazon Picking Challenge: https://www.ijcai.org/proceedings/2017/0676.pdf

IEEE Xplore (great interface for reading paper) : https://ieeexplore.ieee.org/Xplore/home.jsp

Learn about Kuka: http://www.diag.uniroma1.it/~deluca/rob1_en/09_Exercise_DH_KukaLWR4.pdf

Probabilistic Object Tracking (PAPER): https://arxiv.org/abs/1505.00241

Probabilistic Object Tracking (VIDEO): https://www.youtube.com/watch?v=7KNt2L5_jTU

Glossary of Robotics Terms: https://www.motoman.com/en-us/about/company/robotics-glossary

Experiment Video: https://www.youtube.com/watch?v=R9gZF6ihPSk

Object and Arm Tracking / Feedback Video: https://www.youtube.com/watch?v=HpUUqcHzsII

Additional Reference Links: https://www.youtube.com/watch?v=unWnZvXJH2o&t=149s

All References

1. J. Leonard et al. "A perception-driven autonomous urban vehicle" J. Field Robot. vol. 25 no. 10 pp. 727-774 2008.
2. C. Richter A. Bry and N. Roy Polynomial Trajectory Planning for Aggressive Quadrotor Flight in Dense Indoor Environments Cham Switzerland:Springer-Verlag pp. 649-666 2016.
3. C. Atkeson et al. "No falls no resets: Reliable humanoid behavior in the DARPA robotics challenge" Proc. 2015 IEEE-RAS. 15th Int. Conf. Humanoid Robots pp. 623-630 2015.
4. M. Johnson et al. "Team IHMC's lessons learned from the DARPA robotics challenge trials" J. Field Robot. vol. 32 no. 2 pp. 192-208 2015.
5. C. Eppner et al. "Lessons from the amazon picking challenge: Four aspects of robotic systems building" Proc. Robot. Sci. Syst. pp. 4831-4835 2016.
6. N. Correll et al. "Analysis and observations from the first amazon picking challenge" IEEE Trans. Autom. Sci. Eng. vol. 15 no. 1 pp. 172-188 Jan. 2018.
7. M. Wüthrich P. Pastor M. Kalakrishnan J. Bohg and S. Schaal "Probabilistic object tracking using a range camera" Proc. 2013 IEEE/RSJ Int. Conf. Intell. Robots Syst. pp. 3195-3202 2013.
8. C. Garcia Cifuentes J. Issac M. Wüthrich S. Schaal and J. Bohg "Probabilistic articulated real-time tracking for robot manipulation" IEEE Robot. Autom. Lett. vol. 2 no. 2 pp. 577-584 Apr. 2017.
9. N. Ratliff M. Toussaint and S. Schaal "Understanding the geometry of workspace obstacles in motion optimization" Proc. 2015 IEEE Int. Conf. Robot. Autom. pp. 4202-4209 2015.
10. J. Mainprice N. Ratliff and S. Schaal "Warping the workspace geometry with electric potentials for motion optimization of manipulation tasks" Proc. 2016 IEEE/RSJ Int. Conf. Intell. Robots Syst. pp. 3156-3163 2016.
11. D. Kortenkamp R. G. Simmons and D. Brugali "Robotic systems architectures and programming" in Springer Handbook of Robotics New York NY USA:Springer-Verlag pp. 283-306 2016.
12. R. A. Brooks "Elephants don't play chess" Robot. Auton. Syst. vol. 6 no. 1/2 pp. 3-15 1990.
13. L. Righetti et al. "An autonomous manipulation system based on force control and optimization" Auton. Robots vol. 36 no. 1/2 pp. 11-30 2014.
14. J. Kuffner and J. Xiao "Motion for manipulation tasks" in Springer Handbook of Robotics Cham Switzerland:Springer-Verlag pp. 897-930 2016.
15. O. Khatib "Real-time obstacle avoidance for manipulators and mobile robots" Int. J. Robot. Res. vol. 5 no. 1 pp. 90-98 1986.
16. F. Chaumette S. Hutchinson and P. Corke "Visual servoing" in Springer Handbook of Robotics. New York NY USA:Springer-Verlag pp. 841-866 2016.
17. S. Levine C. Finn T. Darrell and P. Abbeel "End-to-end training of deep visuomotor policies" J. Mach. Learn. Res. vol. 17 no. 1 pp. 1334-1373 2016.
18. D. Ferguson and A. Stentz "Using interpolation to improve path planning: The field D* algorithm" J. Field Robot. vol. 23 no. 2 pp. 79-101 2006.
19. Y. Kuwata J. Teo G. Fiore S. Karaman E. Frazzoli and J. P. How "Real-time motion planning with applications to autonomous urban driving" IEEE Trans. Control Syst. Technol. vol. 17 no. 5 pp. 1105-1118 Sep. 2009.
20. O. Brock and O. Khatib "Elastic strips: A framework for motion generation in human environments" Int. J. Robot. Res. vol. 21 no. 12 pp. 1031-1052 2002.
21. R. R. Burridge A. A. Rizzi and D. E. Koditschek "Sequential composition of dynamically dexterous robot behaviors" Int. J. Robot. Res. vol. 18 no. 6 pp. 534-555 1999.
22. A. J. Ijspeert J. Nakanishi H. Hoffmann P. Pastor and S. Schaal "Dynamical movement primitives: Learning attractor models for motor behaviors" Neural Comput. vol. 25 pp. 328-373 Feb. 2013.
23. A. Rai G. Sutanto S. Schaal and F. Meier "Learning feedback terms for reactive planning and control" Proc. 2017 IEEE Int. Conf. Robot. Autom. pp. 2184-2191 2017.
24. D. Kappler P. Pastor M. Kalakrishnan M. Wuthrich and S. Schaal "Data-driven online decision making for autonomous manipulation" Proc. Robot. Sci. Syst. 2015.
25. P. Lehner A. Sieverling and O. Brock "Incremental sensor-based motion generation for mobile manipulators in unknown dynamic environments" Proc. 2015 IEEE Int. Conf. Robot. Autom. pp. 4761-4767 May 2015.
26. Y. Yang V. Ivan Z. Li M. Fallon and S. Vijayakumar "iDRM: Humanoid motion planning with realtime end-pose selection in complex environments" Proc. 2016 IEEE-RAS. 16th Int. Conf. Humanoid Robots pp. 271-278 2016.
27. A. Hornung K. M. Wurm M. Bennewitz C. Stachniss and W. Burgard "OctoMap: An efficient probabilistic 3d mapping framework based on octrees" Auton. Robots vol. 34 pp. 189-206 2013.
28. R. B. Rusu and S. Cousins "3d is here: Point cloud library (PCL)" Proc. 2011 IEEE Int. Conf. Robot. Autom. pp. 1-4 2011.
29. M. Zucker et al. "CHOMP: Covariant hamiltonian optimization for motion planning" Int. J. Robot. Res. vol. 32 pp. 1164-1193 2013.
30. M. Kalakrishnan S. Chitta E. Theodorou P. Pastor and S. Schaal "STOMP: Stochastic trajectory optimization for motion planning" Proc. 2011 IEEE Int. Conf. Robot. Autom. pp. 4569-4574 2011.
31. N. D. Ratliff J. Issac and D. Kappler "Riemannian motion policies" Jan. 2018.
32. M. Toussaint and C. Goerick "A bayesian view on motor control and planning" in From Motor Learning to Interaction Learning in Robots New York NY USA:Springer-Verlag pp. 227-252 2010.
33. M. Kazemi J.-S. Valois J. A. Bagnell and N. S. Pollard "Robust object grasping using force compliant motion primitives" Proc. Robot. Sci. Syst. 2012.
34. R. Deimel C. Eppner J. Álvarez-Ruiz M. Maertens and O. Brock "Exploitation of environmental constraints in human and robotic grasping" Proc. Int. Symp. Robot. Res. pp. 393-409 2013.
35. N. Ratliff F. Meier D. Kappler and S. Schaal "Doomed: Direct online optimization of modeling errors in dynamics" Big Data vol. 4 no. 4 pp. 253-268 2016.
36. F. Meier D. Kappler N. Ratliff and S. Schaal "Towards robust online inverse dynamics learning" Proc. 2016 IEEE/RSJ Int. Conf. Intell. Robots Syst. pp. 4034-4039 2016.

Extended Readings (from earlier slide):

[17]: S. Levine, C. Finn, T. Darrell and P. Abbeel, "End-to-end training of deep visuomotor policies", J. Mach. Learn. Res., vol. 17, no. 1, pp. 1334-1373, 2016: learning motion policies directly from perceptual feedback in form of raw camera images and the system joint state, e.g.,

[5]: C. Eppner et al., "Lessons from the amazon picking challenge: Four aspects of robotic systems building", Proc. Robot. Sci. Syst., pp. 4831-4835, 2016: Amazon Picking Challenge

[7]: M. Wüthrich, P. Pastor, M. Kalakrishnan, J. Bohg and S. Schaal, "Probabilistic object tracking using a range camera", Proc. 2013 IEEE/RSJ Int. Conf. Intell. Robots Syst., pp. 3195-3202, 2013: previous work used in this work.

[8] C. Garcia Cifuentes, J. Issac, M. Wüthrich, S. Schaal and J. Bohg, "Probabilistic articulated real-time tracking for robot manipulation", IEEE Robot. Autom. Lett., vol. 2, no. 2, pp. 577-584, Apr. 2017:

[9] N. Ratliff, M. Toussaint and S. Schaal, "Understanding the geometry of workspace obstacles in motion optimization", Proc. 2015 IEEE Int. Conf. Robot. Autom., pp. 4202-4209, 2015.

[10] J. Mainprice, N. Ratliff and S. Schaal, "Warping the workspace geometry with electric potentials for motion optimization of manipulation tasks", Proc. 2016 IEEE/RSJ Int. Conf. Intell. Robots Syst., pp. 3156-3163, 2016.

[31] N. D. Ratliff, J. Issac and D. Kappler, "Riemannian motion policies", Jan. 2018: previous work used in this work.

Thank you all for your attention today!

Insights from Jeannette Bohg (for Discussion)

"If you work with a complex robotic system, you need a large team of team players who are experts in the large variety of components that are needed to make a robot work (perception, learning, control, planning, hardware, infrastructure, ...). The reason for this paper to be one of my favourite papers is that the authors were a great team and worked really hard together to make this large quantitative set of experiments on the system level work. We had daily scrums during the time we build this and were all sitting in one big lab space which made communication easy.

Sometimes I'm wondering if the reason we don't yet have really great advances in (manipulation) robotics is that in order to build such a complex system you need a lot of people who get along well. And this is hard. Robotics might be a people problem in that sense."