Phvsics-Based Manipulation under Uncertainty

Michael Koval
mkoval@cs.cmu.edu

February 9, 2016
Hands: Design and Control for Dexterous Manipulation

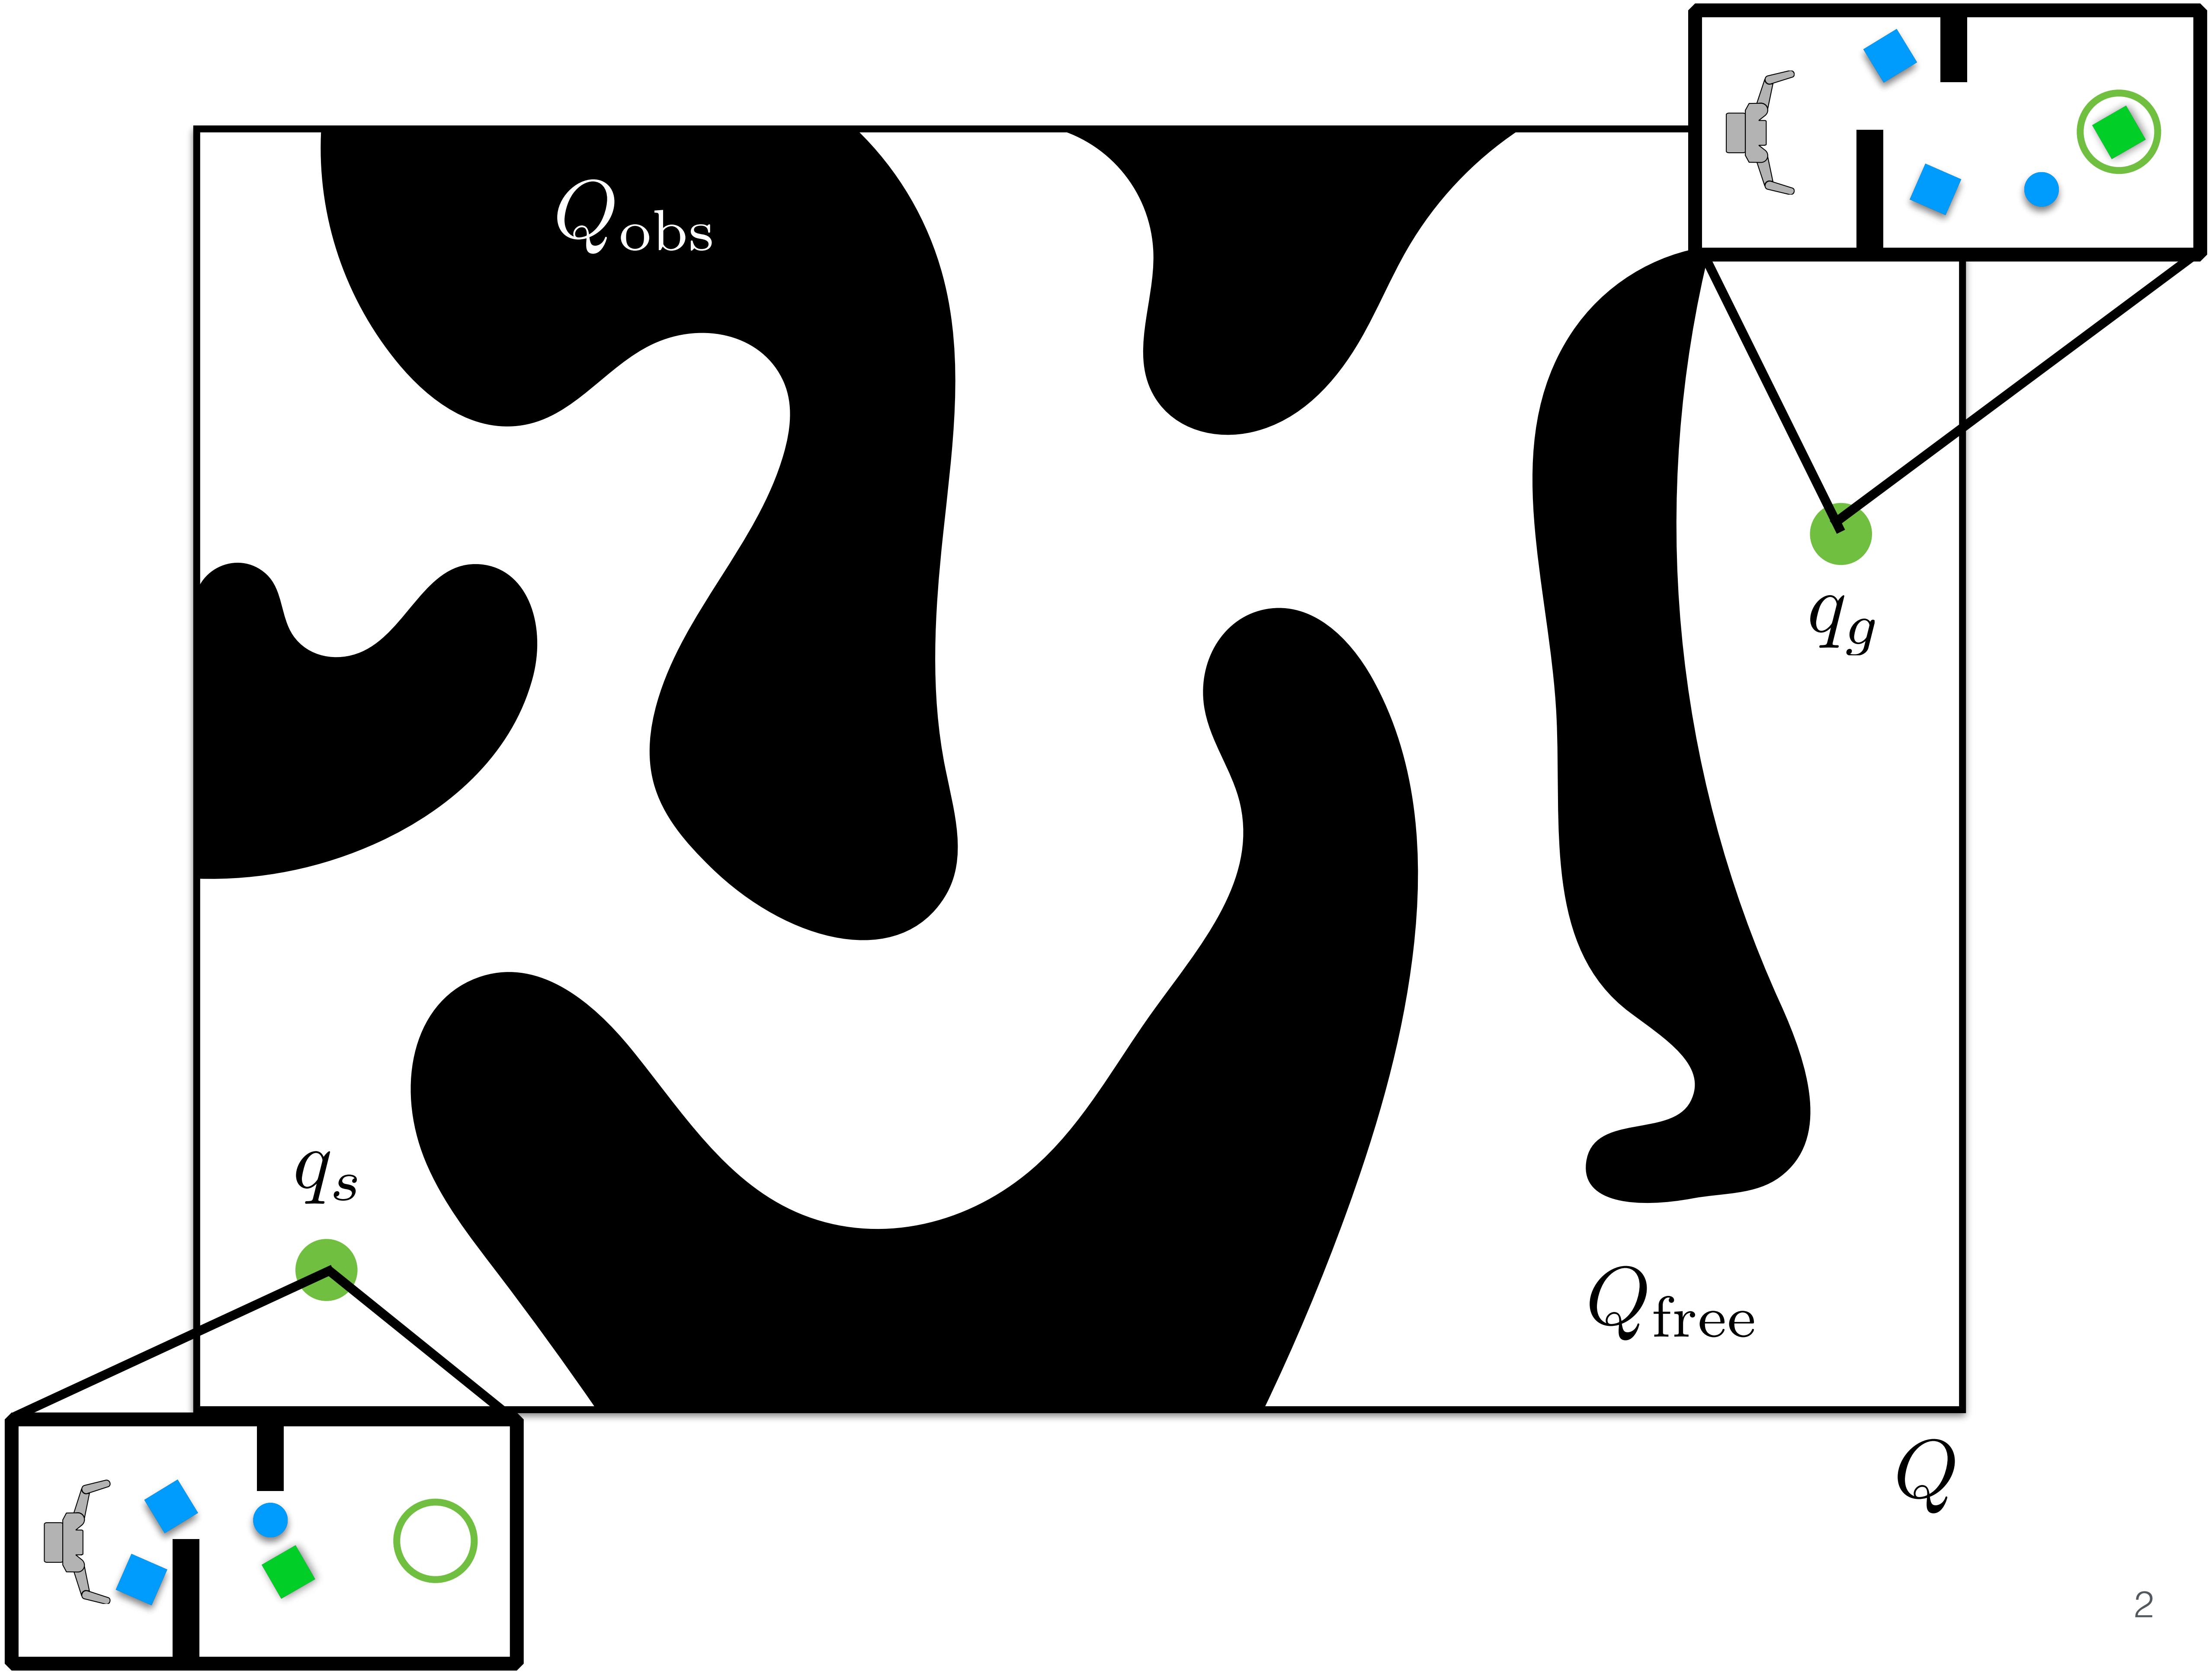
Q_{obs}
q_g
q_s
Q_{free}
Q

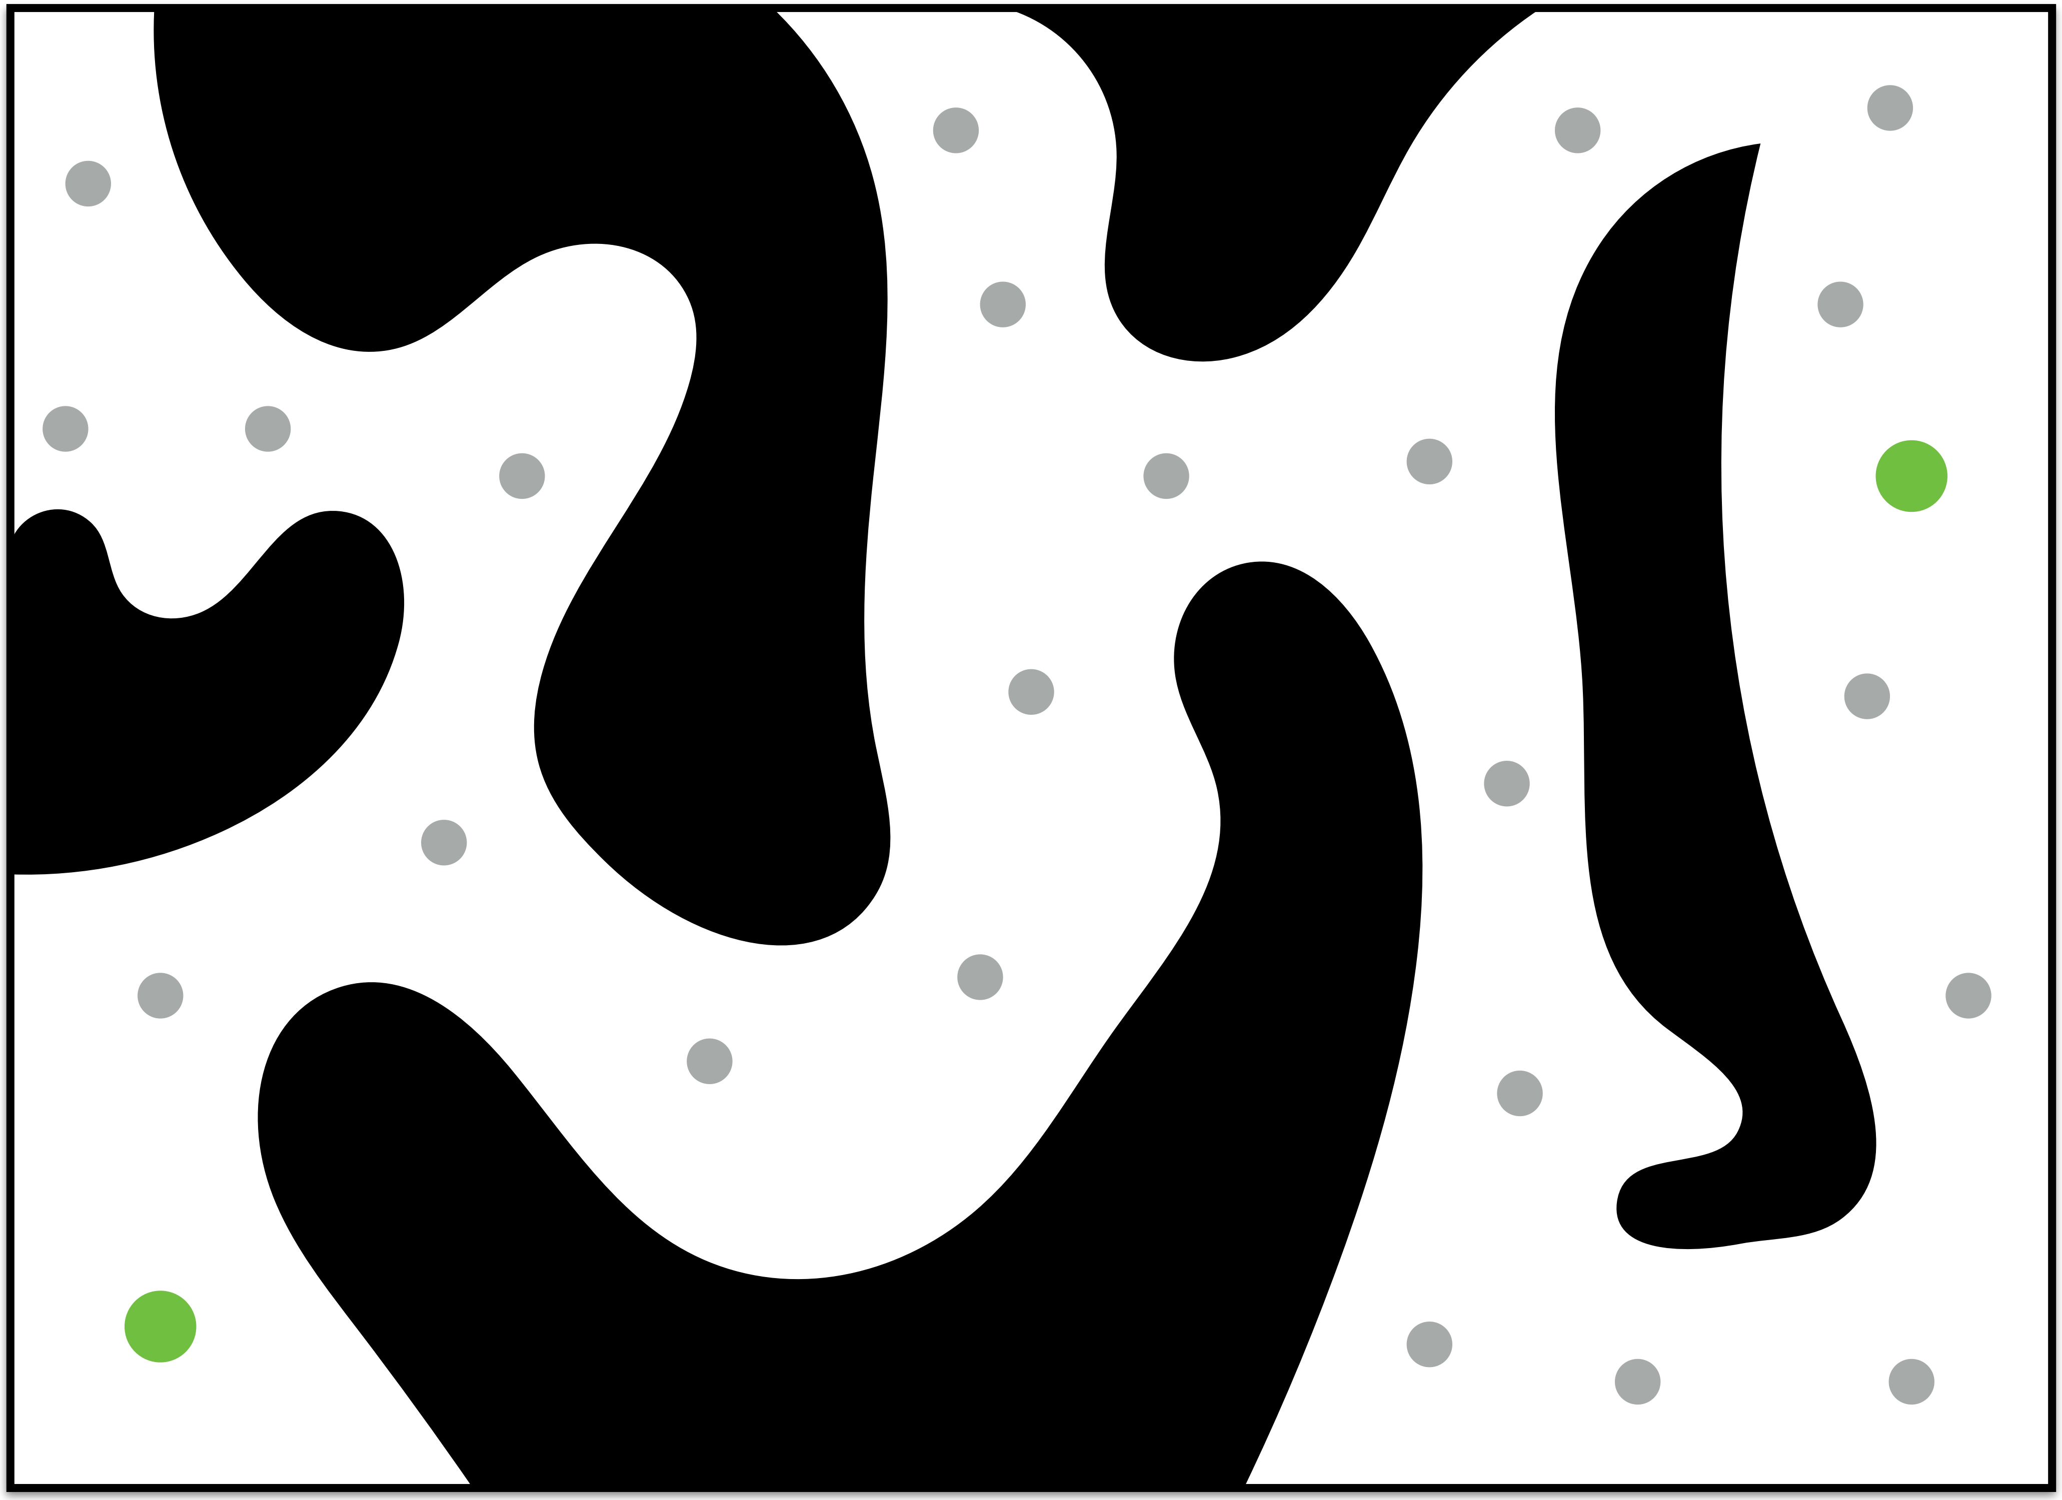

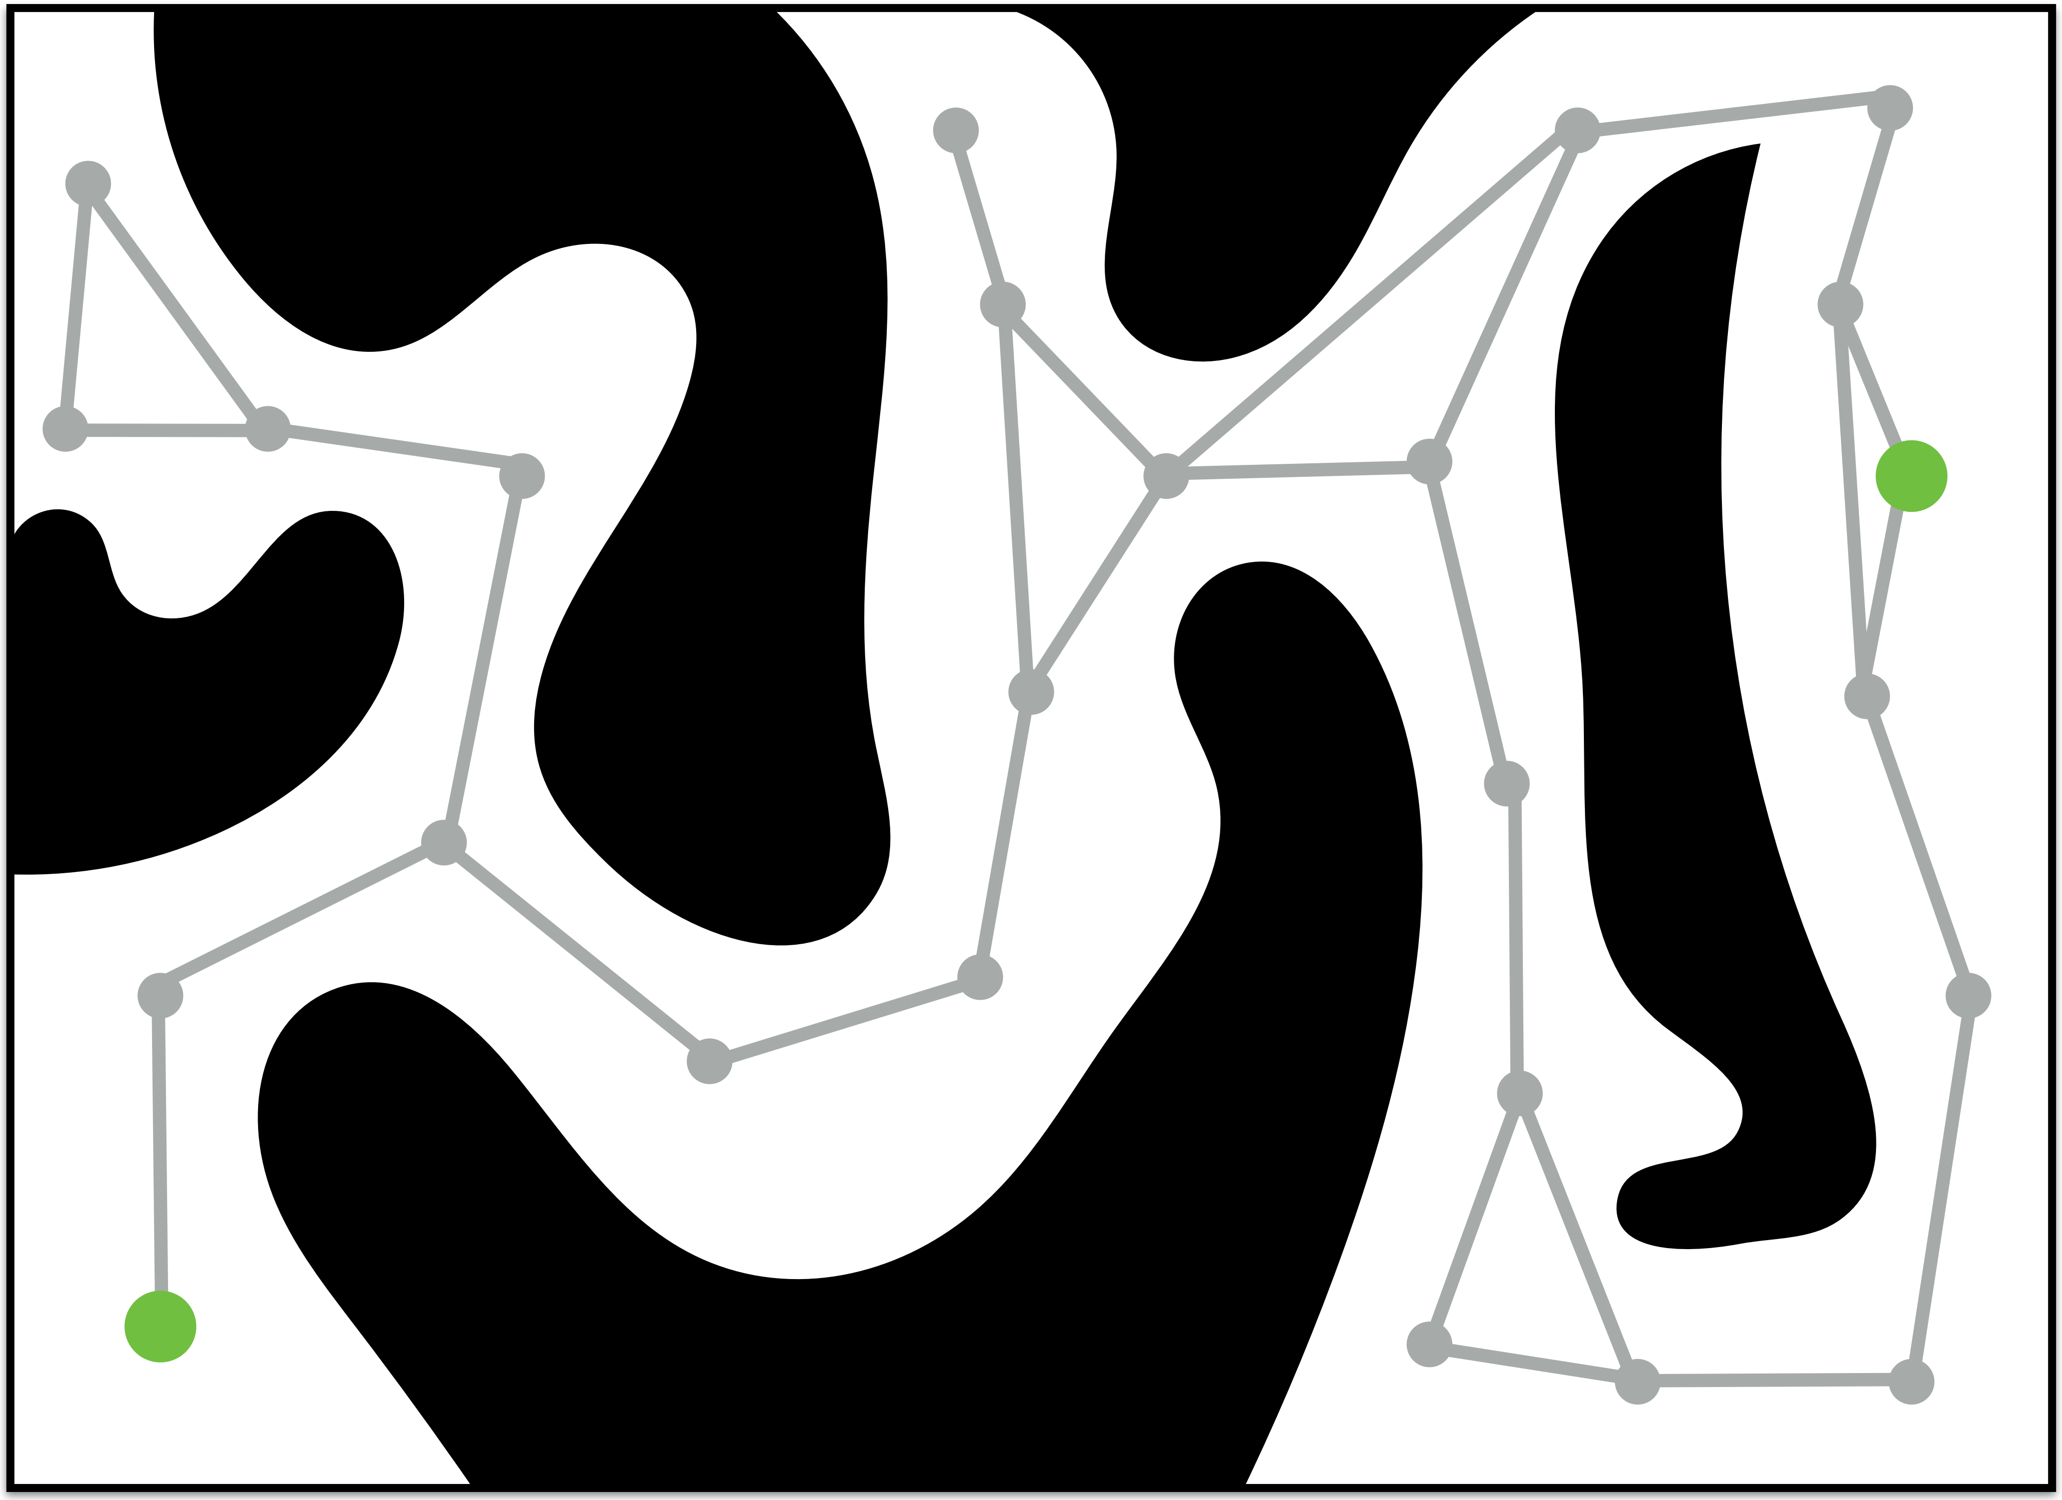

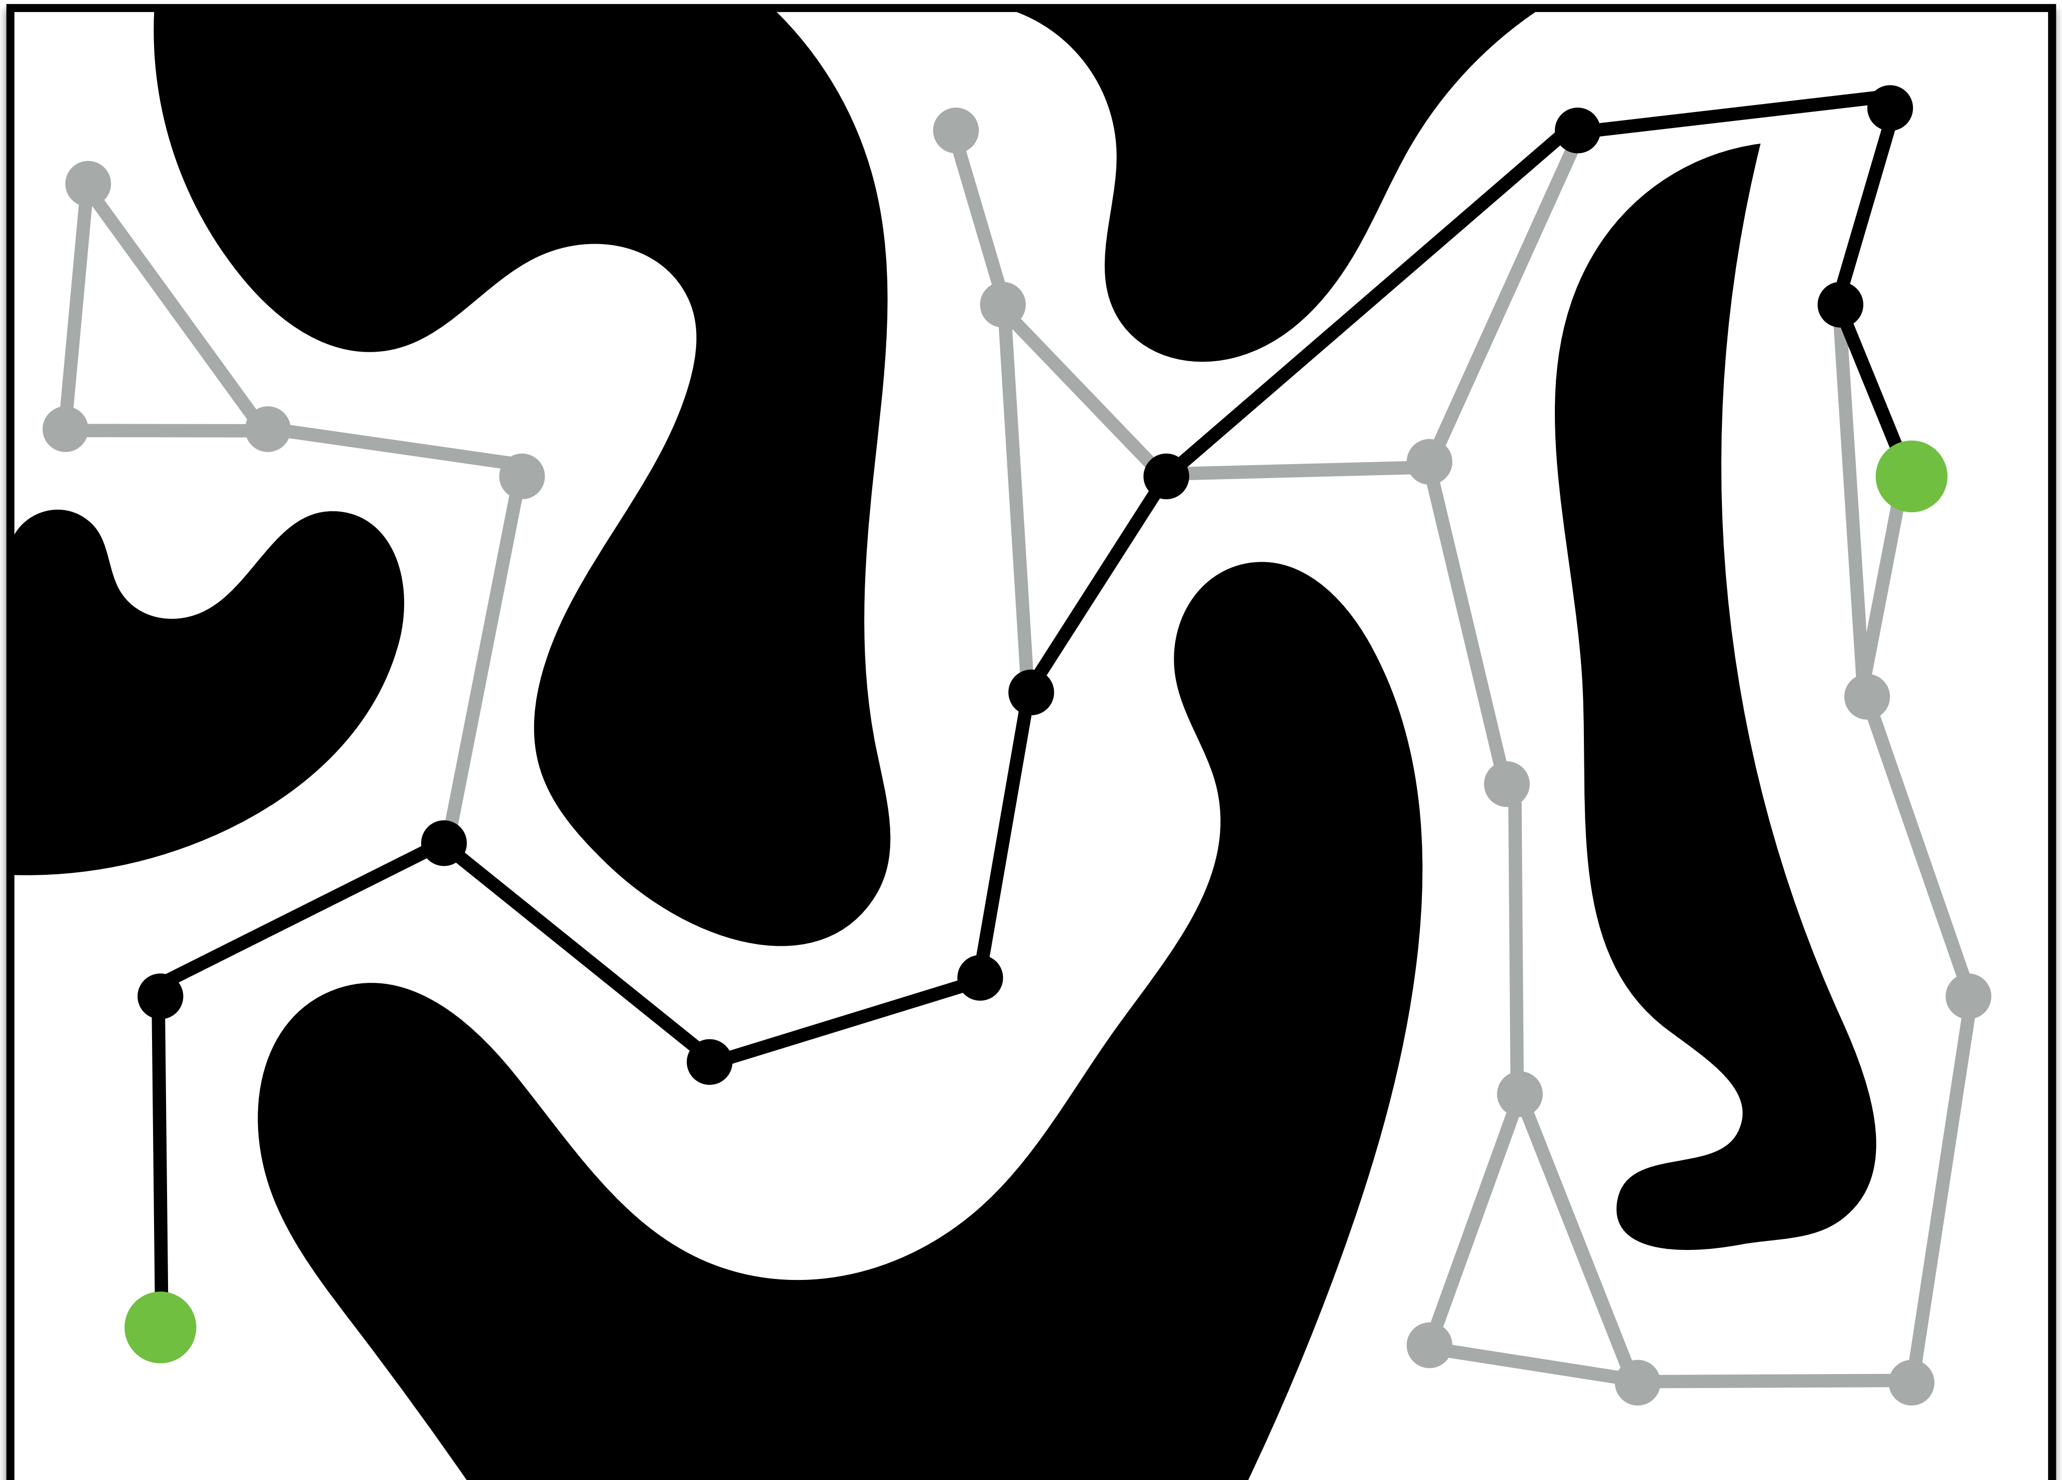

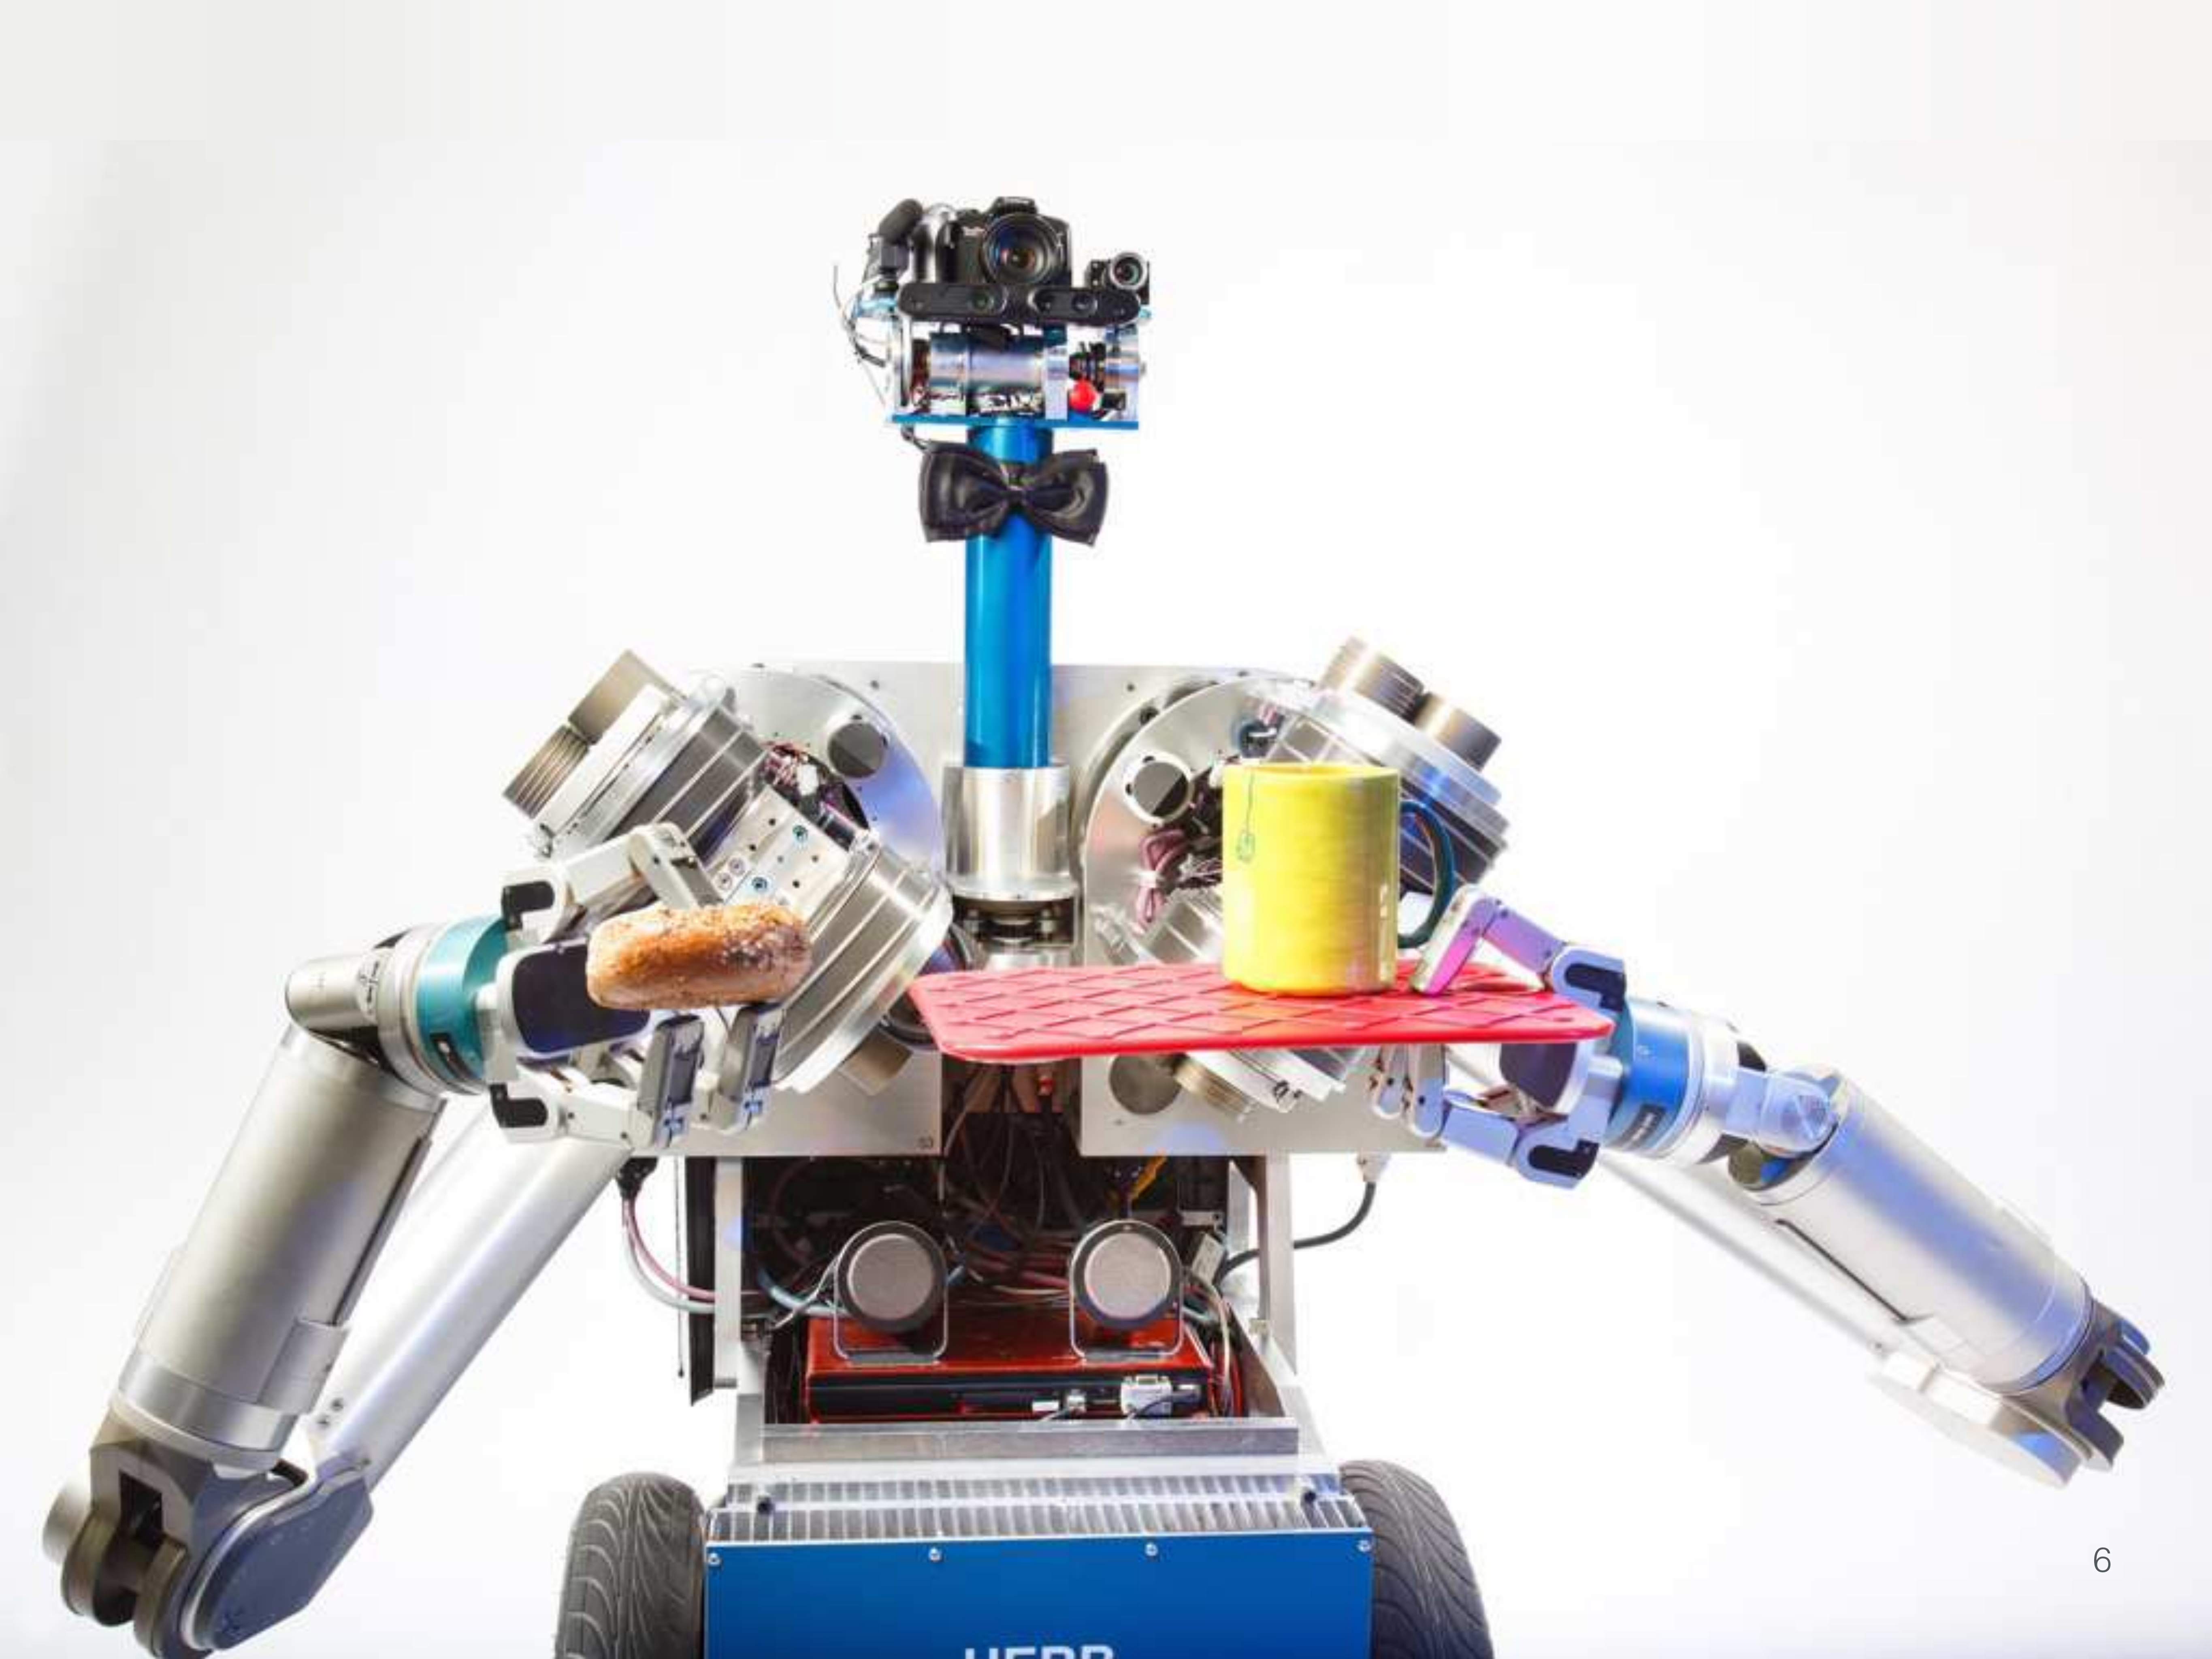

(4×)

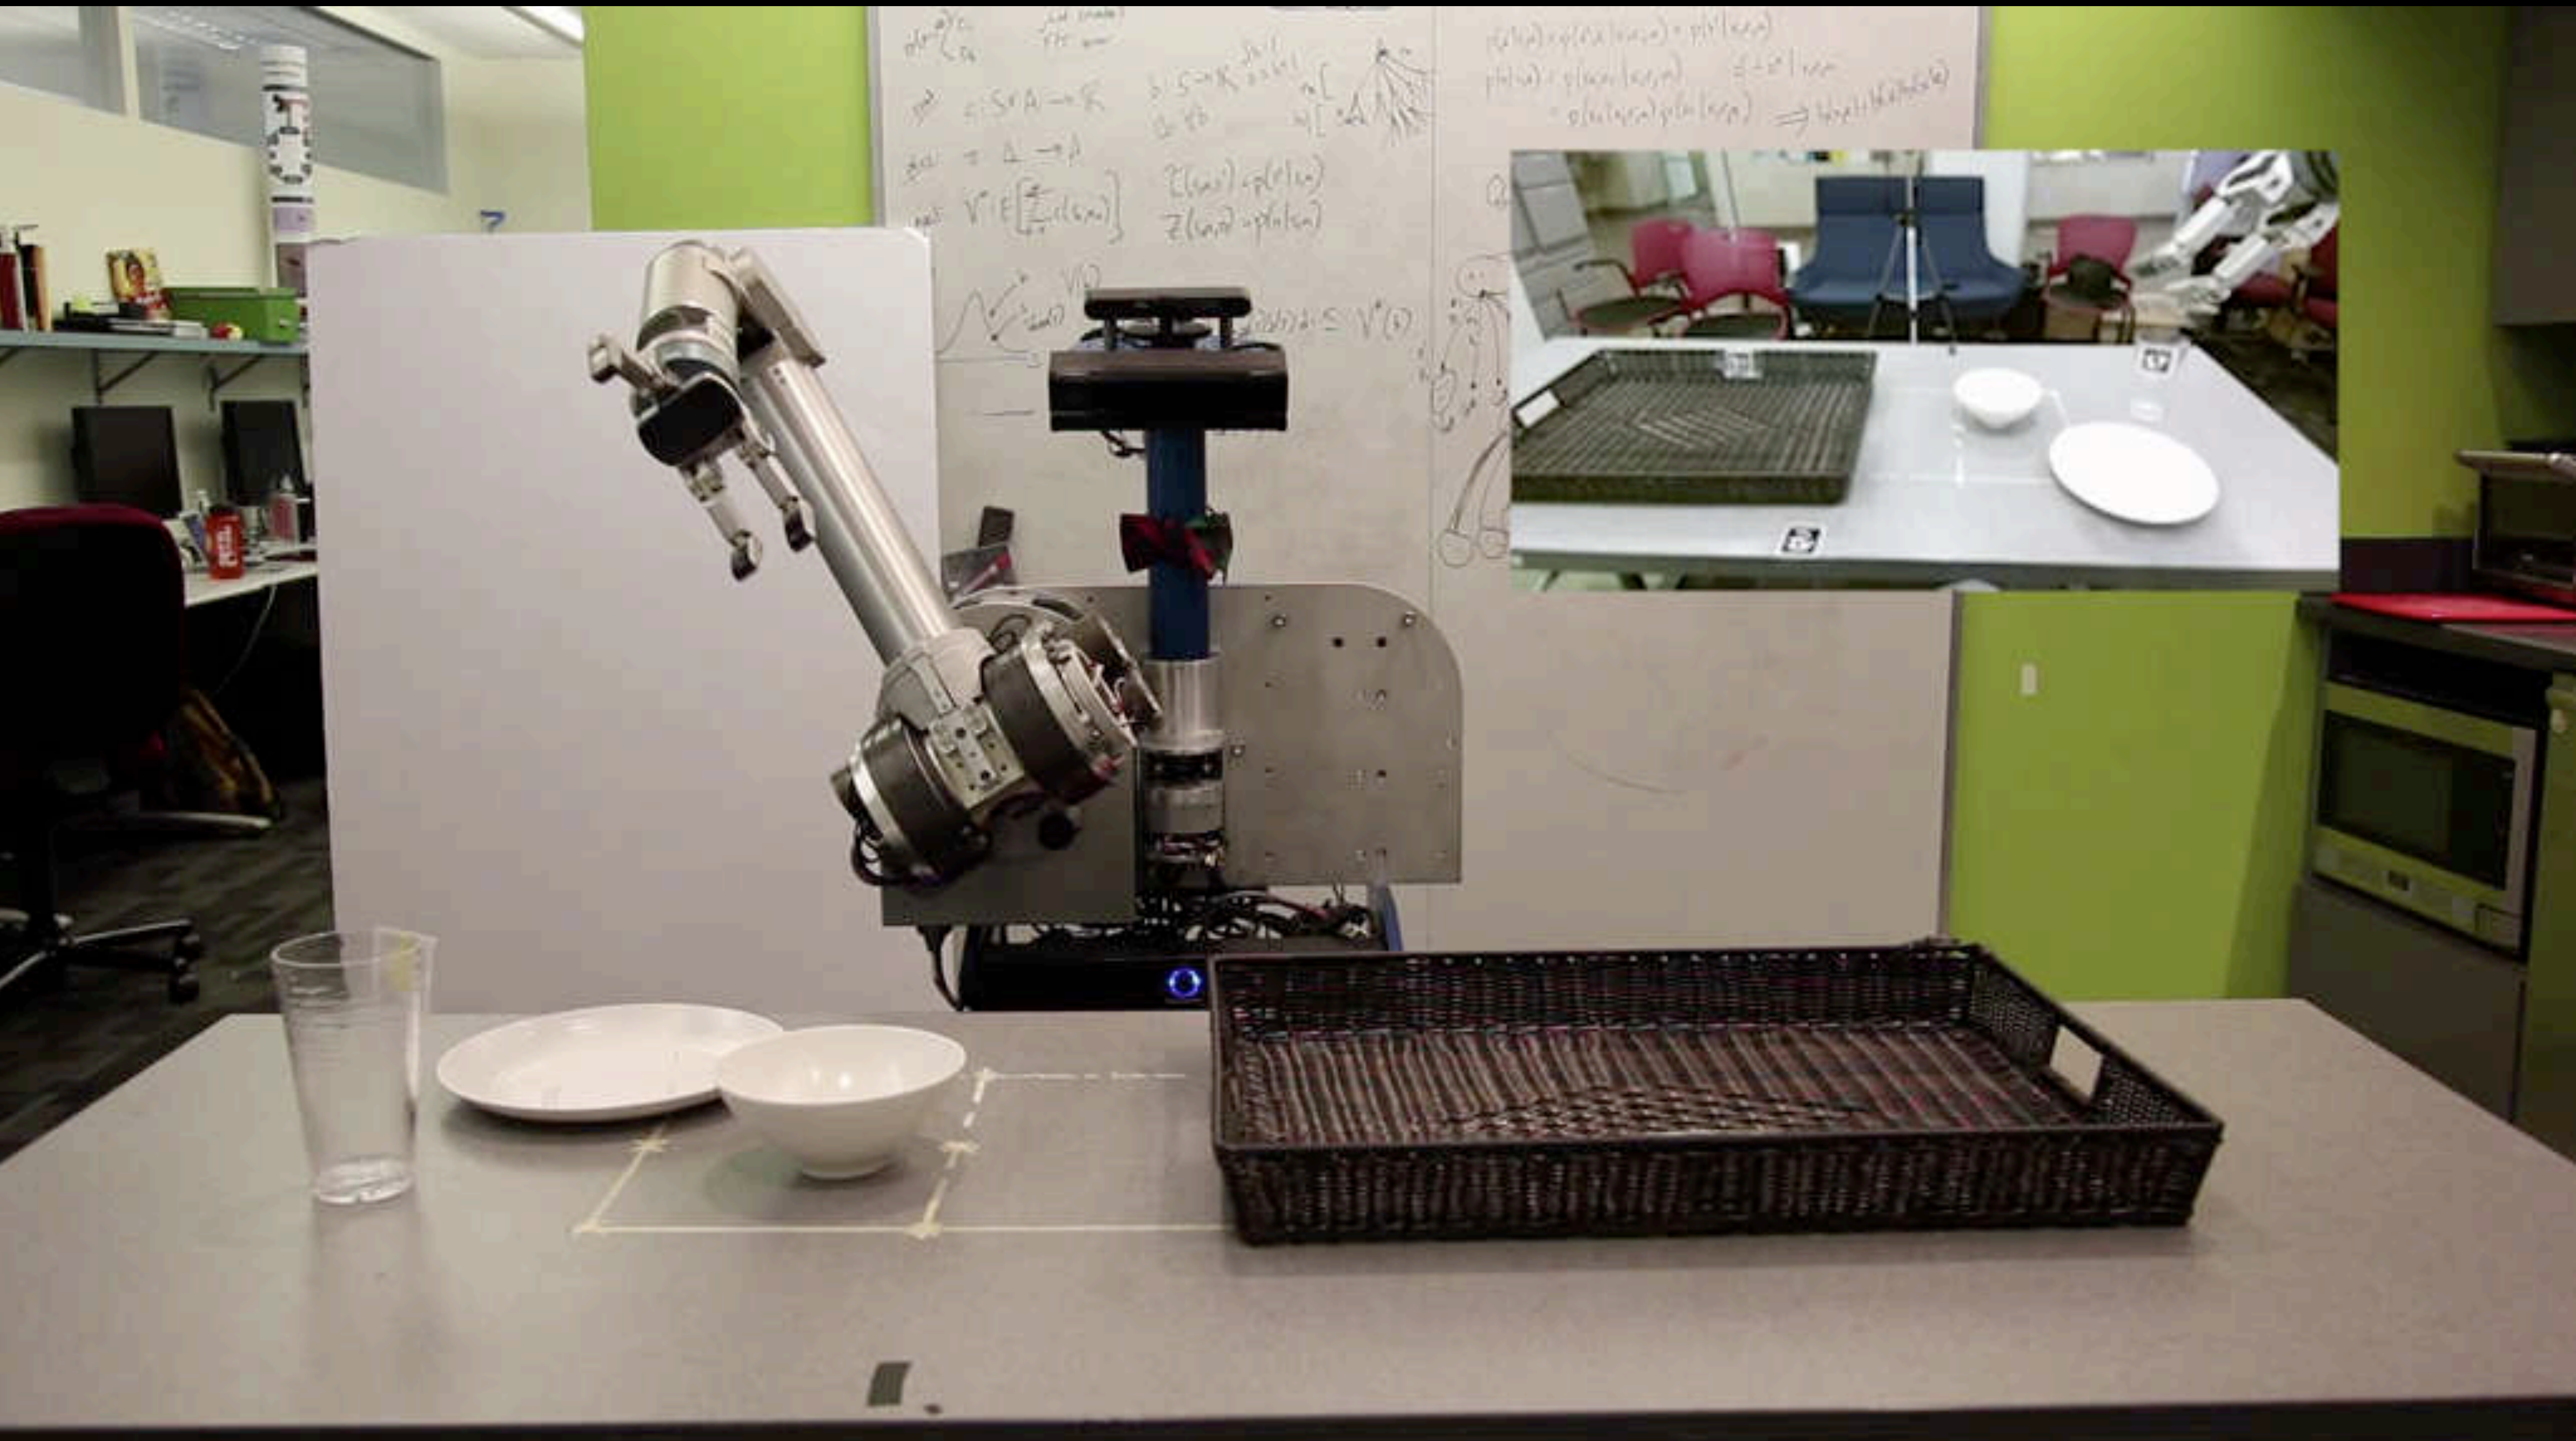

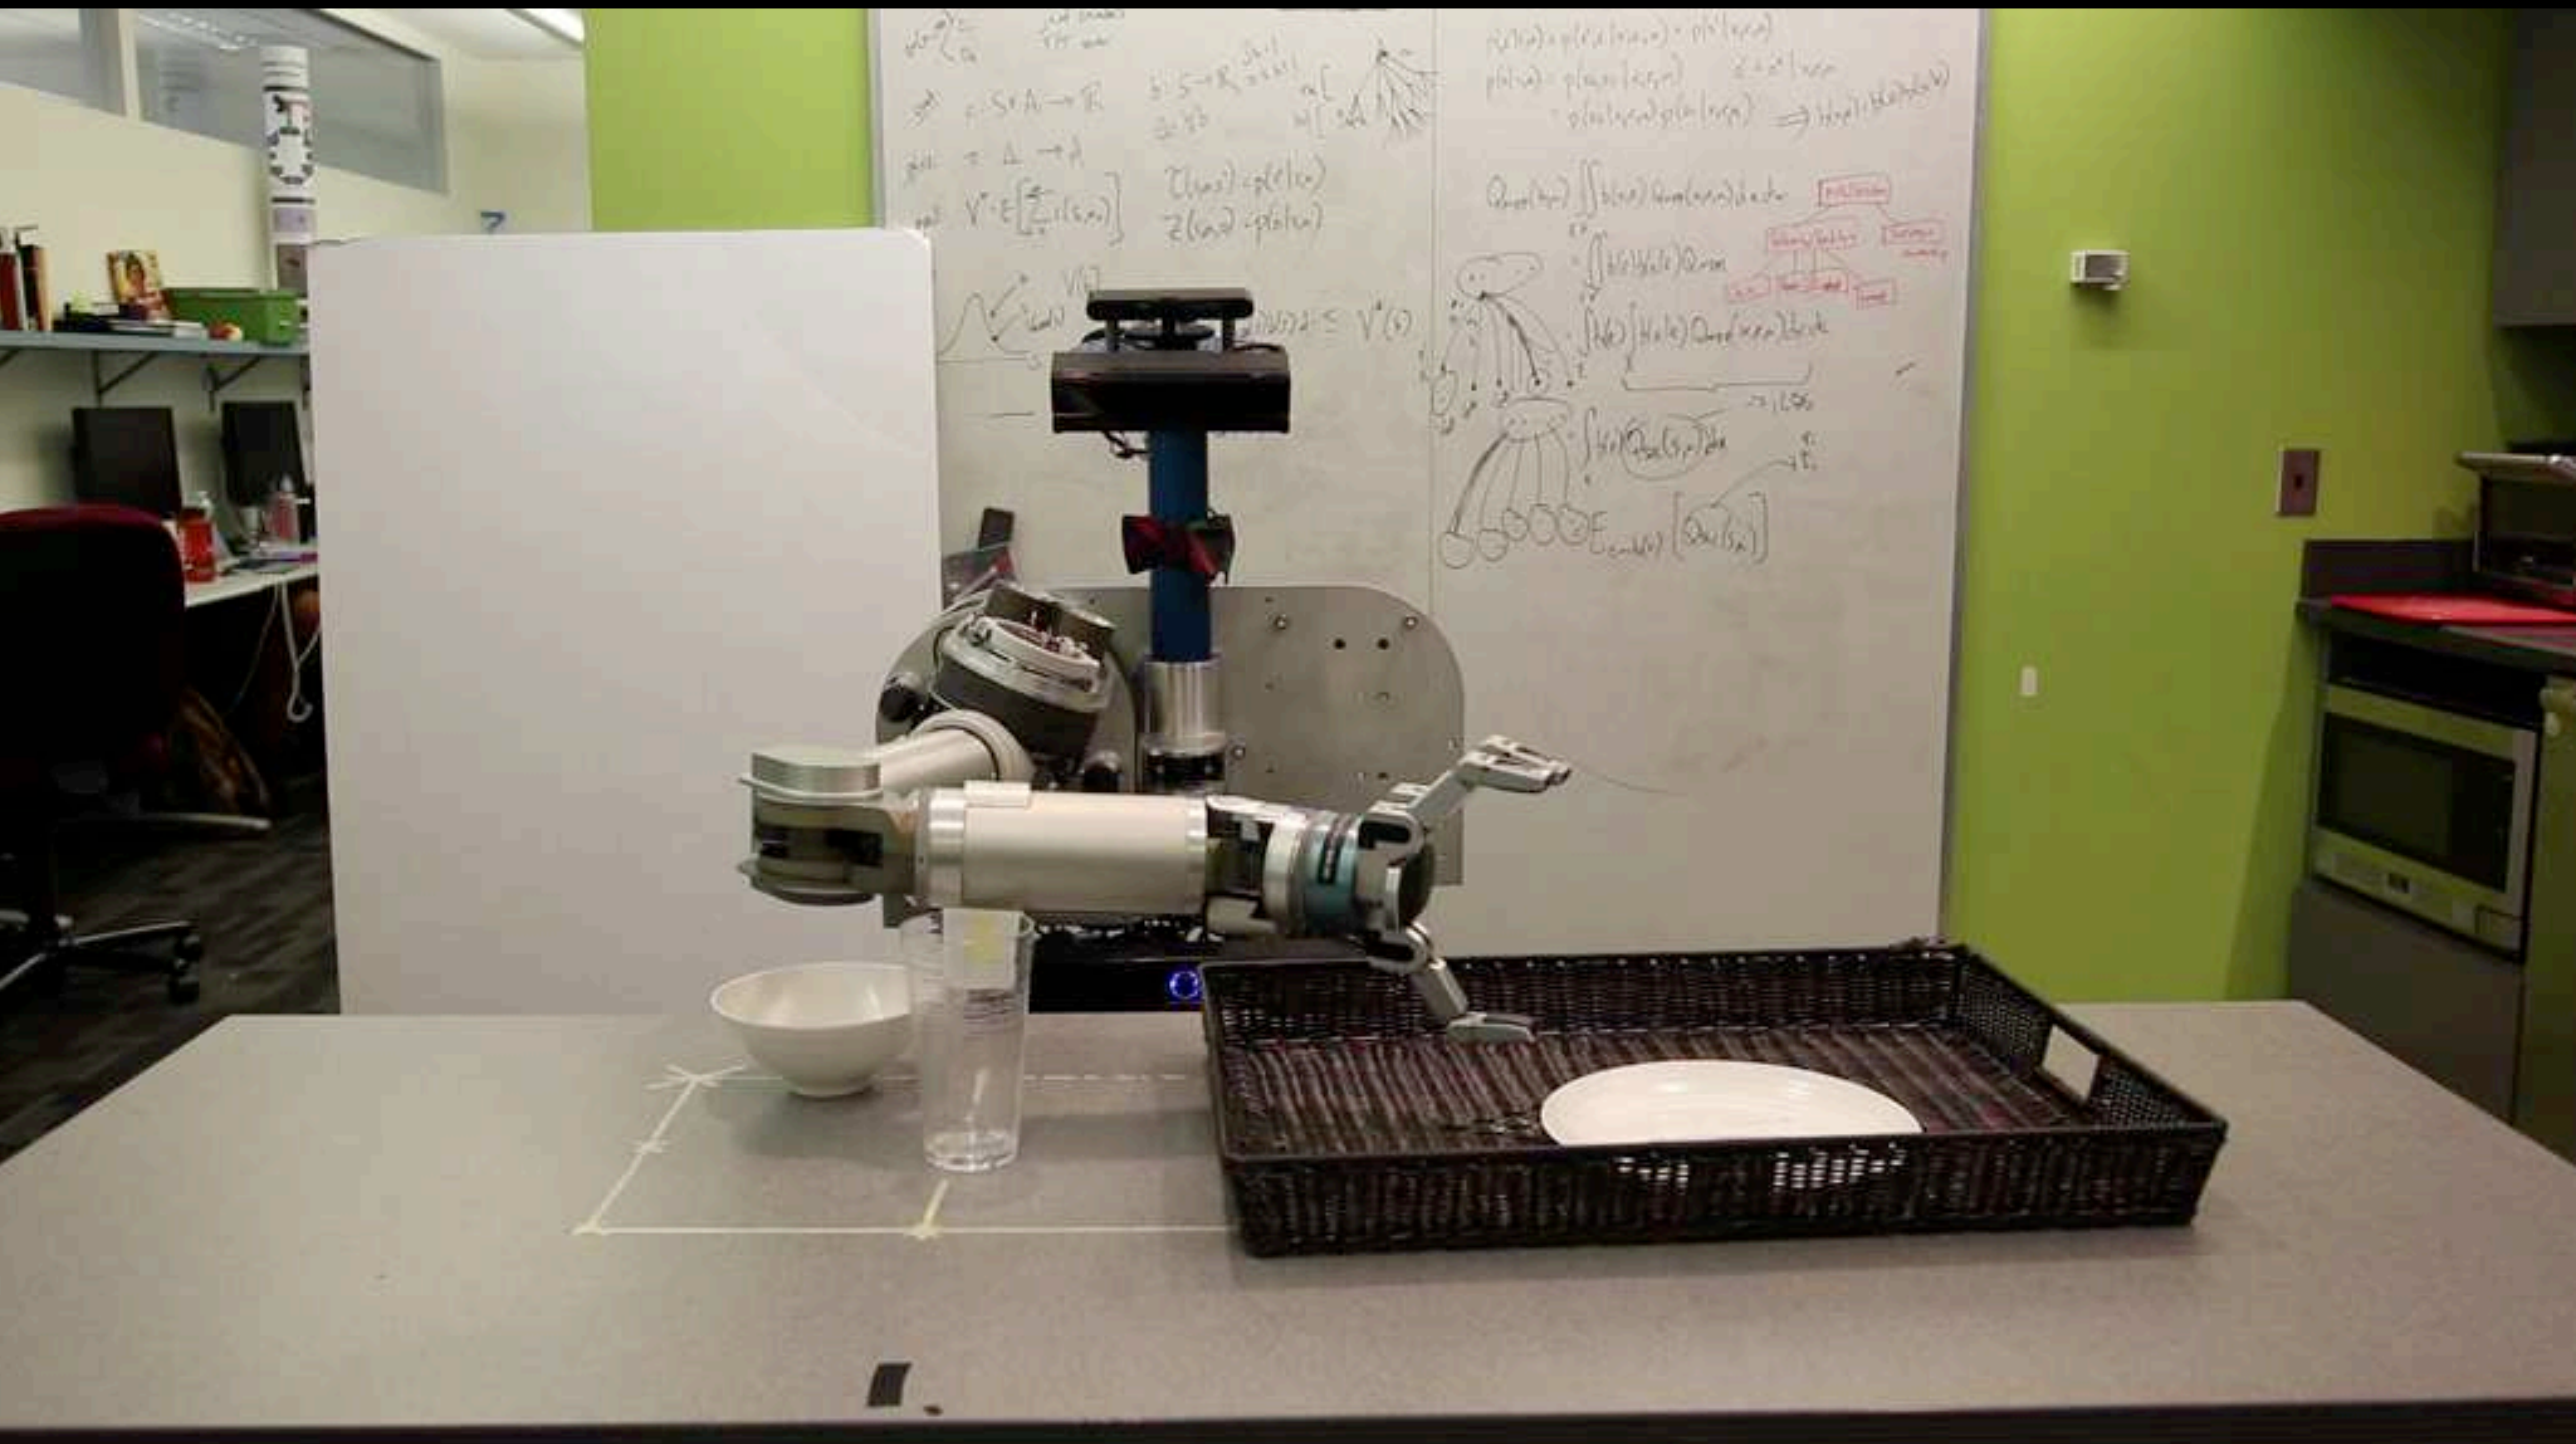


LOW CALORIE


DEWALT
18V

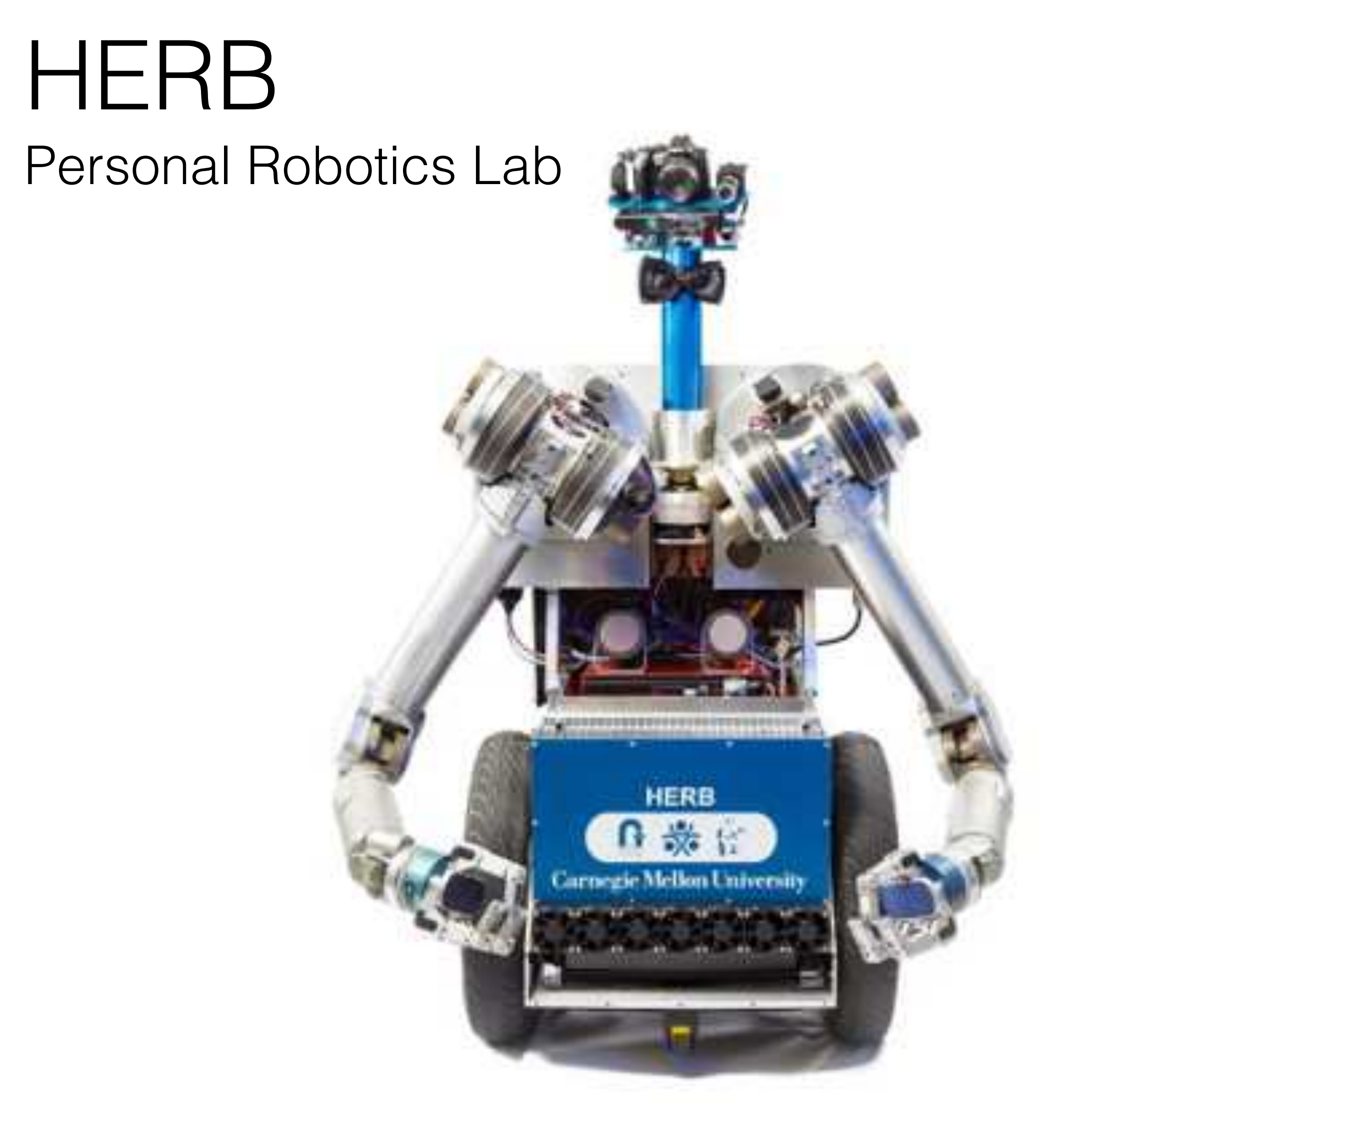
HERB
Personal Robotics Lab
HERB
Carnegie Mellon University


Andy
CMU/NREC

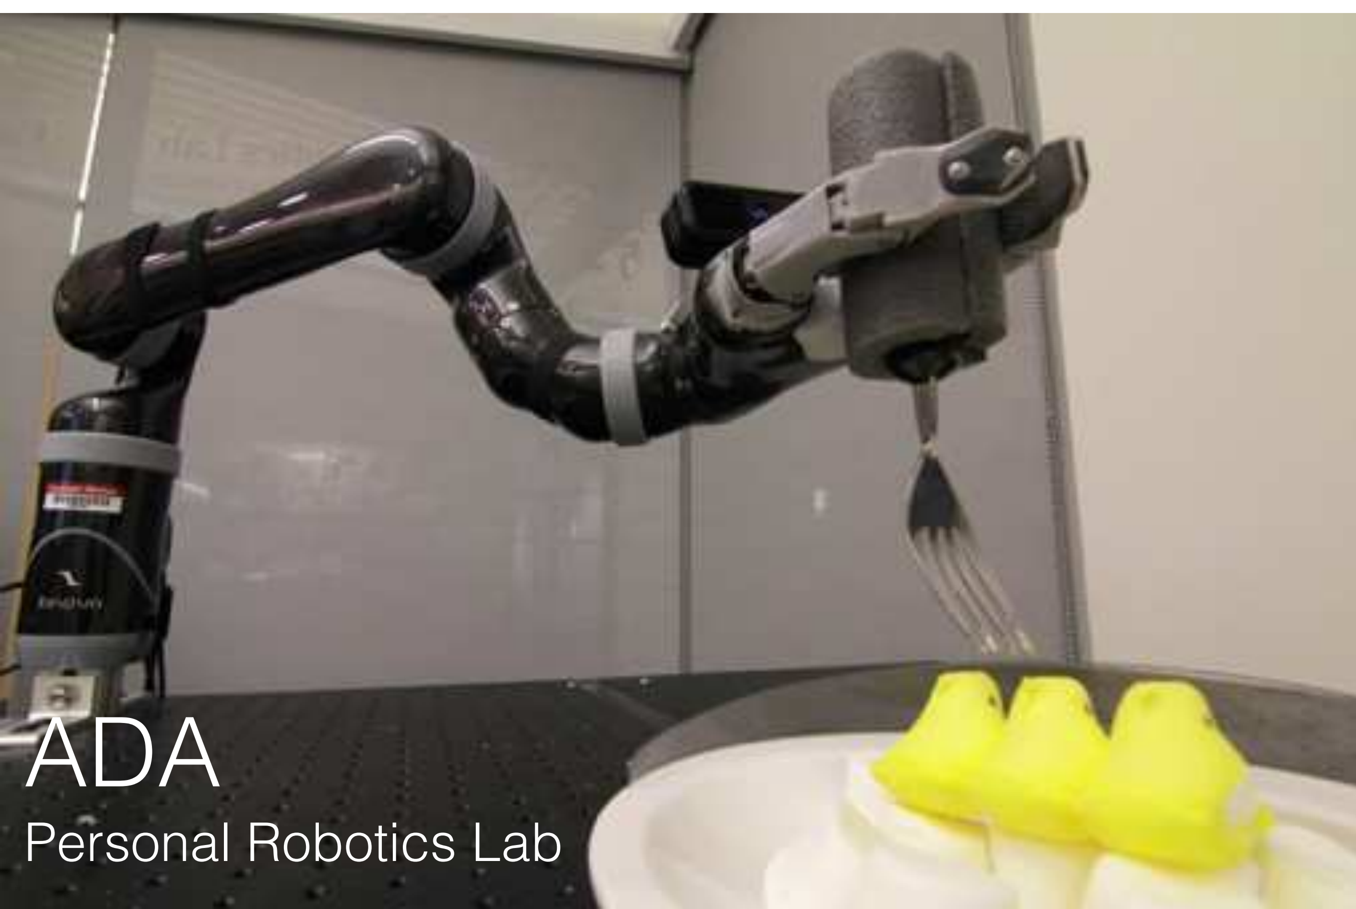
ADA
Personal Robotics Lab


GM
Robonaut 2
NASA/JSC

What caused these failures?

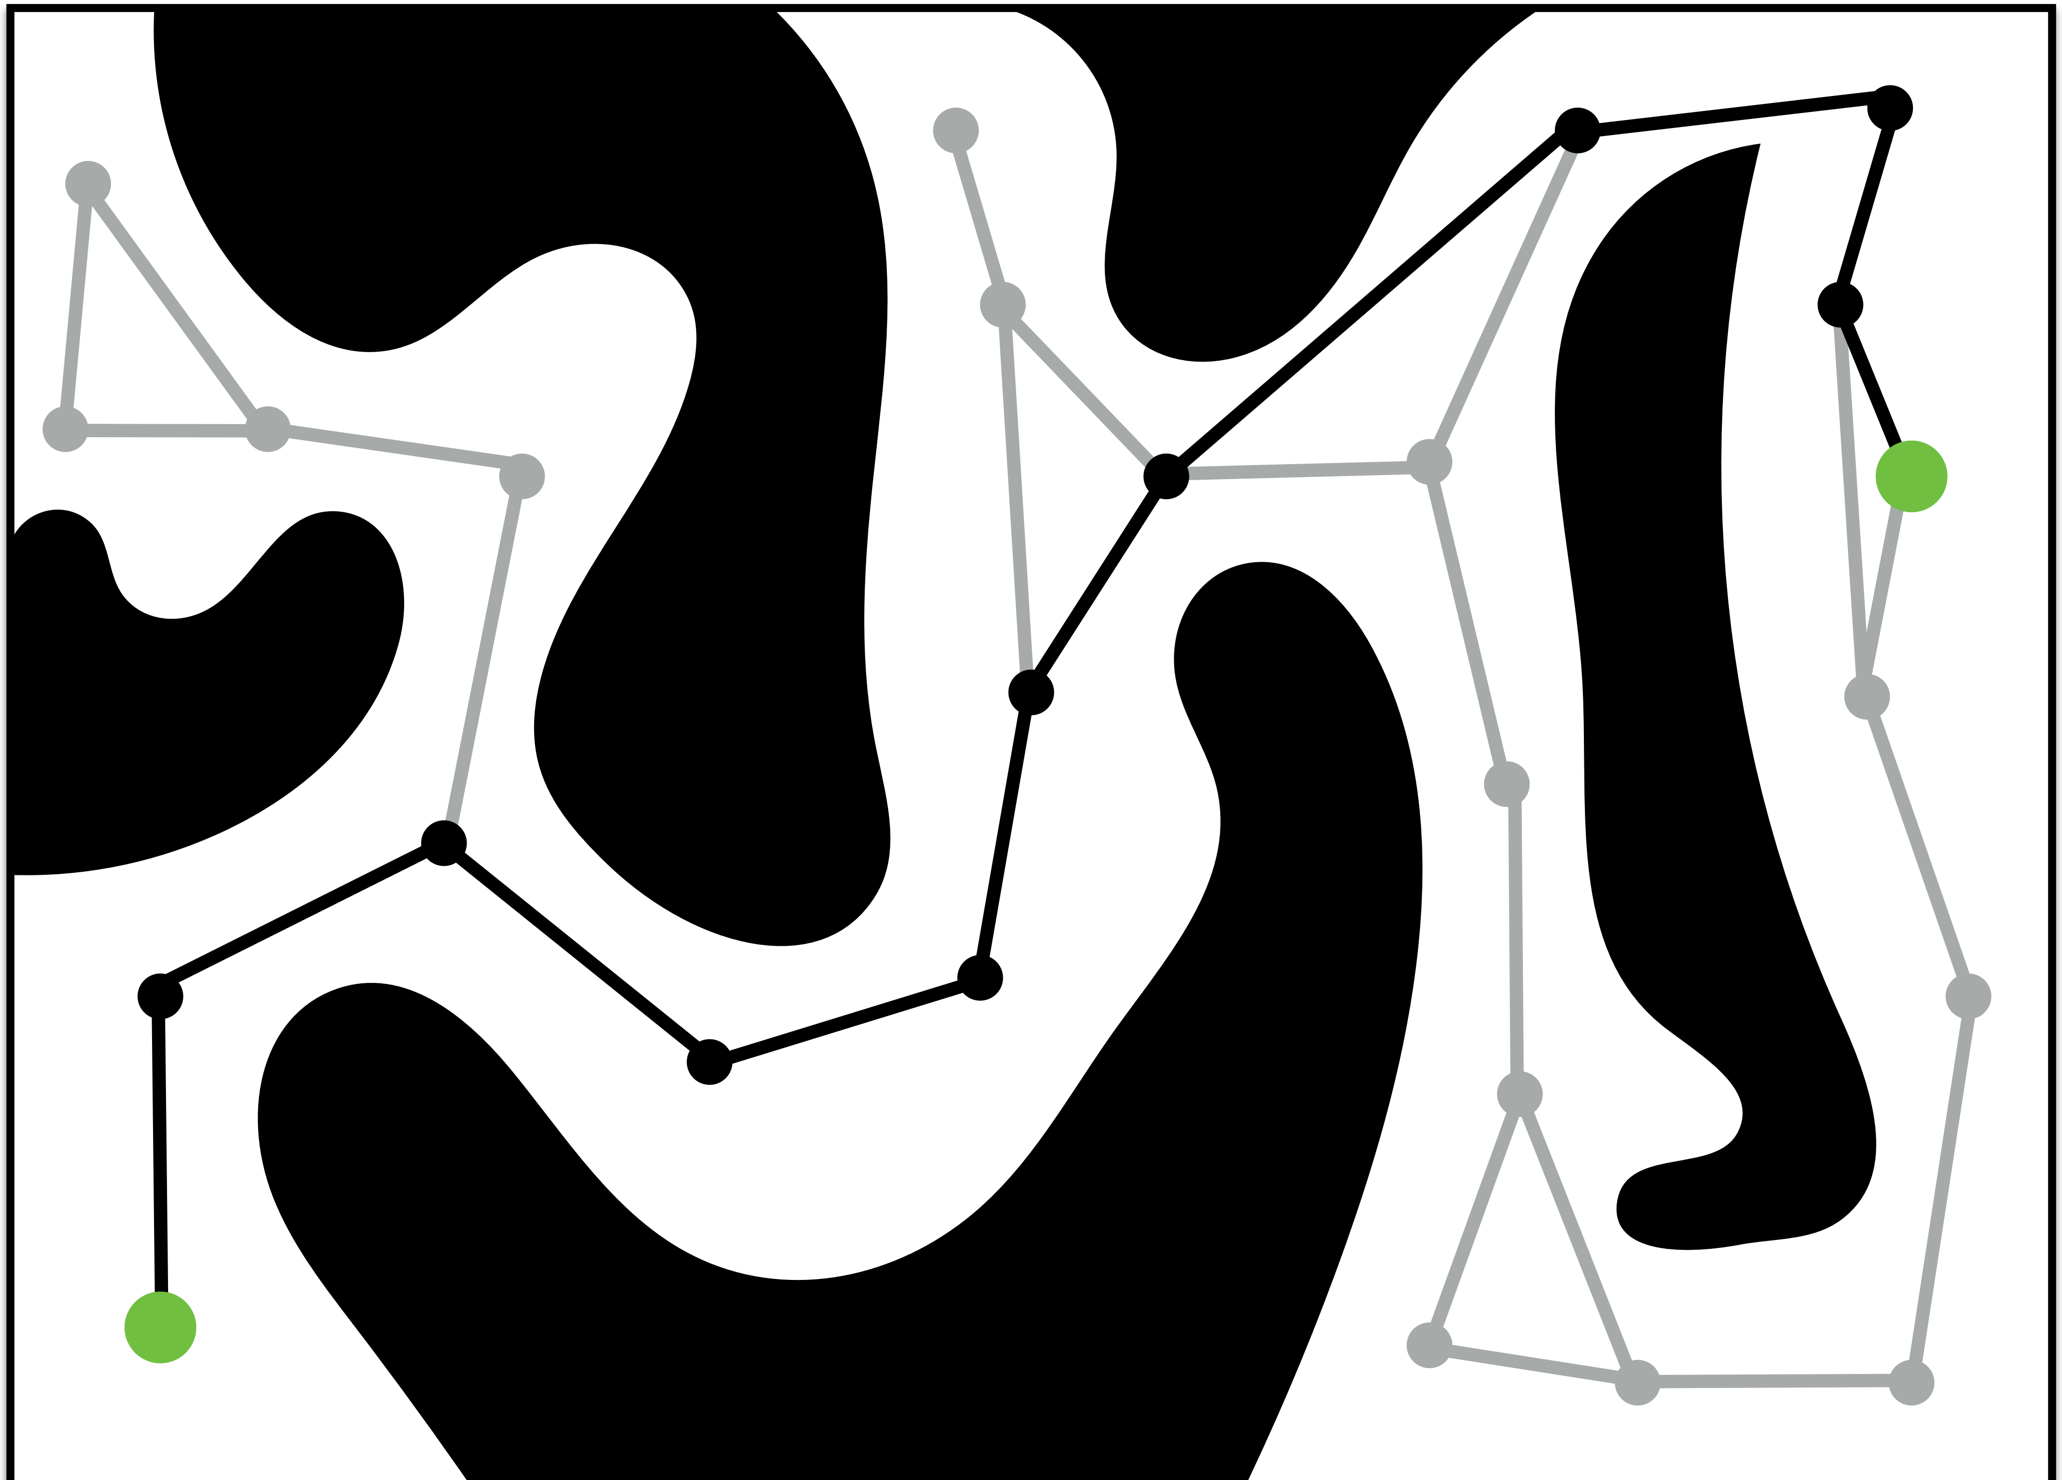

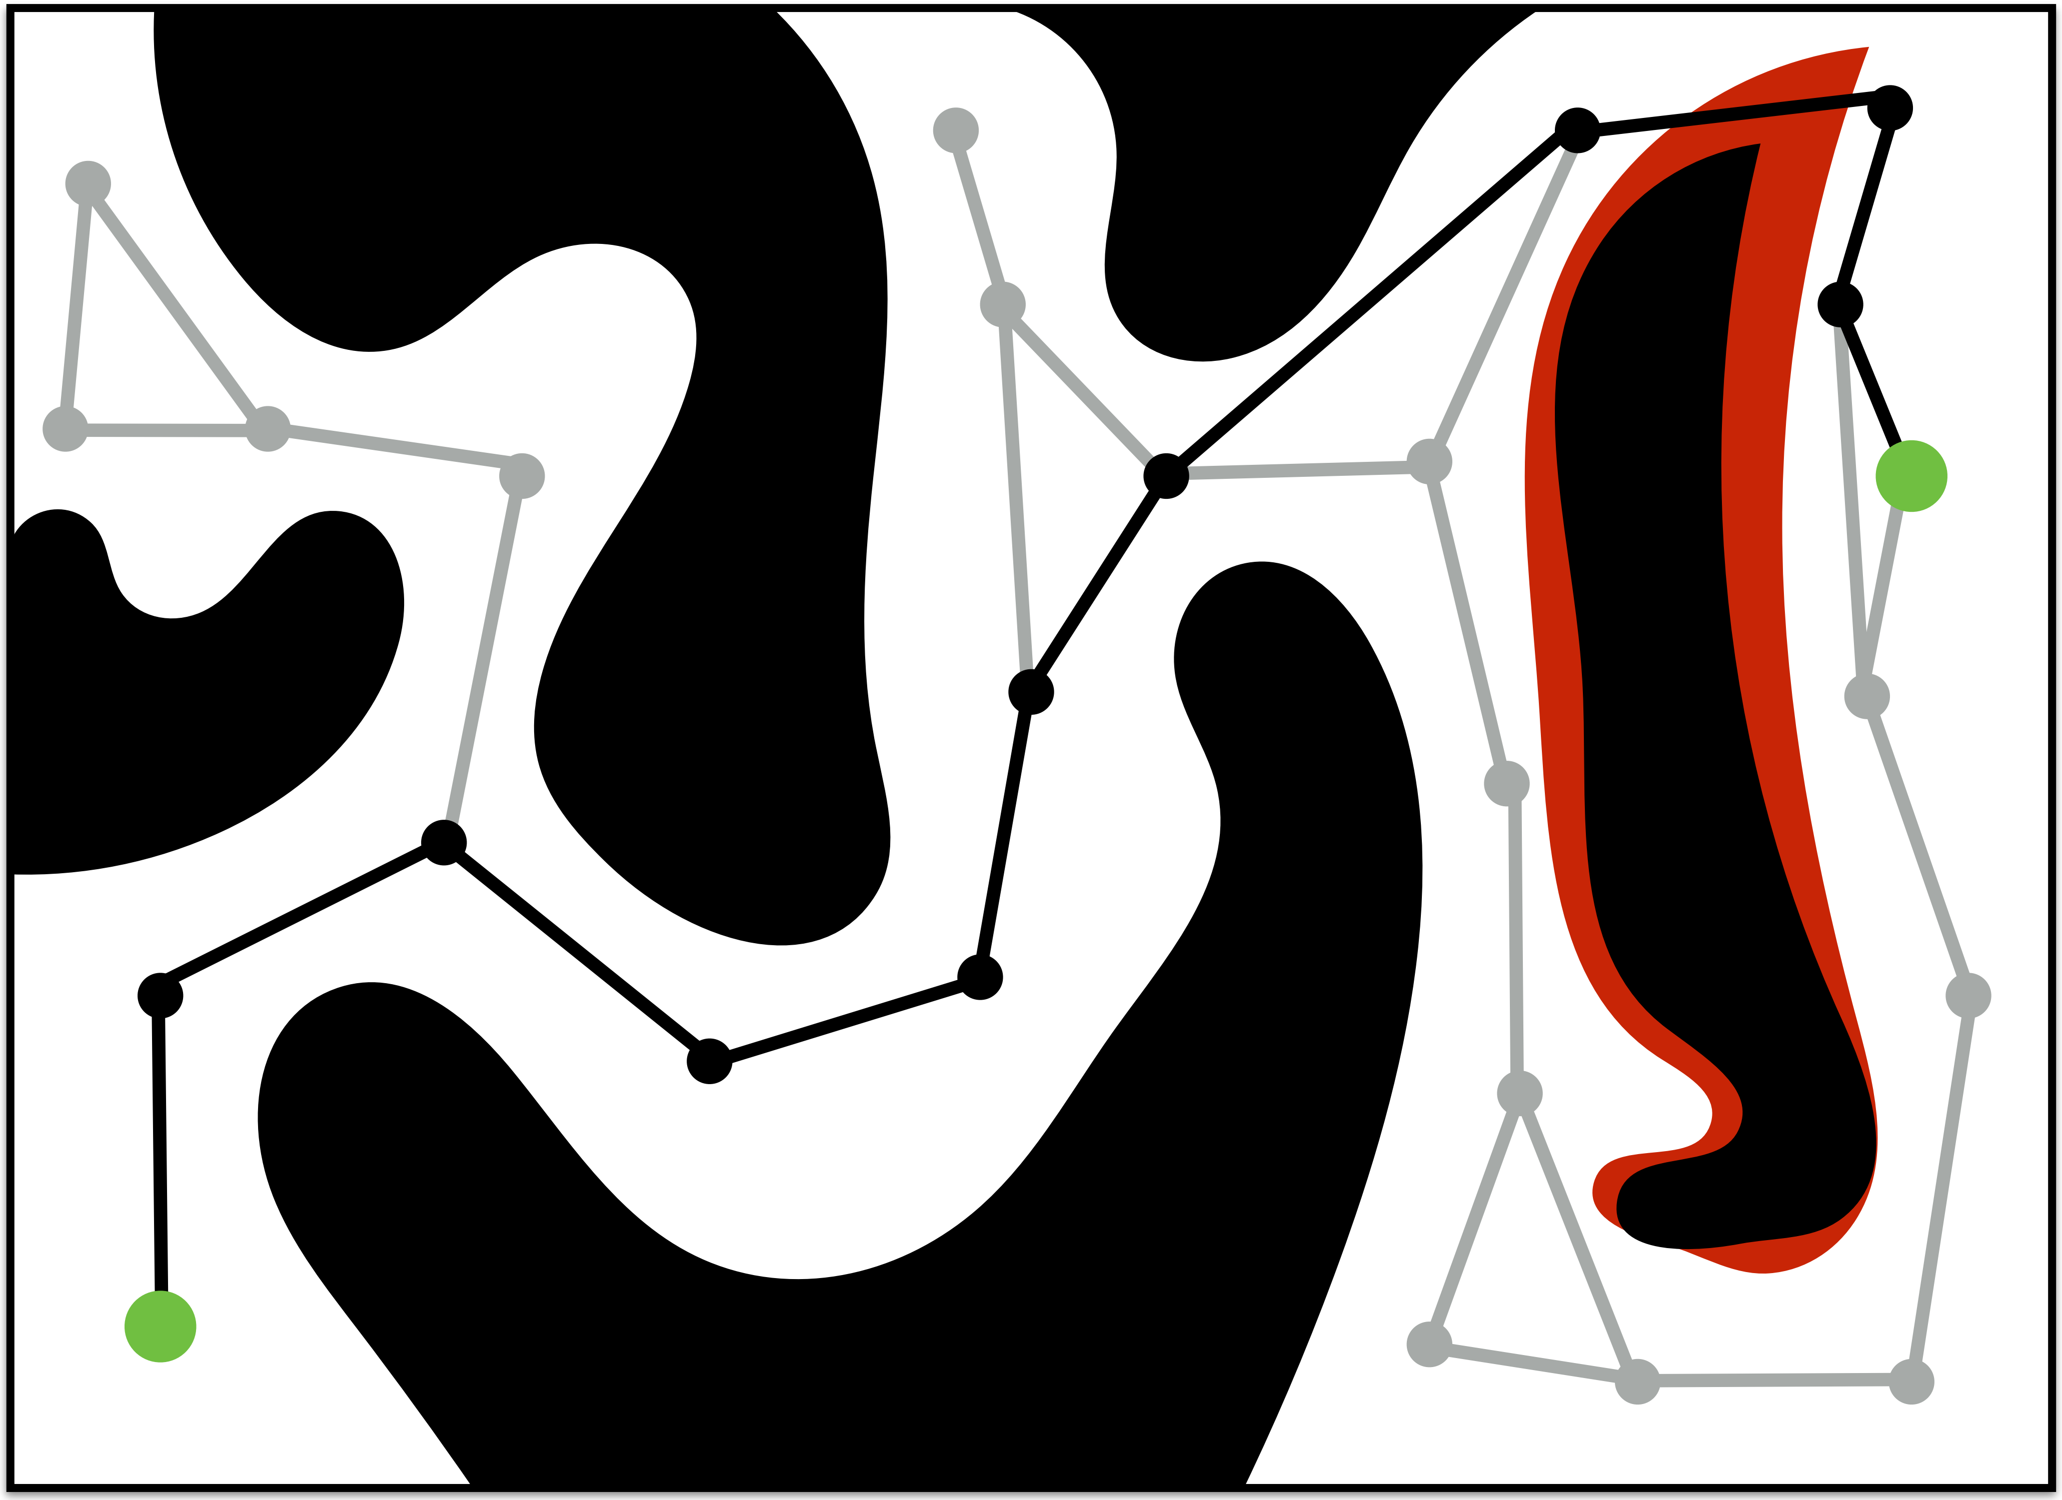

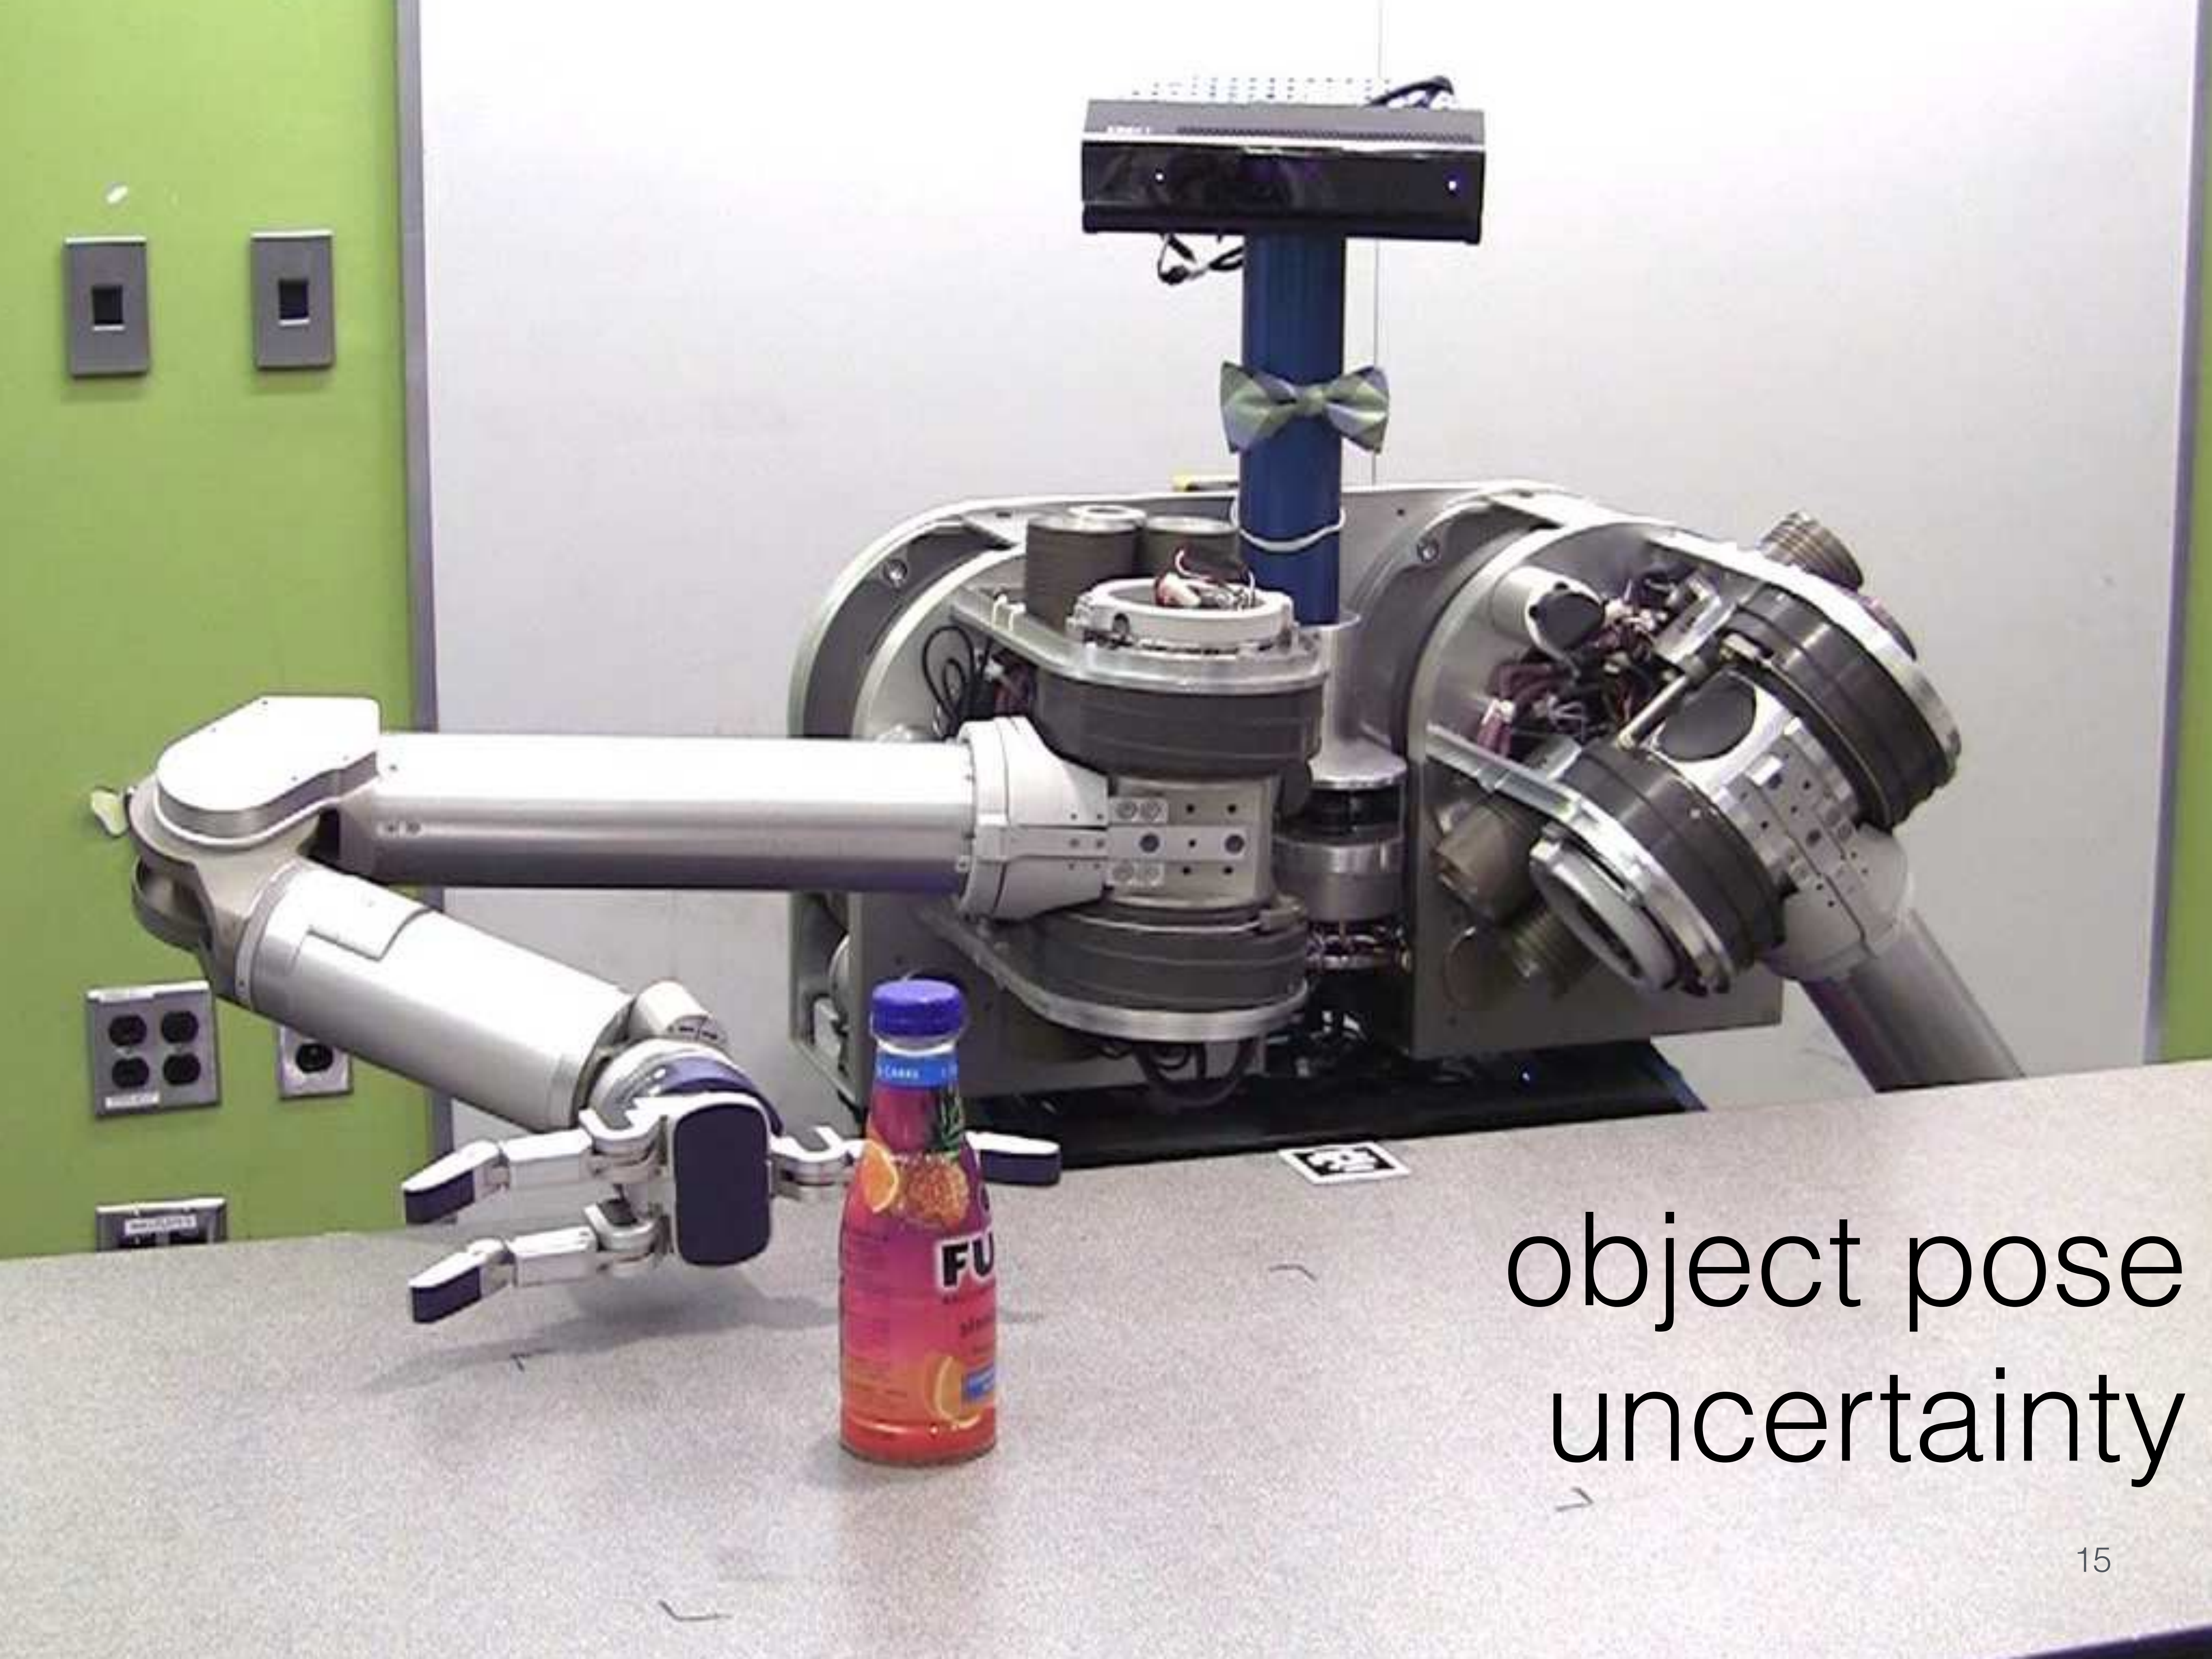

object pose uncertainty

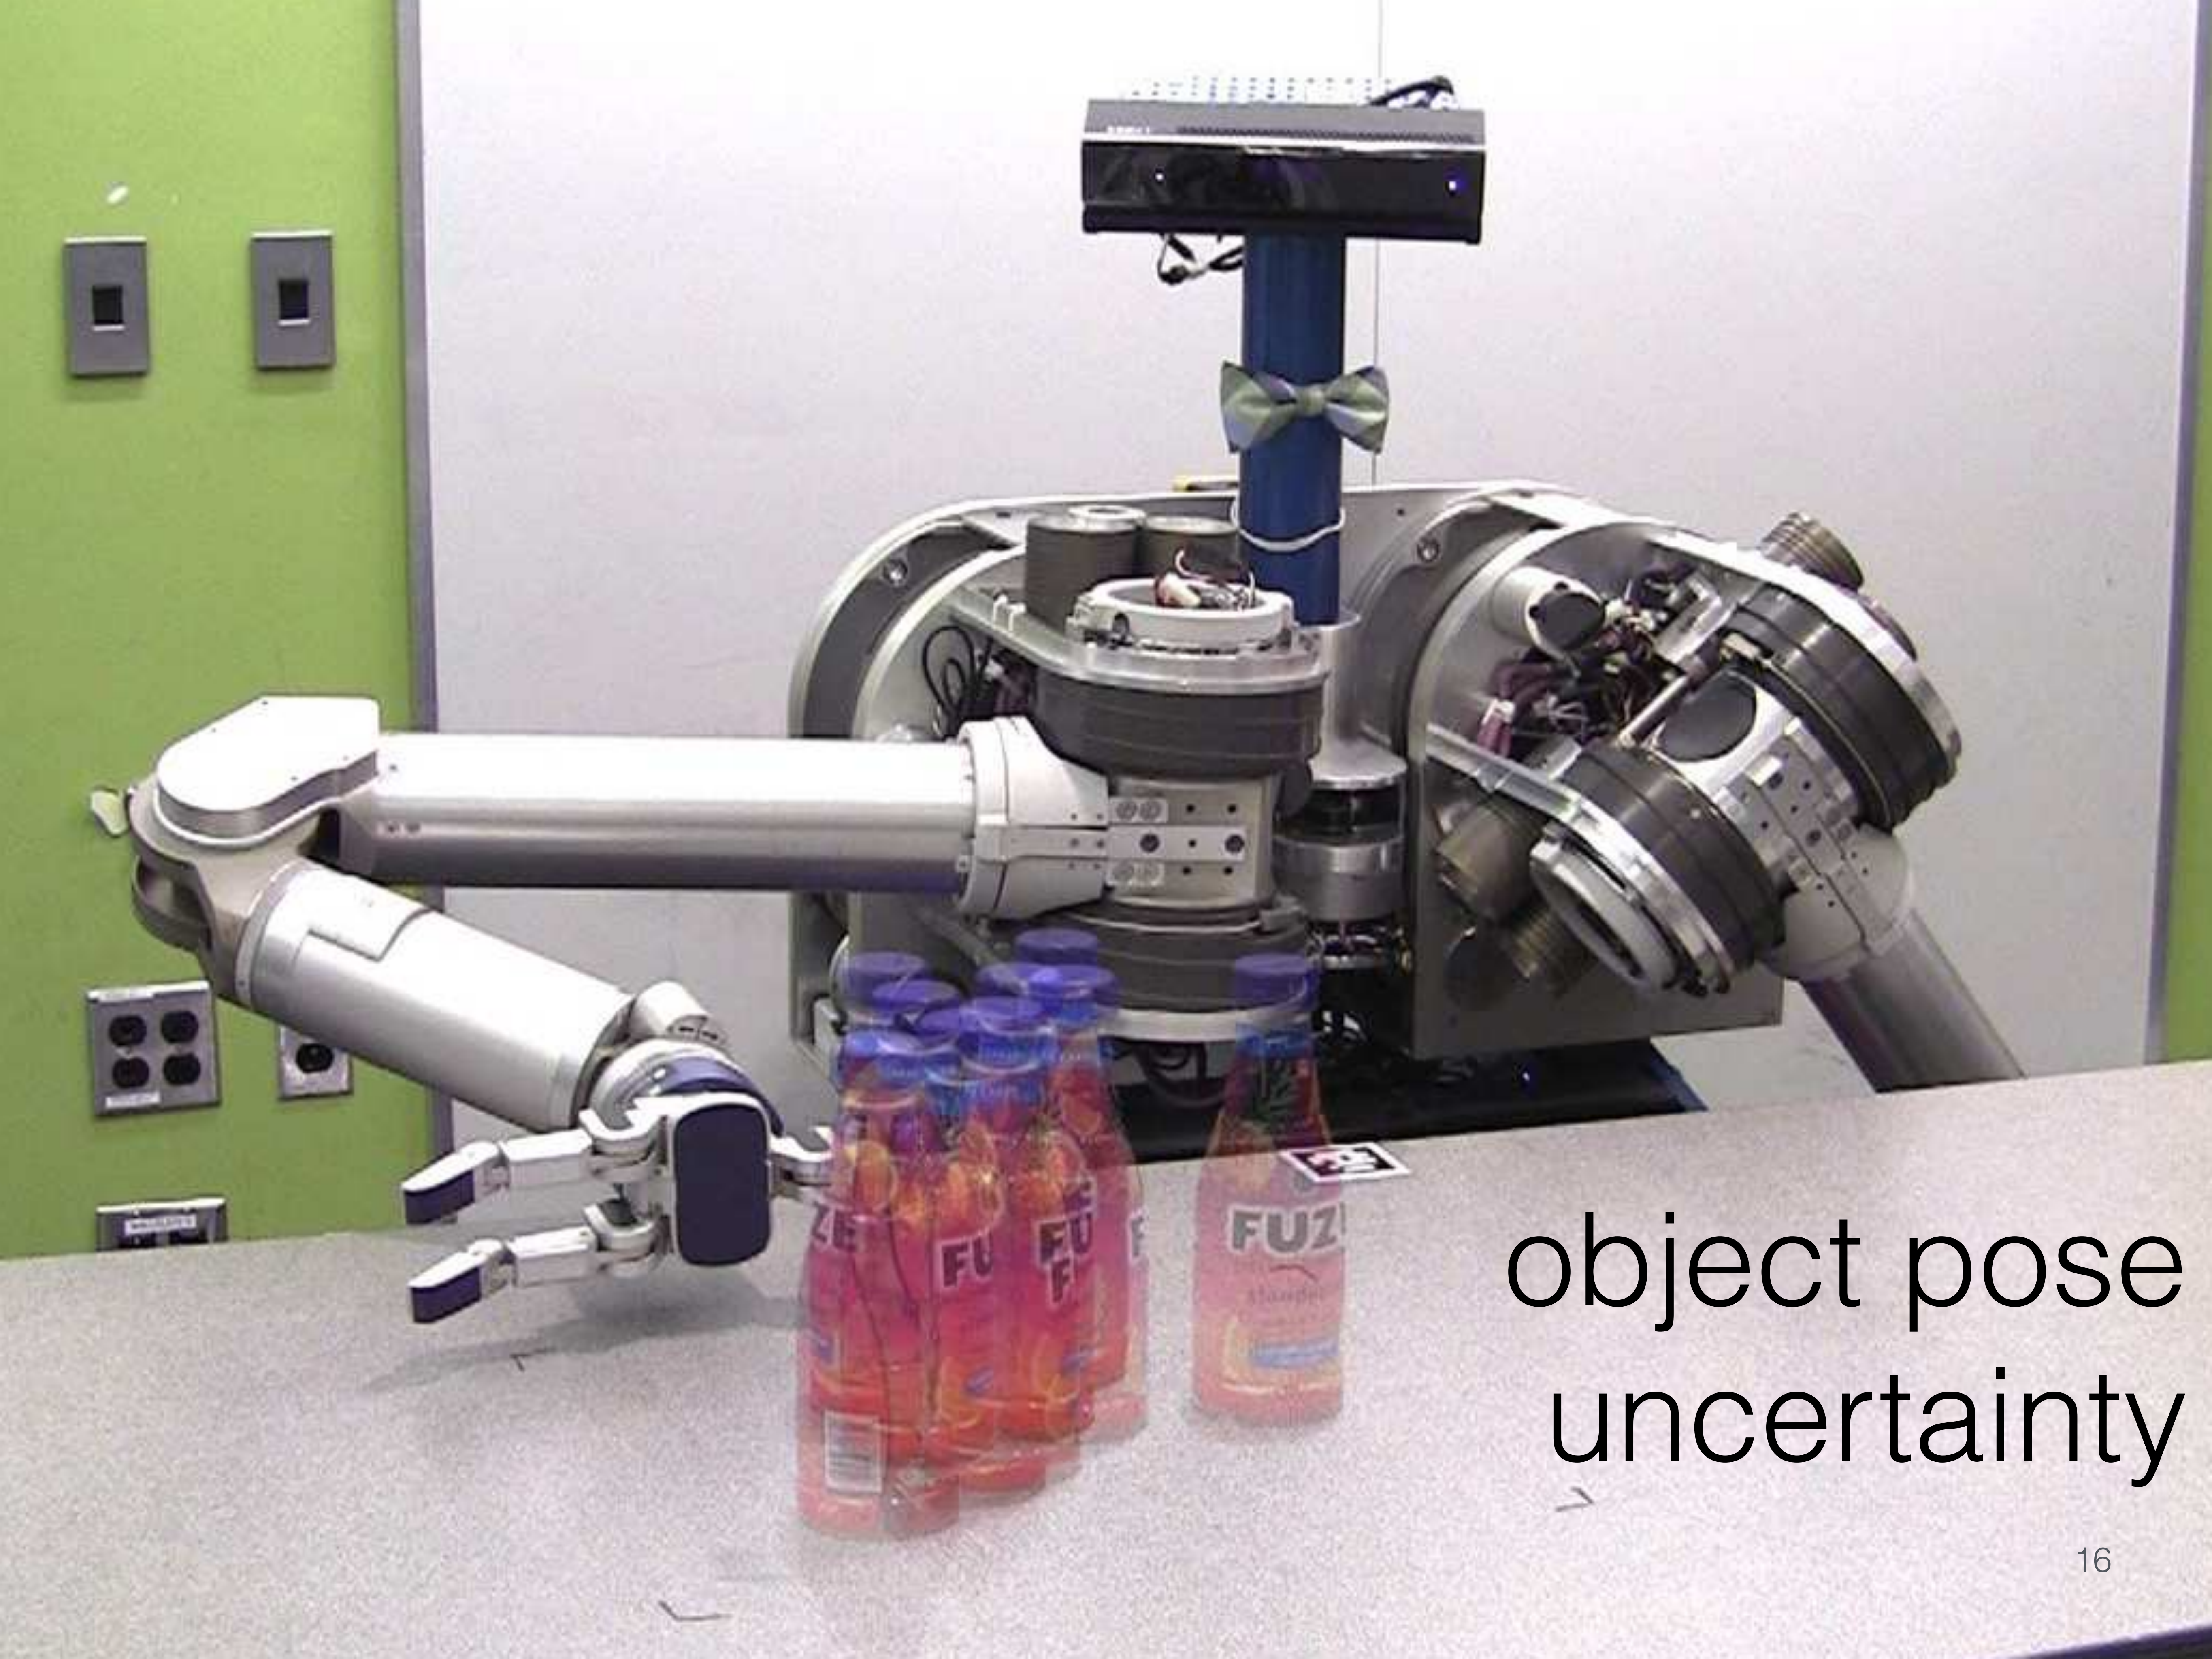

object pose uncertainty

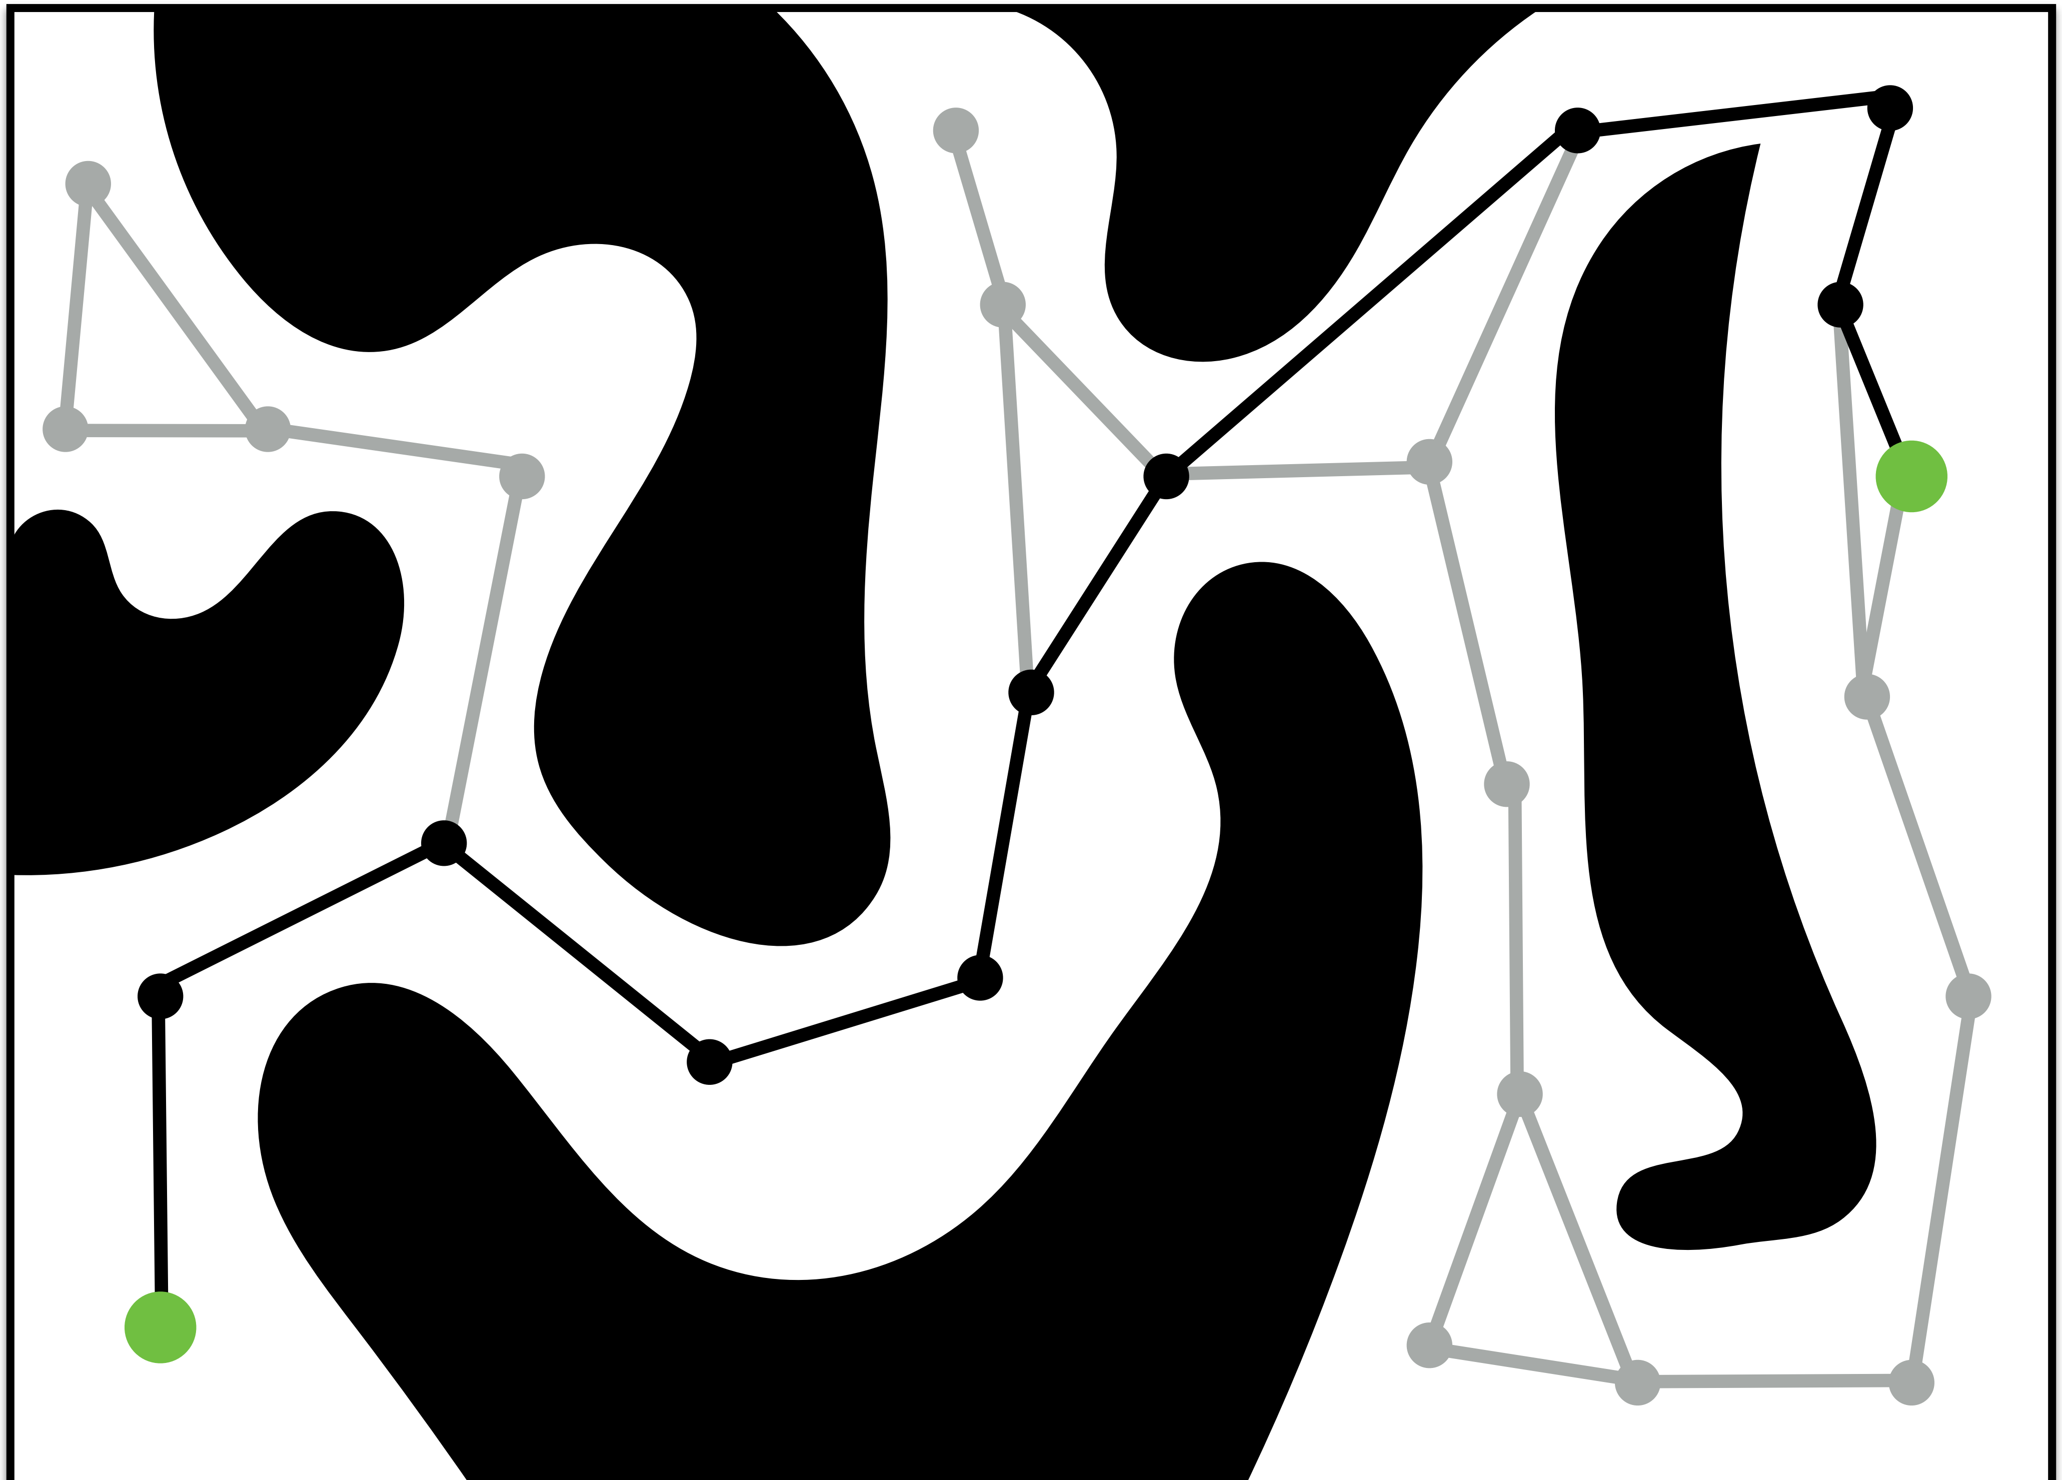

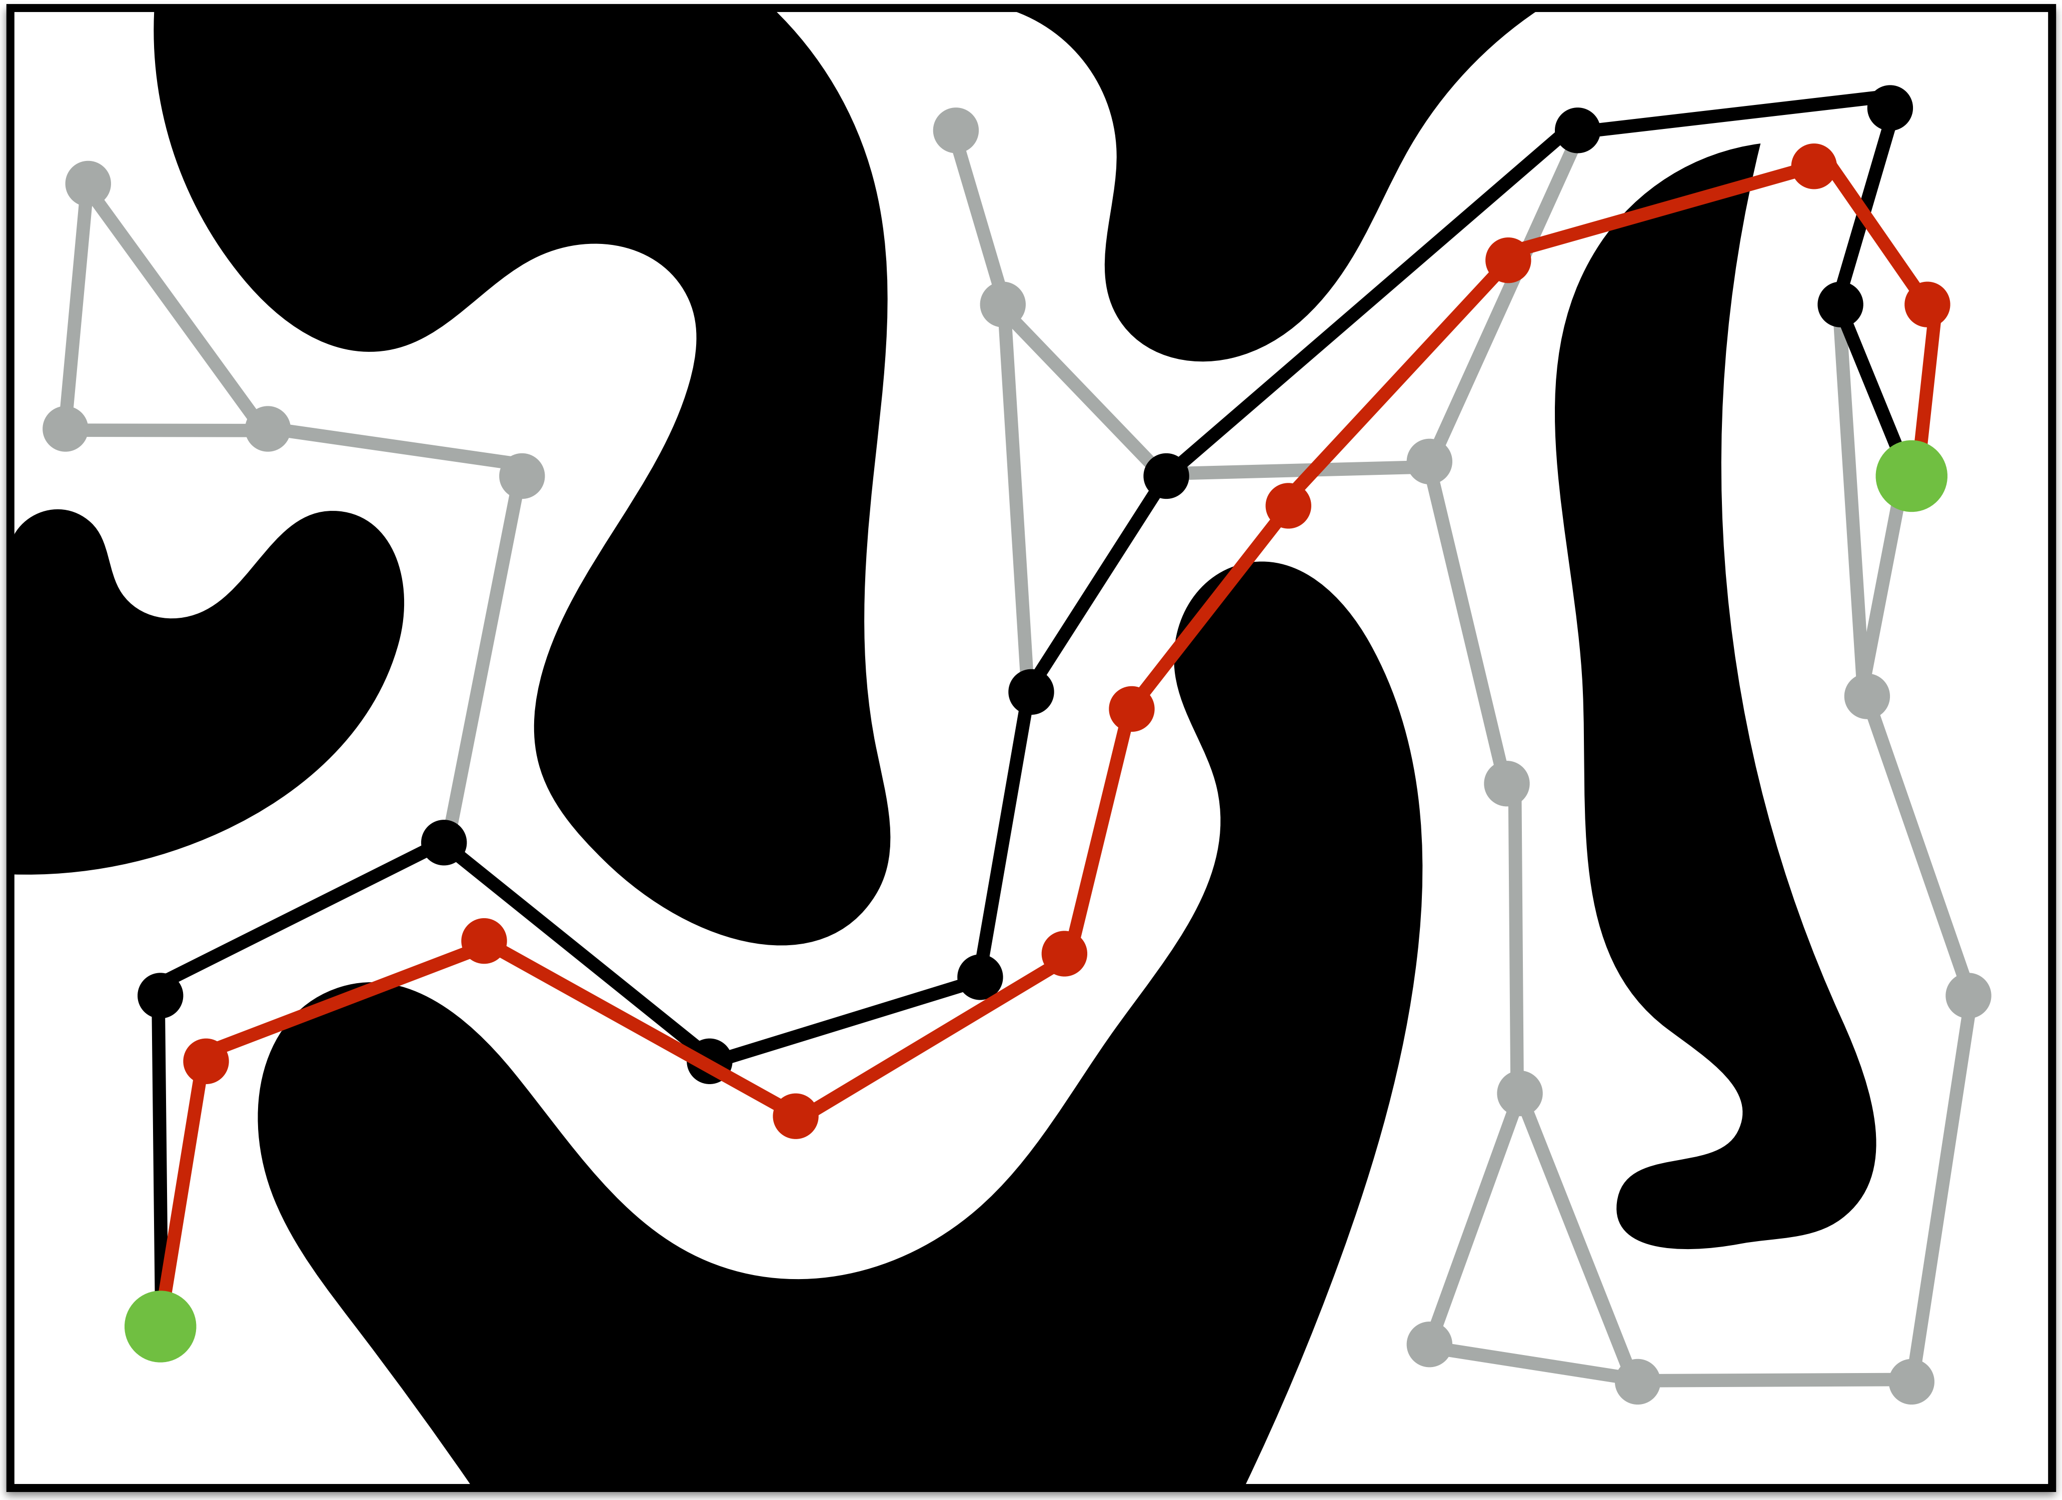

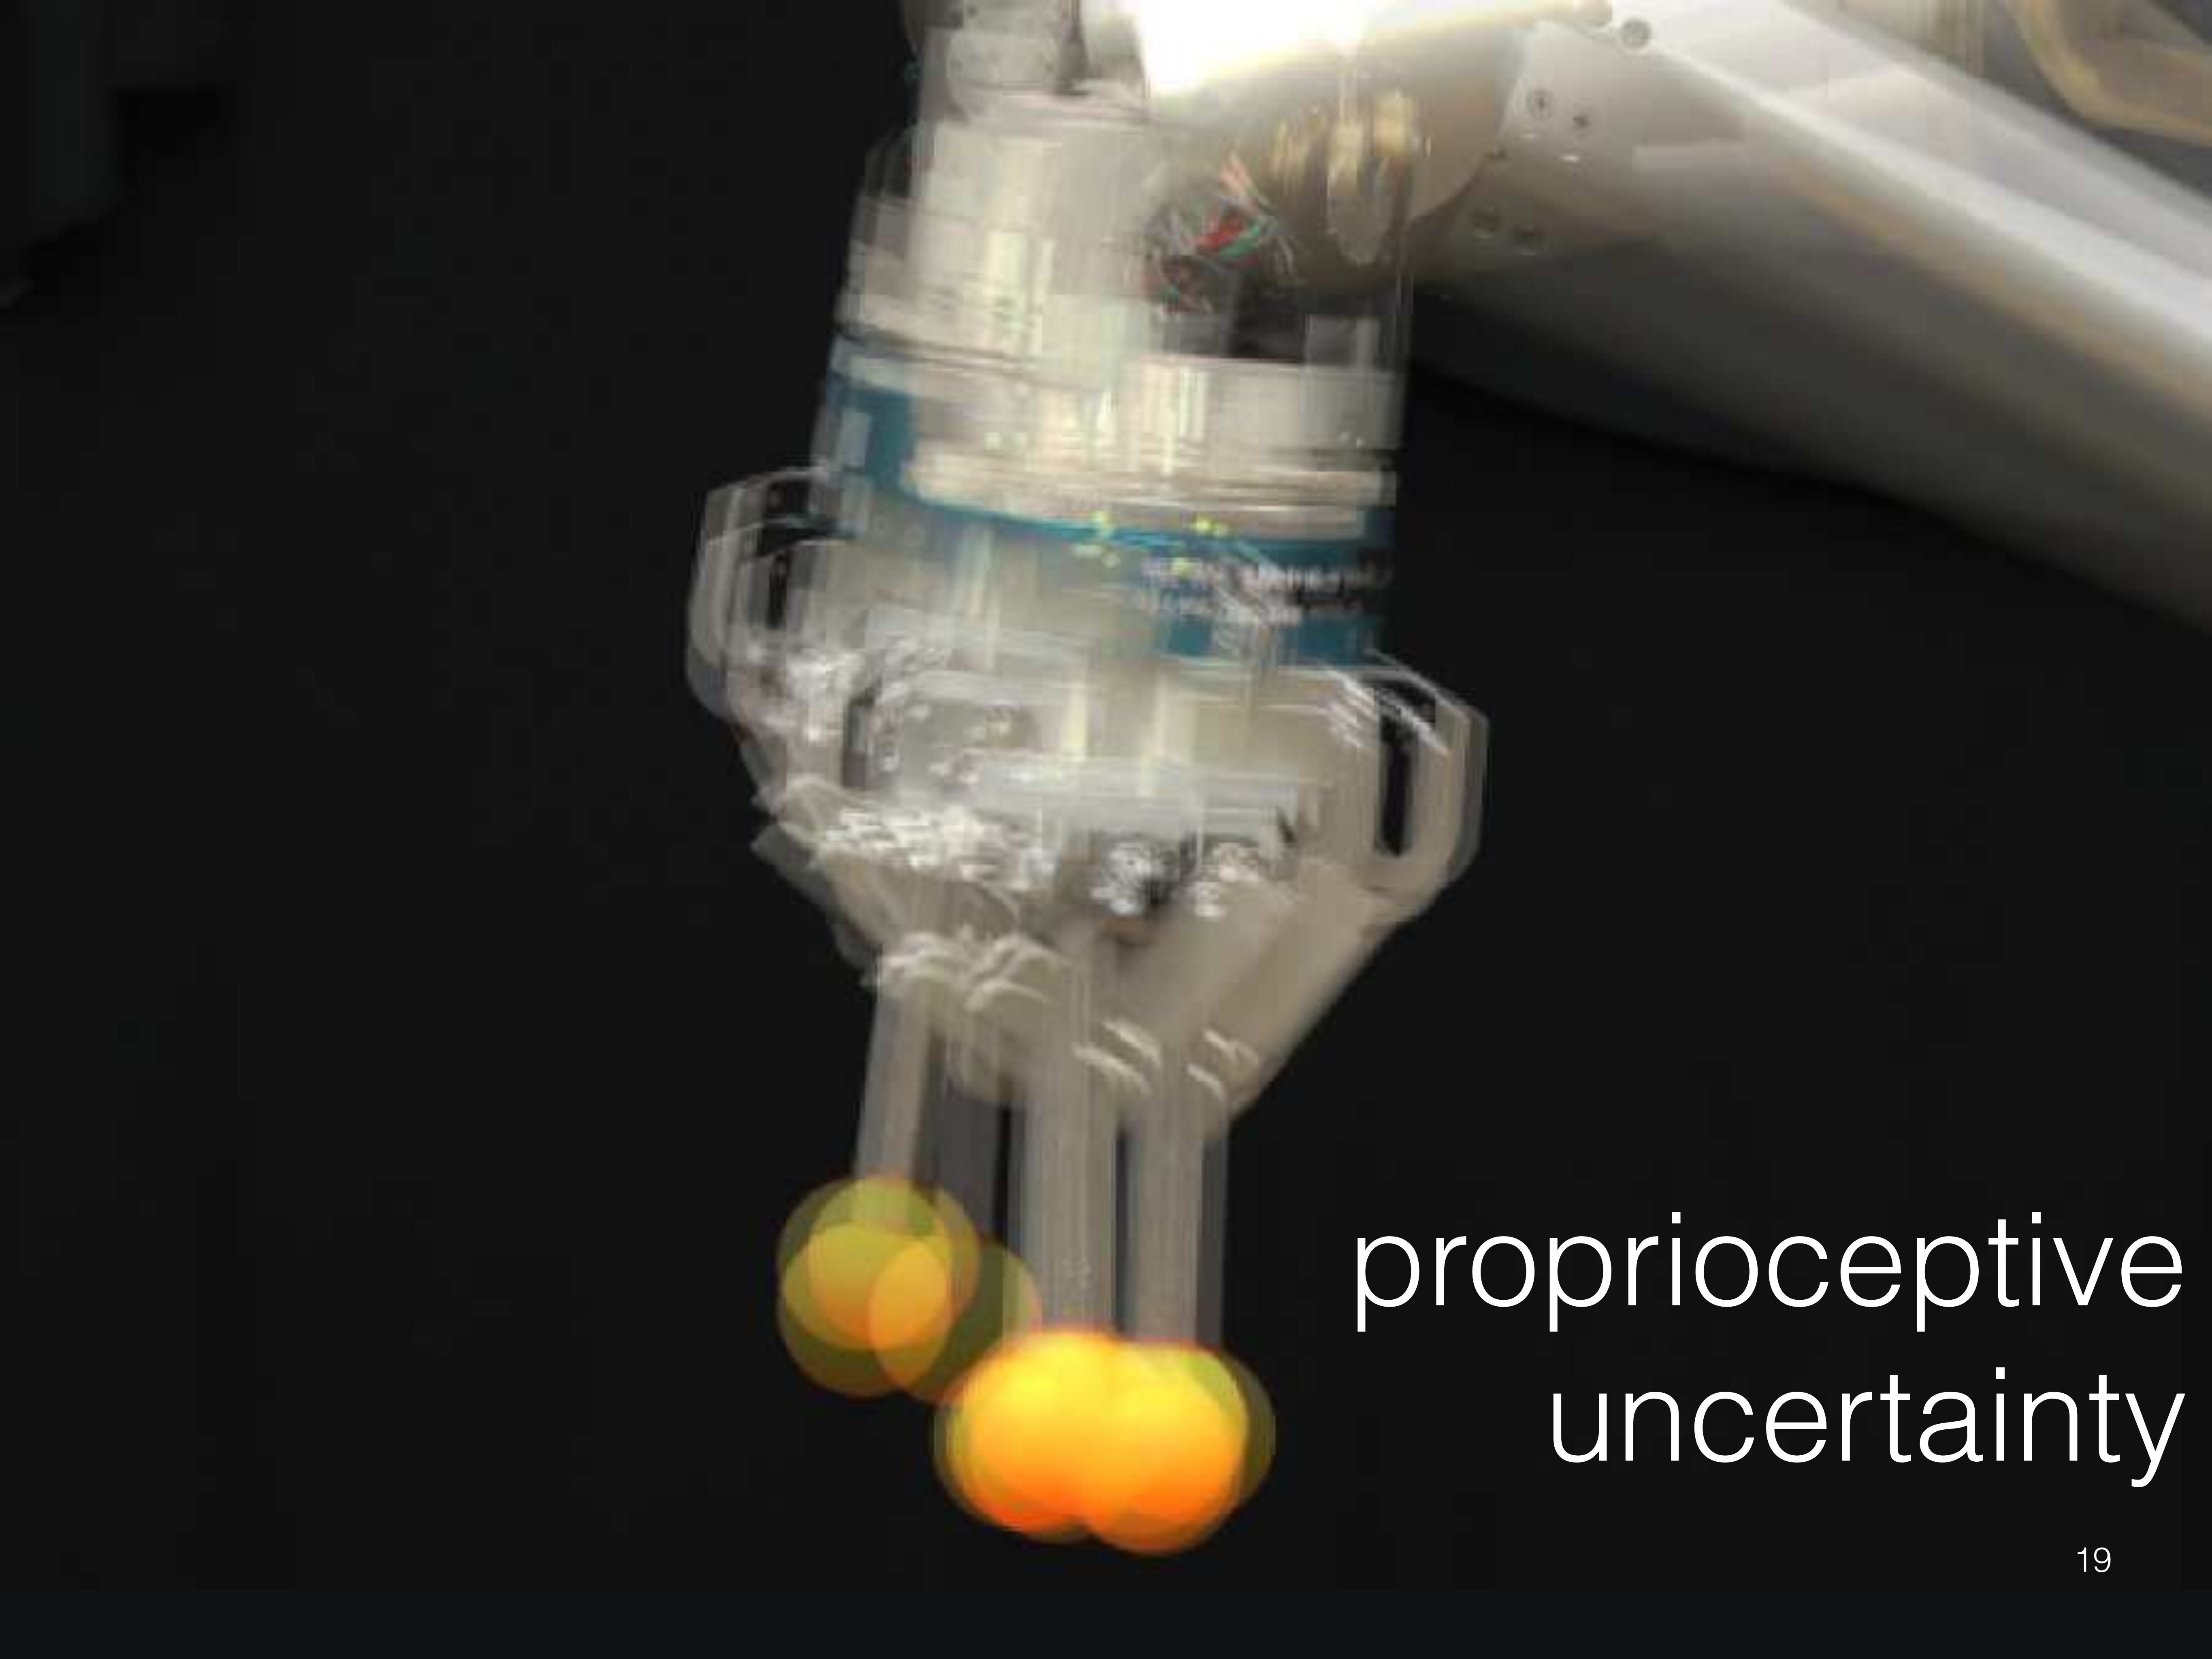
proprioceptive
uncertainty


proprioceptive
uncertainty

model uncertainty

K. Hauser. “Robust contact generation for robot simulation with unstructured meshes.” ISRR, 2013.

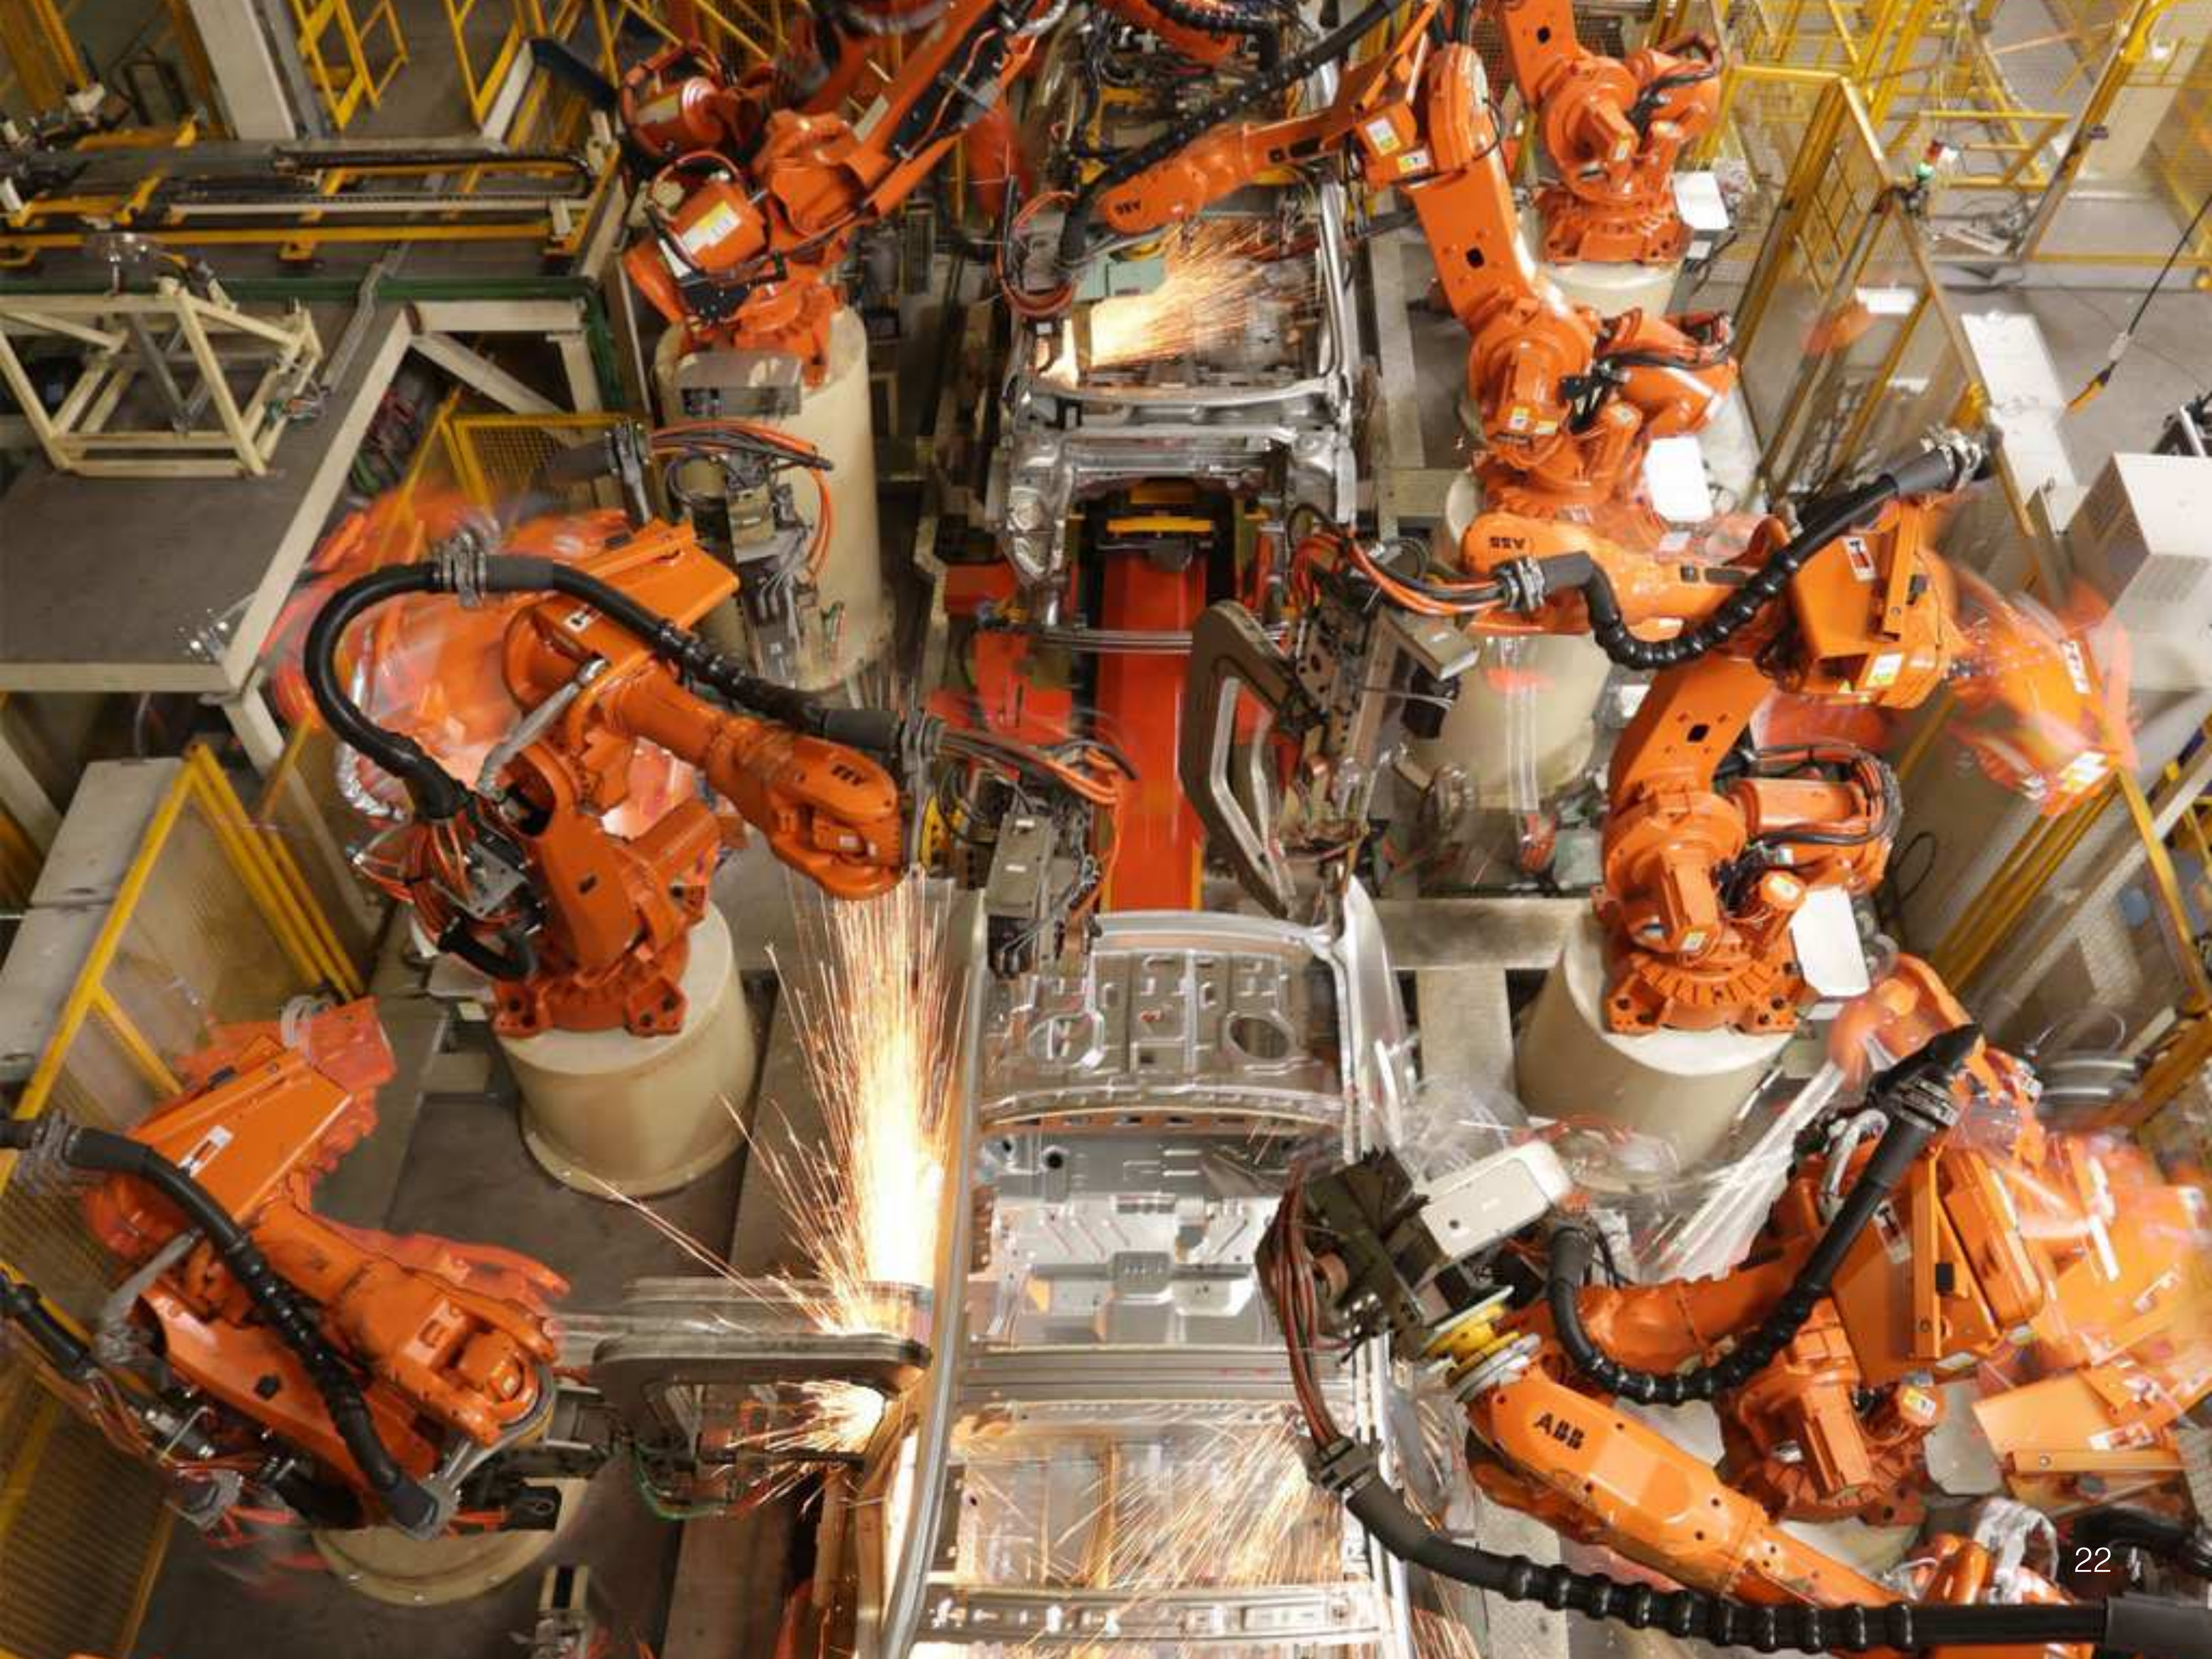
ABB

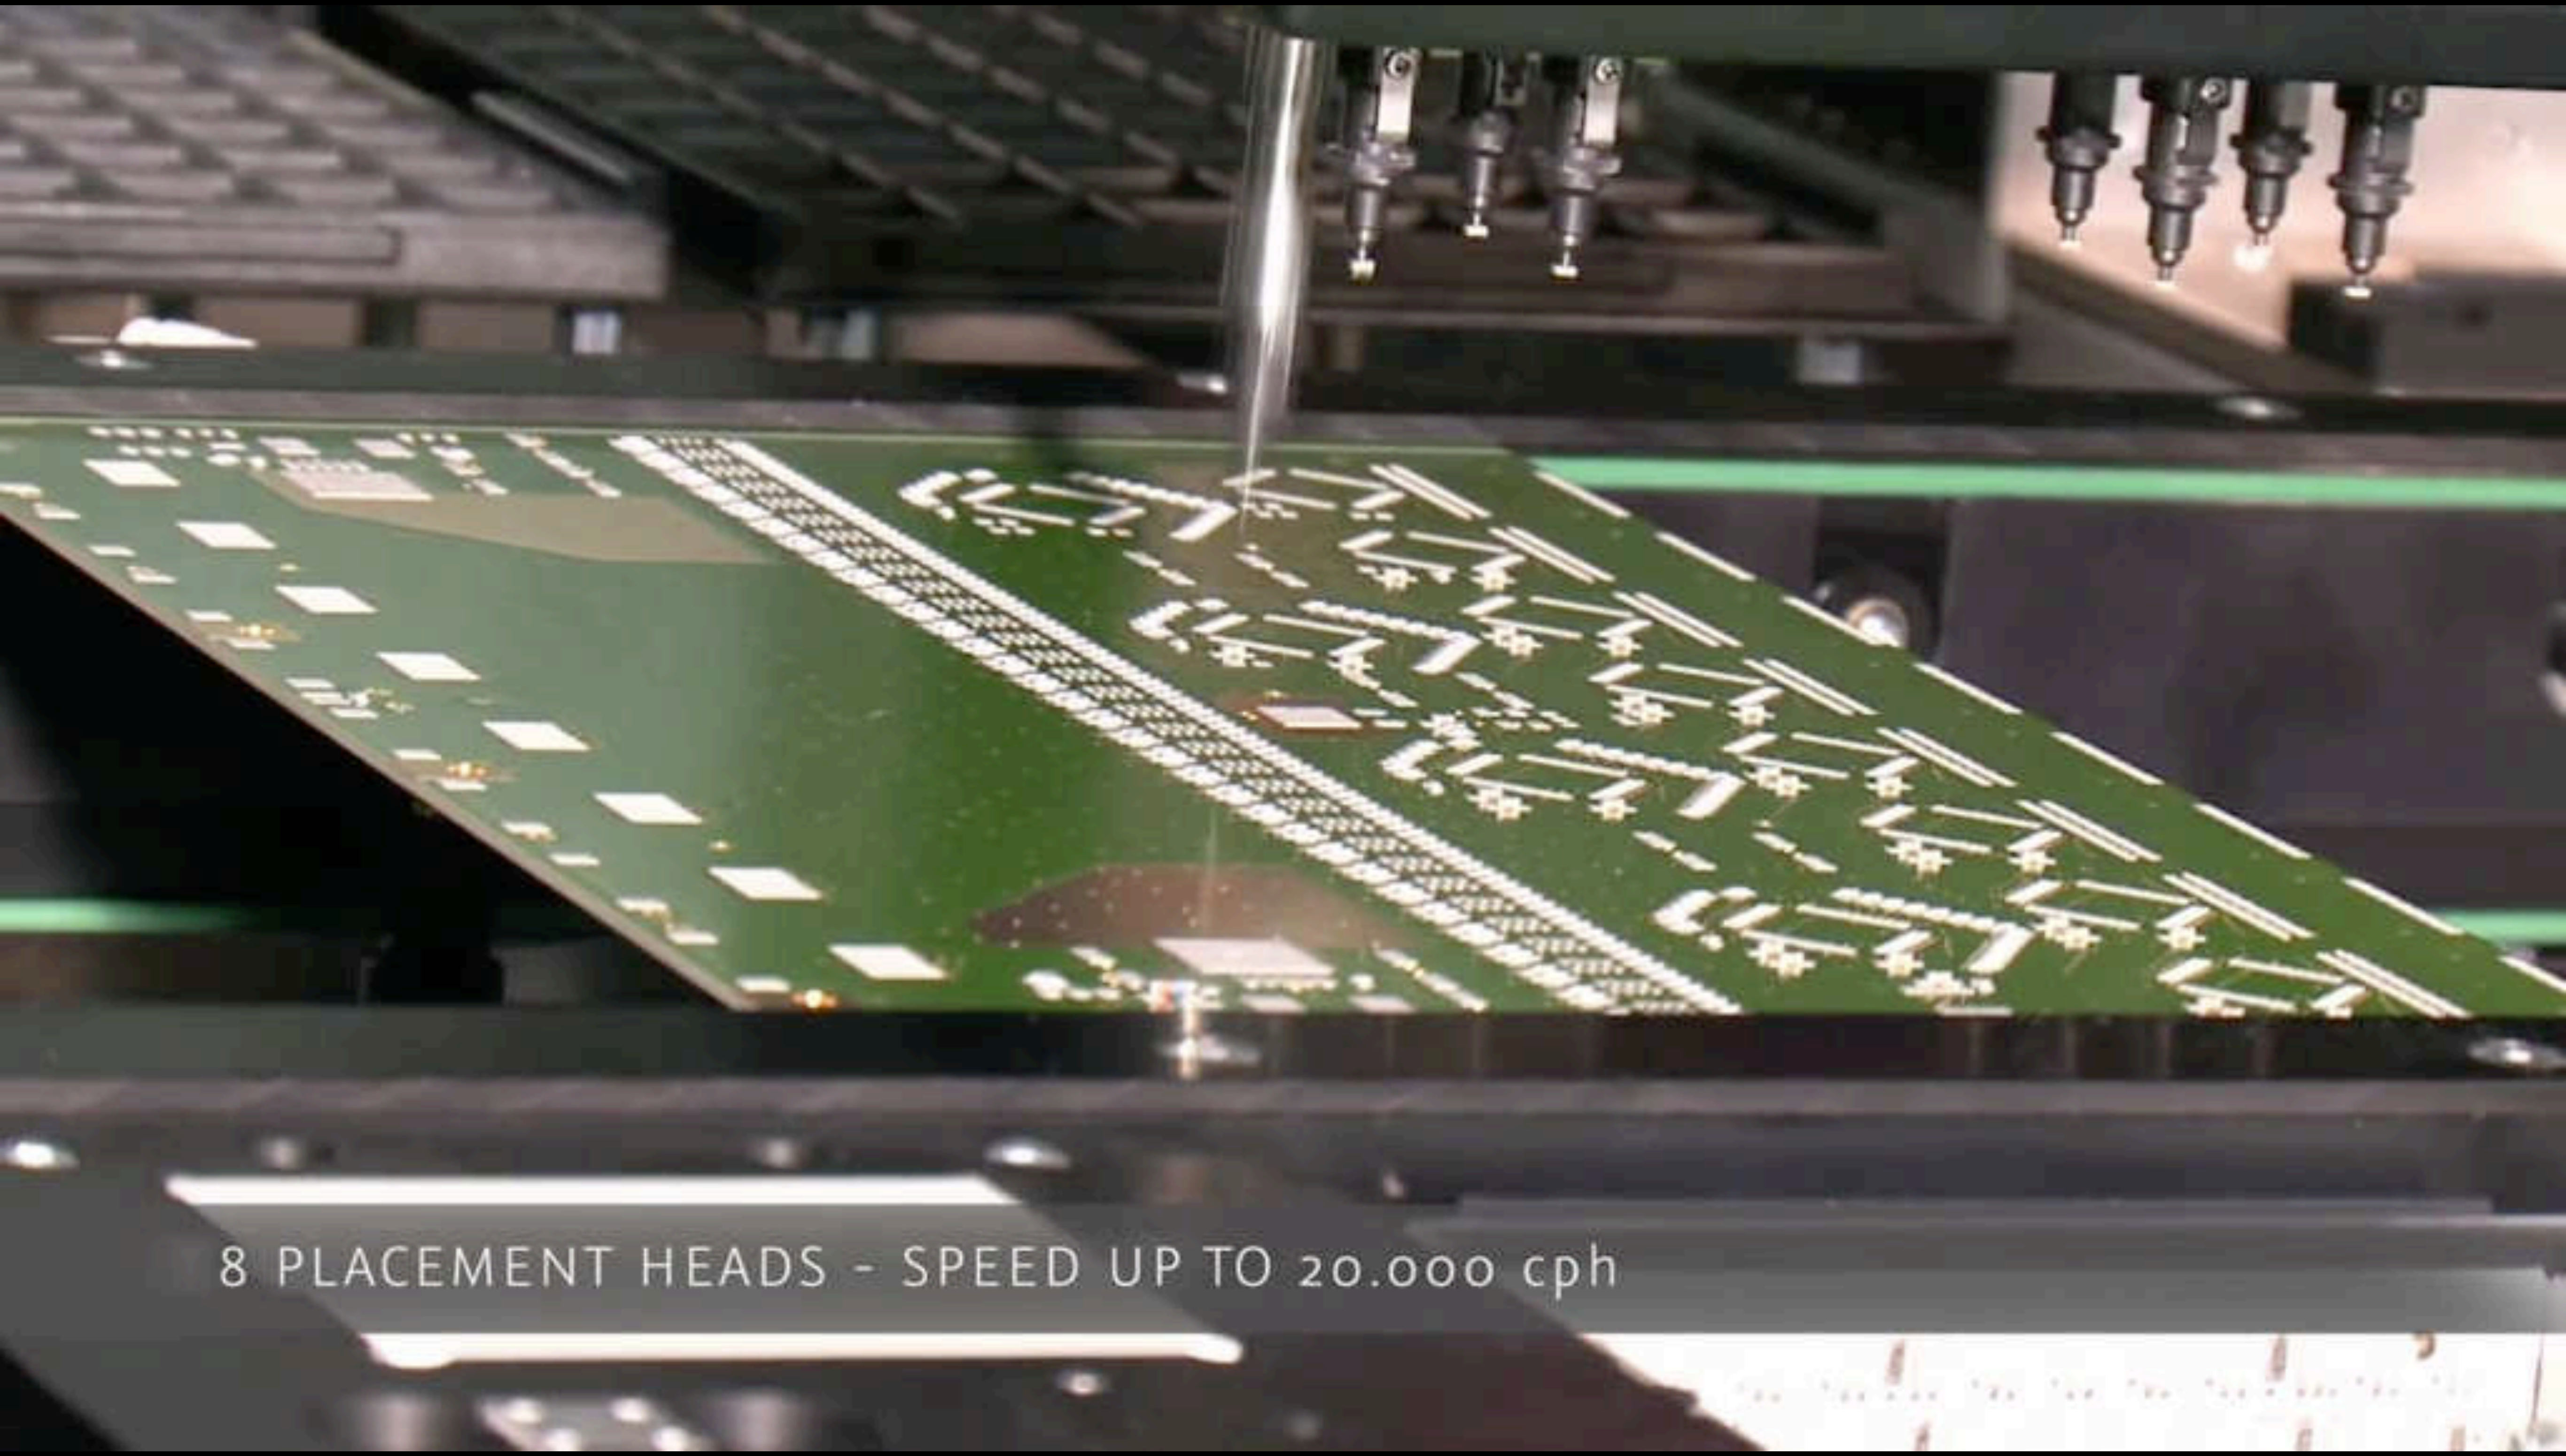

Source: https://youtu.be/S8qkaTsr2_o - ESSEMTEC pick and place machine

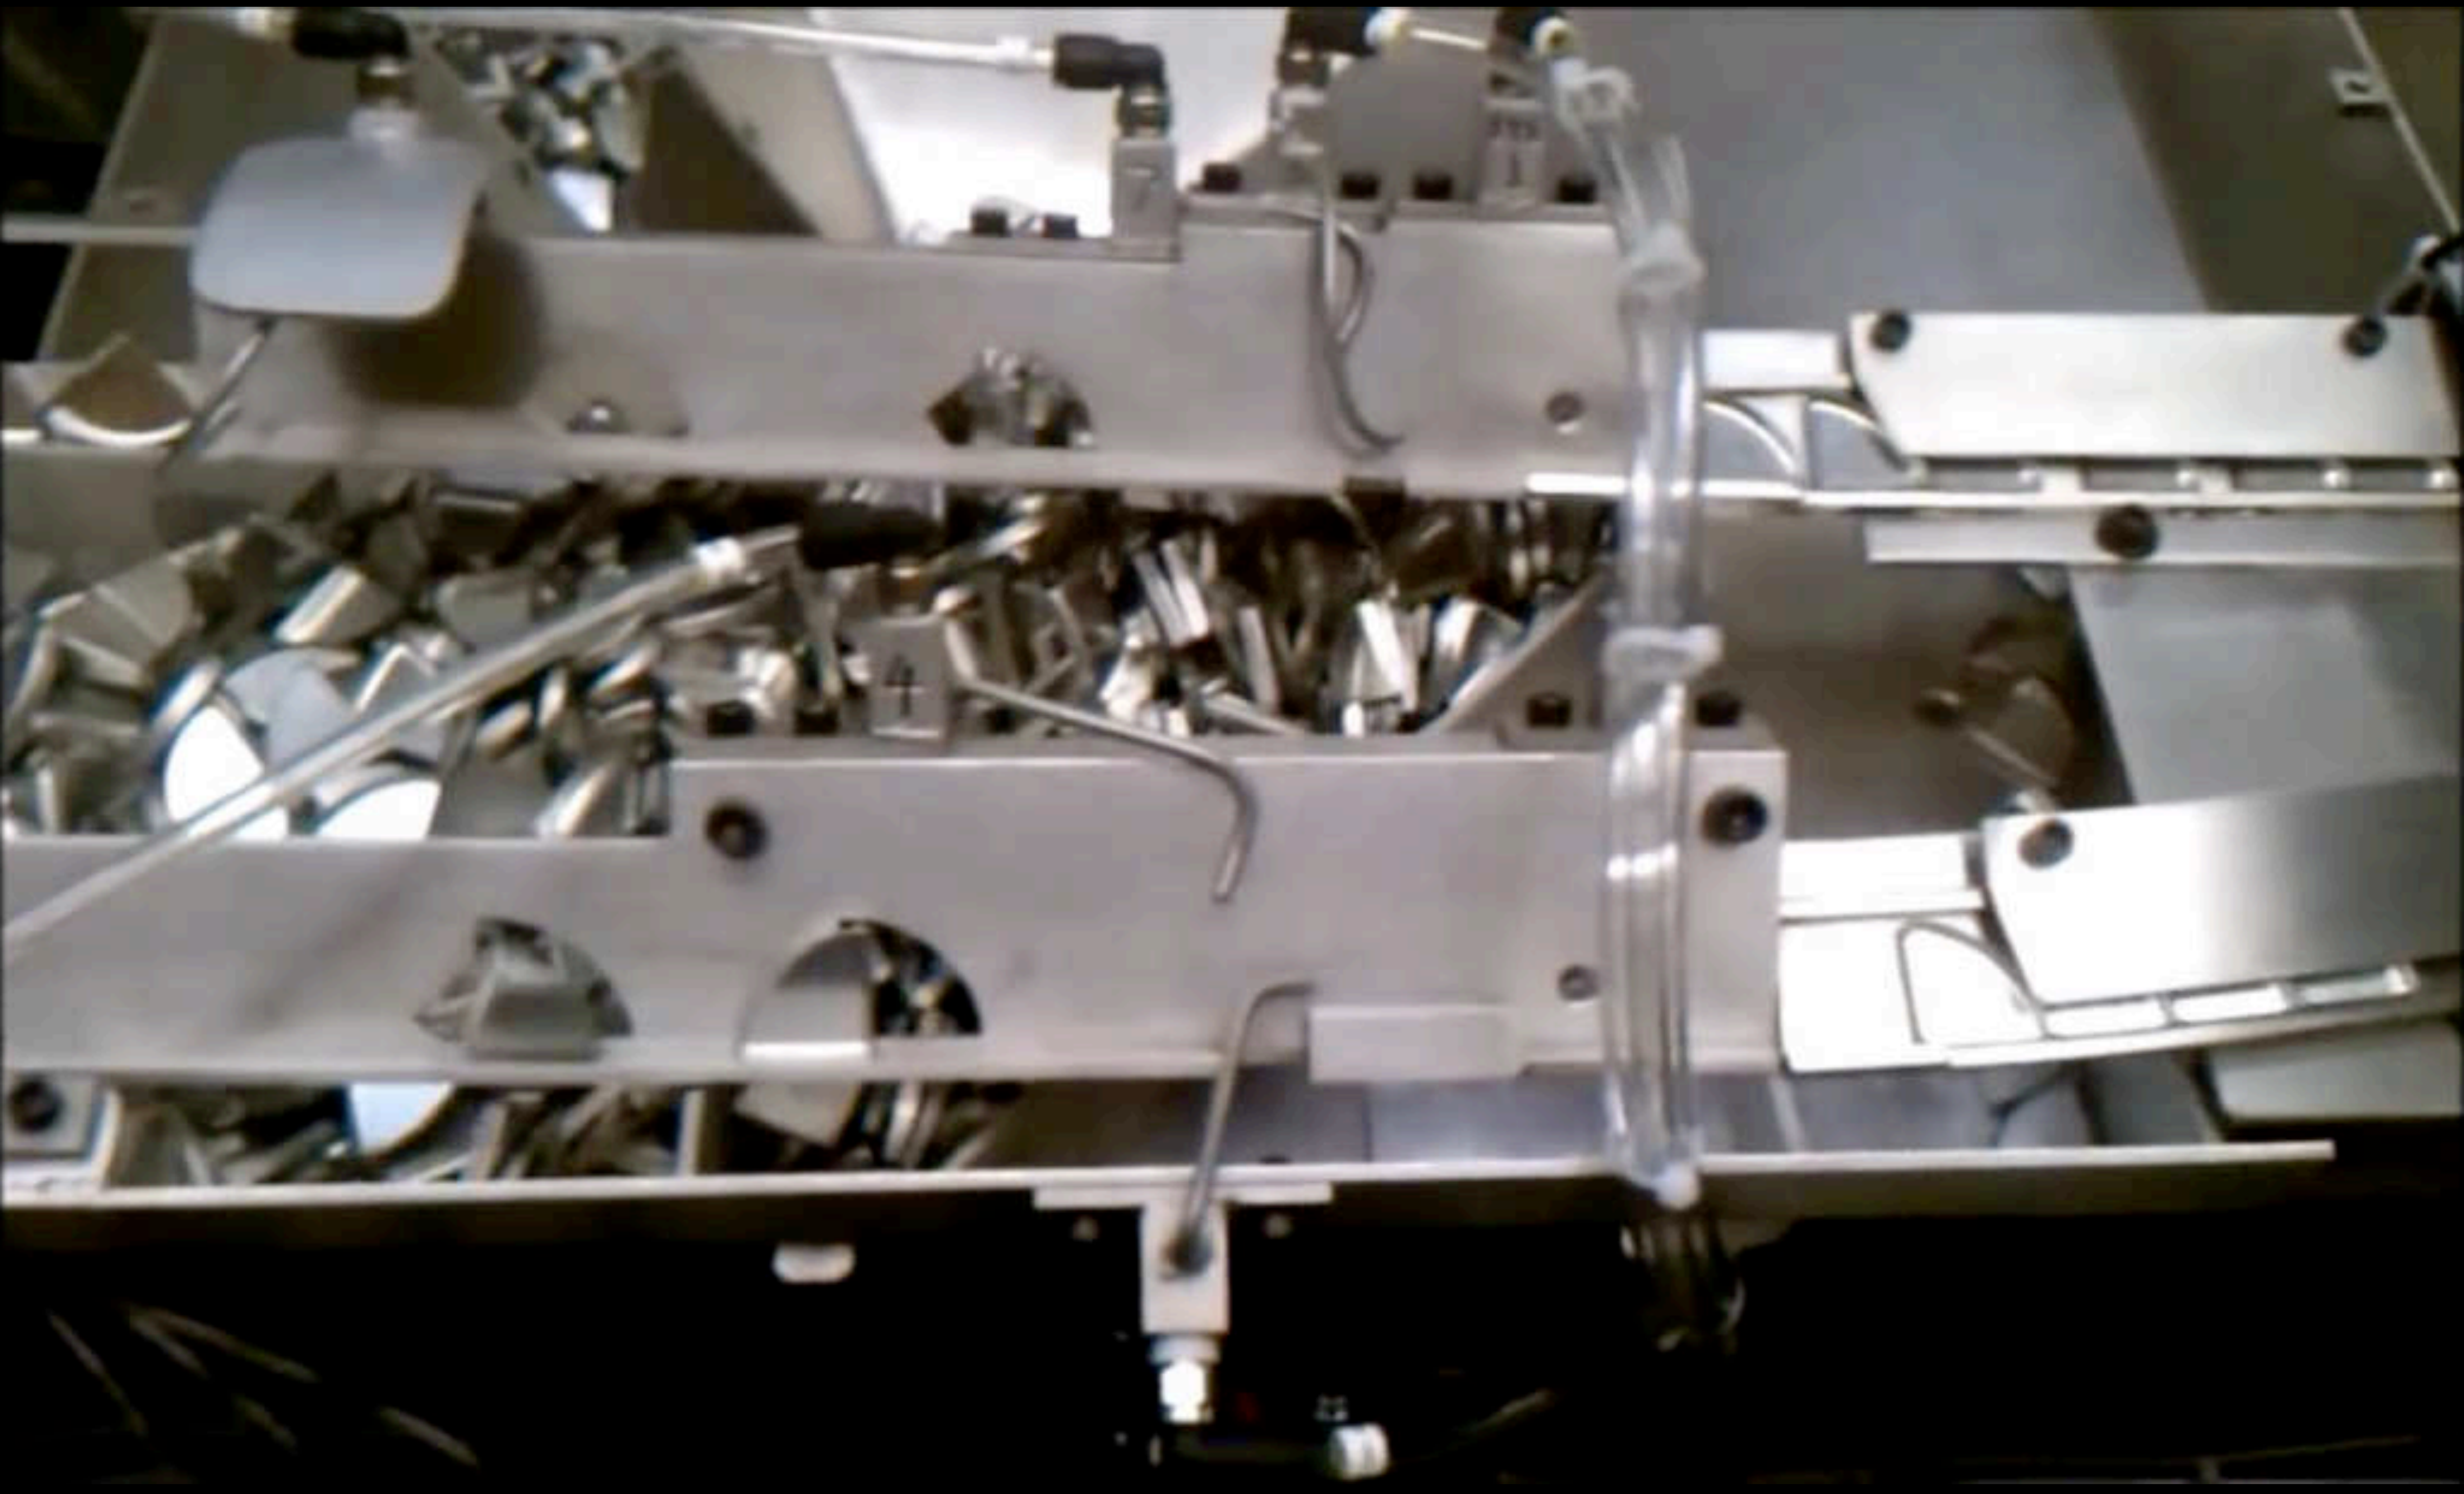

Source: https://youtu.be/yygM-MSxvew - ELCON part feeder machine

Source: https://youtu.be/nkLd45Ftfhc - ABB bottle packing robot

How can we manipulate under uncertainty?

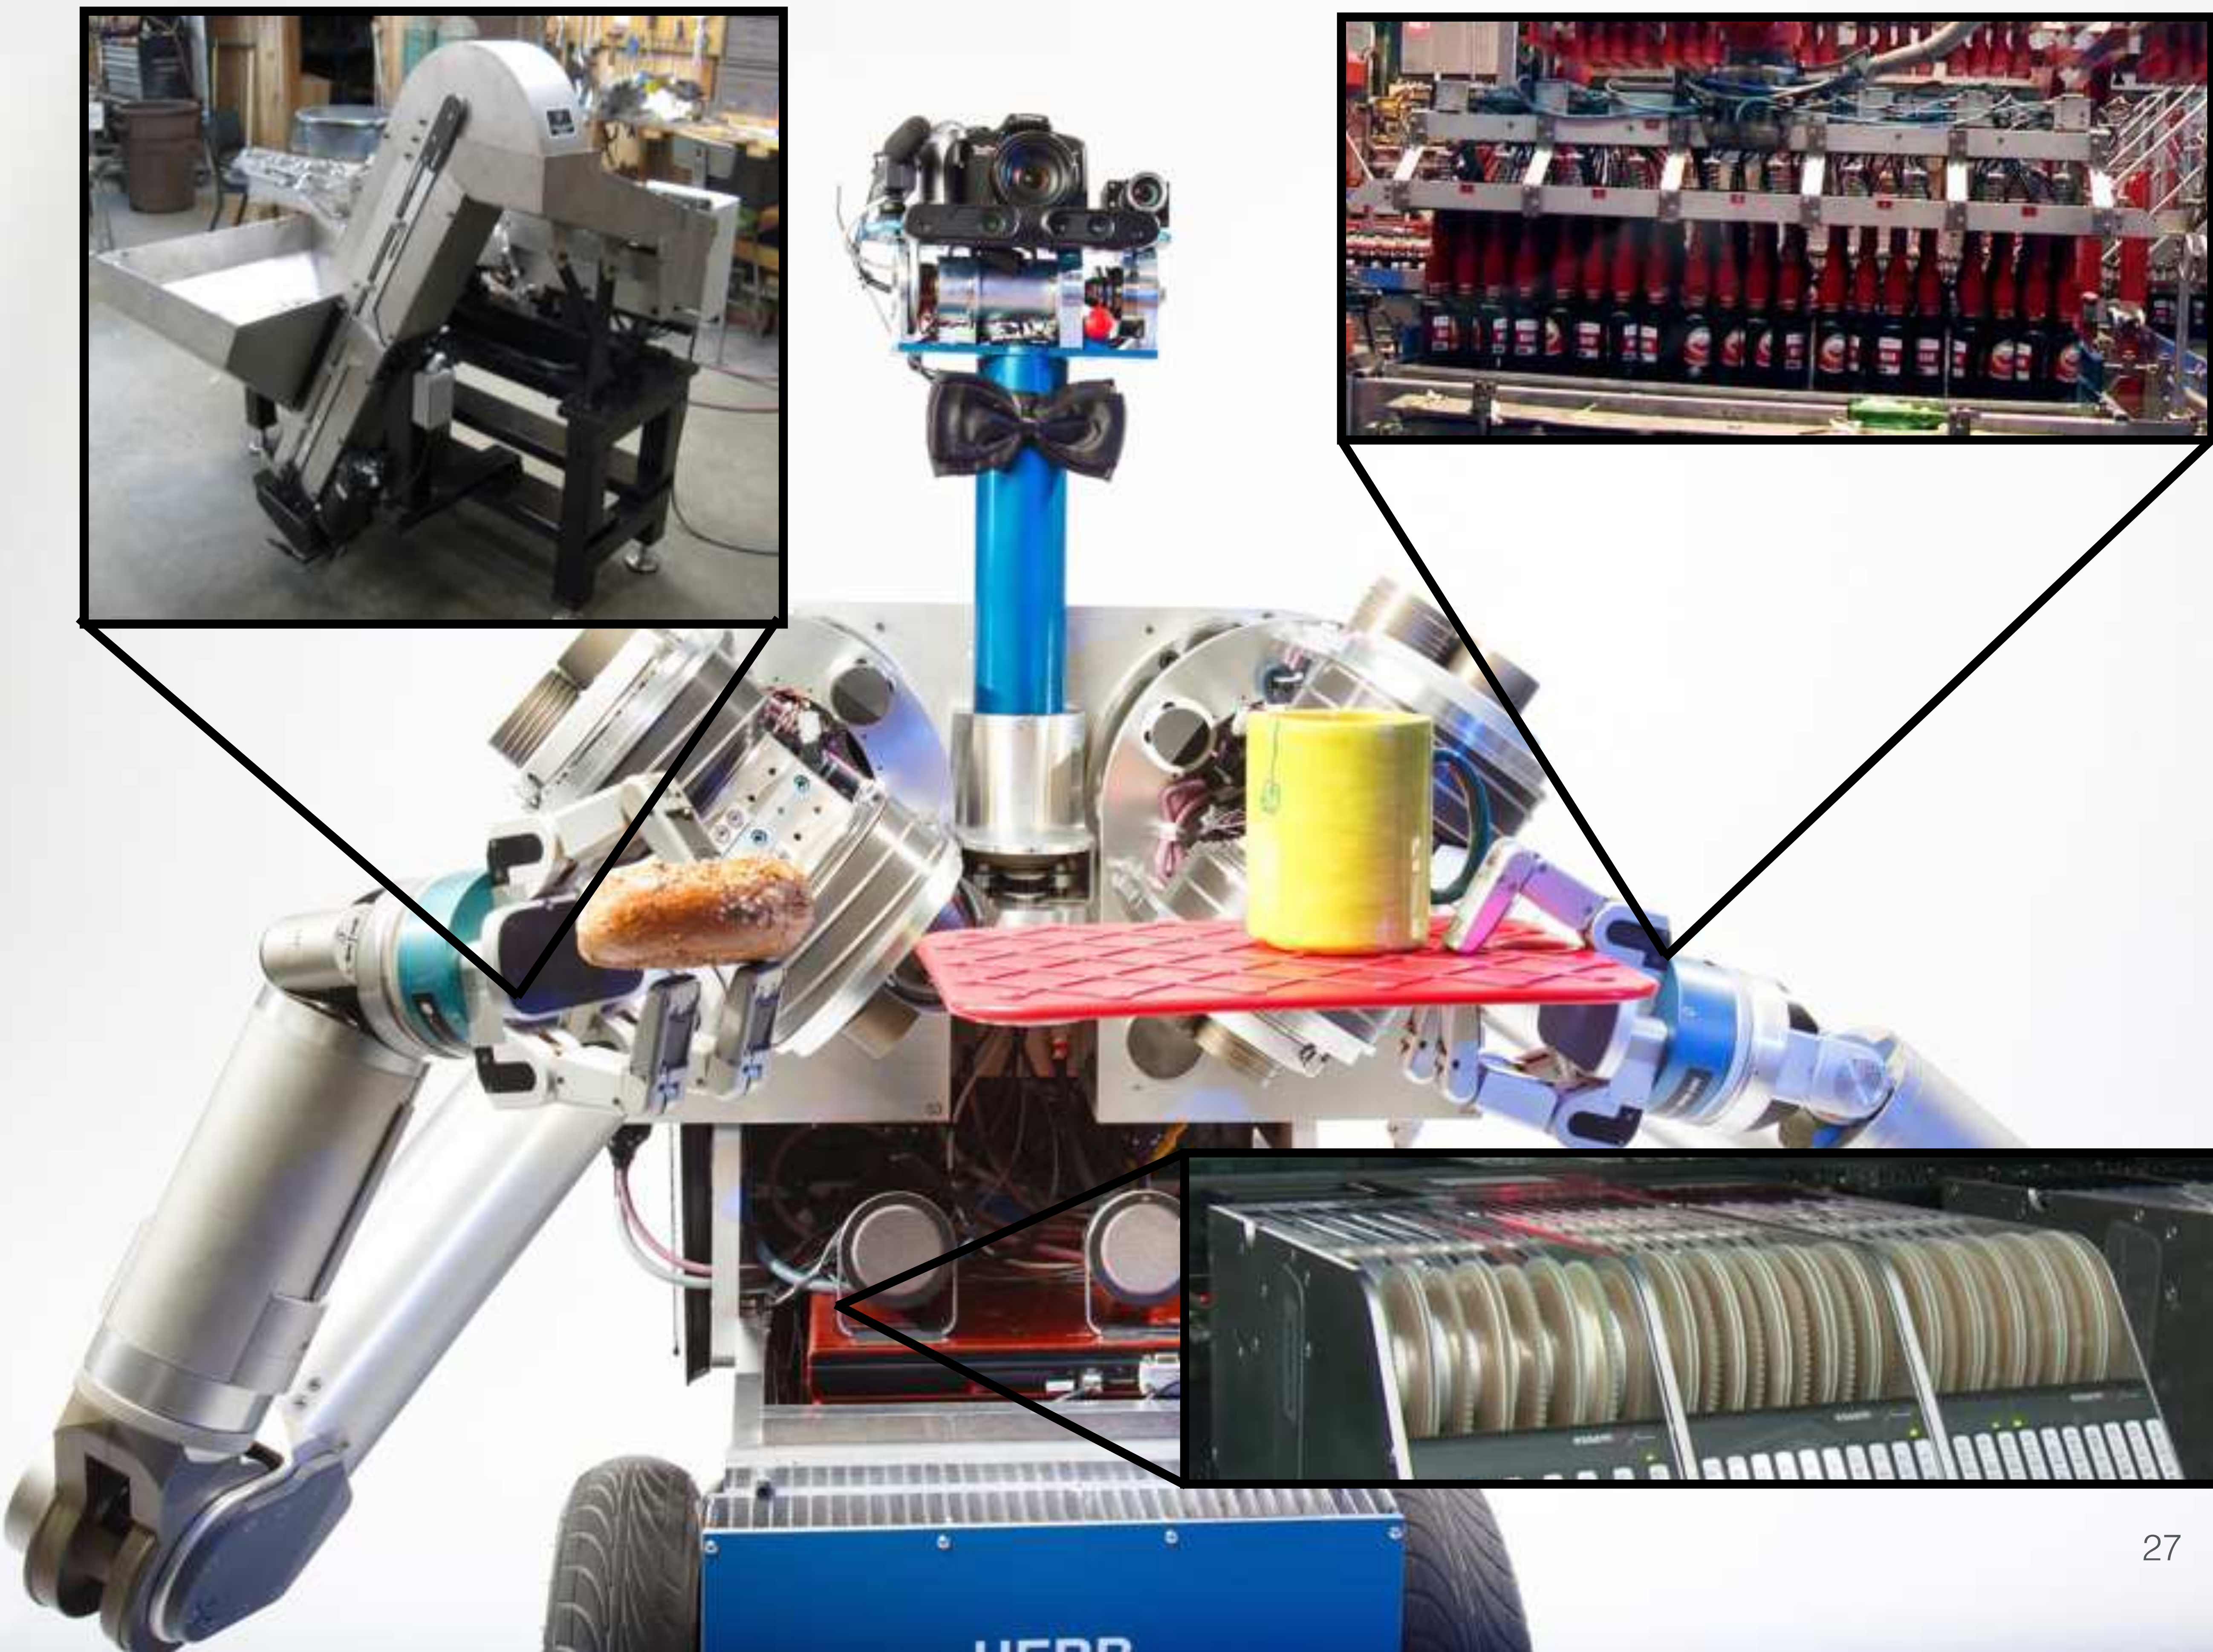

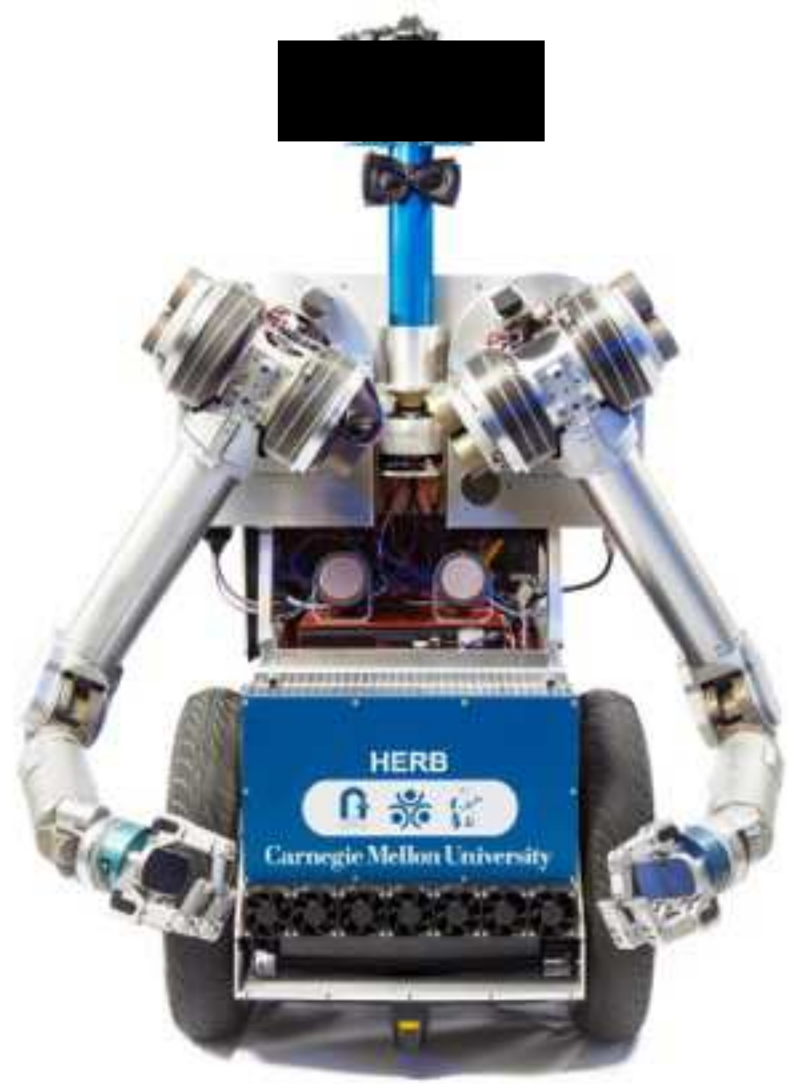

Closed-loop or open-loop?

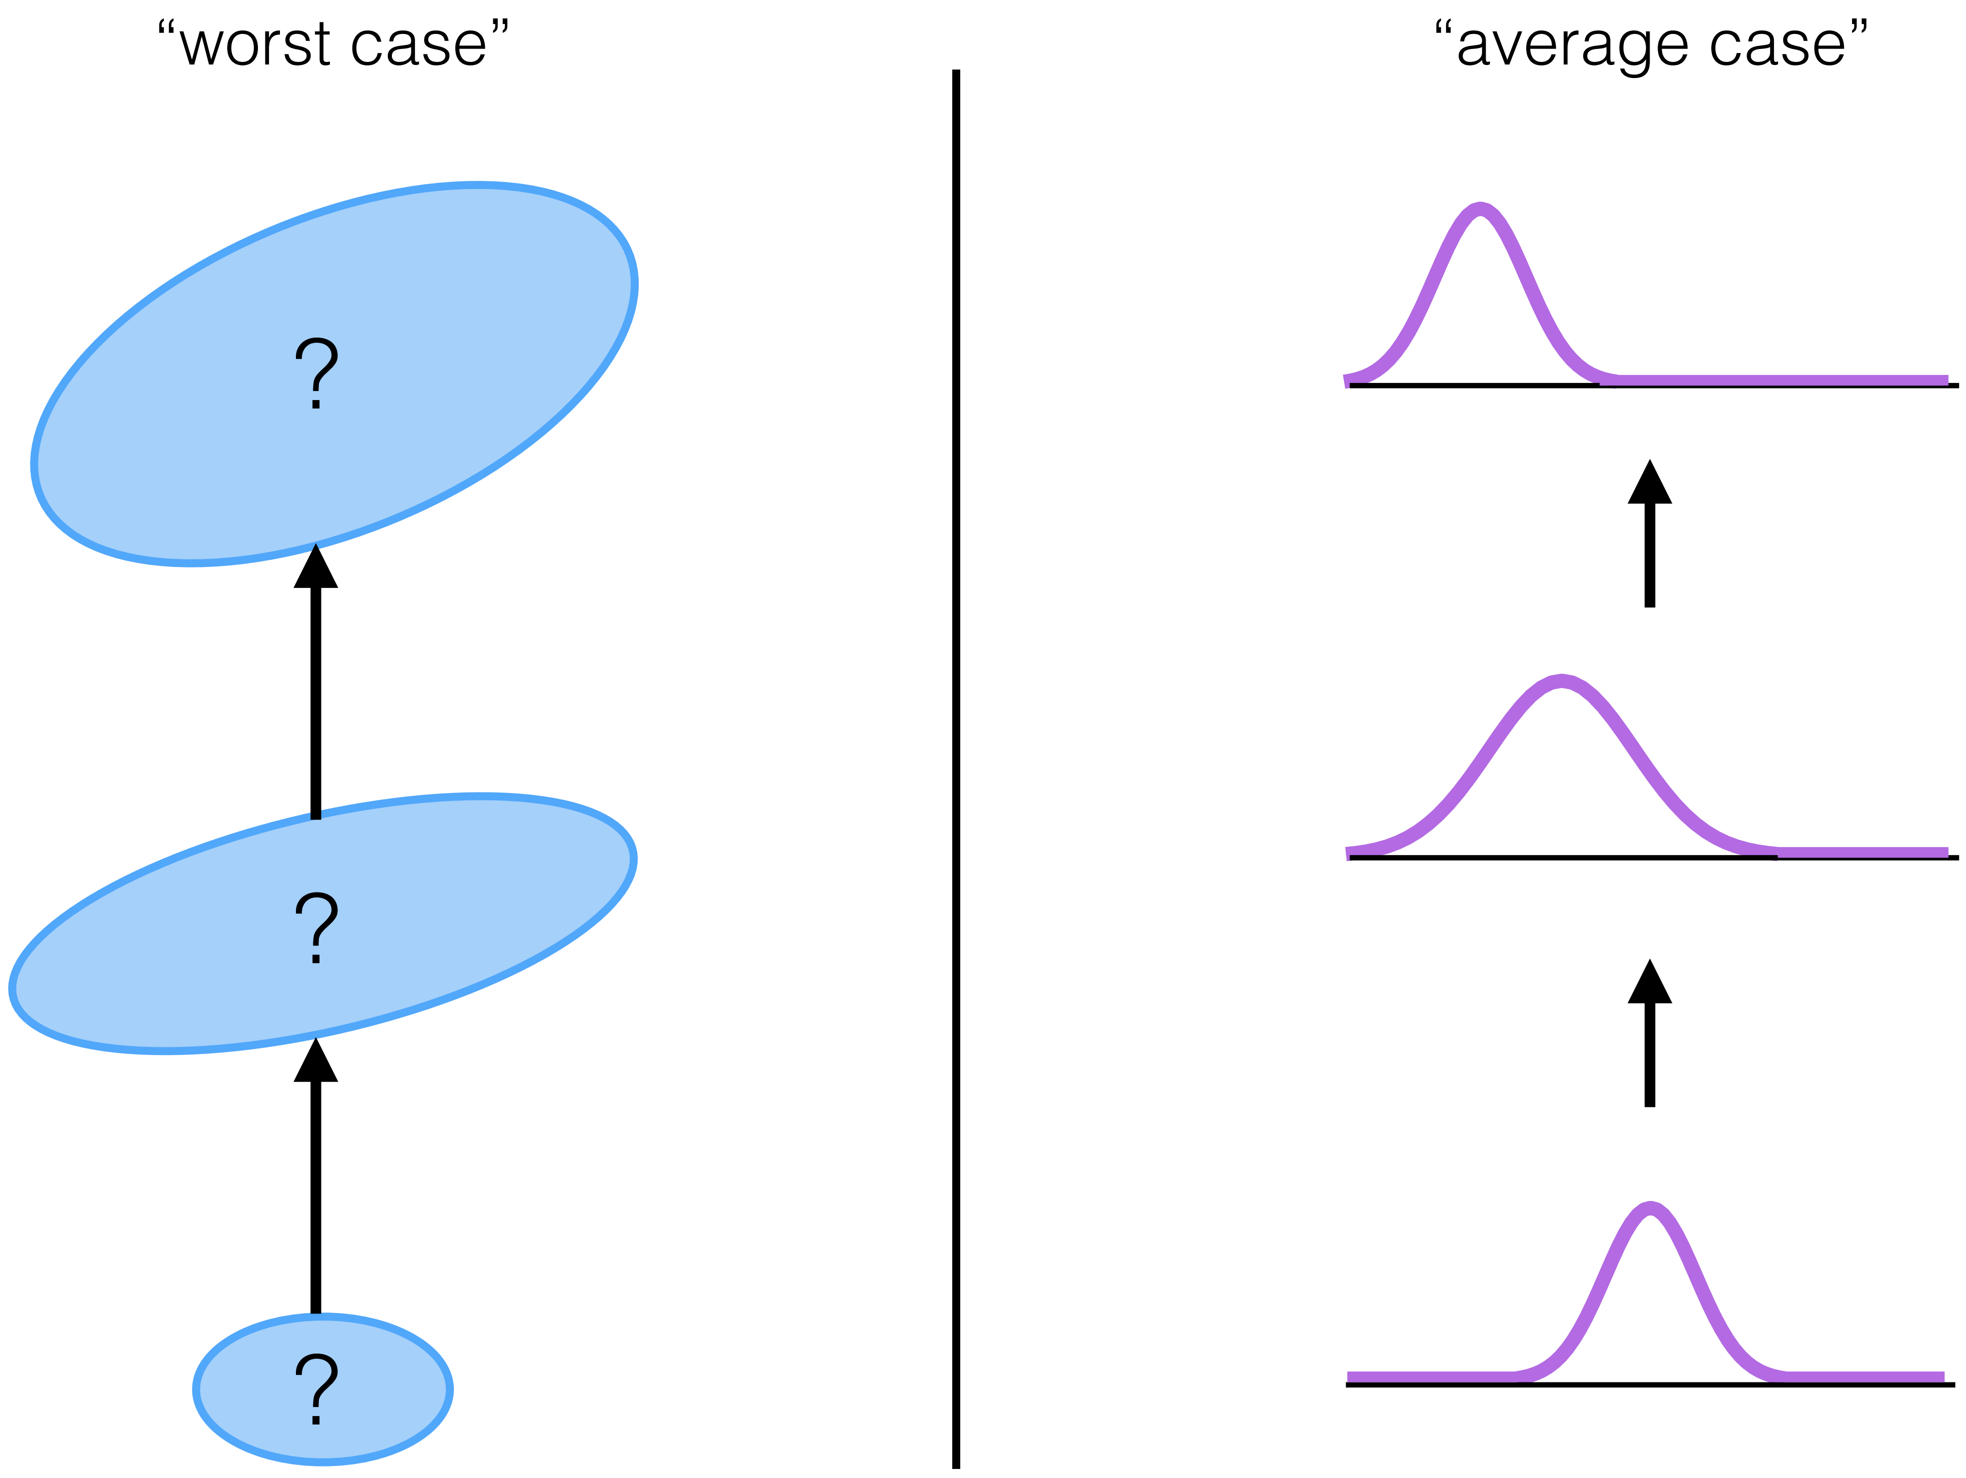

Non-deterministic or probabilistic uncertainty?

"Kalman filter"

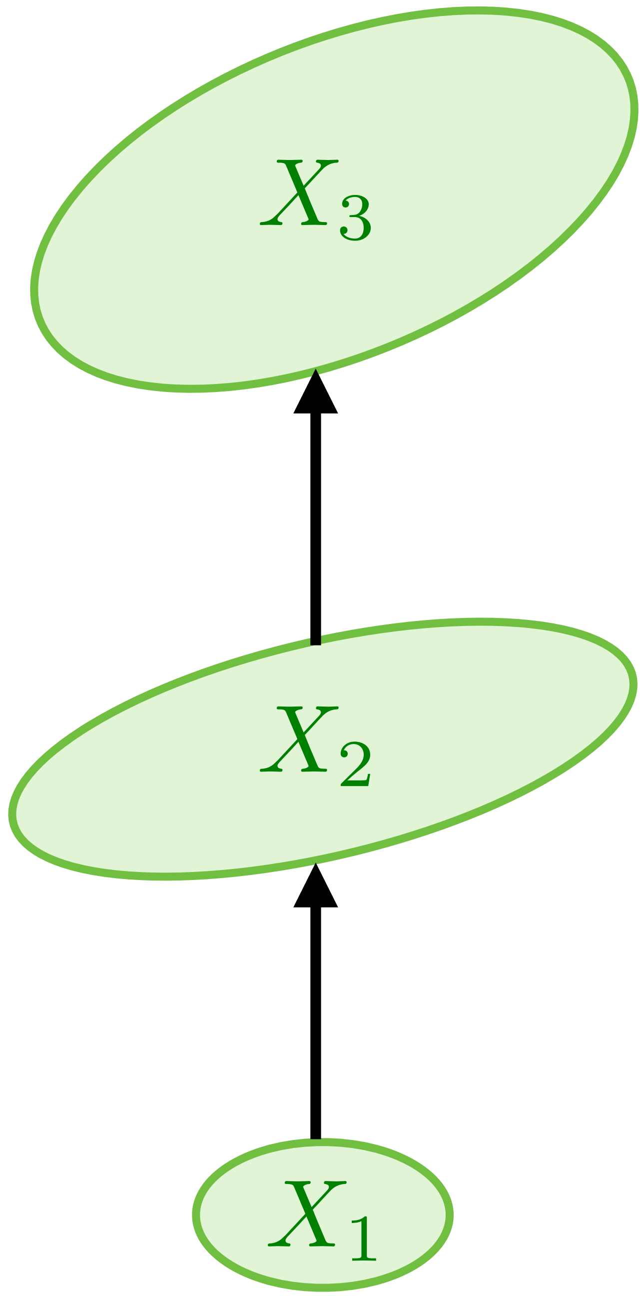

"particle filter"

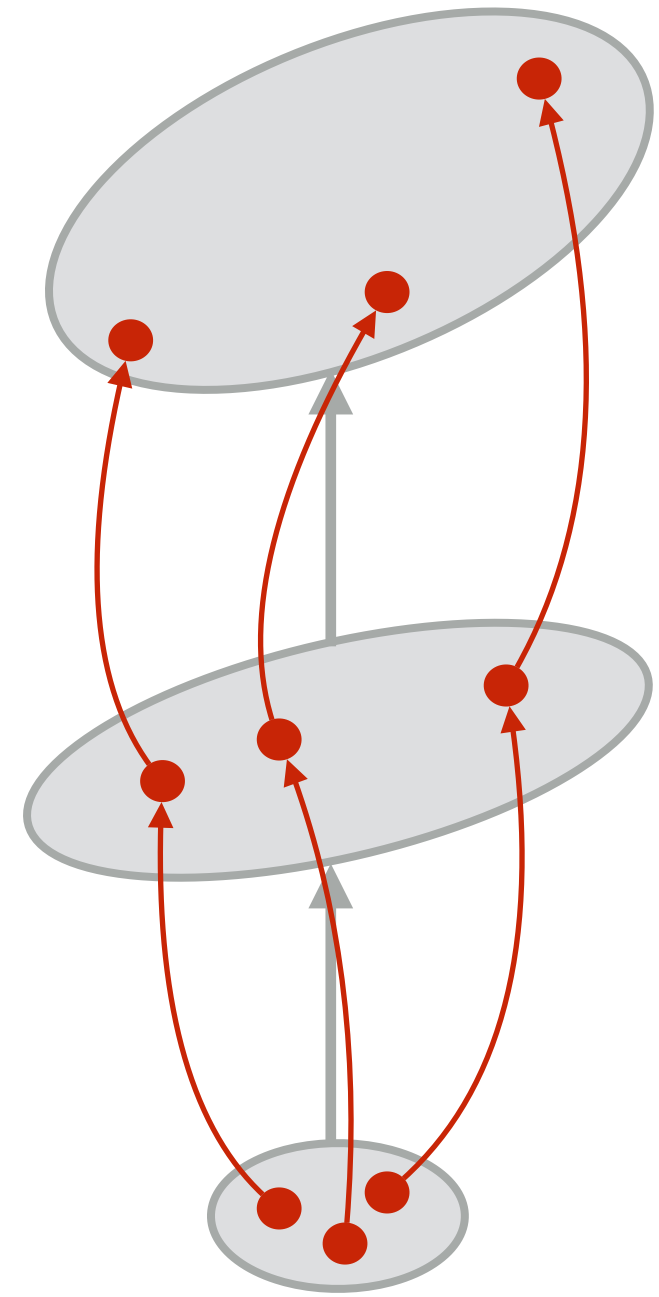

Closed-form or sample-based representation?

Estimate, react, or plan?

Visual Feedback

Tactile Feedback

No Feedback

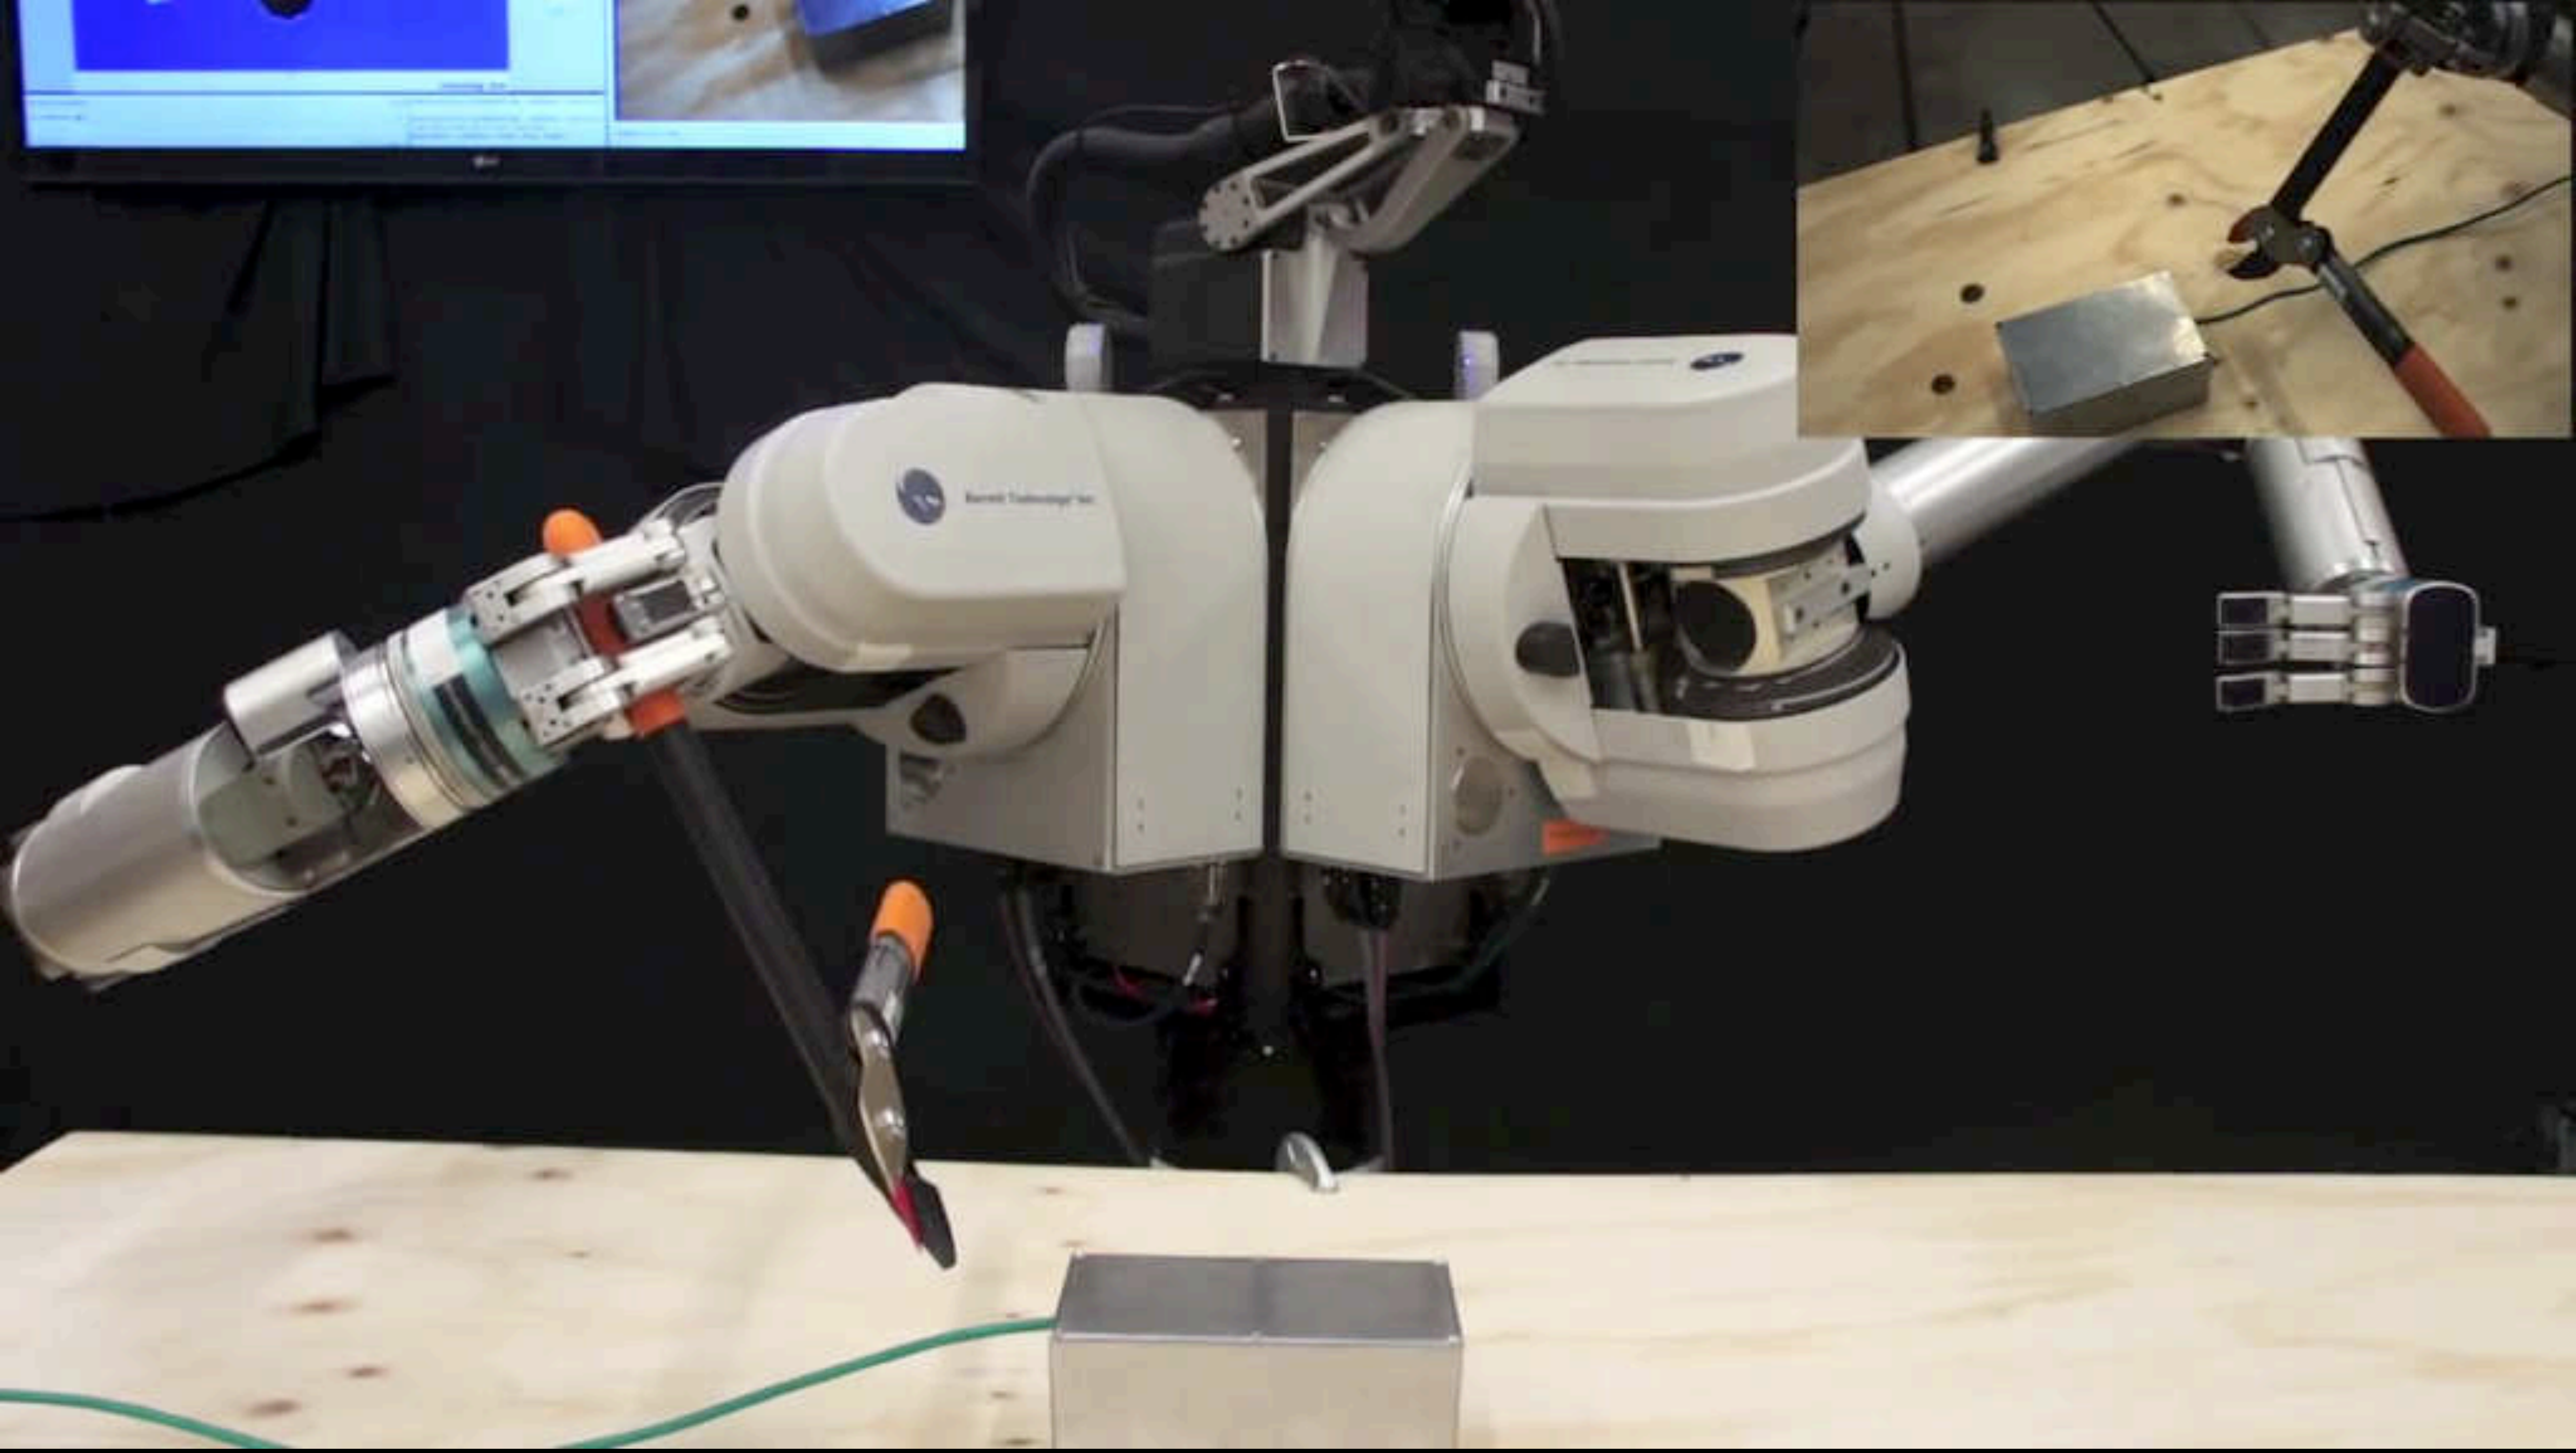

Image-space visual servoing

S. Hutchinson, G.D. Hager, and P. Corke. "A tutorial on visual servo control." T-RA, 1996.

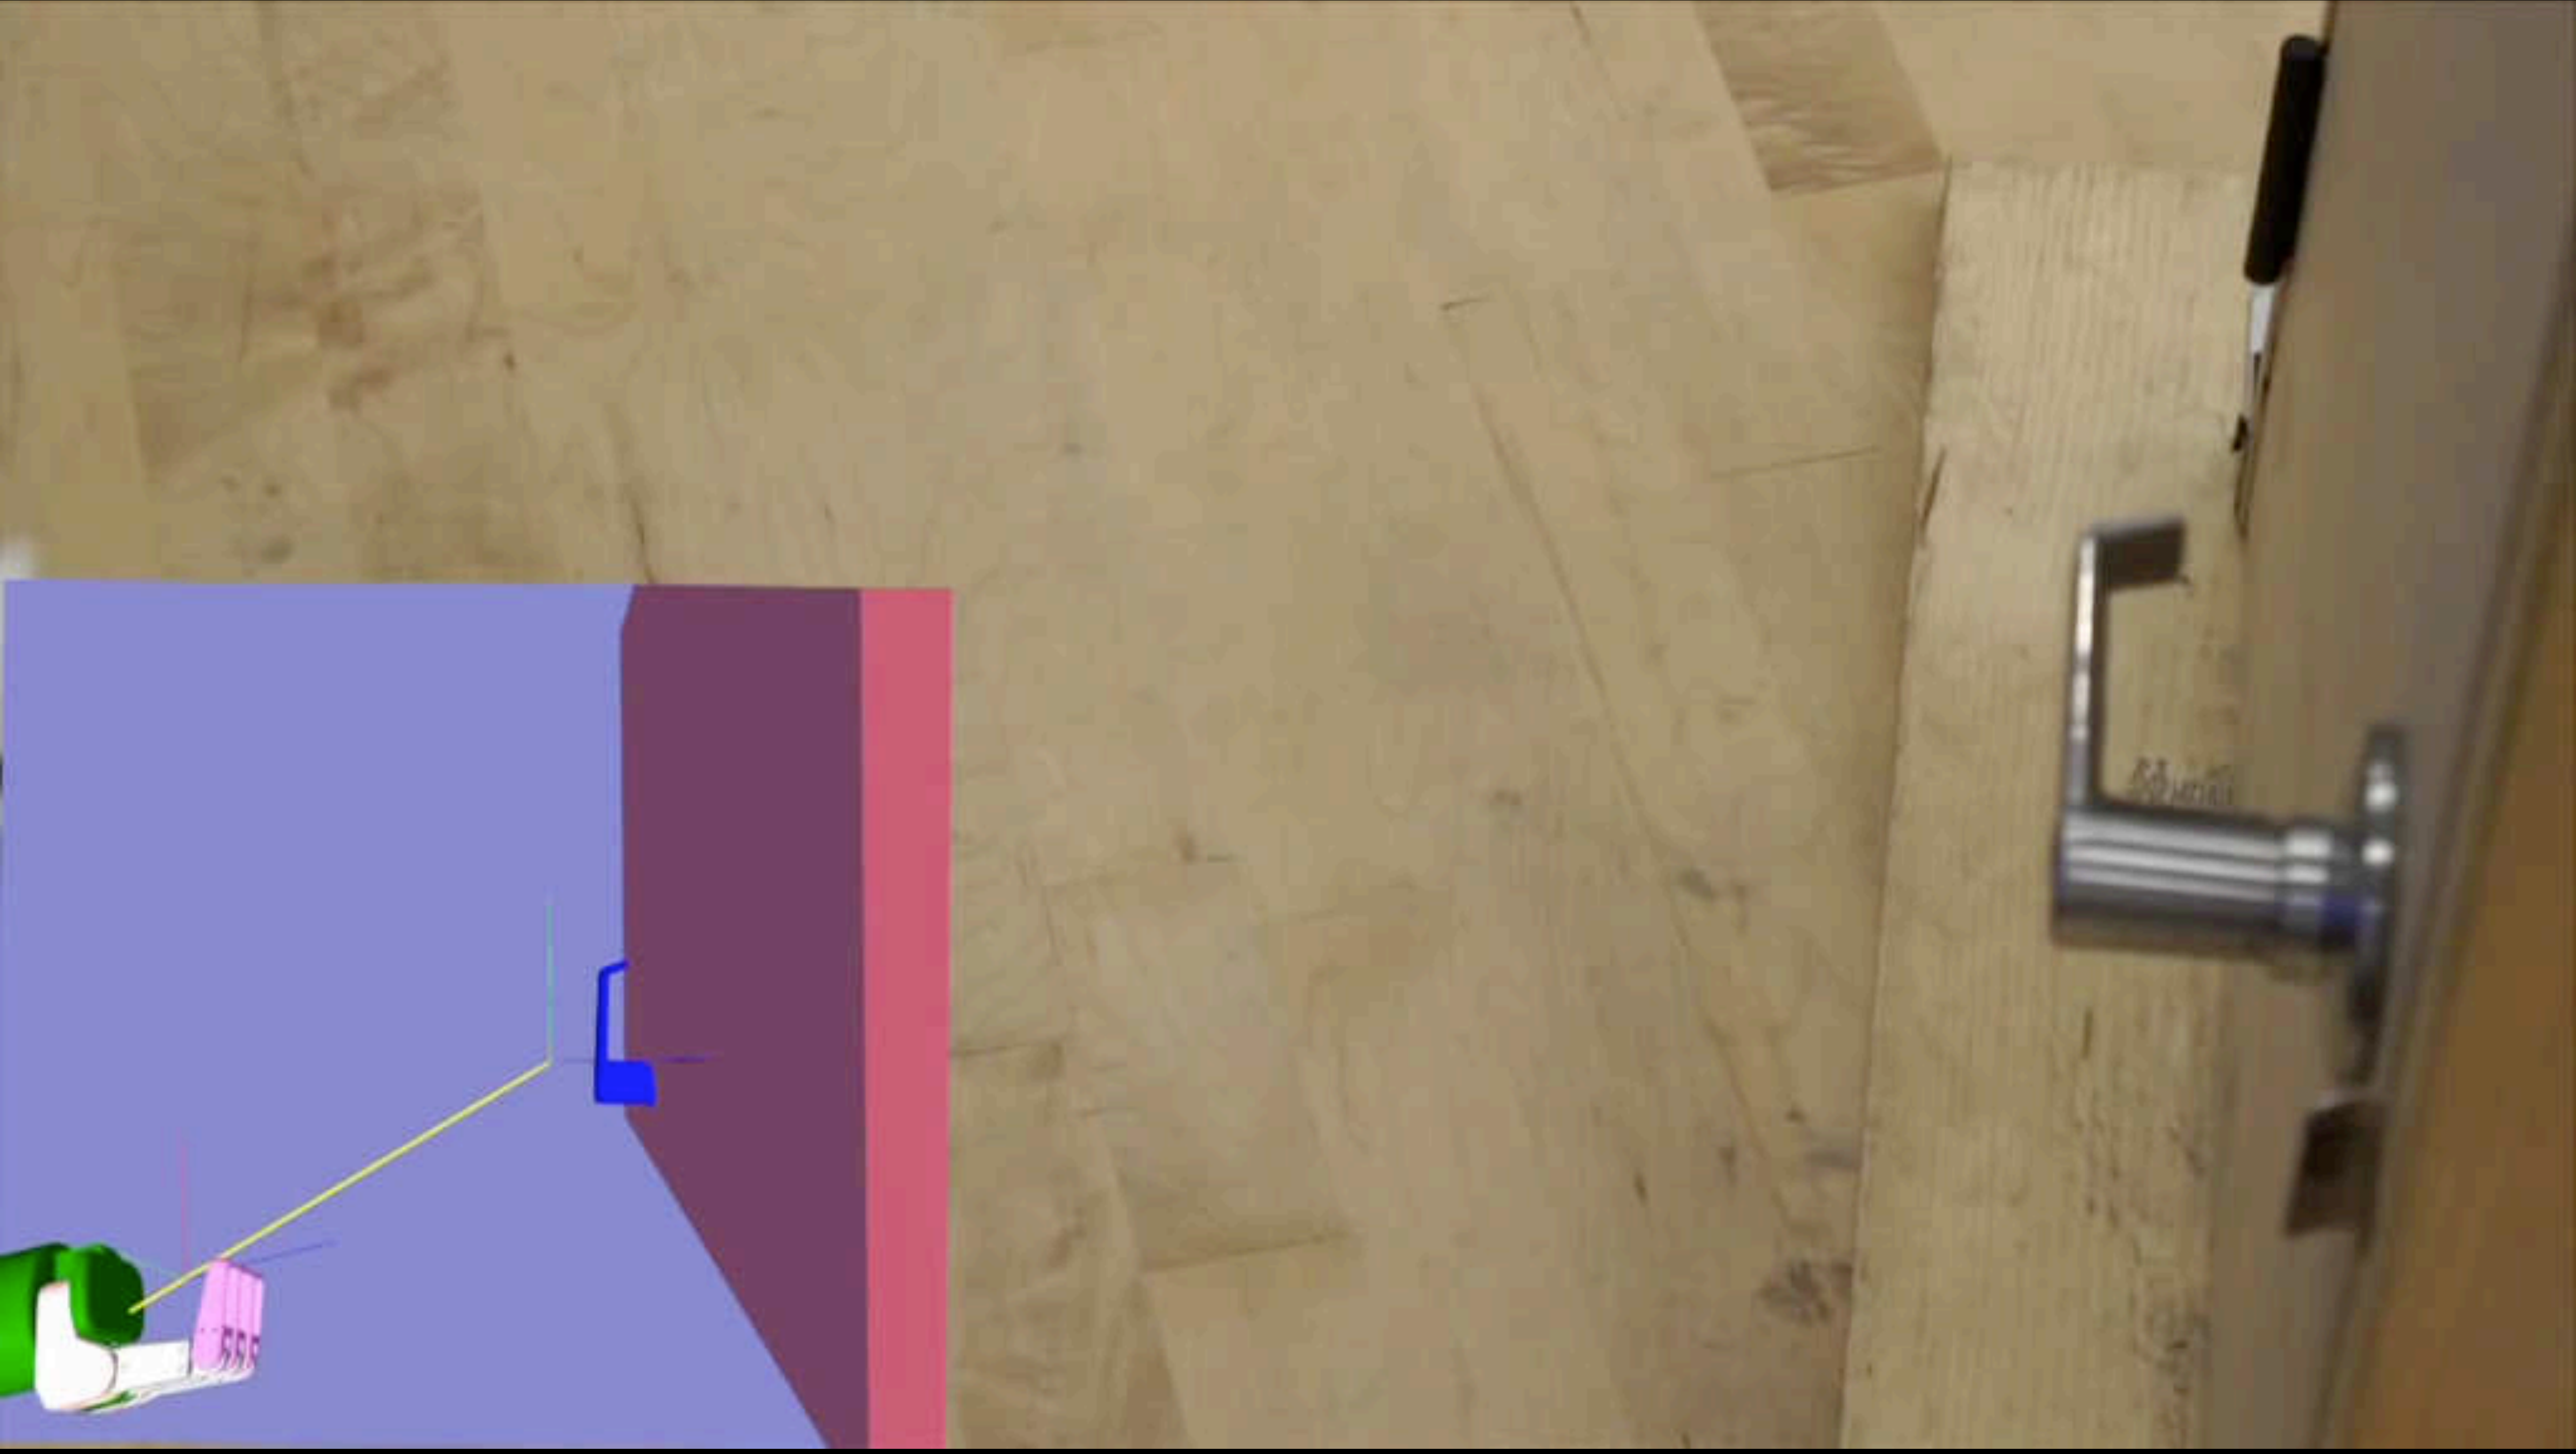

Markerless real-time articulated tracking

M. Klingensmith et al. "Closed-loop servoing using real-time markerless arm tracking." ICRA, 2013. (Workshop)

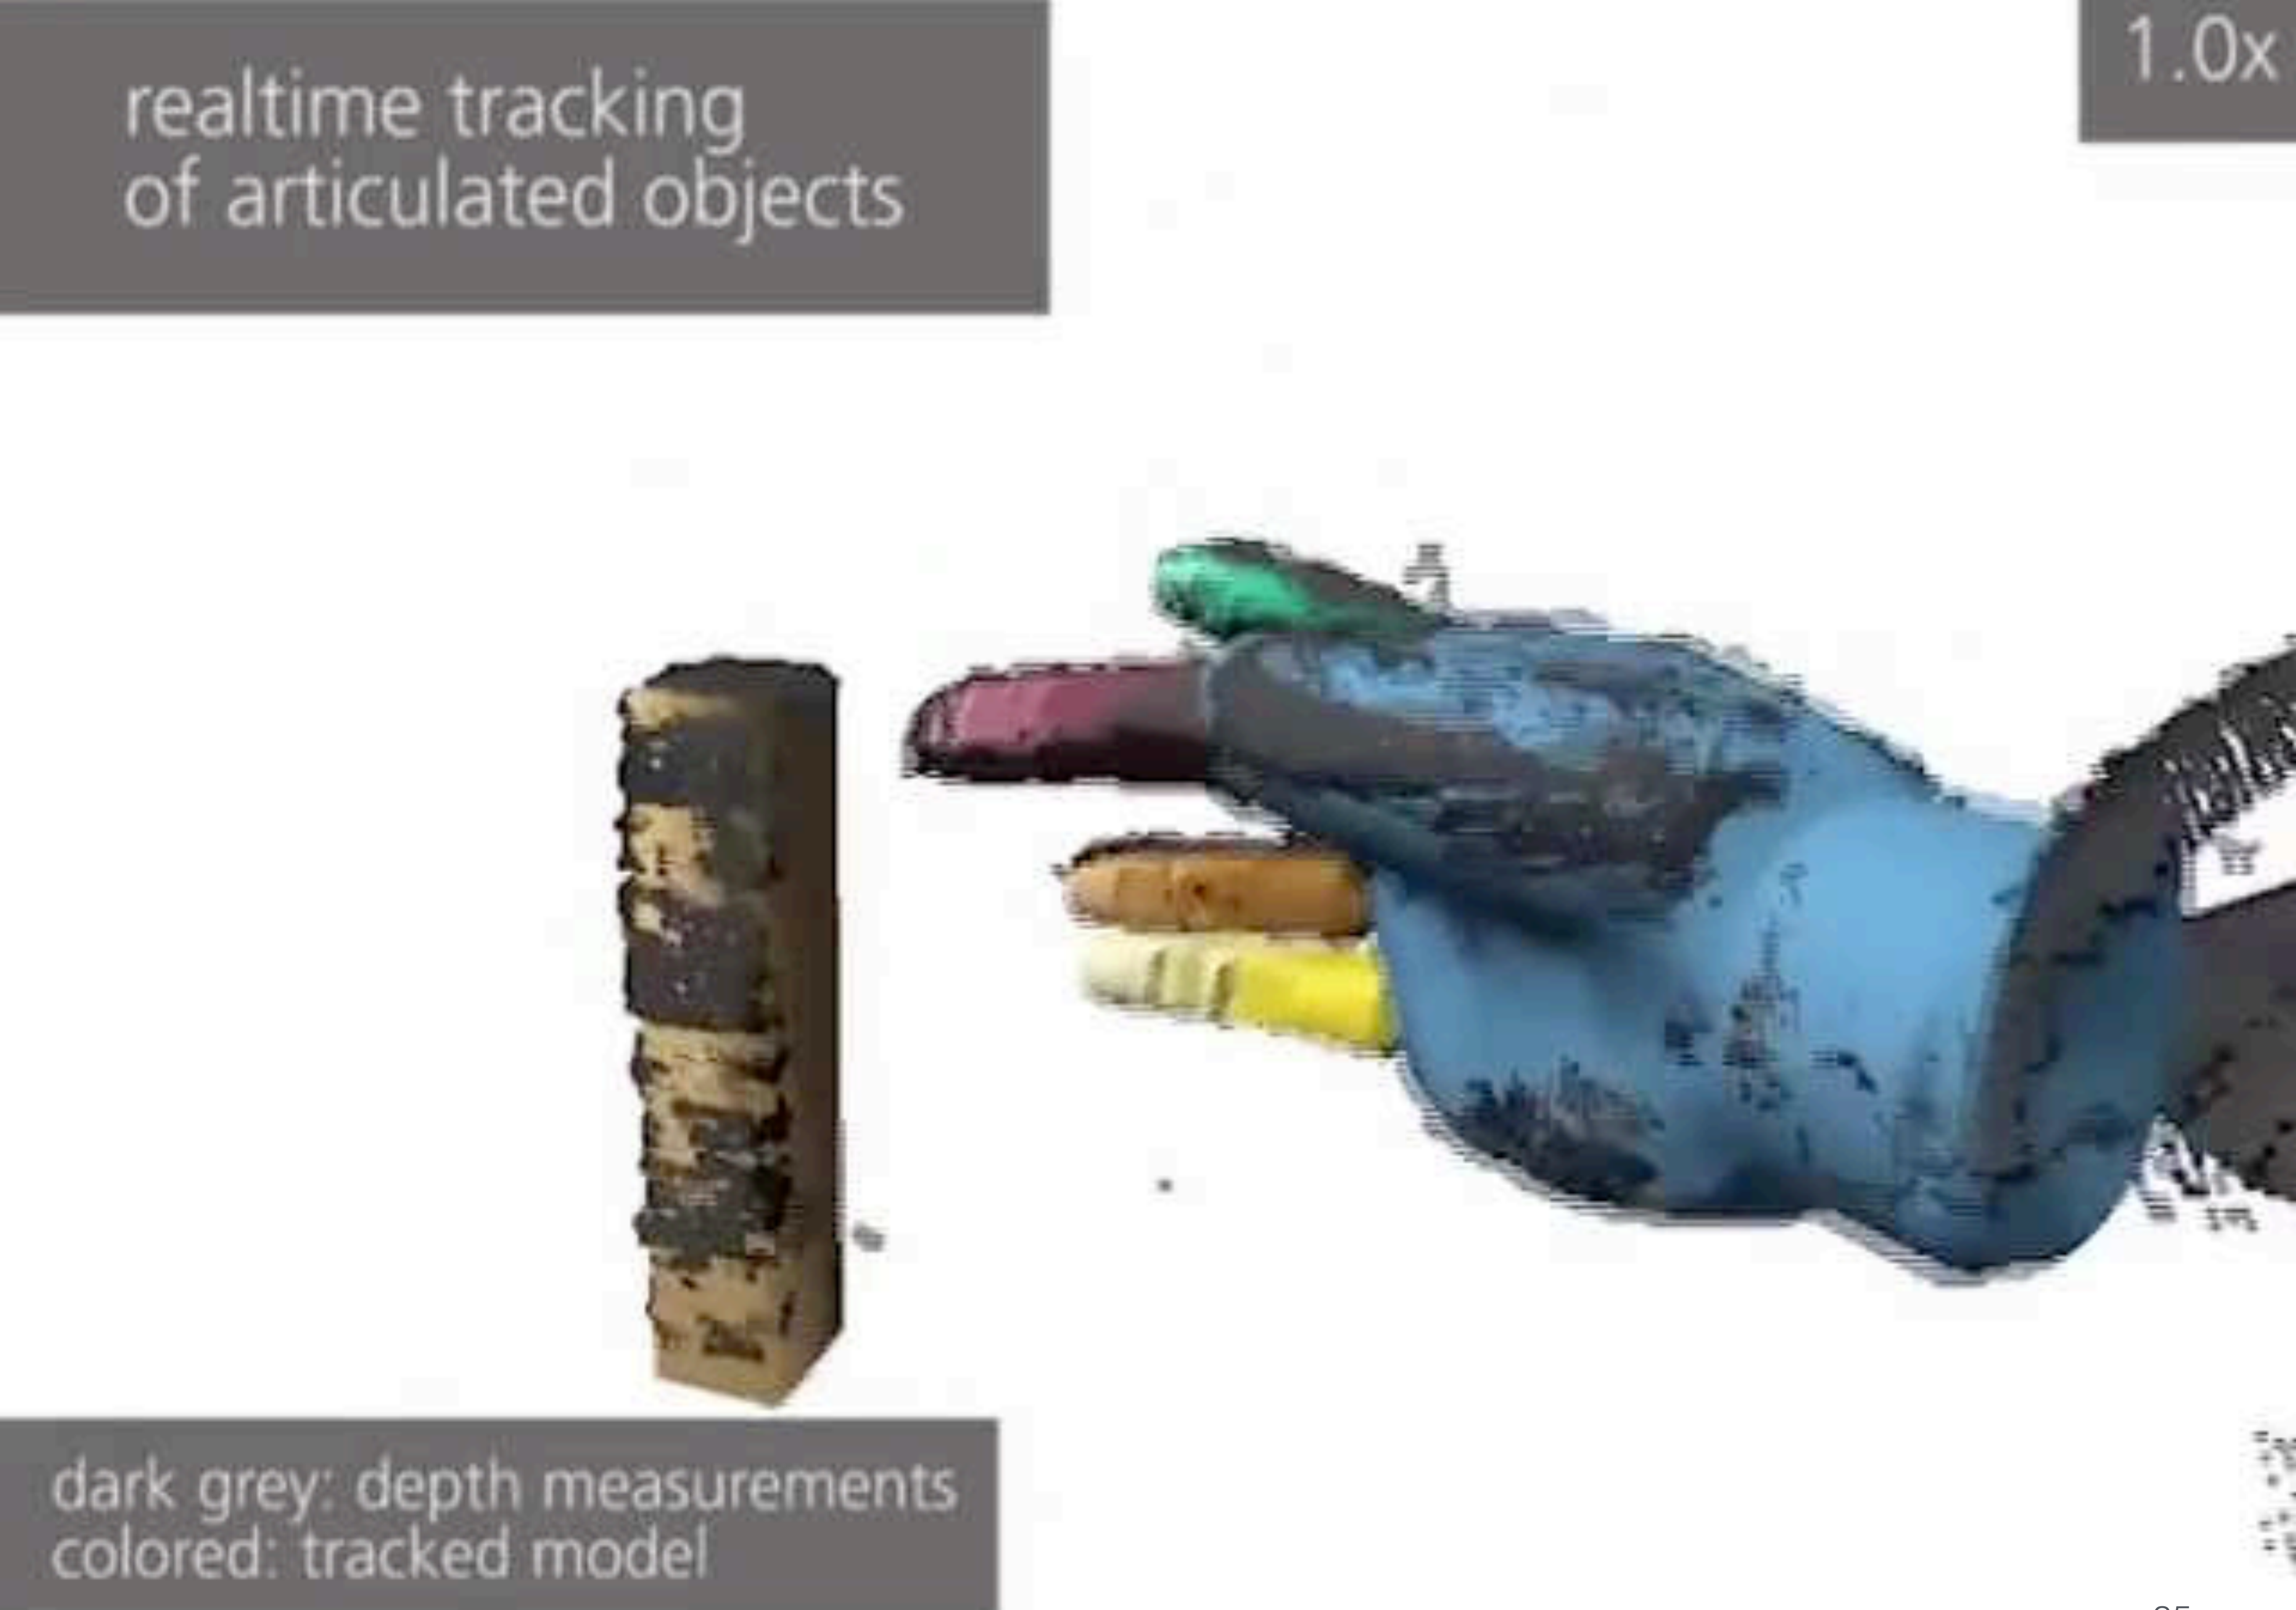

T. Schmidt, R.A. Newcombe, and D. Fox. “DART: Dense Articulated Real-Time Tracking.” RSS, 2014.
T. Schmidt, K. Hertkorn, R.A. Newcombe, Z. Marton, M. Suppa, and D. Fox. "Depth-based tracking with physical constraints for robotic manipulation.” ICRA, 2015.

Visual Feedback

Tactile Feedback

No Feedback

Visual Feedback

Tactile Feedback

No Feedback

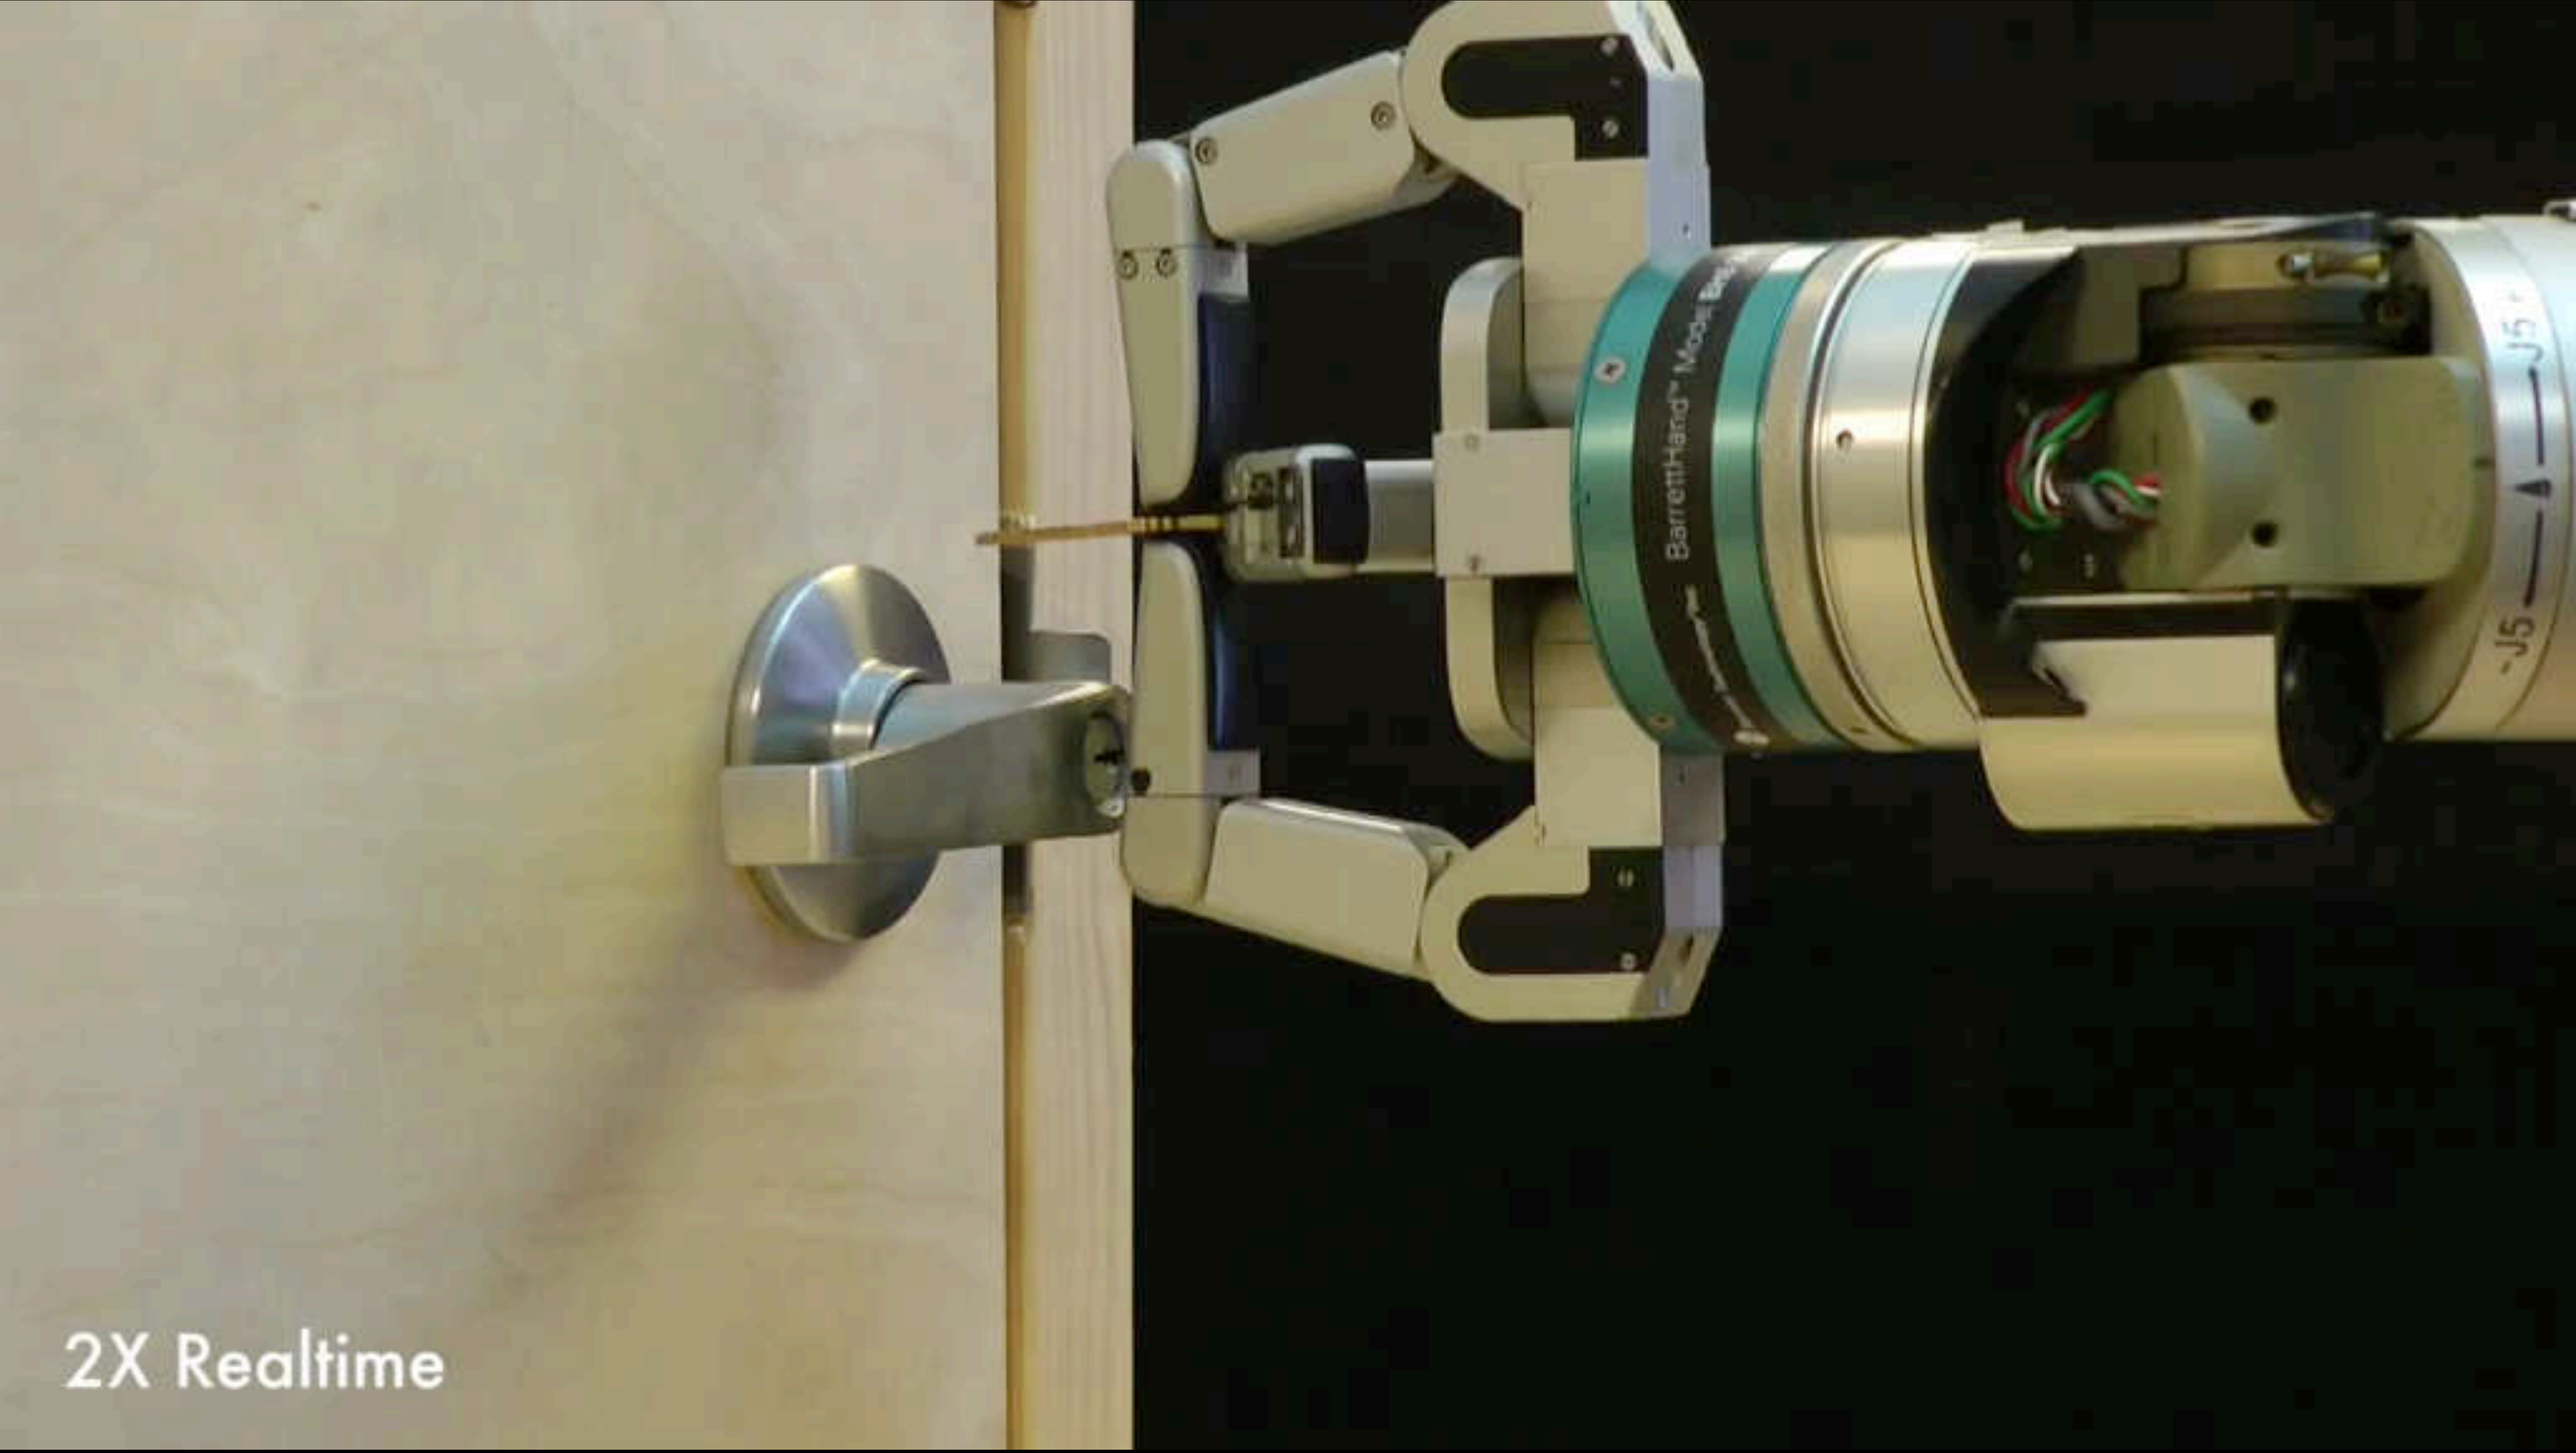

Use “guarded moves” to reduce uncertainty

Plan a sequence that maximizes information gain

S. Javdani, M. Klingsmith, J.A. Bagnell, N.S. Pollard, S.S. Srinivasa. "Efficient touch based localization through submodularity." ICRA, 2013.

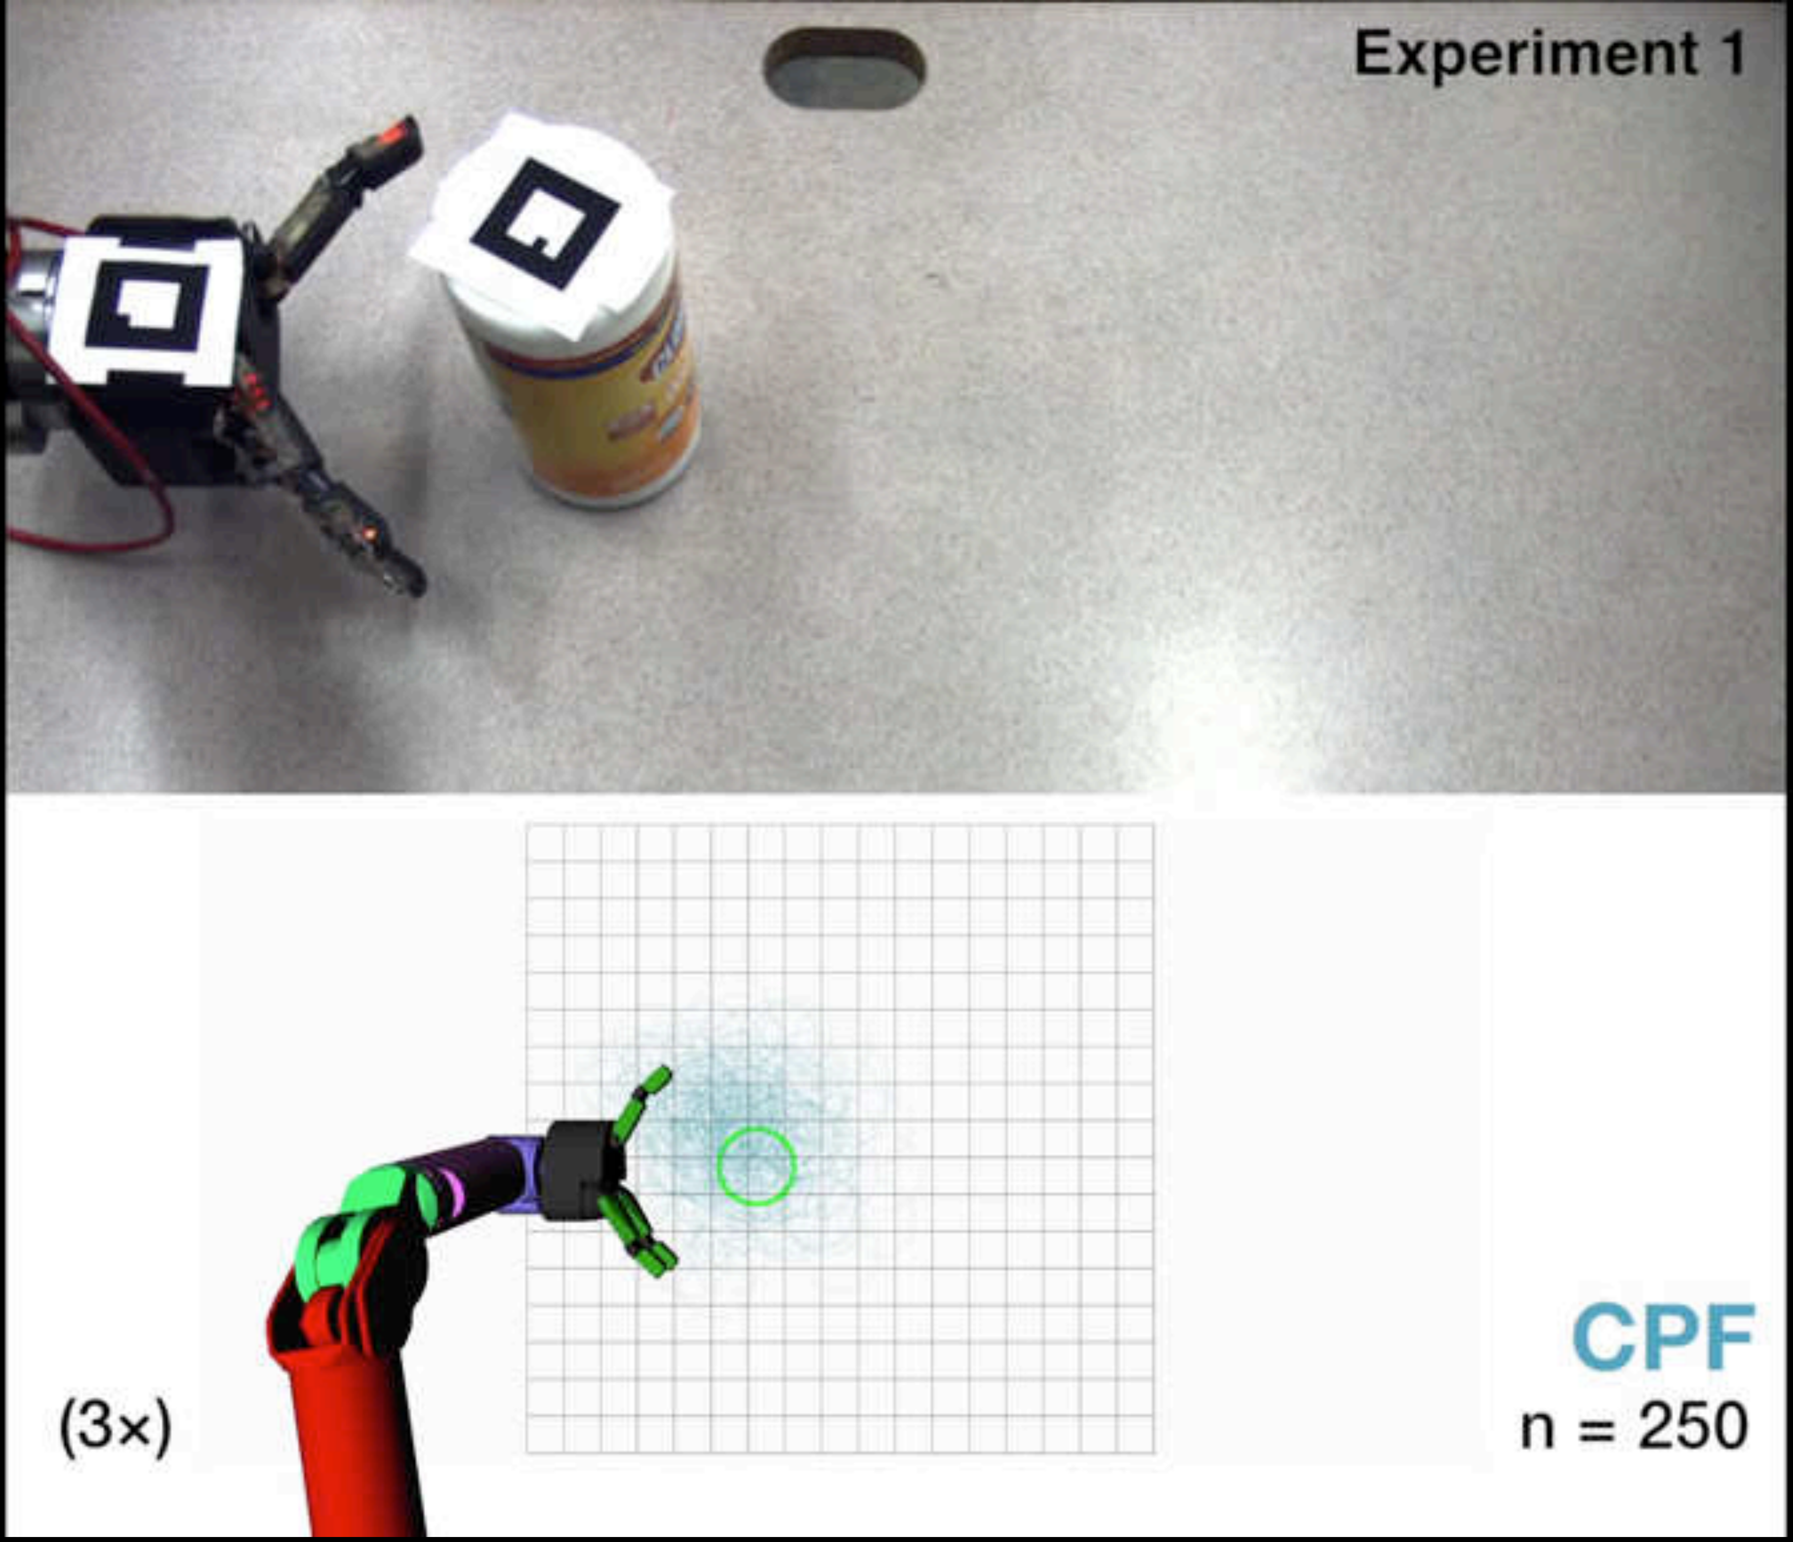

Estimate the pose of the object using tactile sensing

M.C. Koval, M.R. Dogar, N.S. Pollard, and S.S. Srinivasa “Pose estimation for contact manipulation using manifold particle filters." IROS, 2013.
M.C. Koval, N.S. Pollard, and S.S. Srinivasa. “Manifold representations for state estimation in contact manipulation.” ISRR, 2013.
M.C. Koval, N.S. Pollard, and S.S. Srinivasa. “Pose estimation for planar contact manipulation with manifold particle filters.” IJRR, 2015.

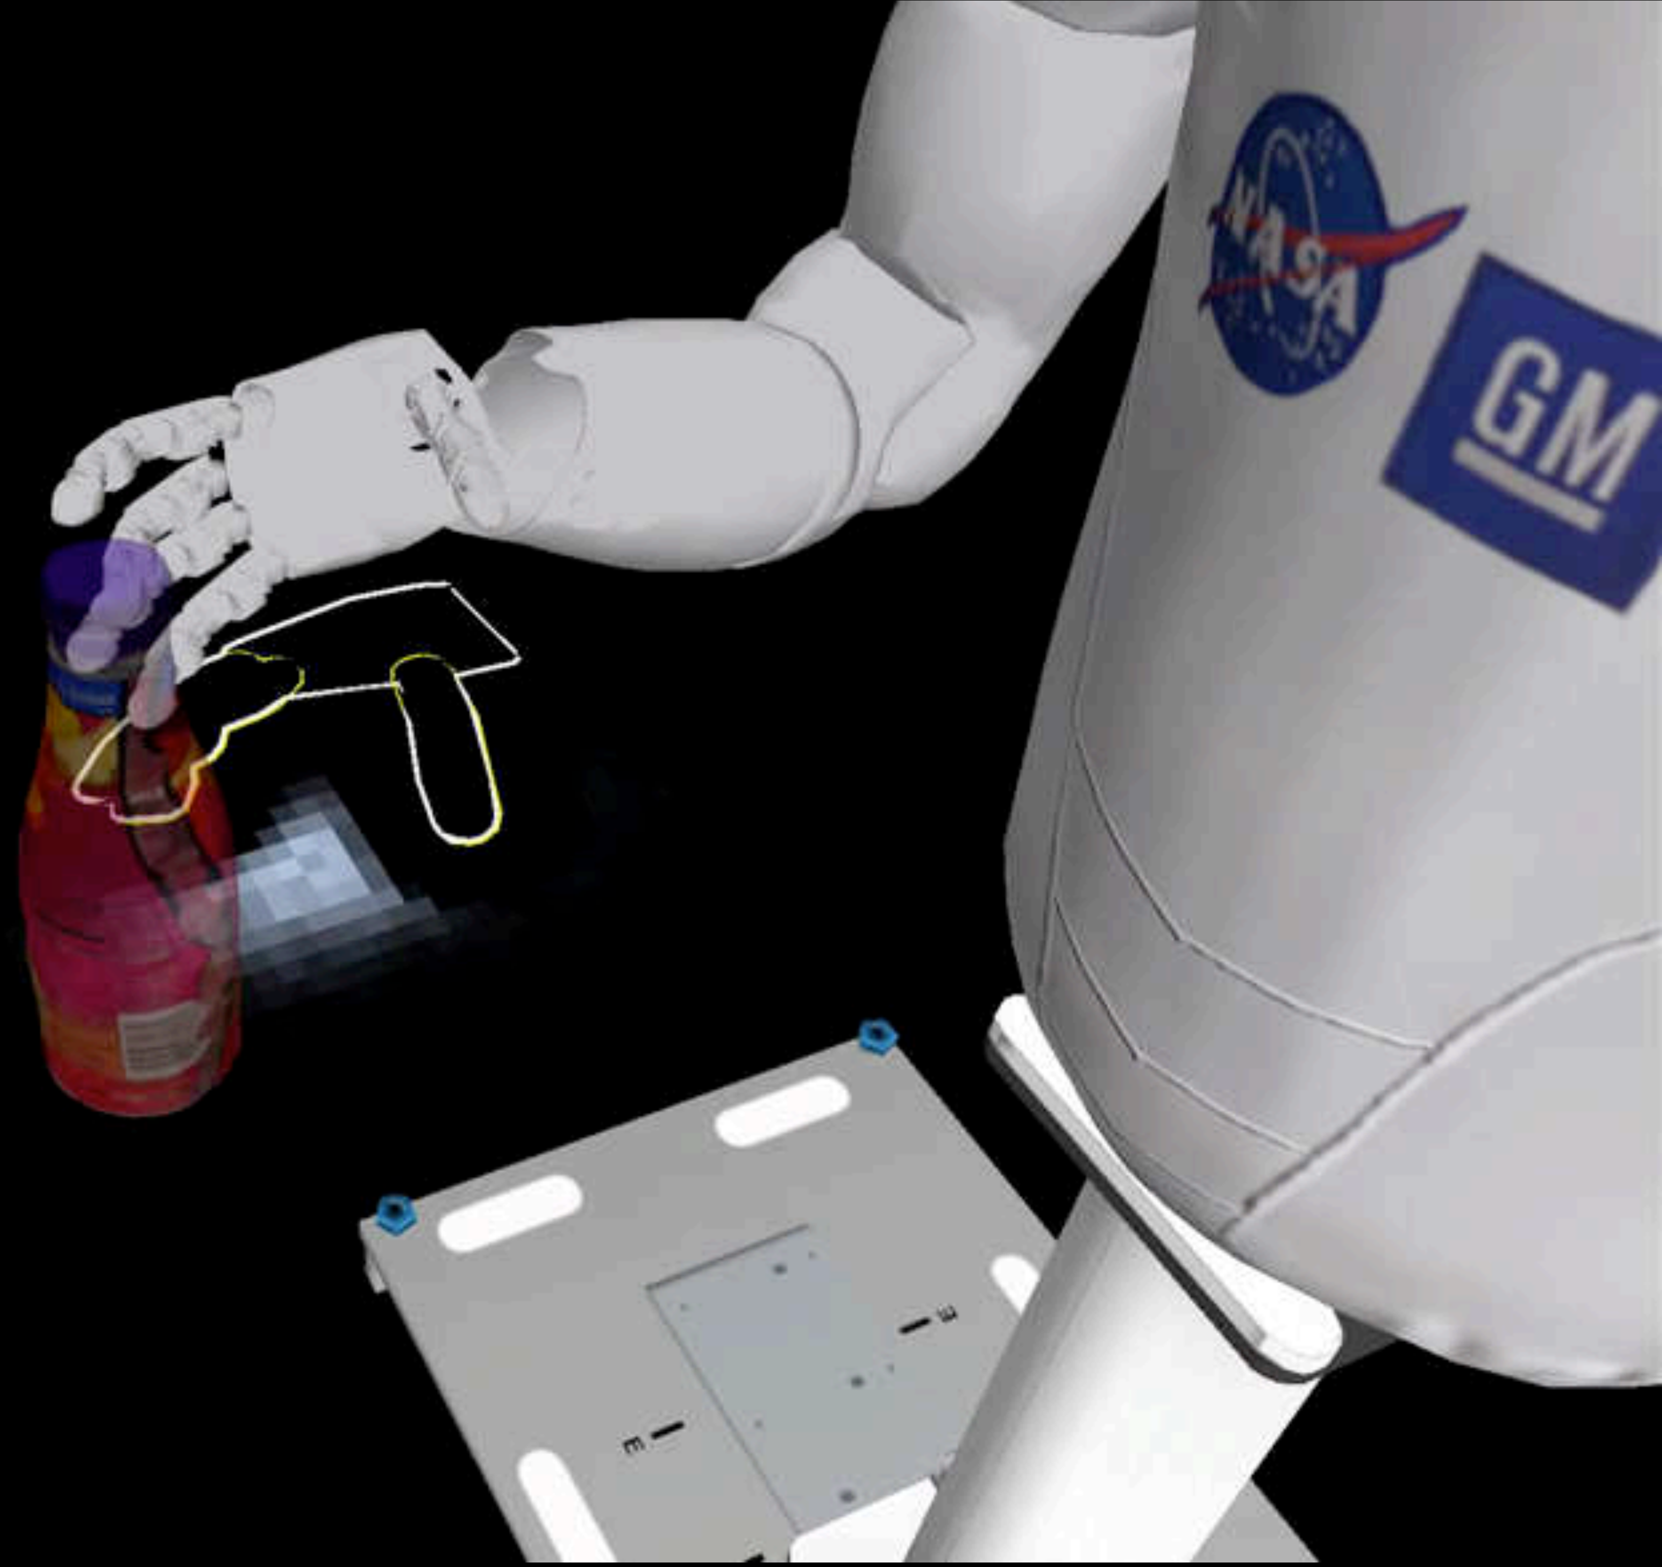

Closed-loop grasping using contact sensing

M.C. Koval, N.S. Pollard, S.S. Srinivasa. “Pre- and post-contact policy decomposition for planar contact manipulation under uncertainty." RSS, 2014.
M.C. Koval, N.S. Pollard, S.S. Srinivasa. “Pre- and post-contact policy decomposition for planar contact manipulation under uncertainty.” IJRR, 2015.

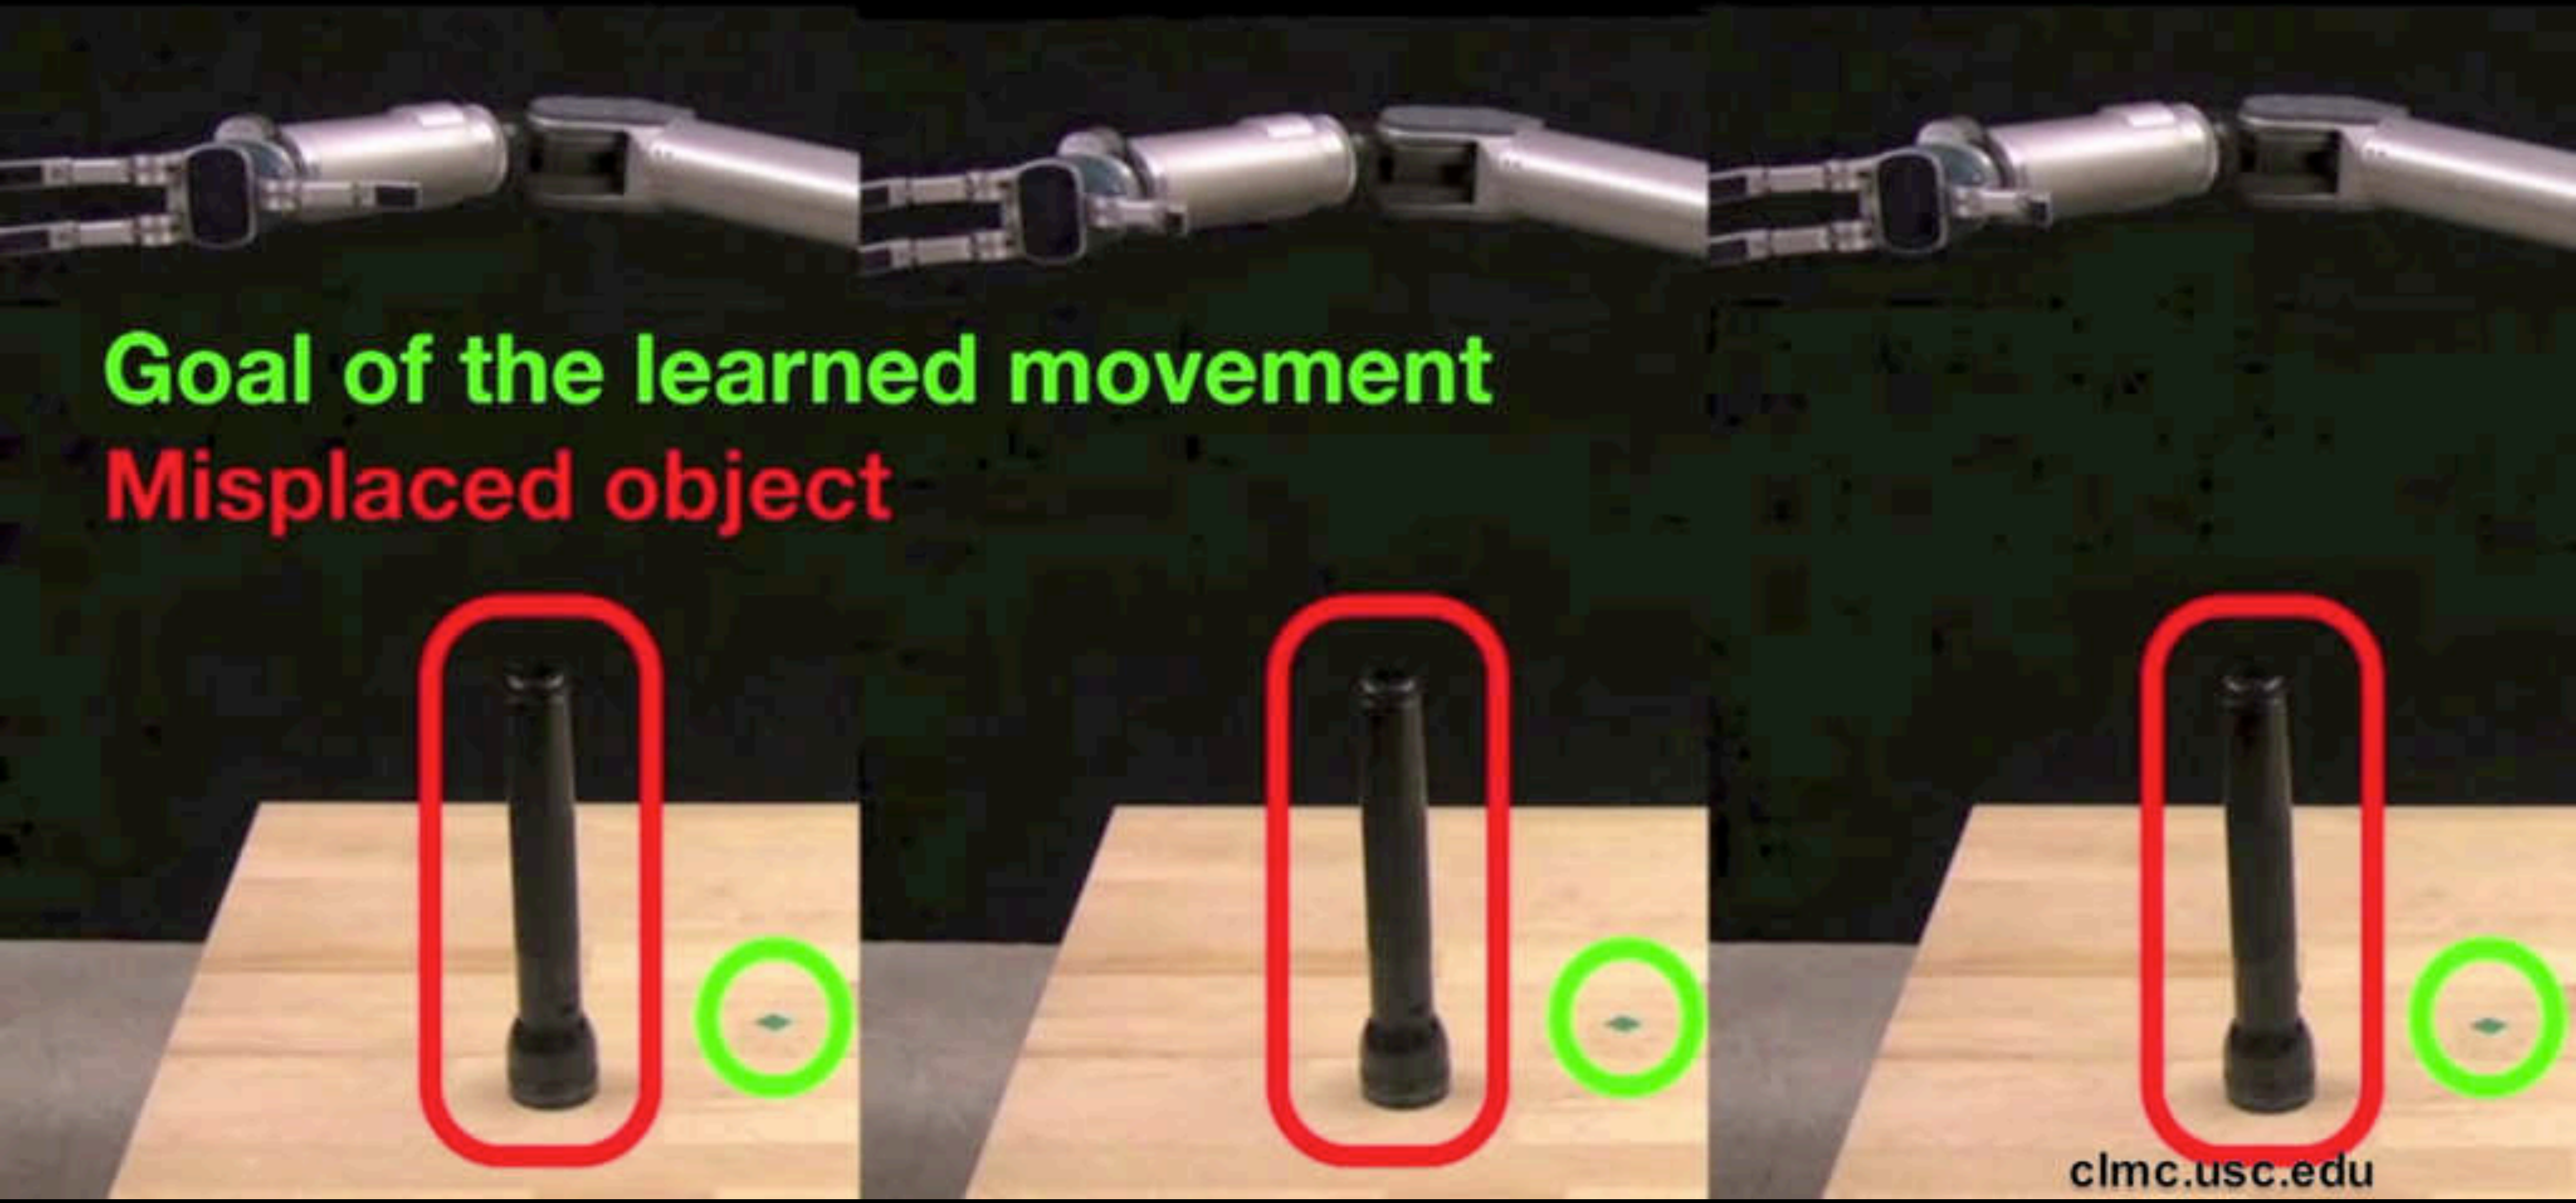

Learn feedback policies that use sensor feedback

P. Pastor, L. Righetti, M. Kalakrishnan, and S. Schaal. "Online movement adaptation based on previous sensor measurements." IROS, 2011.

Policy Input

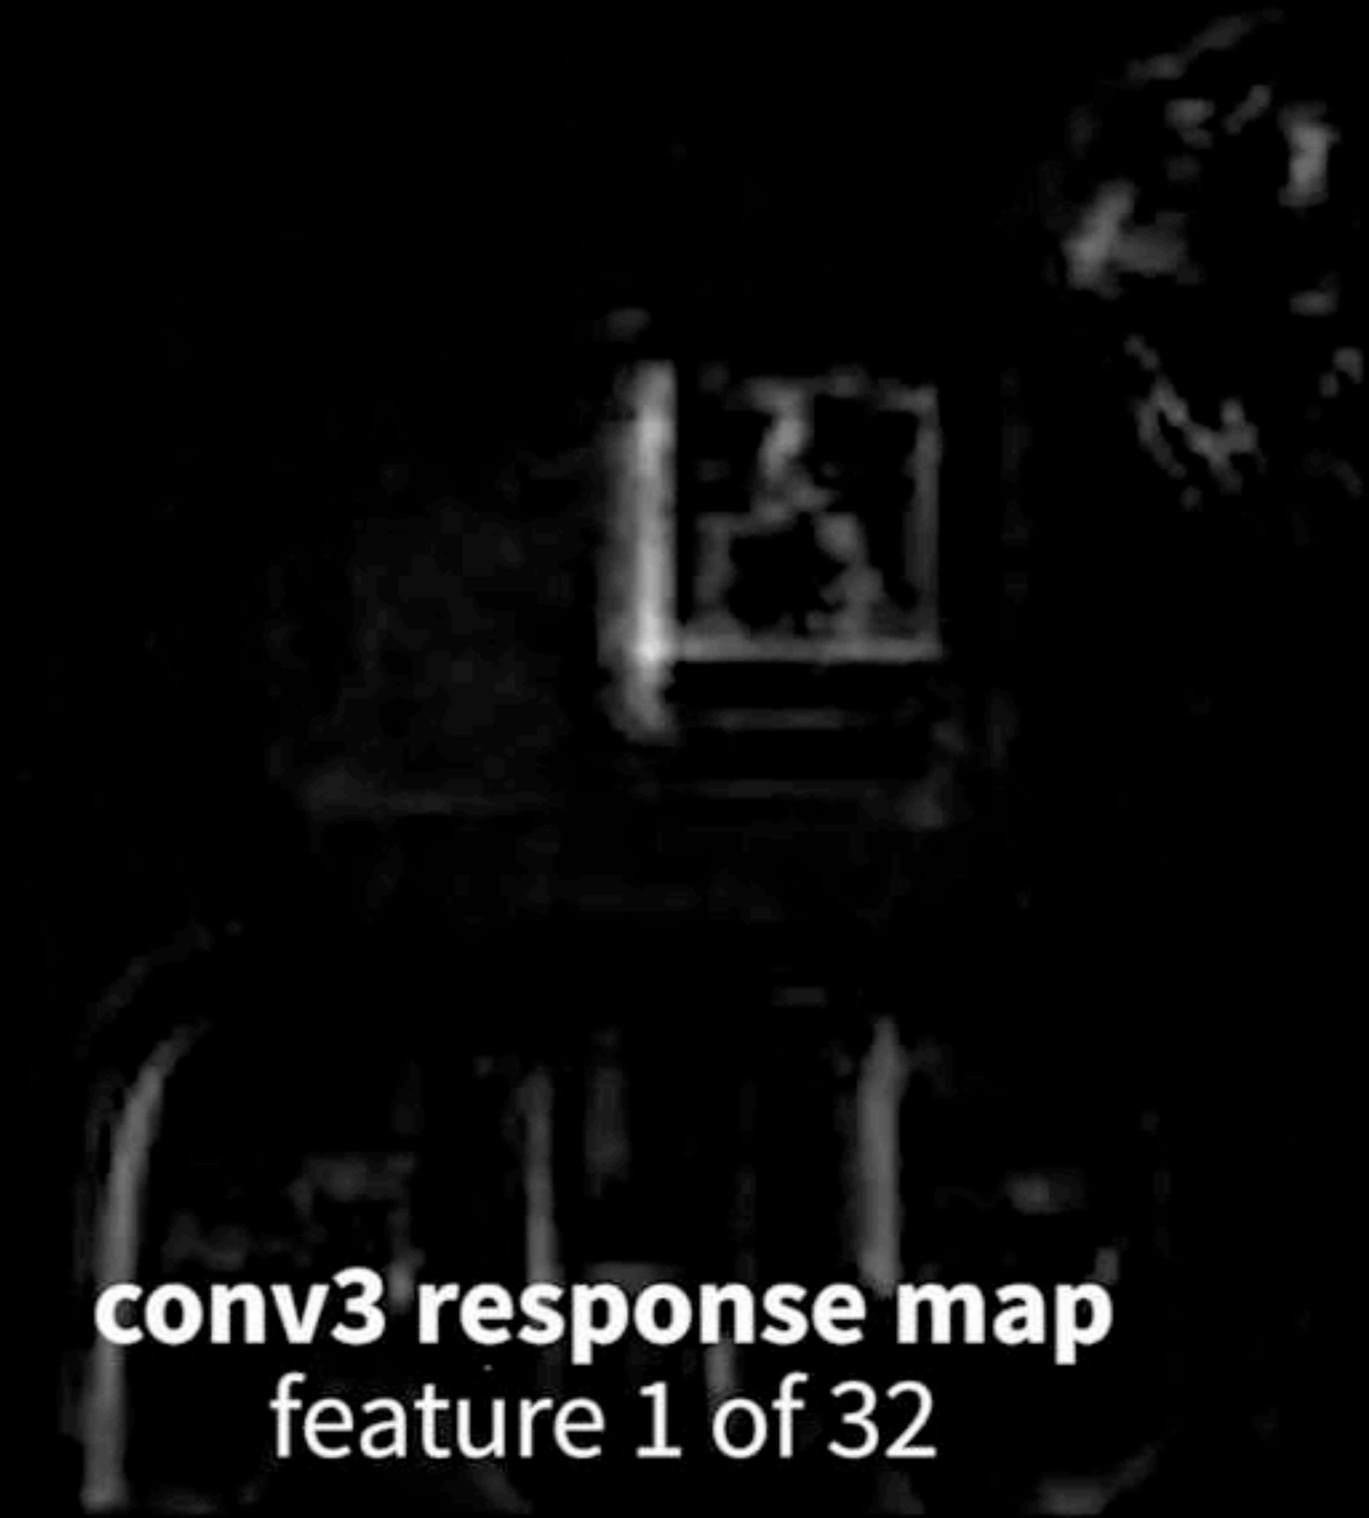

Learn feedback policies that use sensor feedback

J. Fu, S. Levine, P. Abbeel. “One-Shot Learning of Manipulation Skills with Online Dynamics Adaptation and Neural Network Priors.” arXiv, 2015.

Visual Feedback

Tactile Feedback

No Feedback

Visual Feedback

Tactile Feedback

No Feedback

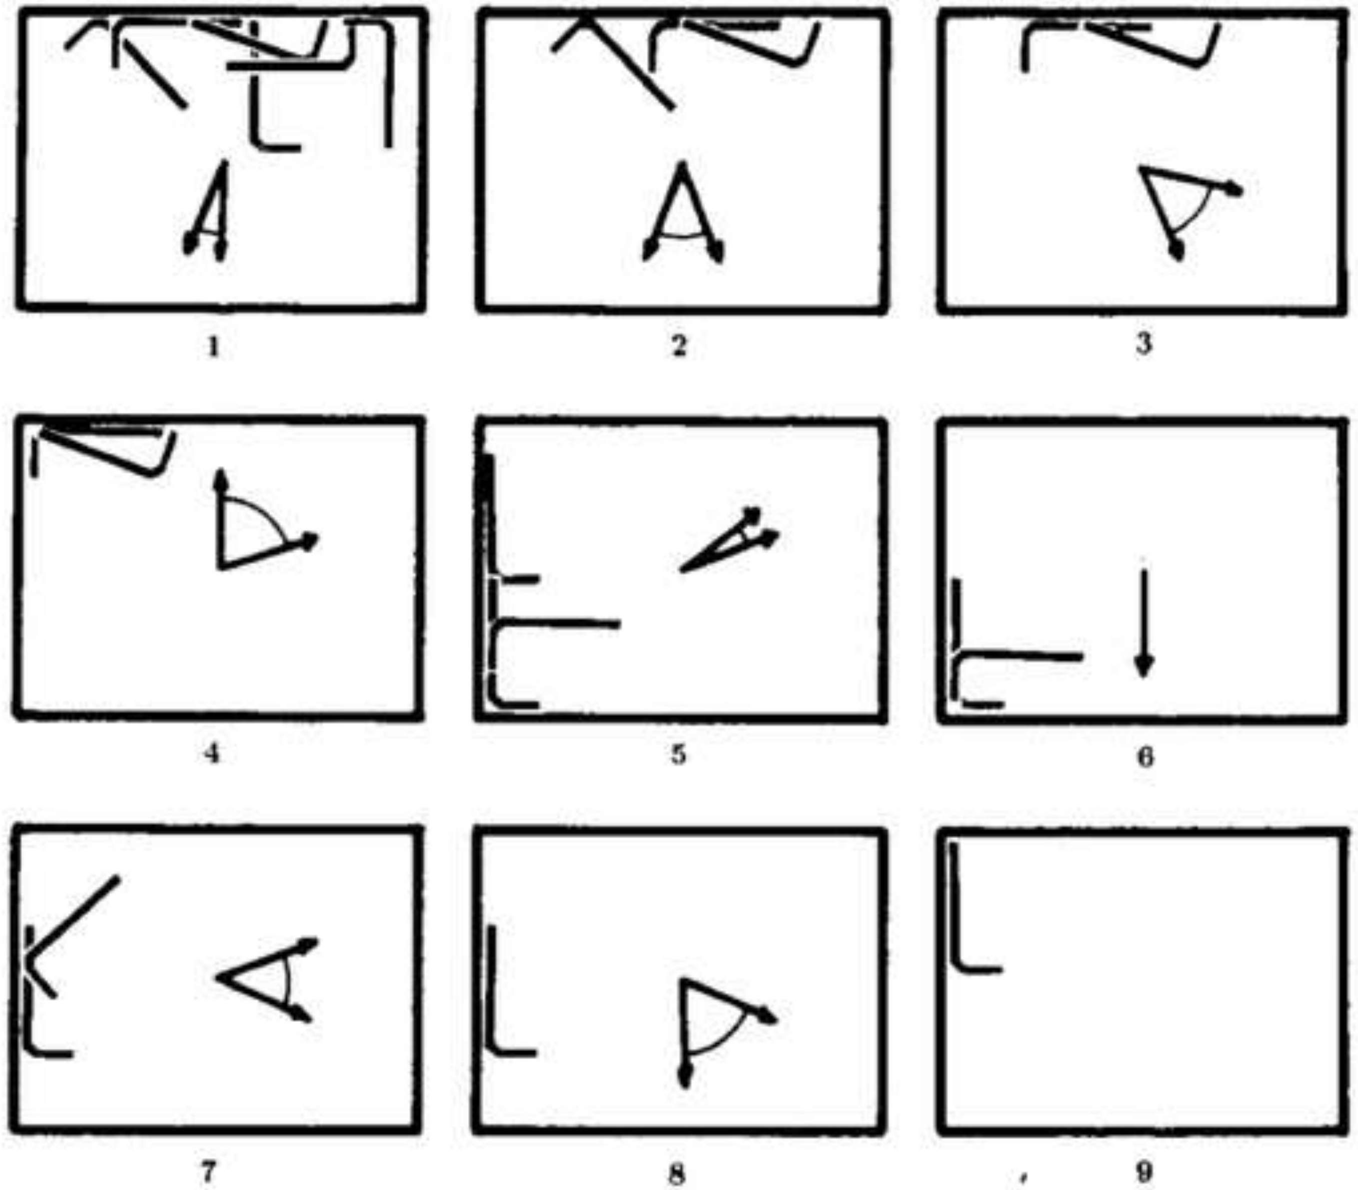

Open-loop robotic part alignment

M. Erdmann and M. Mason. “An exploration of sensorless manipulation.” IEEE Journal of Robotics and Automation, 1988.

Push Grasping

M. Dogar and S. Srinivasa. “Push-grasping with dexterous hands: Mechanics and a method.” IROS, 2010.
M. Dogar and S. Srinivasa. “A framework for push-grasping in clutter.” RSS, 2011.
M. Dogar, K. Hsiao, M. Ciocarlie, and S. Srinivasa “Physics-based grasp planning through clutter.” RSS, 2012.

Rearrangement Planning

M. Dogar and S. Srinivasa. “A Planning Framework for Non-Prehensile Manipulation under Clutter and Uncertainty.” AuRo, 2012.

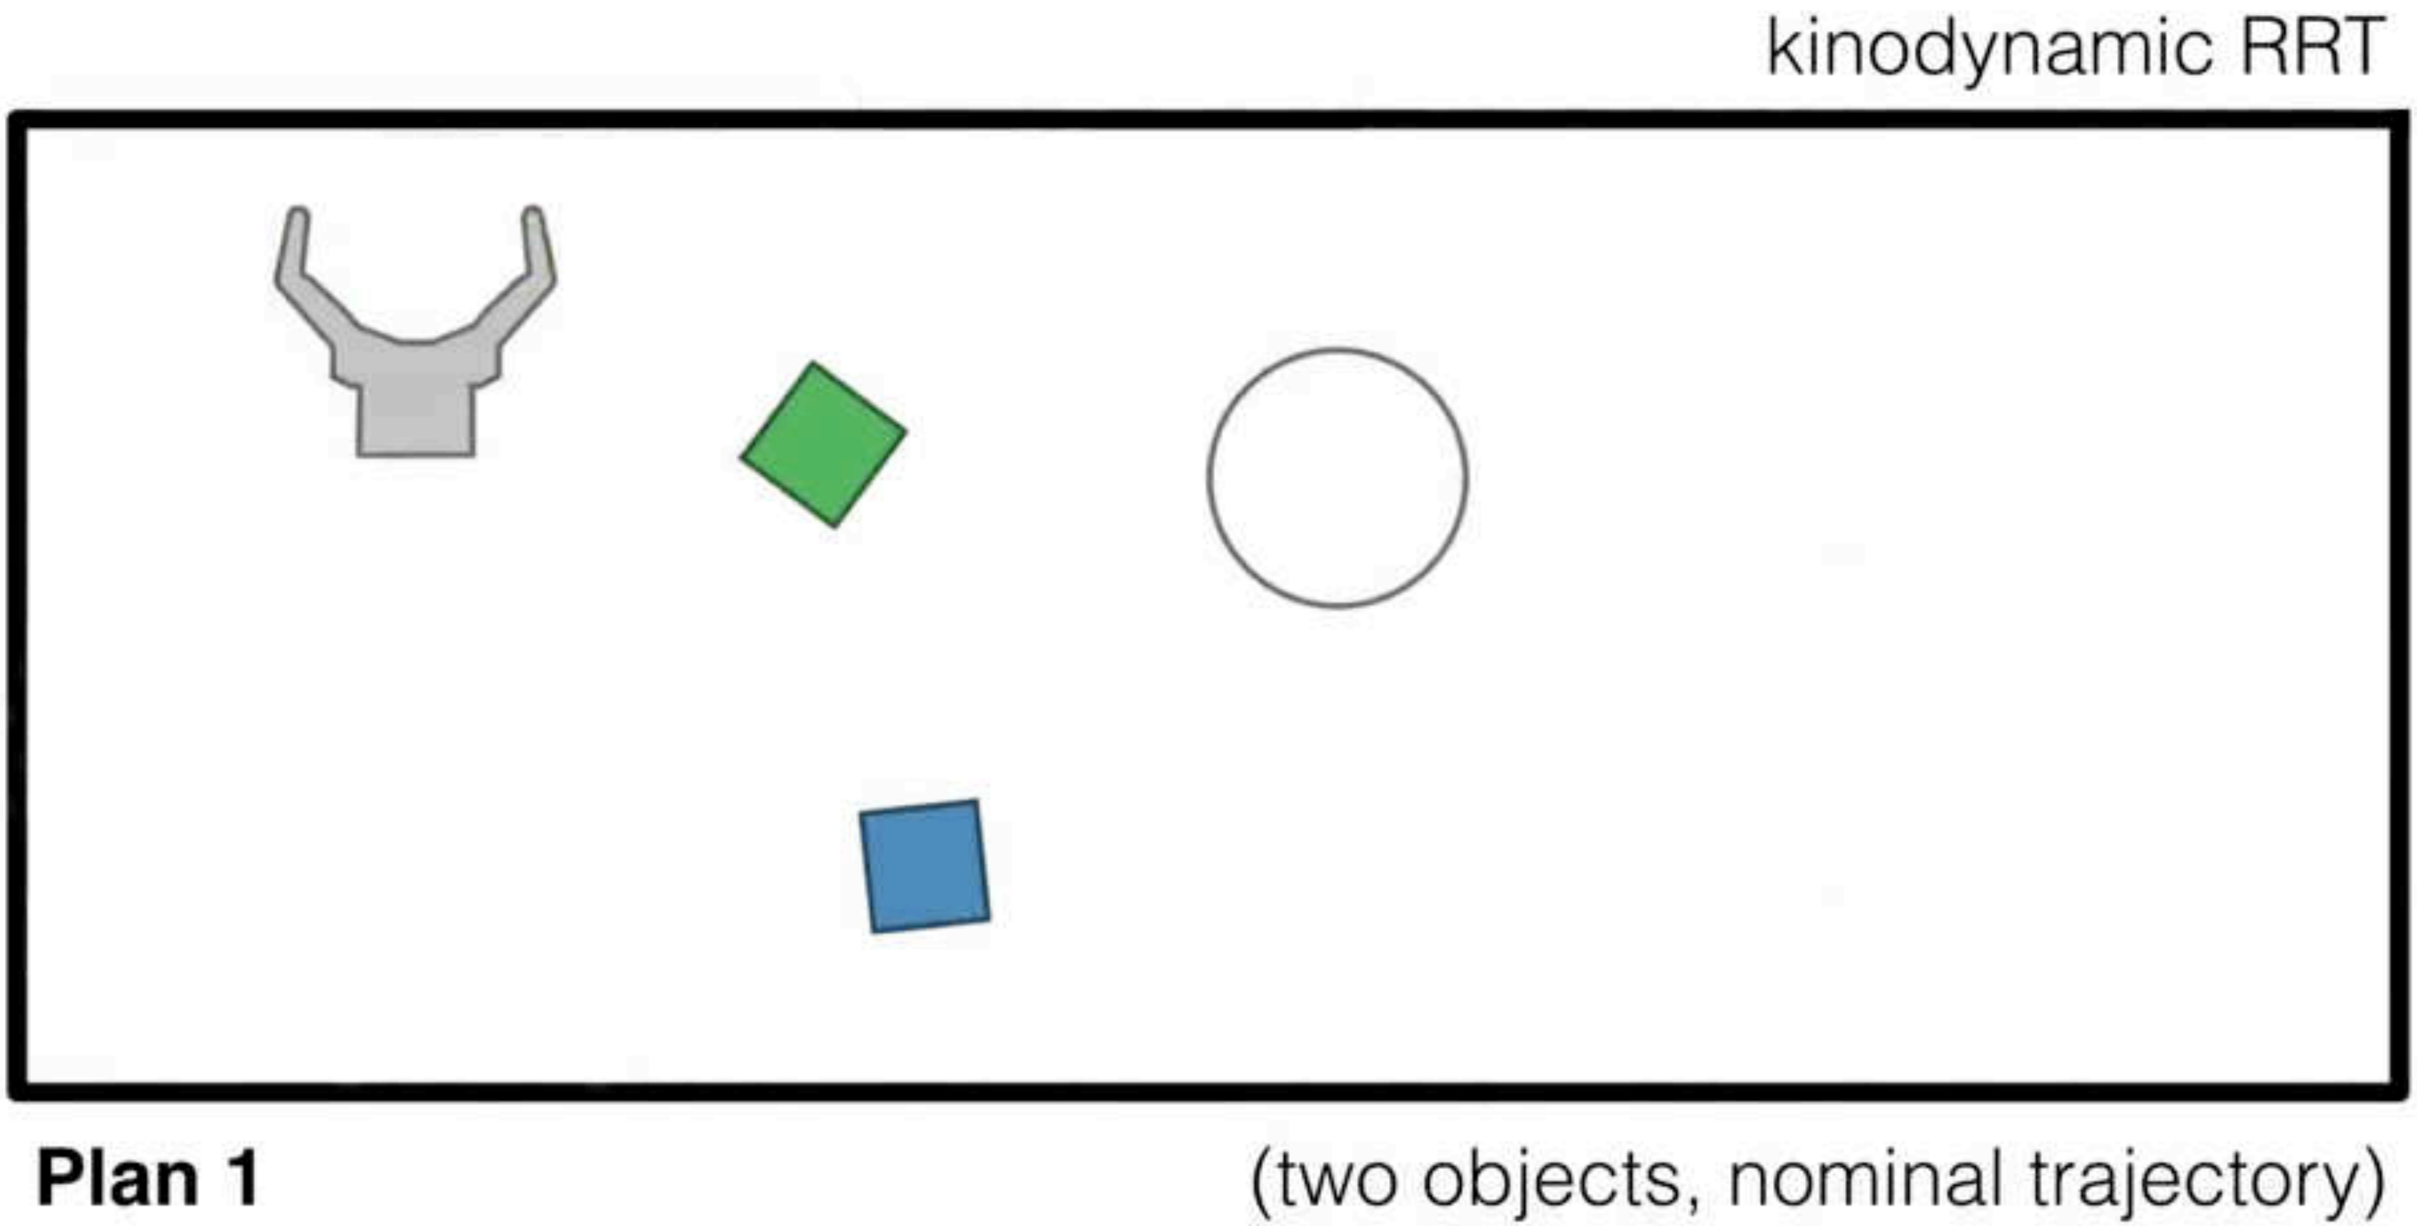

Robust Trajectory Selection

M.C. Koval, J.E. King, N.S. Pollard, and S.S. Srinivasa. “Robust trajectory selection for rearrangement planning as a multi-armed bandit problem.” IROS, 2015.

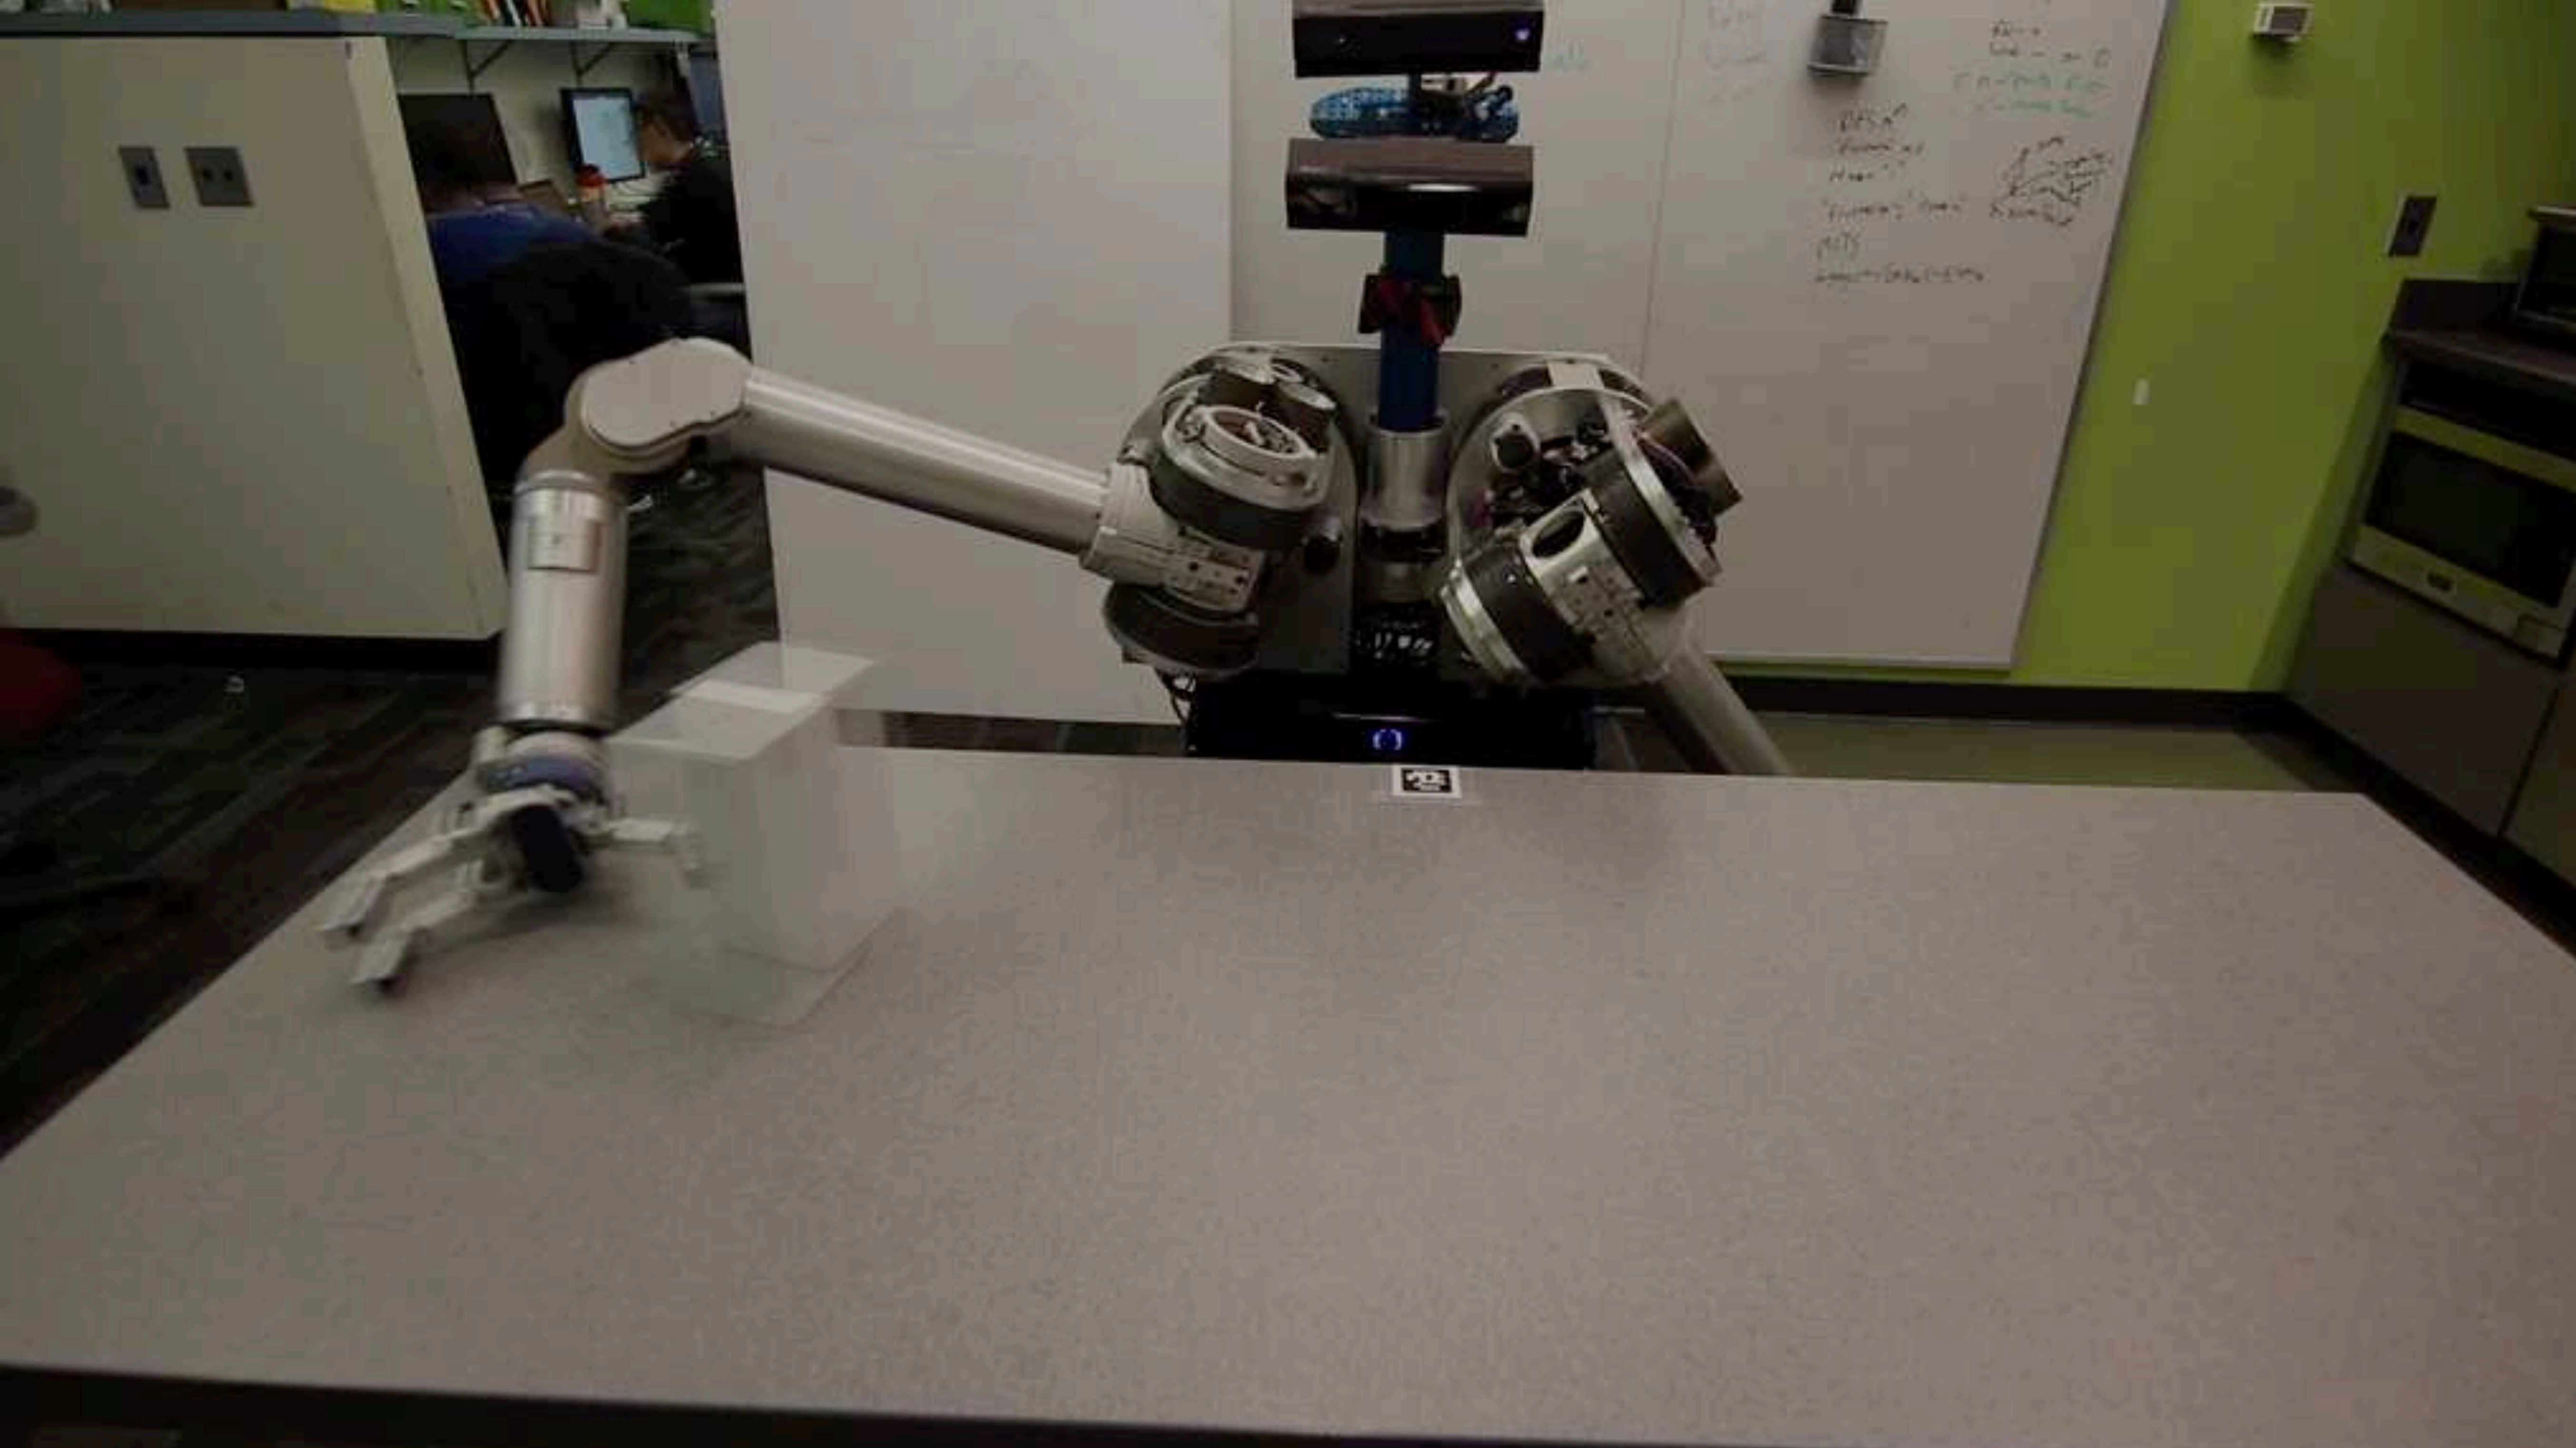

Convergent Planning

A.M. Johnson, J.E. King, and S.S. Srinivasa. "Convergent Planning." RAL, 2016. In press.

A brief introduction to POMDPs.

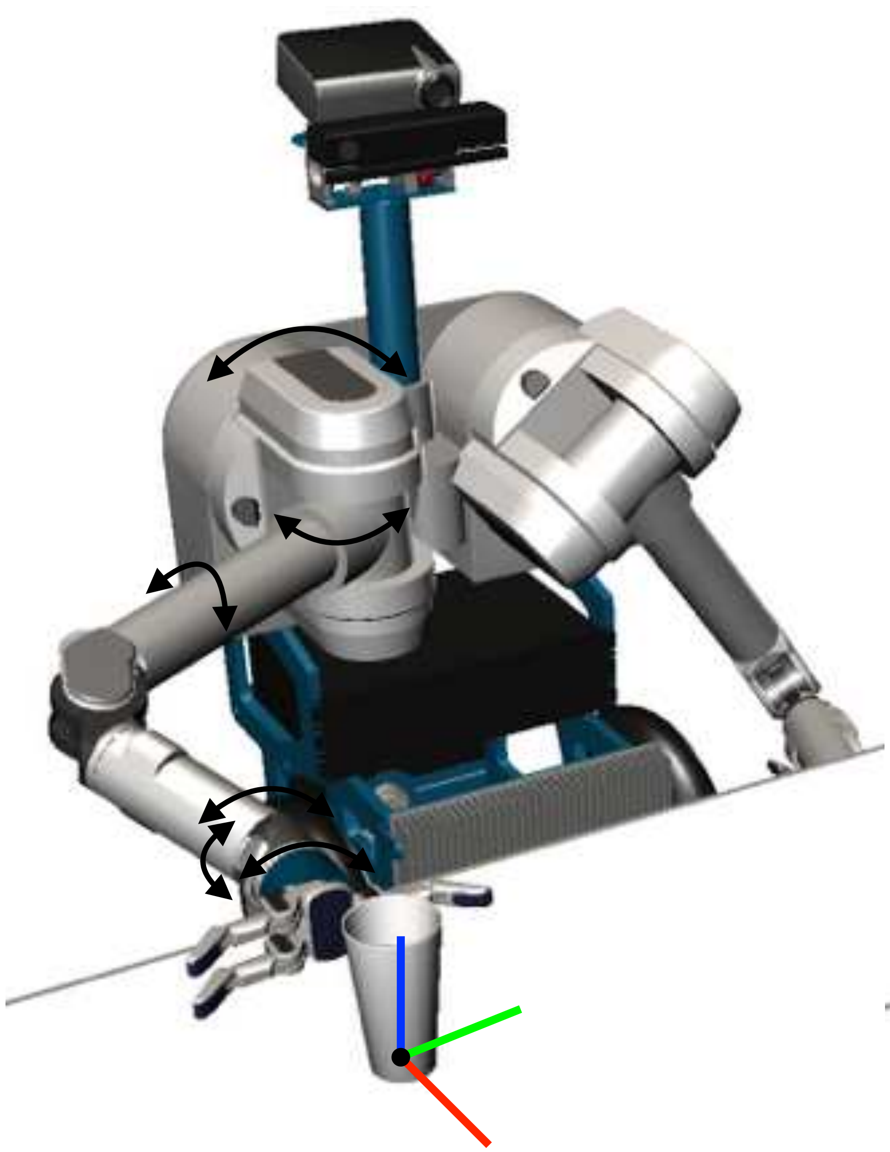

state

$s = (q, x)$

action

$a = (\dot{q}, \Delta t)$

$T = p(s'|s, a)$

observation

$o = (o_{\mathrm{q}}, o_{\mathrm{c}})$

$\Omega = p(o|s, a)$

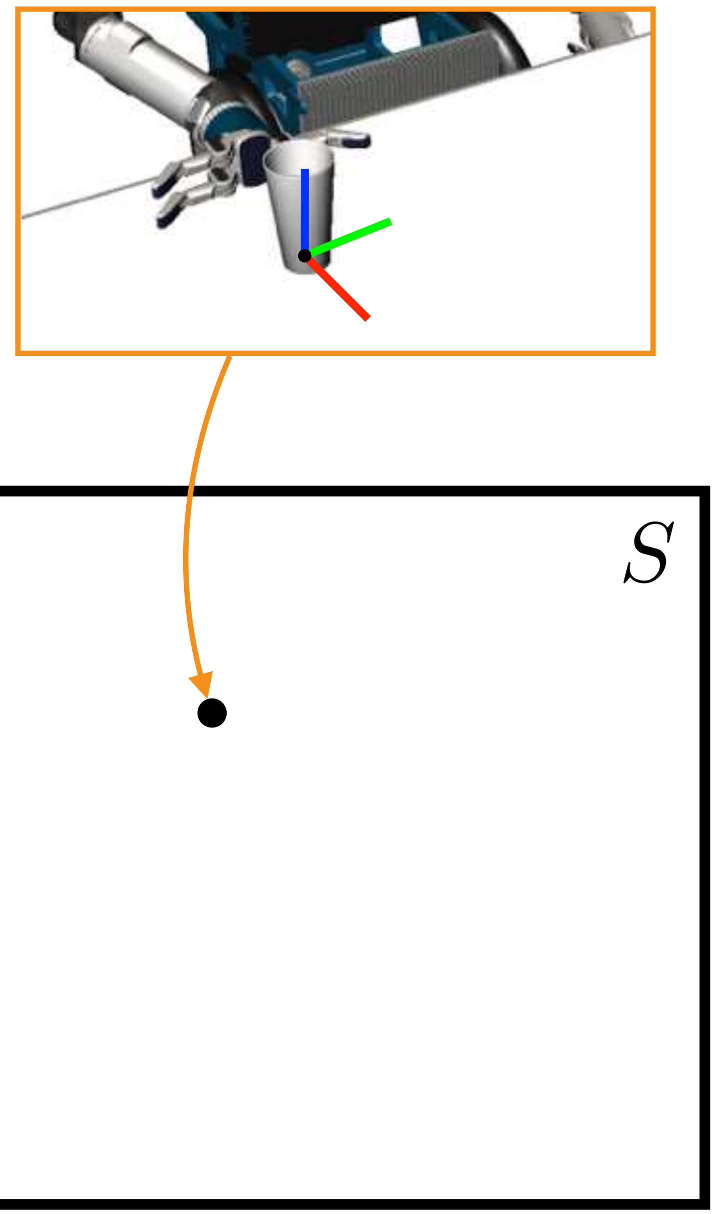

state space

$$\dim(S) = n$$

Planning in Belief Space

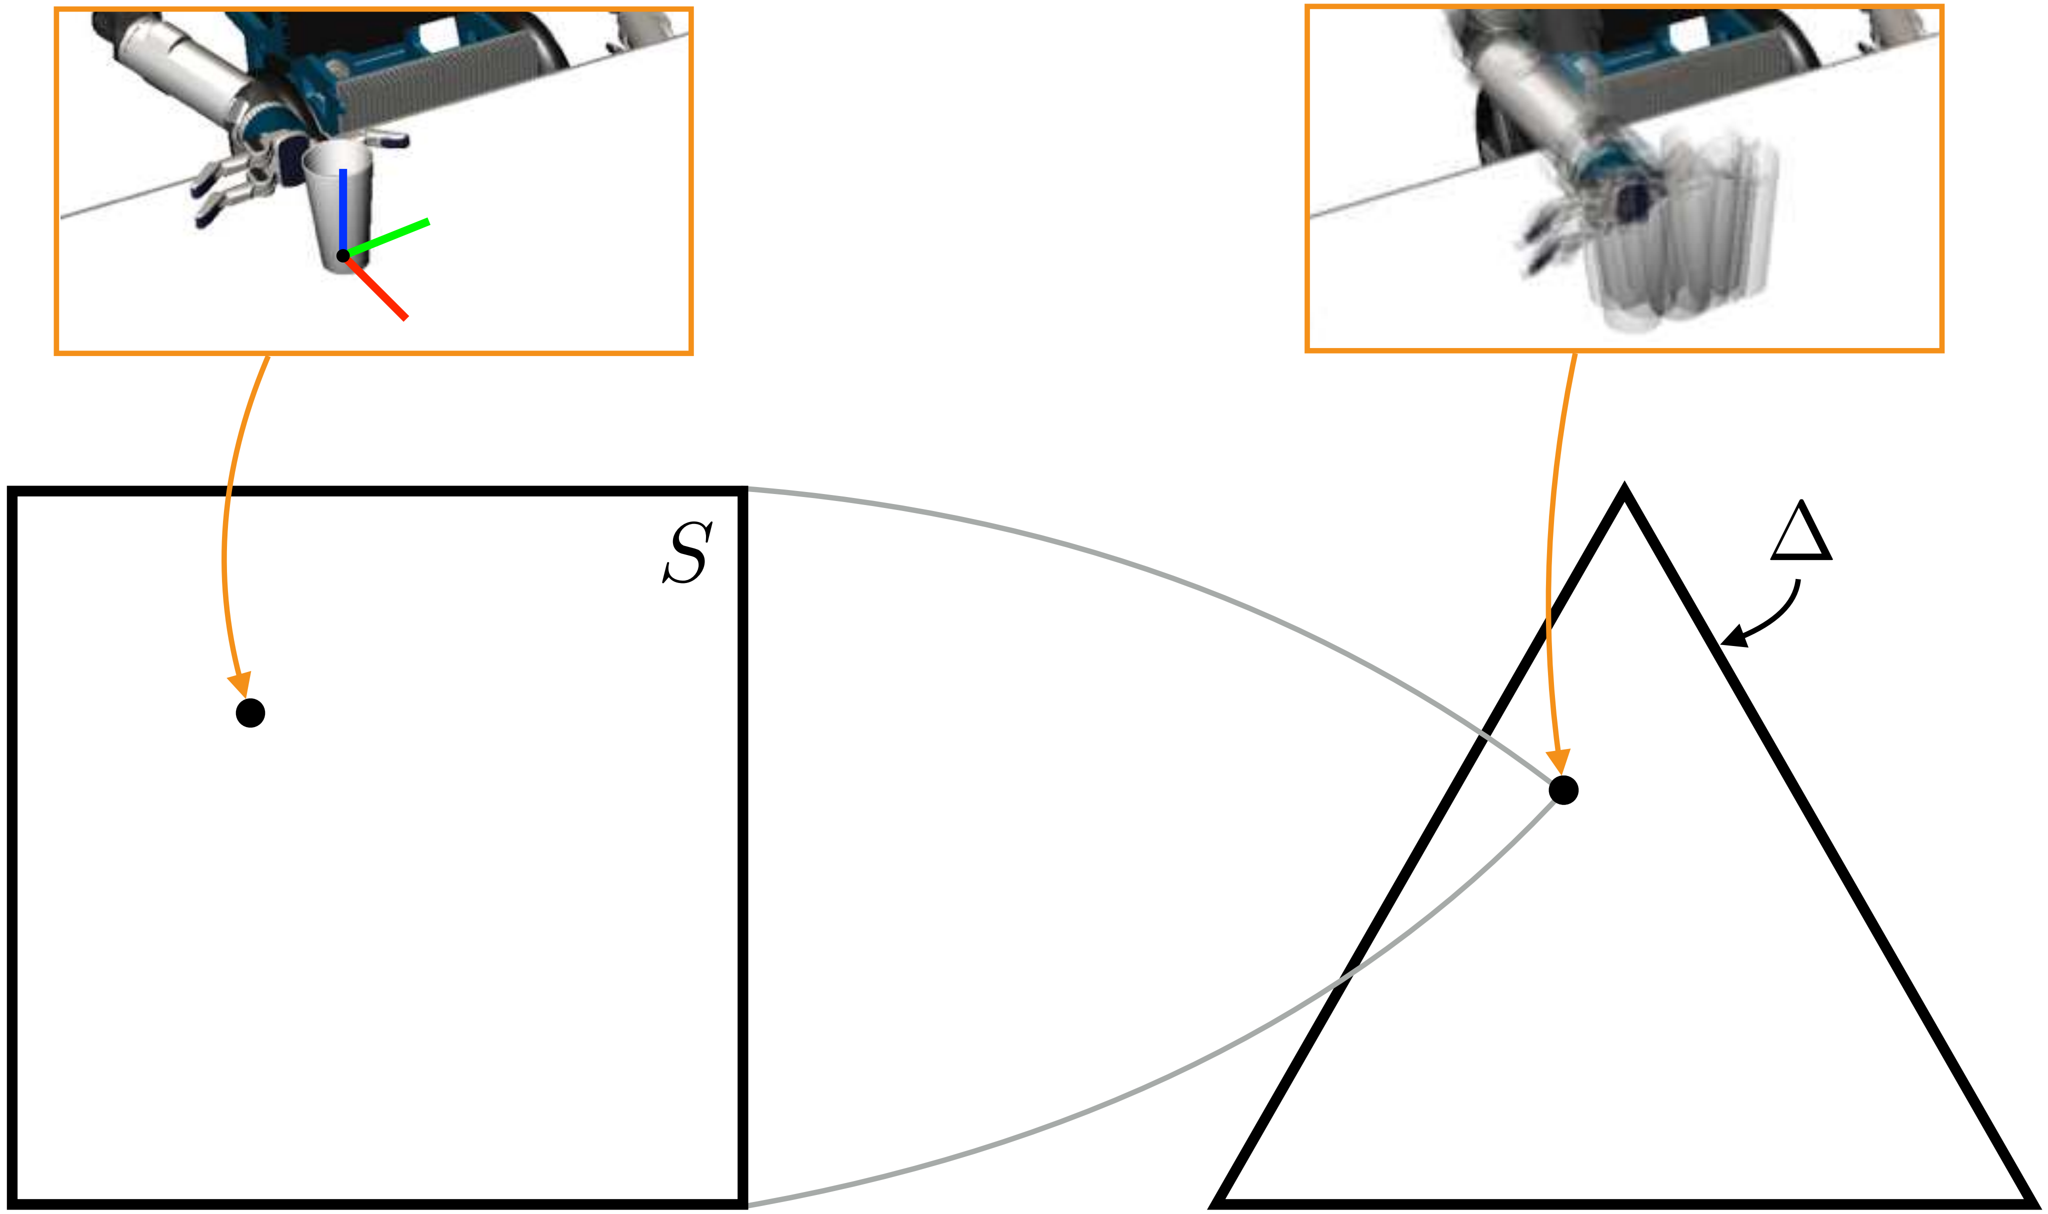

state space

$$\dim(S) = n$$

belief space

$$\dim(\Delta) = \infty$$

Planning in Belief Space

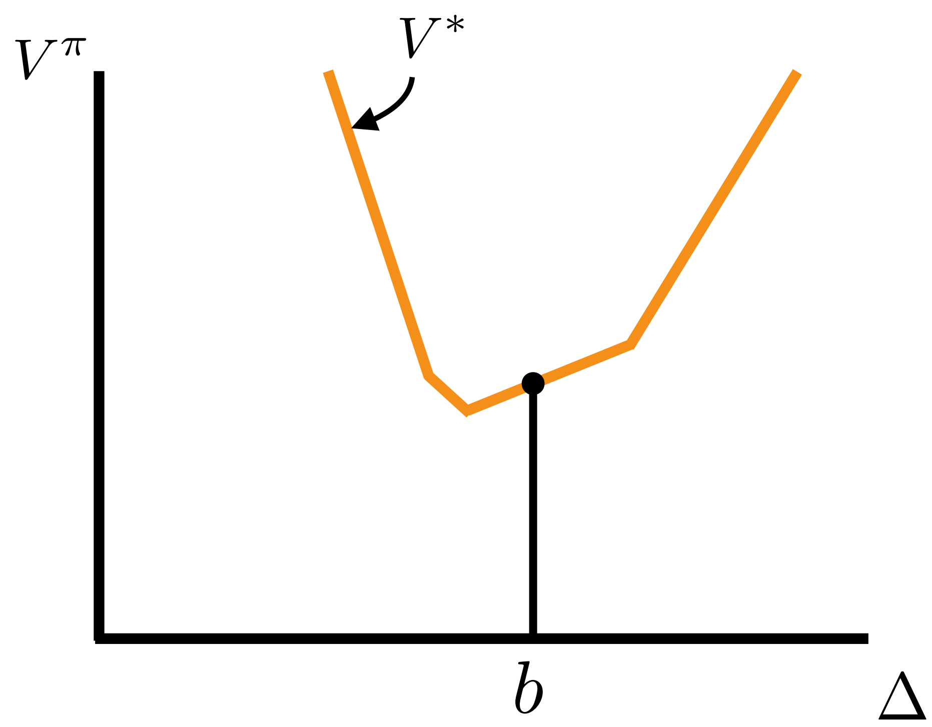

Offline Planning

Point-Based Methods

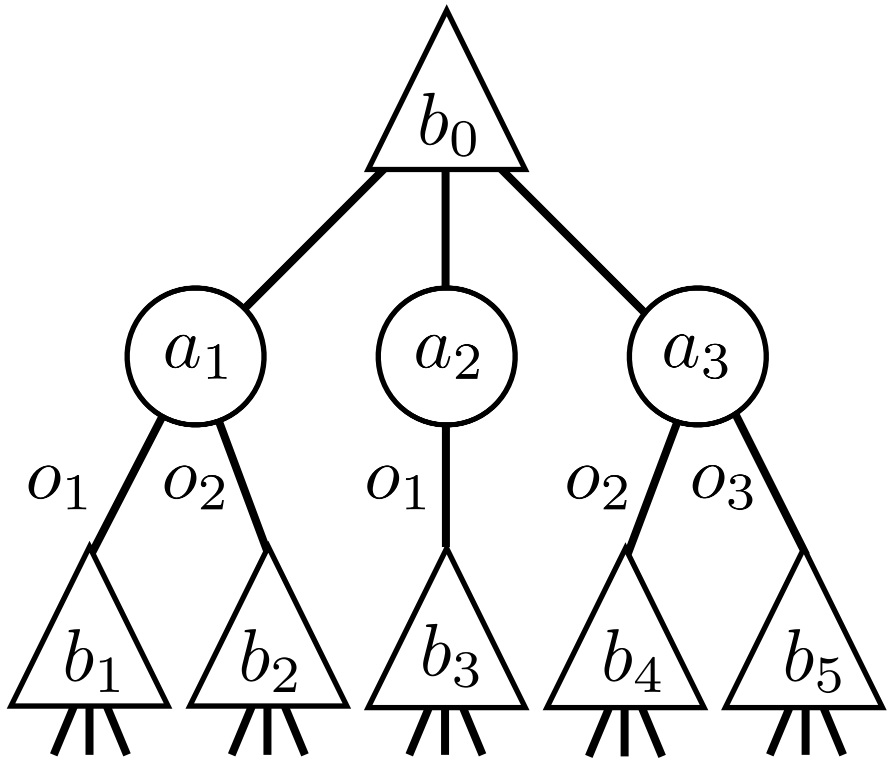

Online Planning

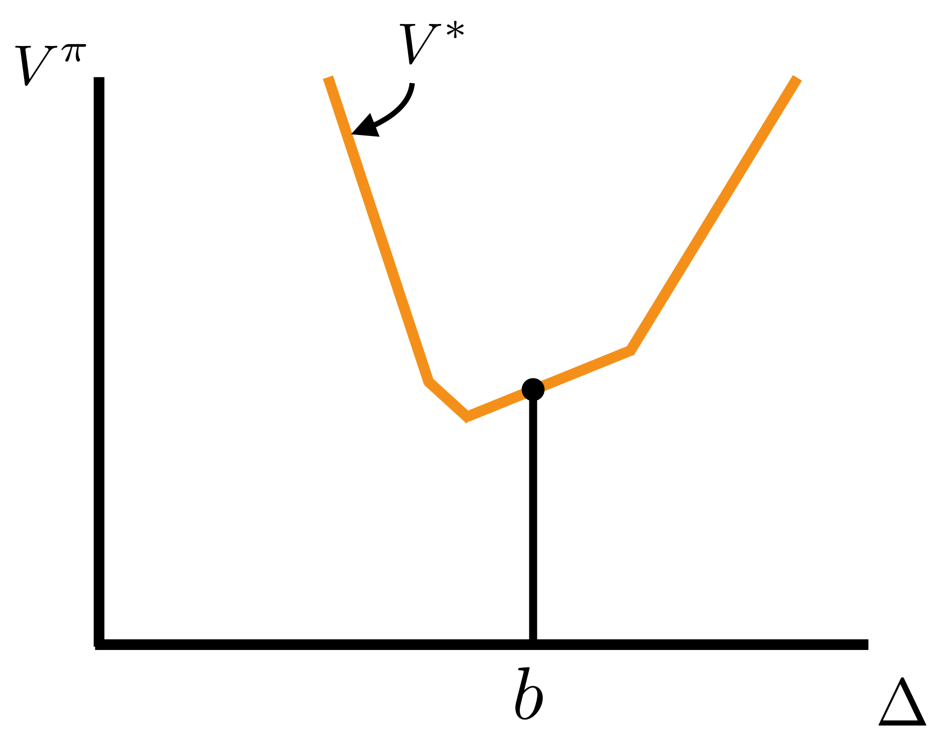

Offline Planning

Point-Based Methods

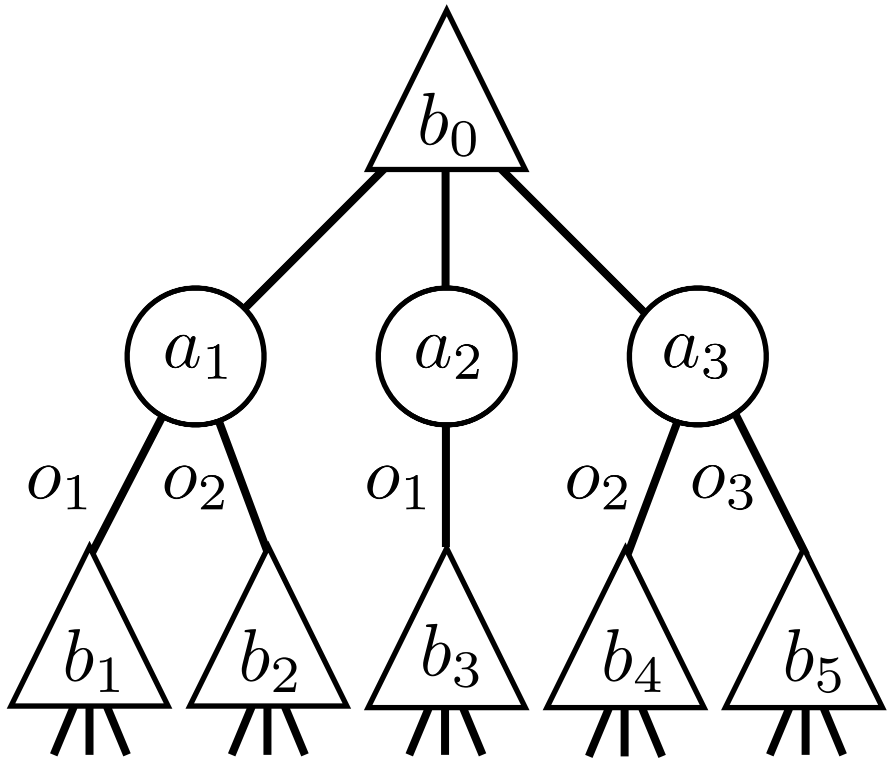

Online Planning

$$V^{\pi} = \sum_{t=1}^{\infty} \gamma^{t} R(s_t, a_t)$$

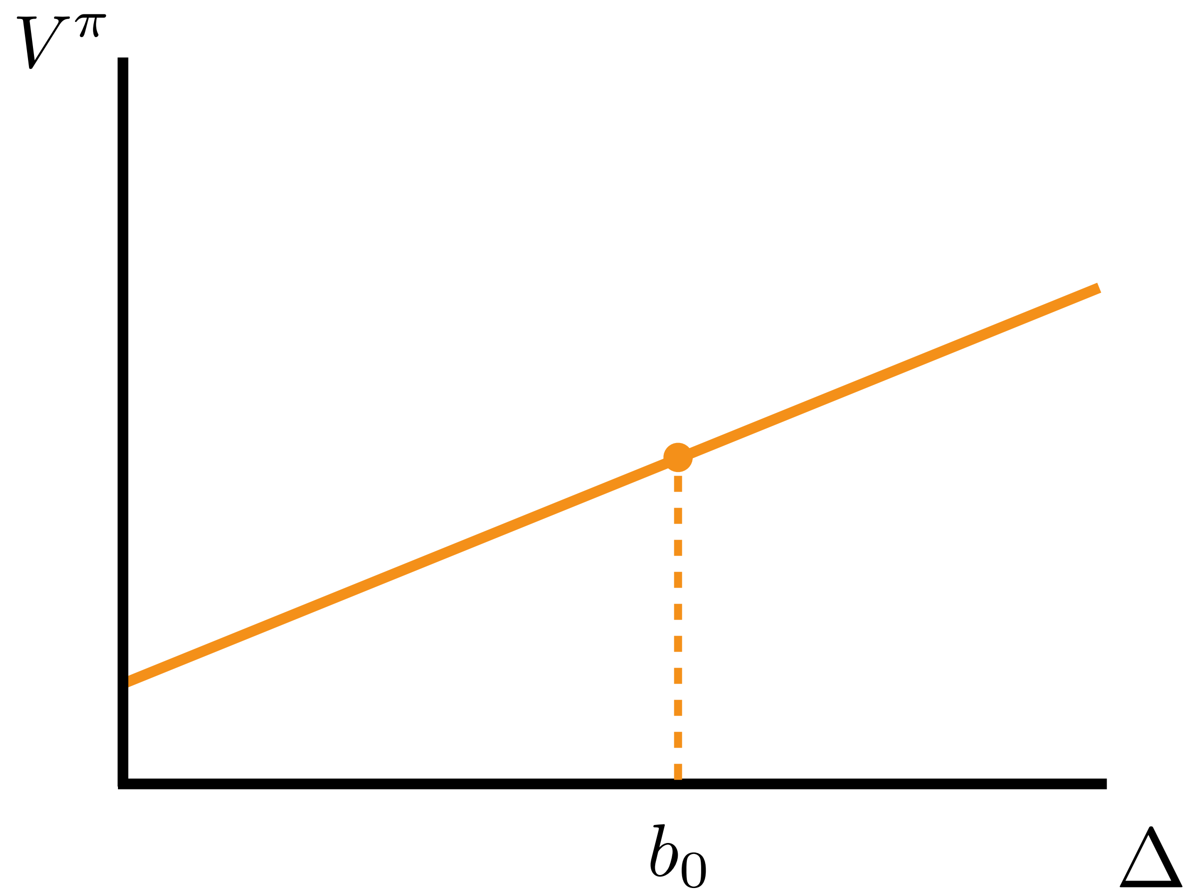

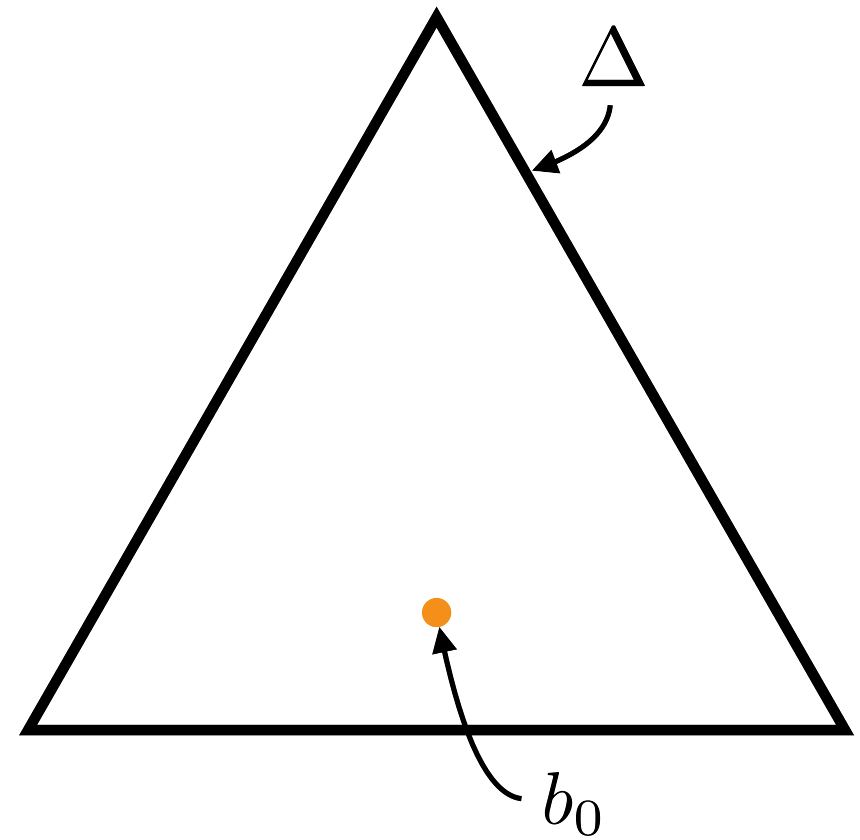

Point-based solvers

J. Pineau, G. Gordon, and S. Thrun. Point-based value iteration: An anytime algorithm for POMDPs. IJCAI, 2003.
H. Kurniawati, D. Hsu, and W.S. Lee. SARSOP: Efficient point-based POMDP planning by approximating optimally reachable belief spaces. RSS, 2008.

$$V^\pi = \sum_{t=1}^{\infty} \gamma^t R(s_t, a_t)$$

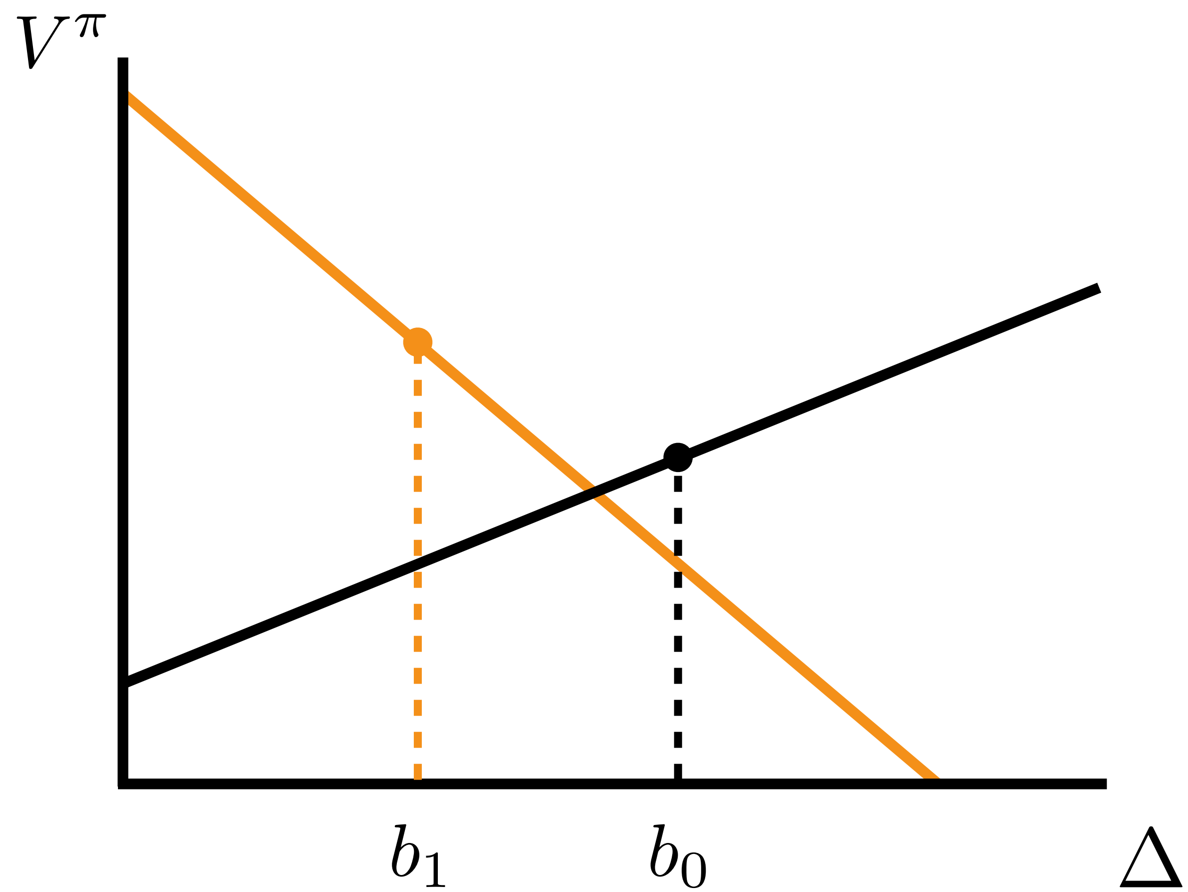

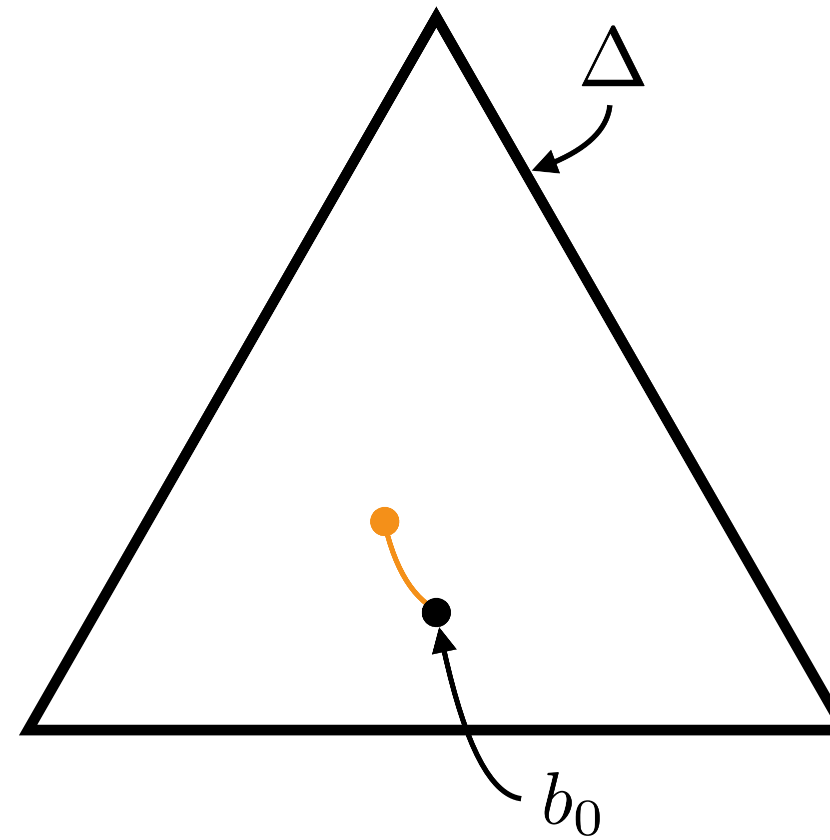

Point-based solvers

J. Pineau, G. Gordon, and S. Thrun. Point-based value iteration: An anytime algorithm for POMDPs. IJCAI, 2003.
H. Kurniawati, D. Hsu, and W.S. Lee. SARSOP: Efficient point-based POMDP planning by approximating optimally reachable belief spaces. RSS, 2008.

$$V^{\pi} = \sum_{t=1}^{\infty} \gamma^t R(s_t, a_t)$$

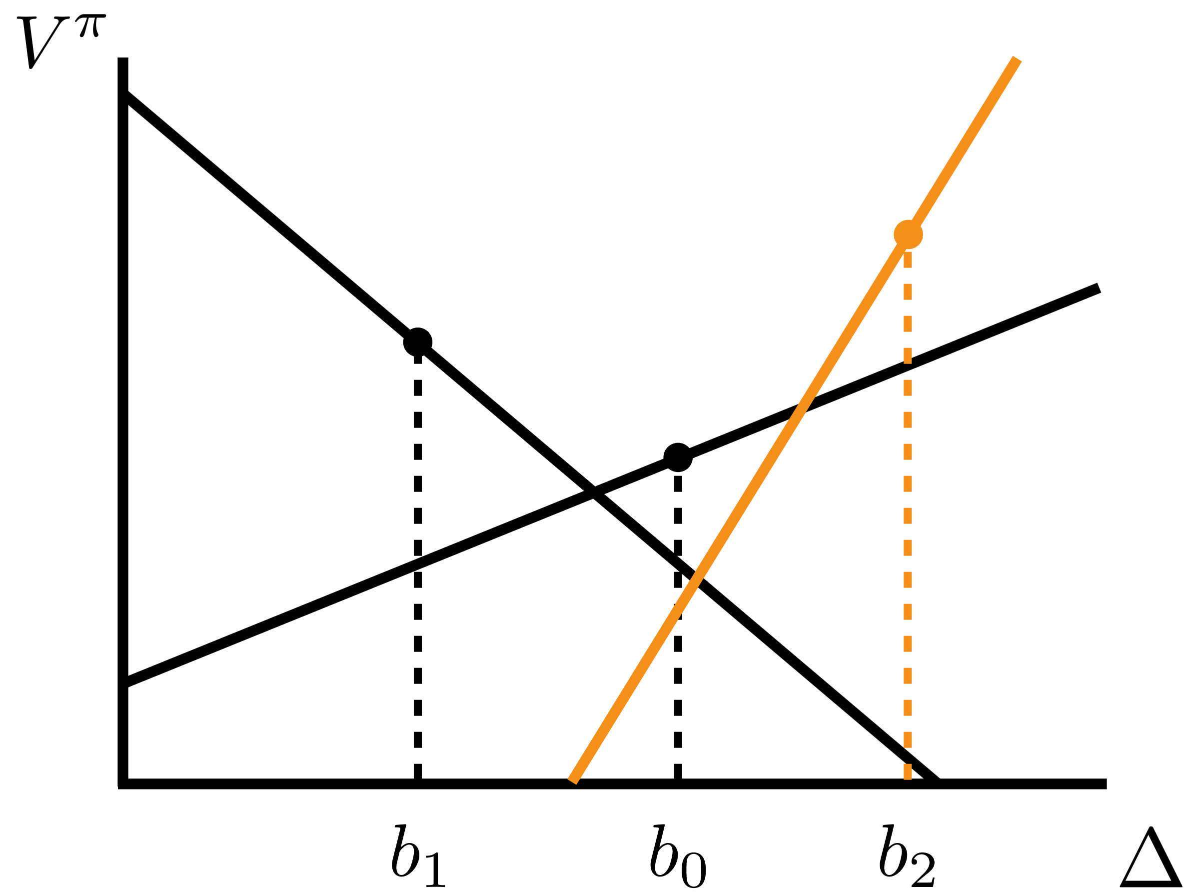

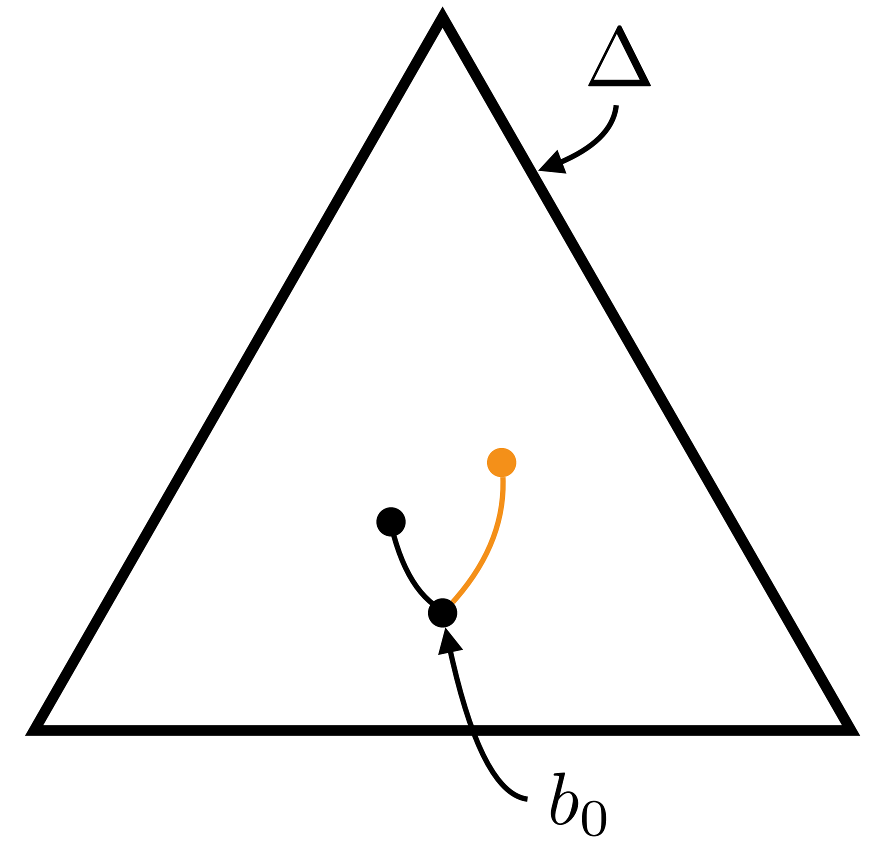

Point-based solvers

J. Pineau, G. Gordon, and S. Thrun. Point-based value iteration: An anytime algorithm for POMDPs. IJCAI, 2003.
H. Kurniawati, D. Hsu, and W.S. Lee. SARSOP: Efficient point-based POMDP planning by approximating optimally reachable belief spaces. RSS, 2008.

$$V^\pi = \sum_{t=1}^{\infty} \gamma^t R(s_t, a_t)$$

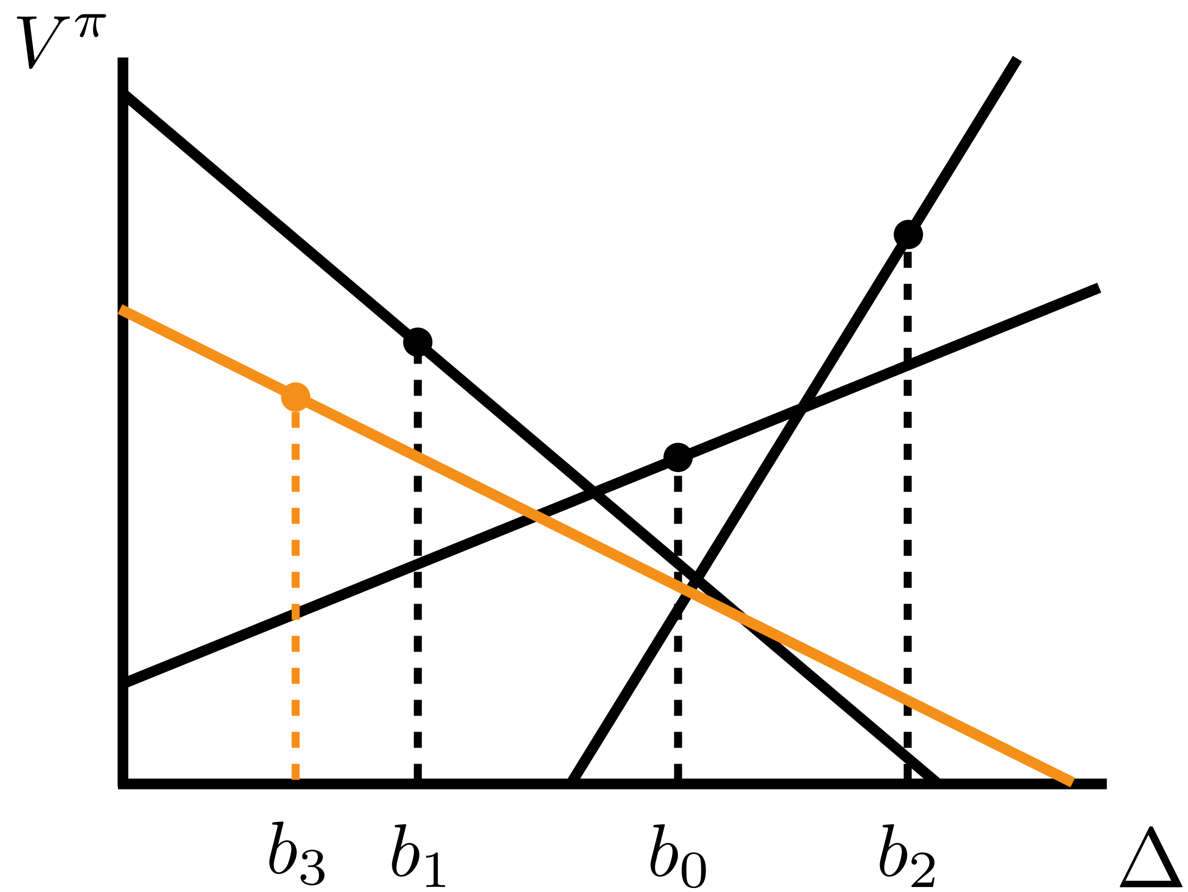

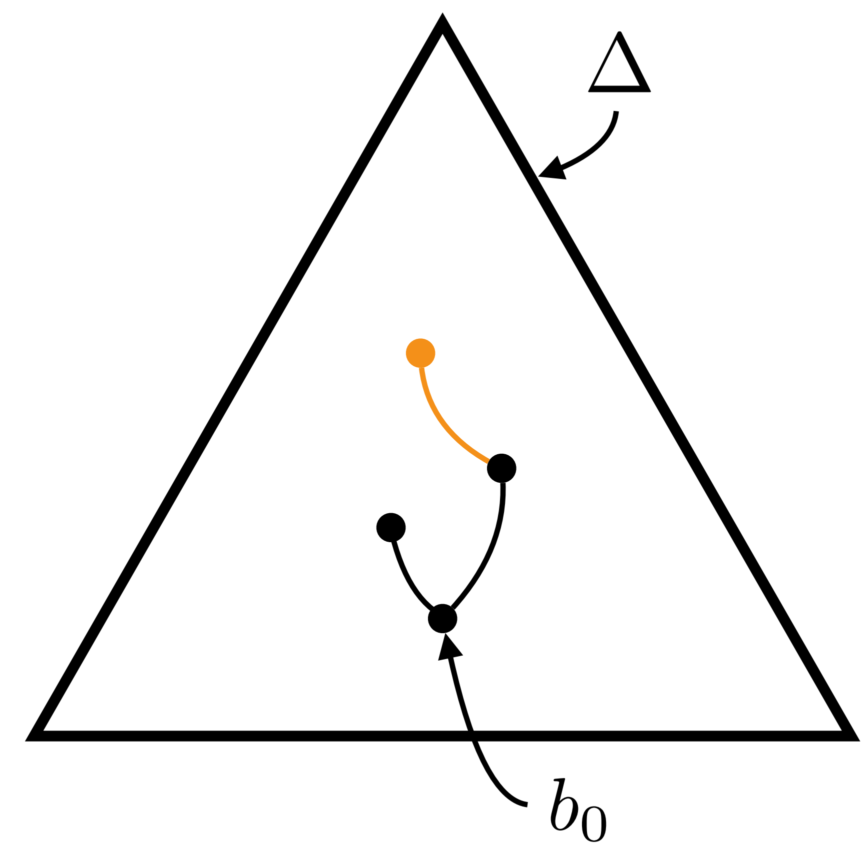

Point-based solvers

J. Pineau, G. Gordon, and S. Thrun. Point-based value iteration: An anytime algorithm for POMDPs. IJCAI, 2003.
H. Kurniawati, D. Hsu, and W.S. Lee. SARSOP: Efficient point-based POMDP planning by approximating optimally reachable belief spaces. RSS, 2008.

$$V^{\pi} = \sum_{t=1}^{\infty} \gamma^t R(s_t, a_t)$$

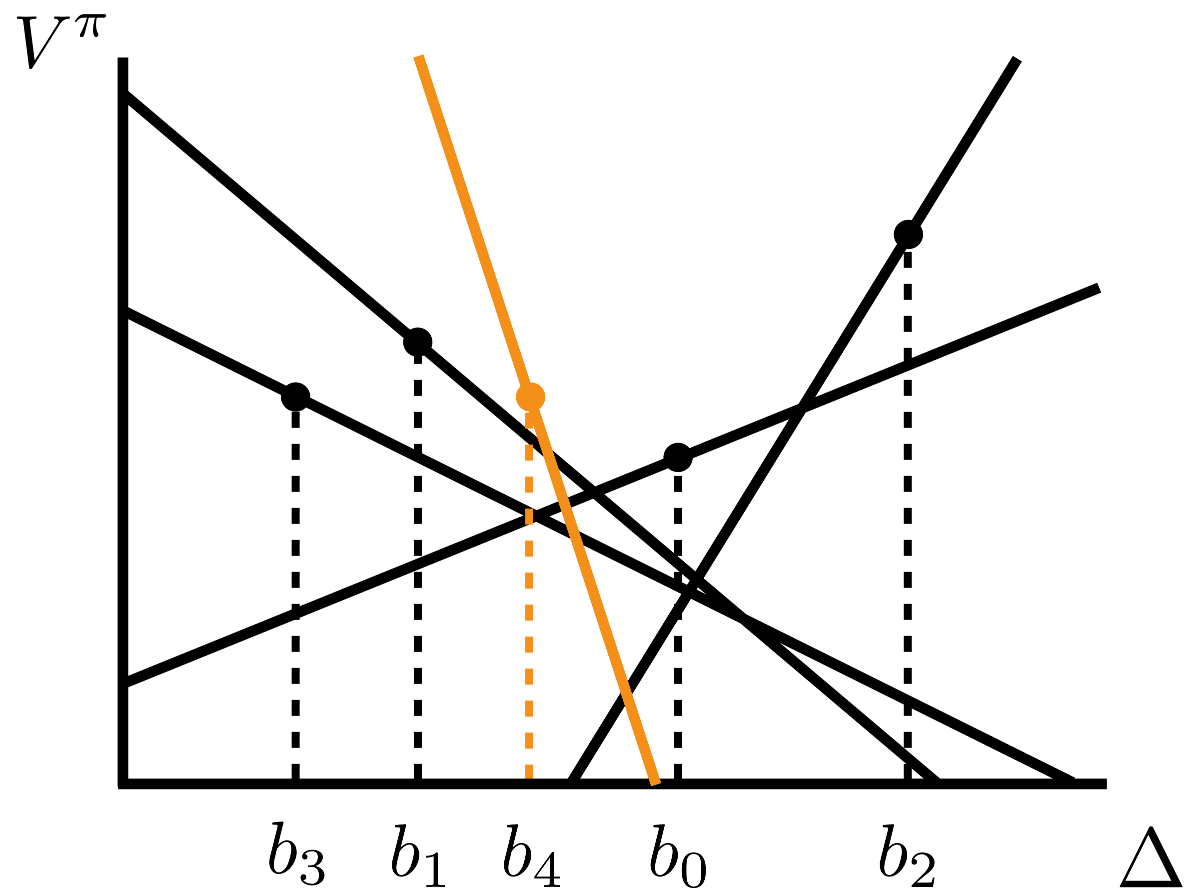

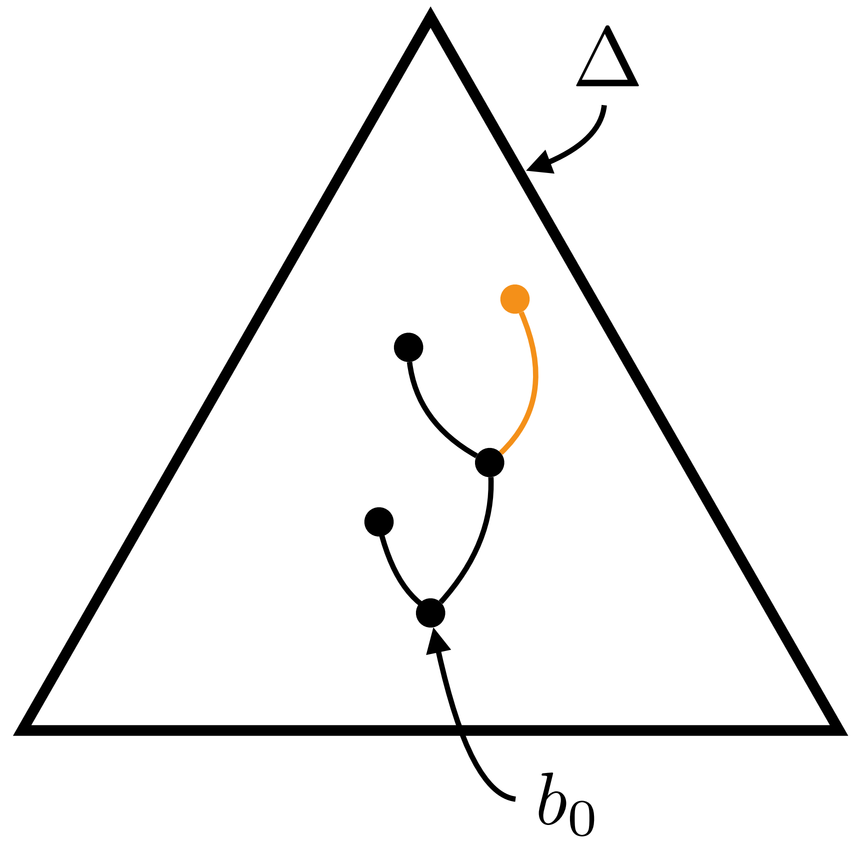

Point-based solvers

J. Pineau, G. Gordon, and S. Thrun. Point-based value iteration: An anytime algorithm for POMDPs. IJCAI, 2003.
H. Kurniawati, D. Hsu, and W.S. Lee. SARSOP: Efficient point-based POMDP planning by approximating optimally reachable belief spaces. RSS, 2008.

$$\pi^* = \arg\max_\pi V^\pi [b(s_0)]$$

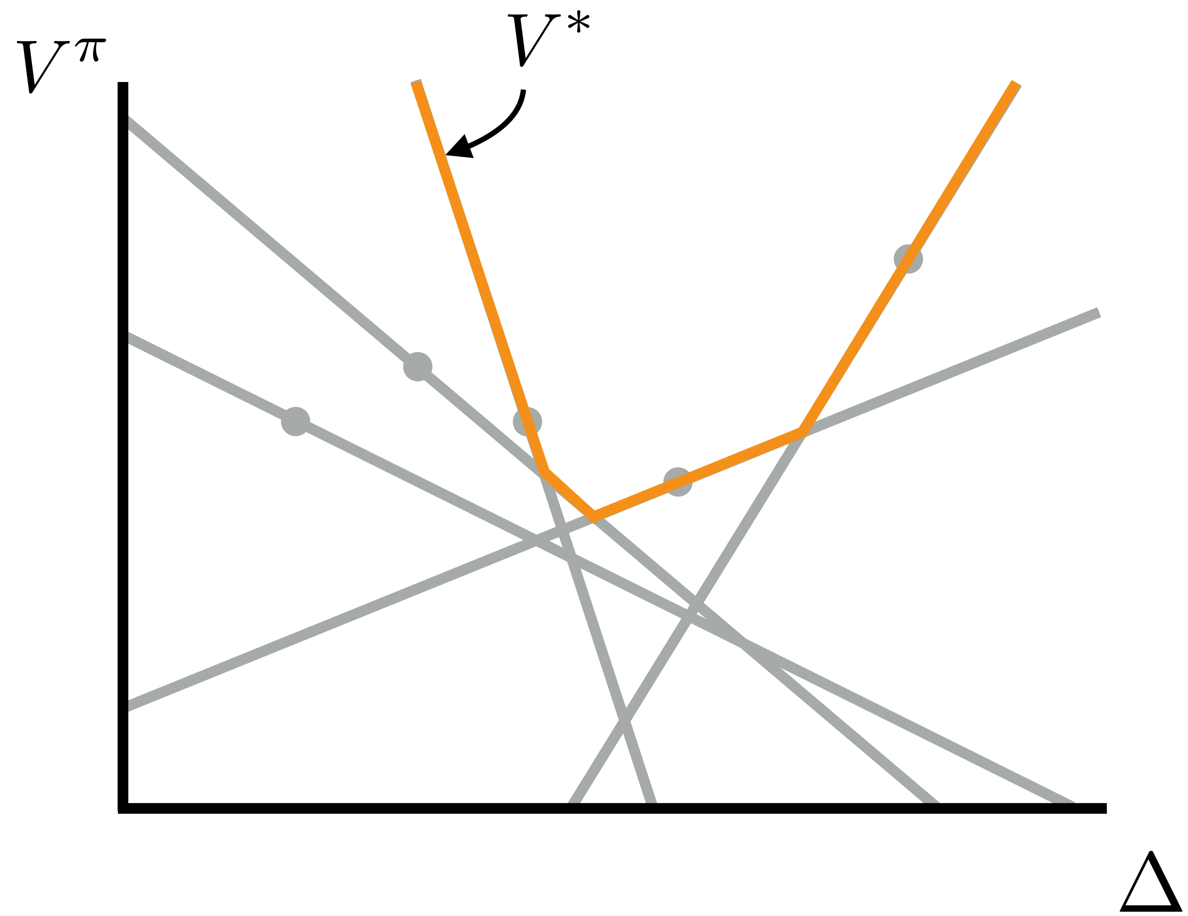

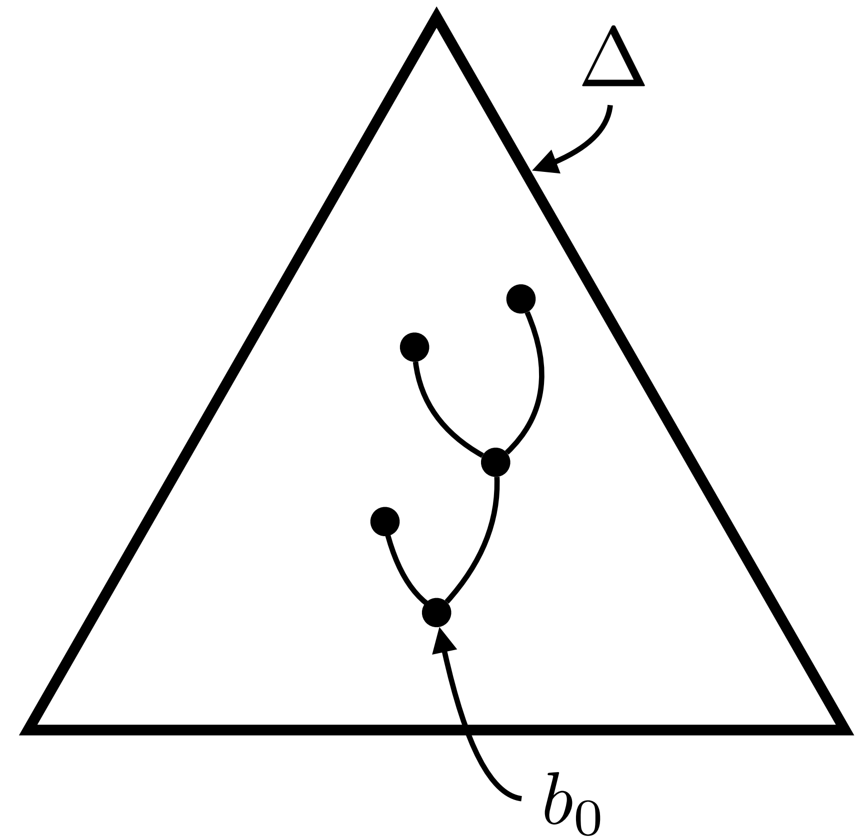

Point-based solvers

J. Pineau, G. Gordon, and S. Thrun. Point-based value iteration: An anytime algorithm for POMDPs. IJCAI, 2003.
H. Kurniawati, D. Hsu, and W.S. Lee. SARSOP: Efficient point-based POMDP planning by approximating optimally reachable belief spaces. RSS, 2008.

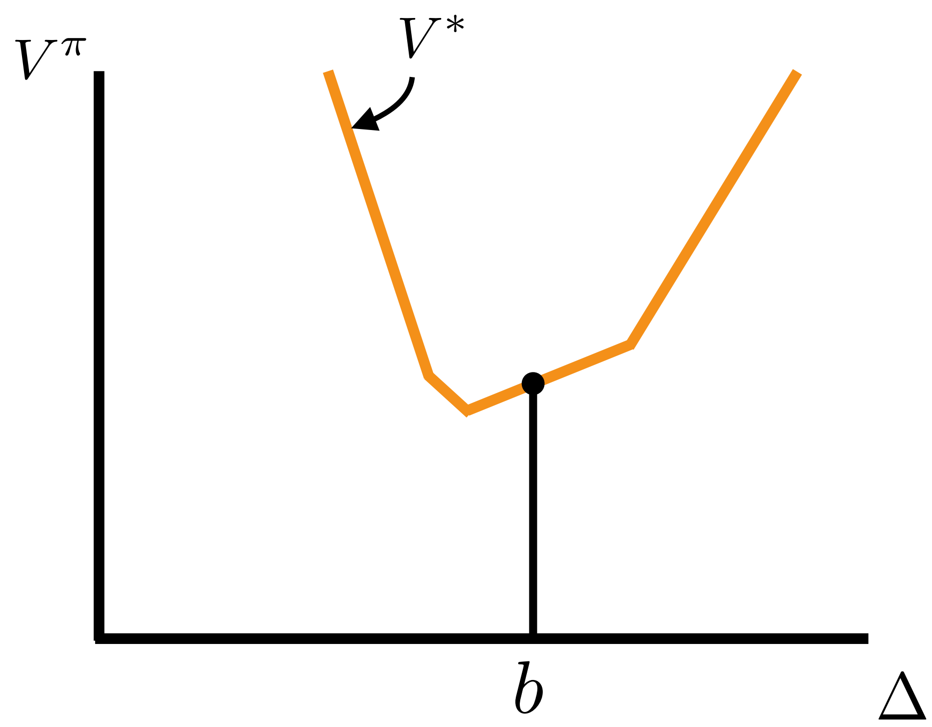

Offline Planning

Point-Based Methods

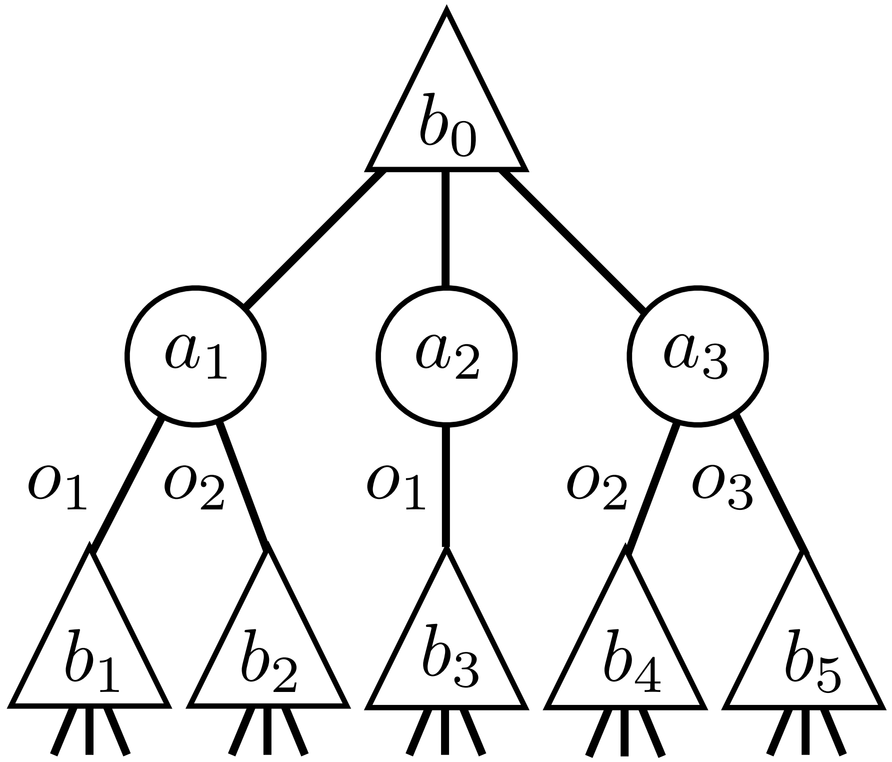

Online Planning

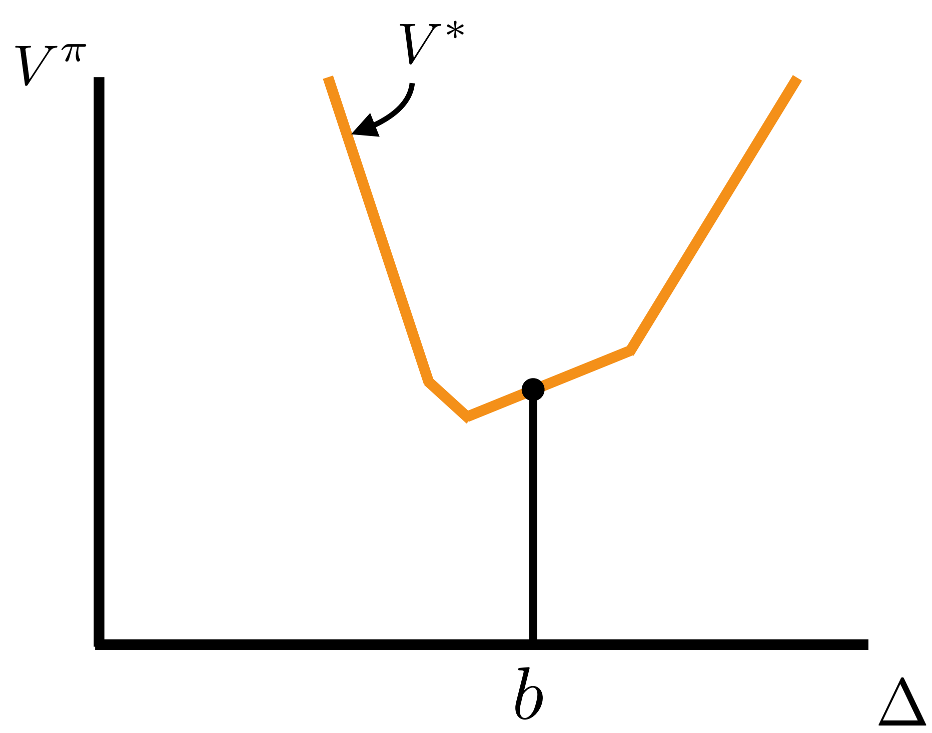

Offline Planning

Point-Based Methods

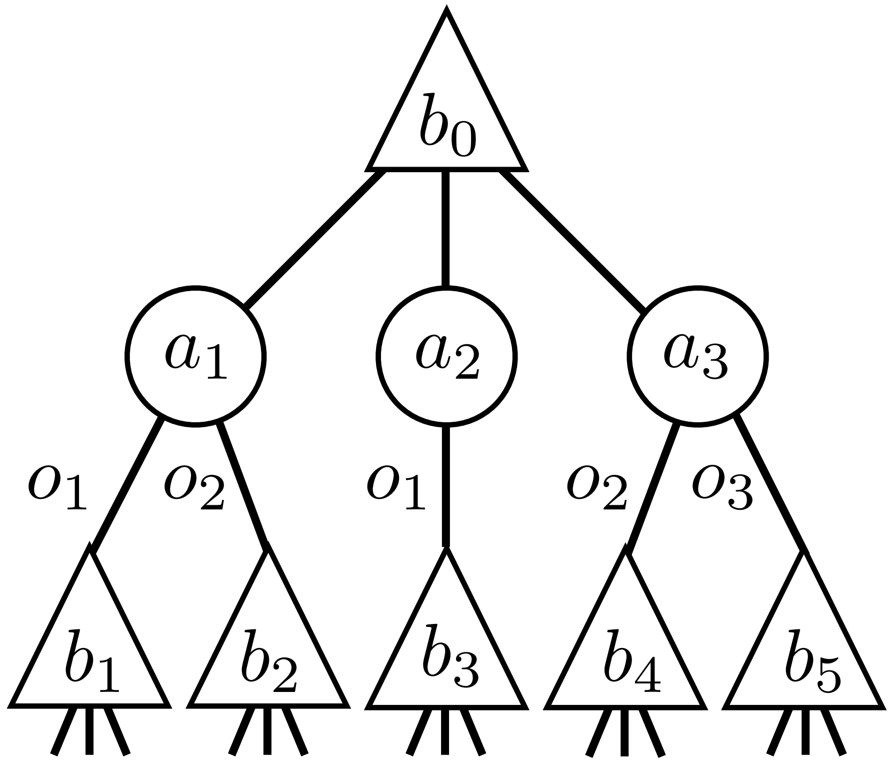

Online Planning

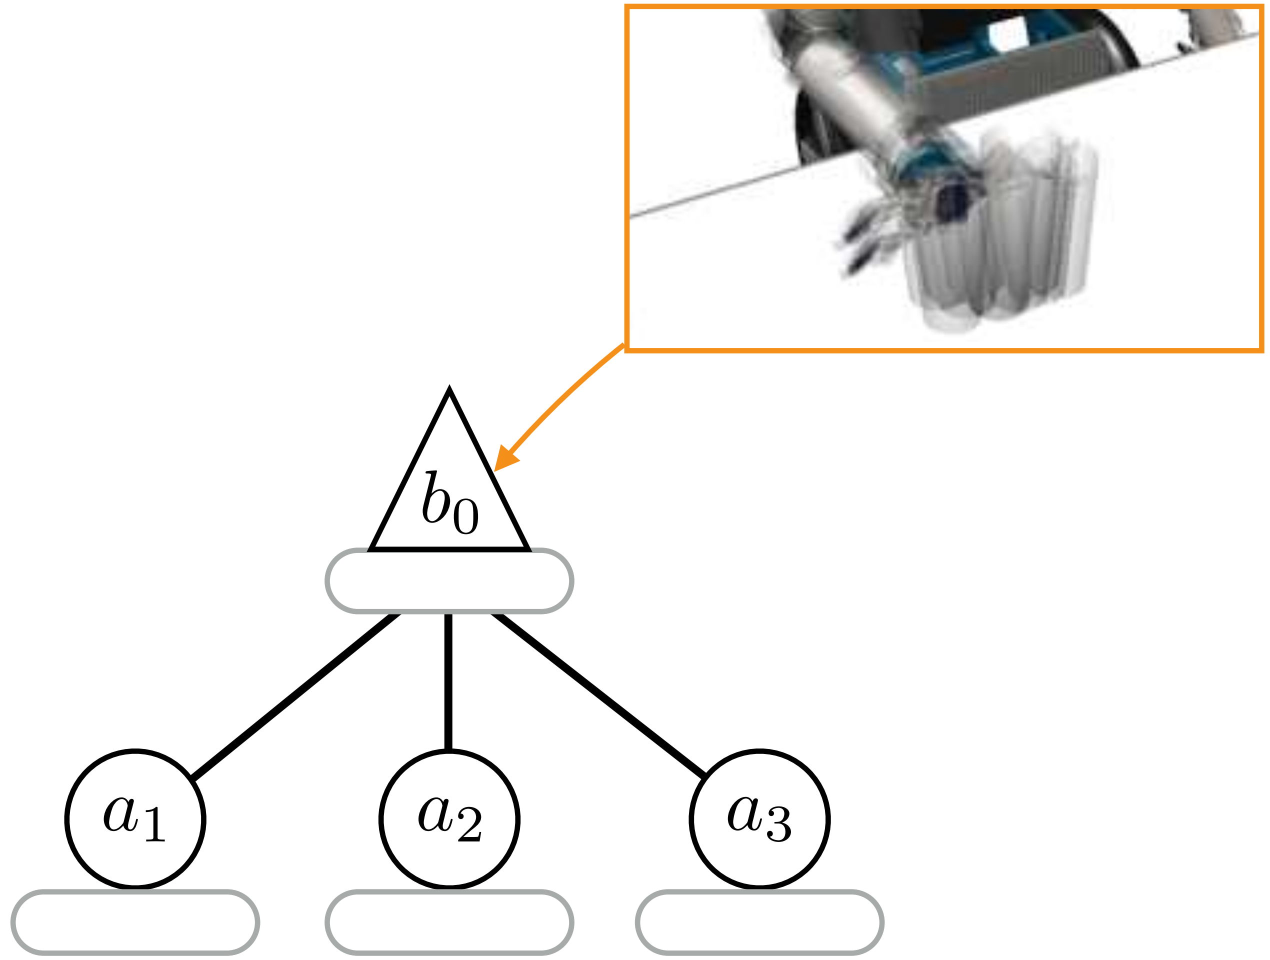

D. Silver and J. Veness. "Monte-Carlo planning in large POMDPs." NIPS, 2010.
A. Somani, N. Yi, D. Hsu, and W.S. Lee. "DESPOT: Online POMDP planning with regularization." NIPS, 2013.

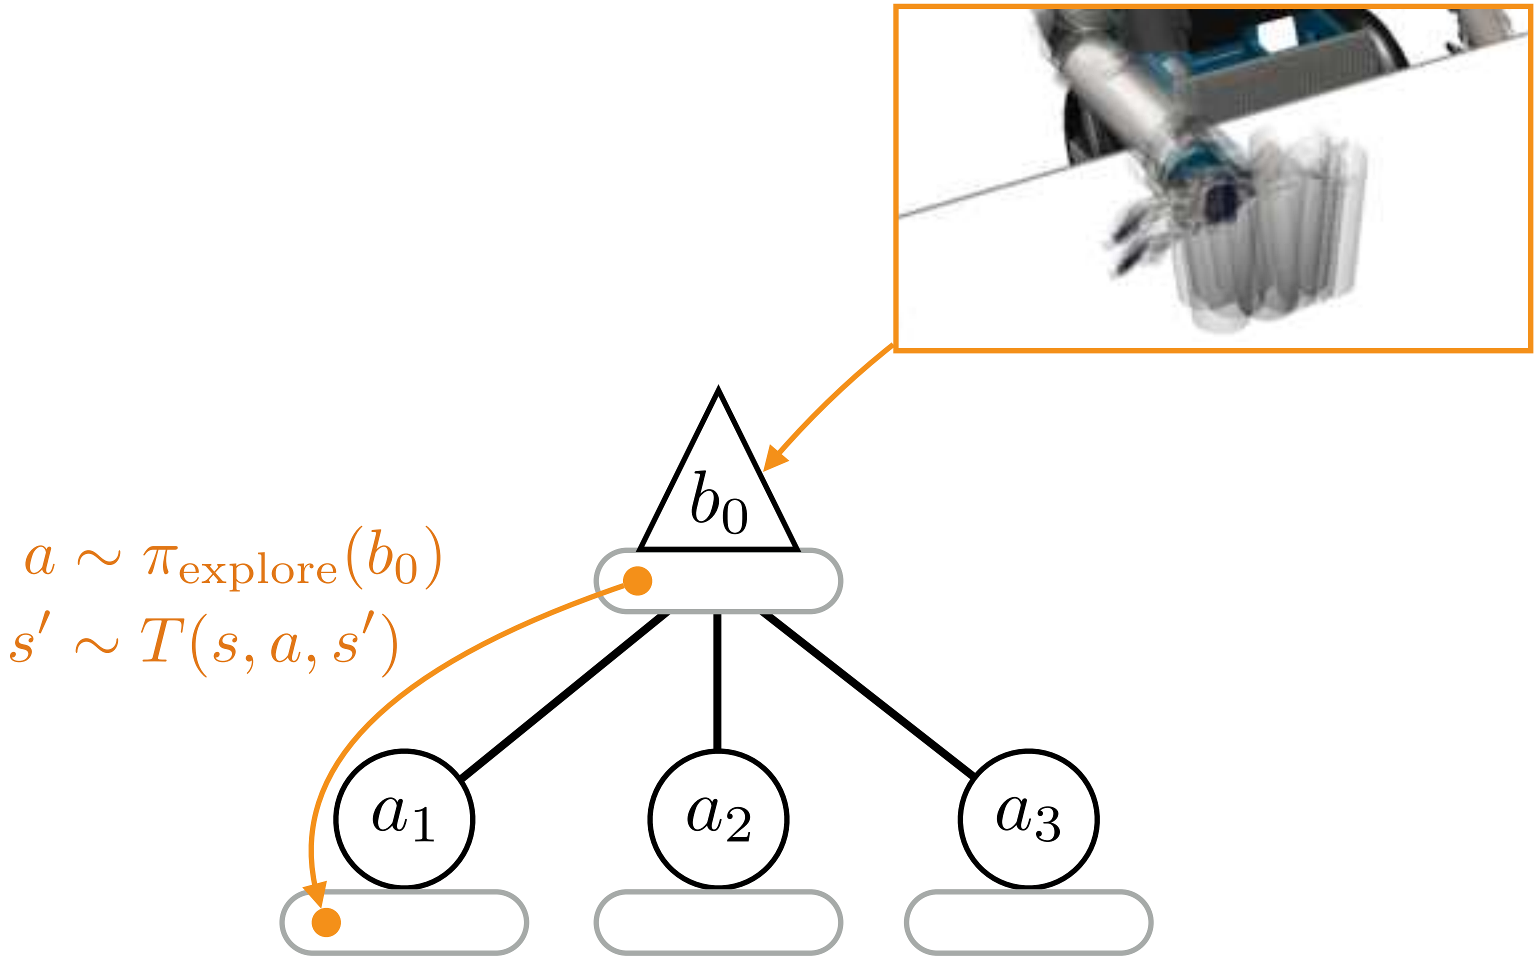

D. Silver and J. Veness. "Monte-Carlo planning in large POMDPs." NIPS, 2010.
A. Somani, N. Yi, D. Hsu, and W.S. Lee. "DESPOT: Online POMDP planning with regularization." NIPS, 2013.

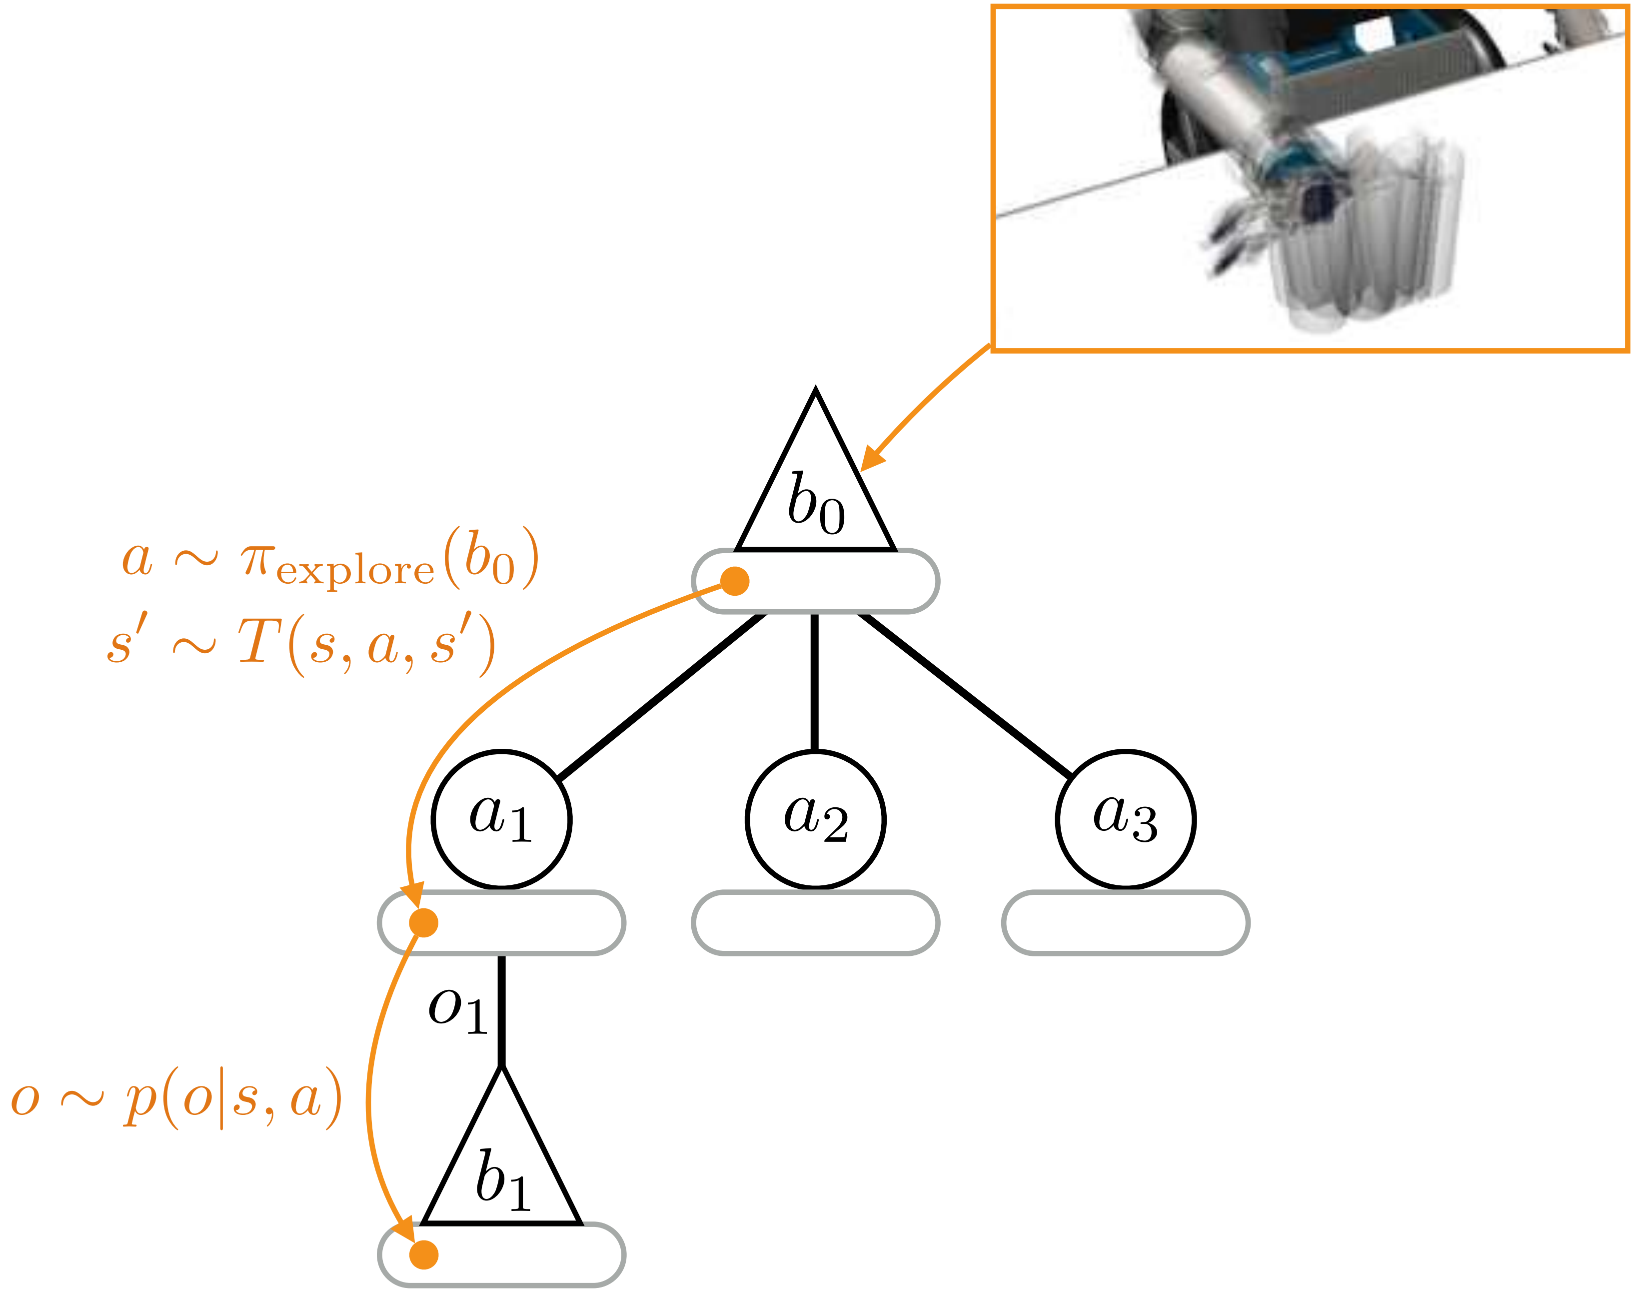

D. Silver and J. Veness. "Monte-Carlo planning in large POMDPs." NIPS, 2010.
A. Somani, N. Yi, D. Hsu, and W.S. Lee. "DESPOT: Online POMDP planning with regularization." NIPS, 2013.

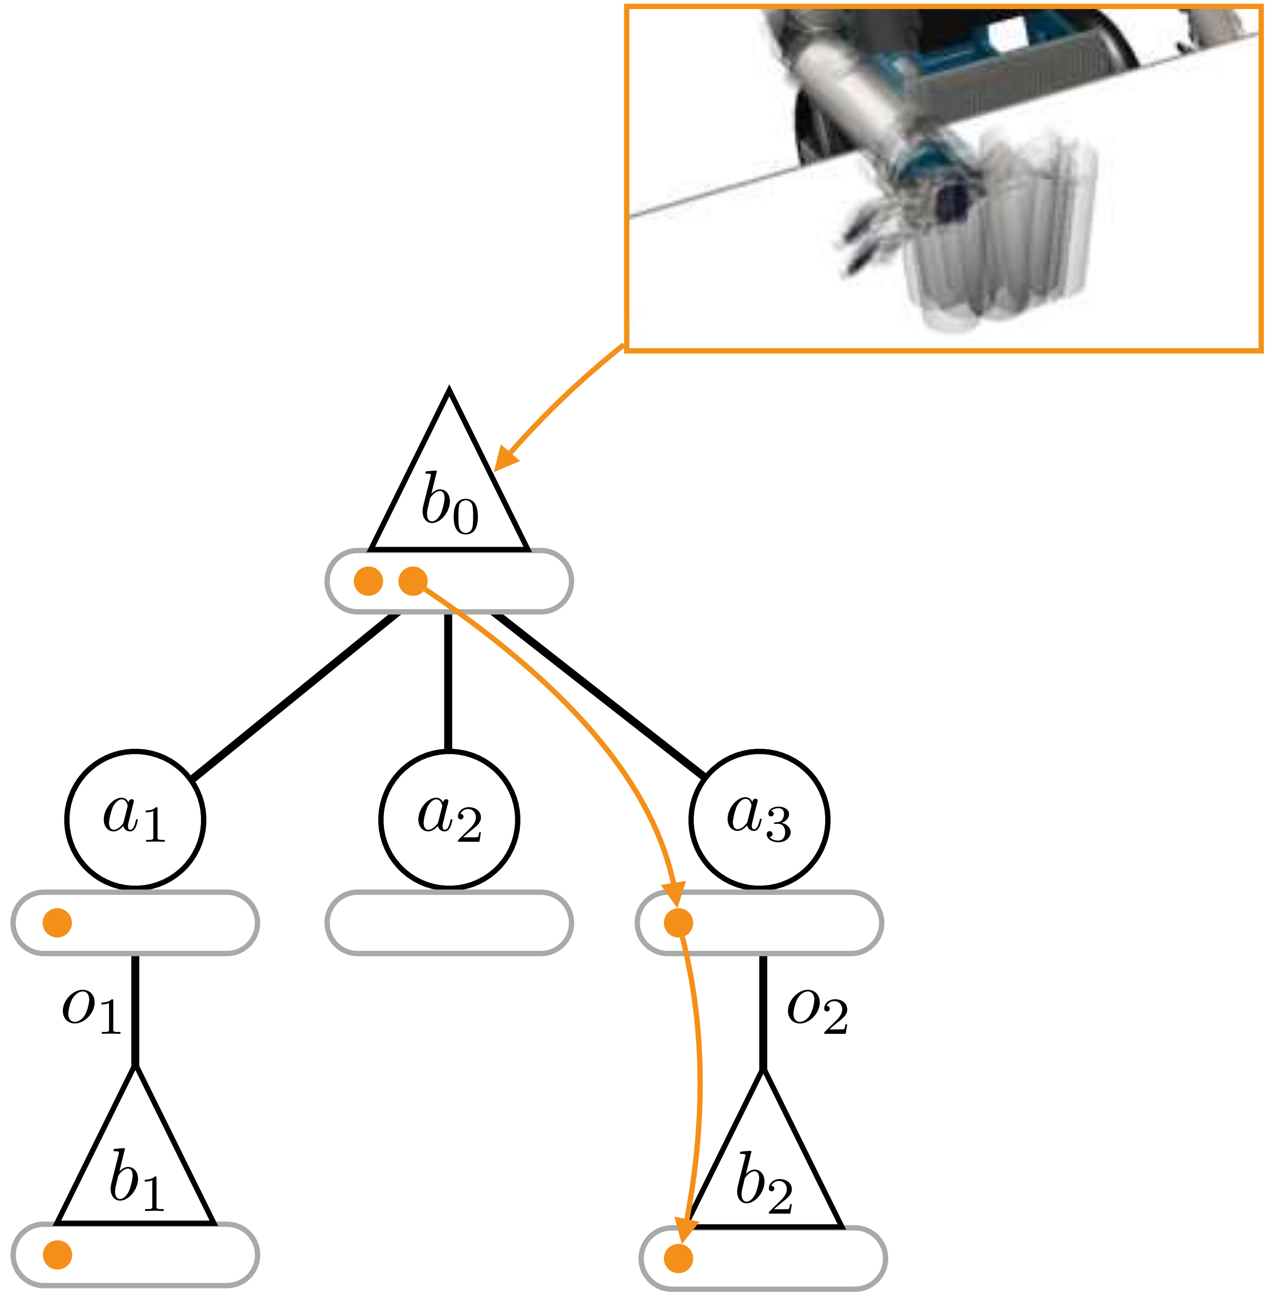

D. Silver and J. Veness. "Monte-Carlo planning in large POMDPs." NIPS, 2010.
A. Somani, N. Yi, D. Hsu, and W.S. Lee. "DESPOT: Online POMDP planning with regularization." NIPS, 2013.

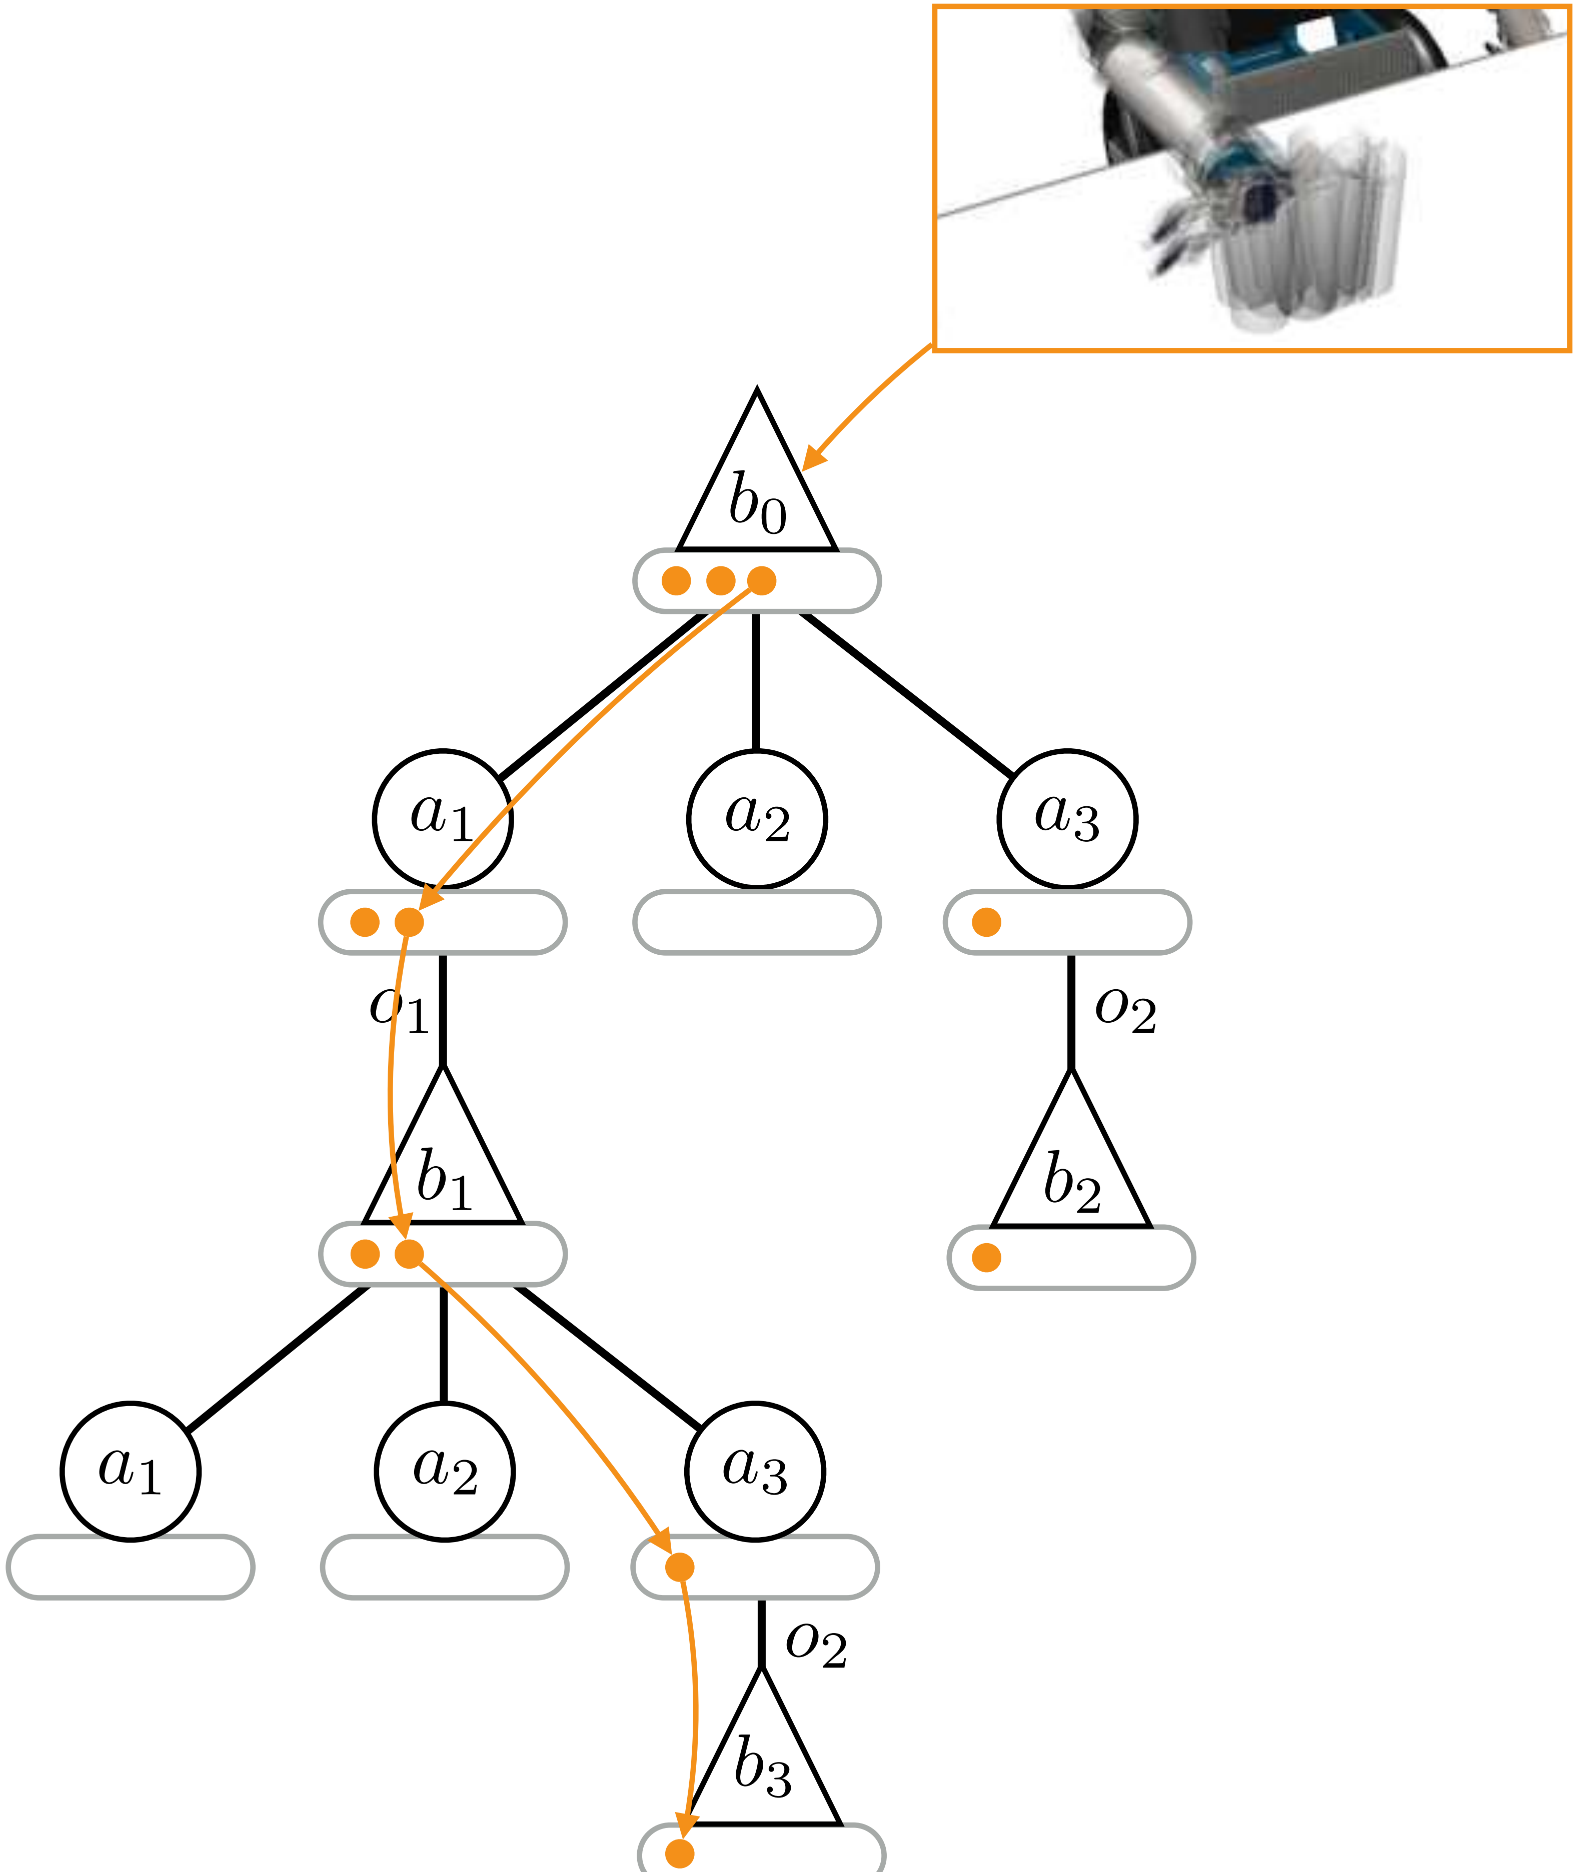
b_0
a_1
a_2
a_3
o_1
o_2
b_1
b_2
a_1
a_2
a_3
o_2
b_3

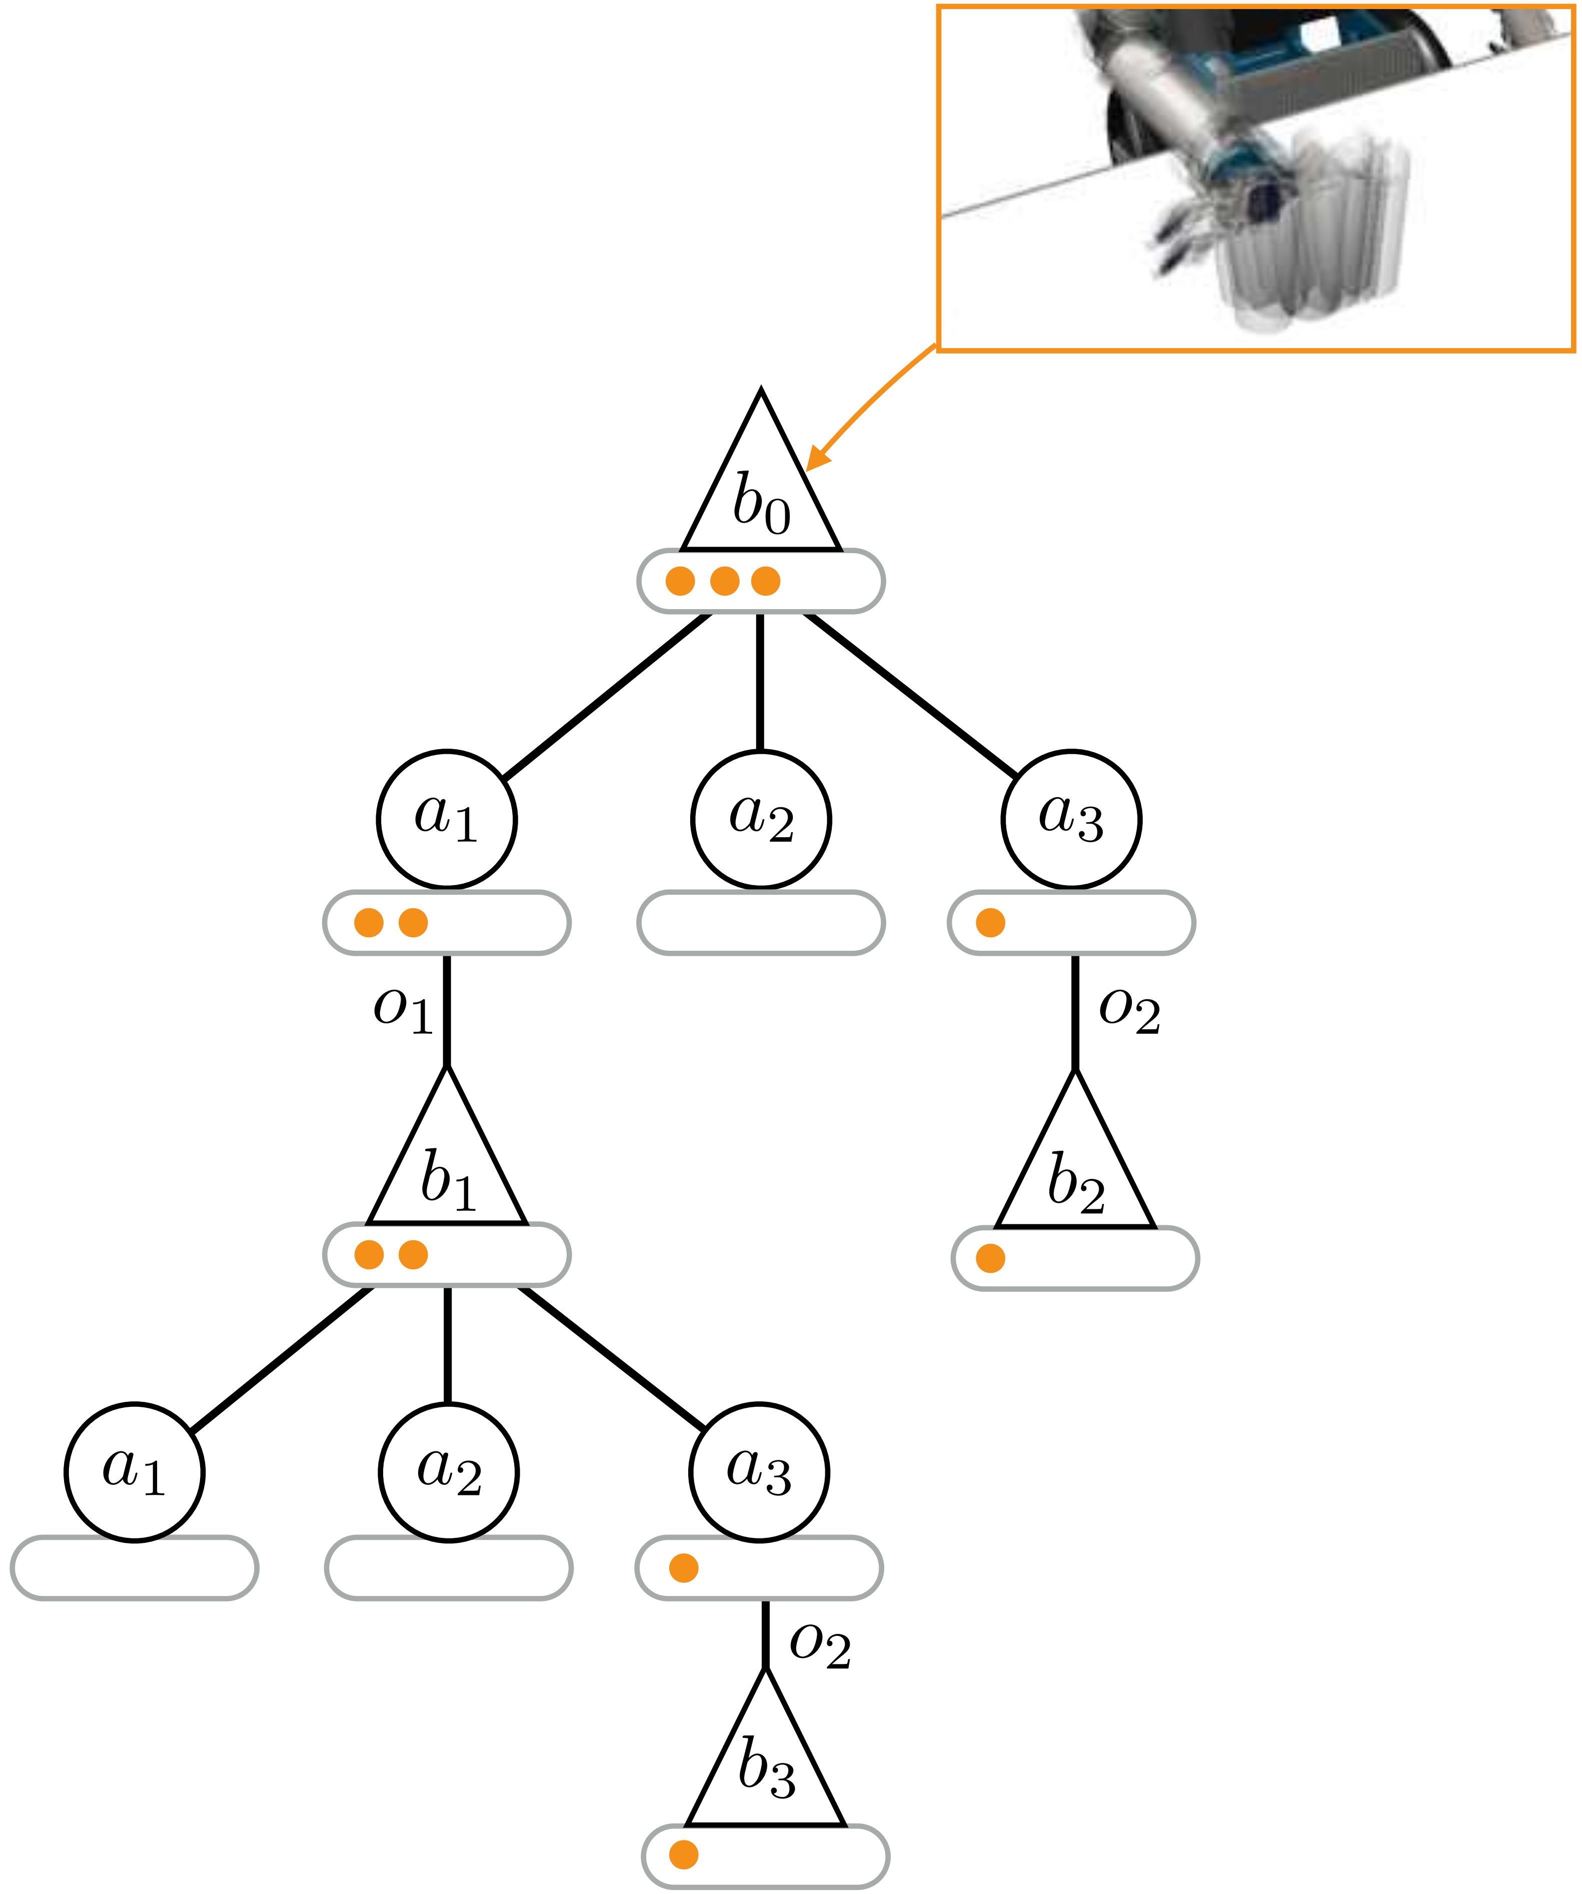
b_0
a_1
a_2
a_3
o_1
o_2
b_1
b_2
a_1
a_2
a_3
o_2
b_3

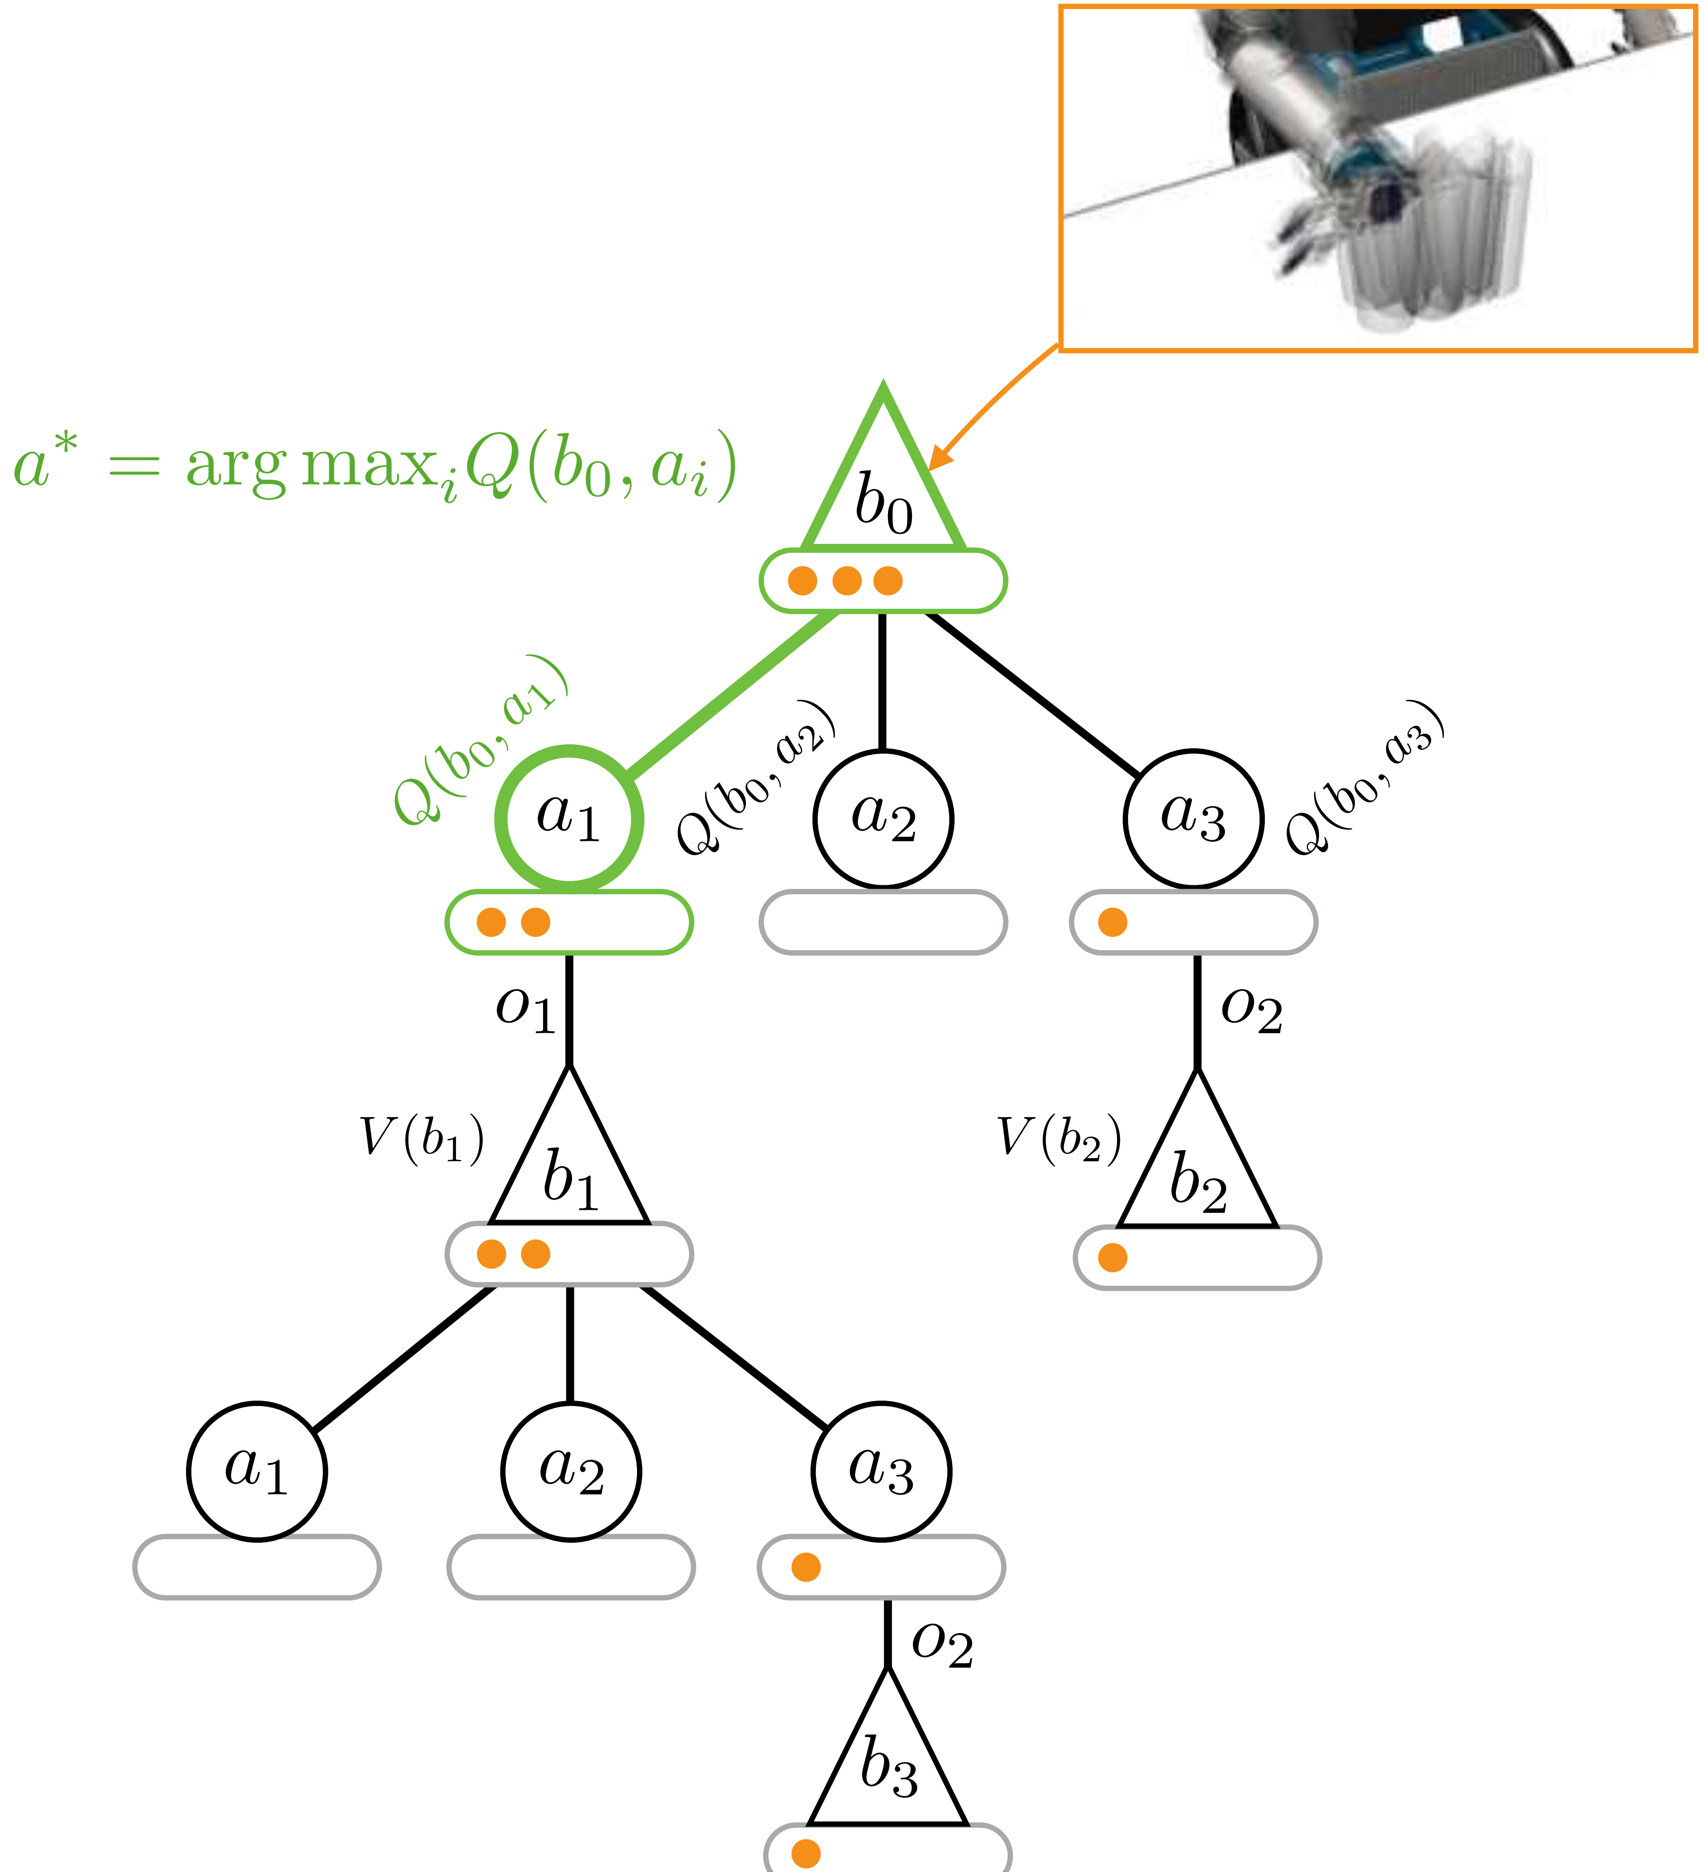

$a^* = \arg\max_i Q(b_0, a_i)$
b_0
$Q(b_0, a_1)$
$Q(b_0, a_2)$
$Q(b_0, a_3)$
a_1
a_2
a_3
o_1
o_2
$V(b_1)$
b_1
$V(b_2)$
b_2
a_1
a_2
a_3
o_2
b_3

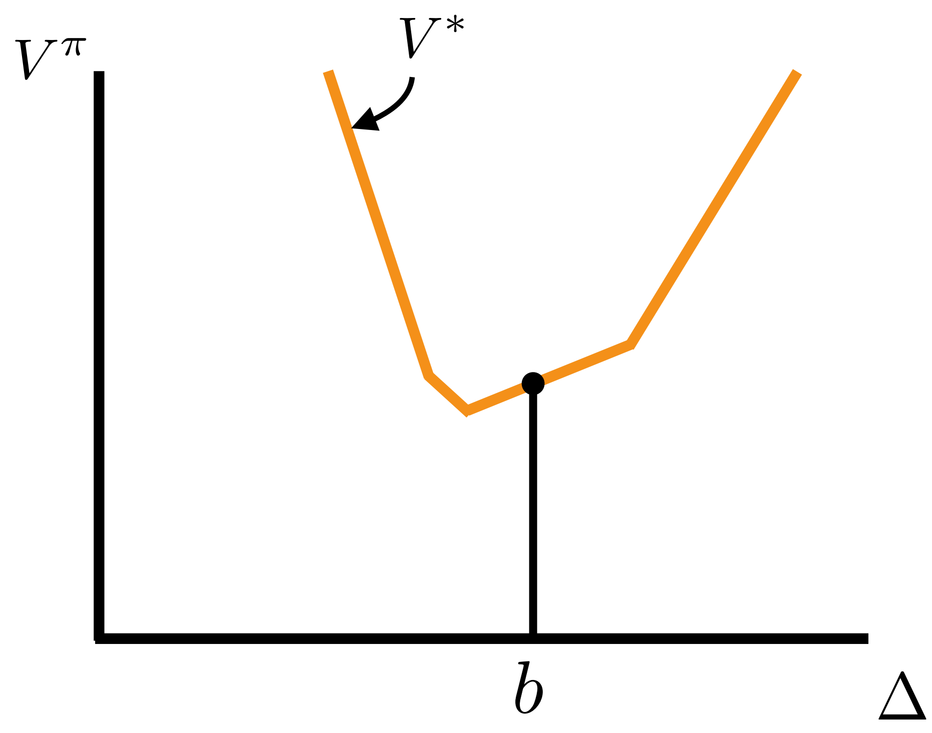

Offline Planning

Point-Based Methods

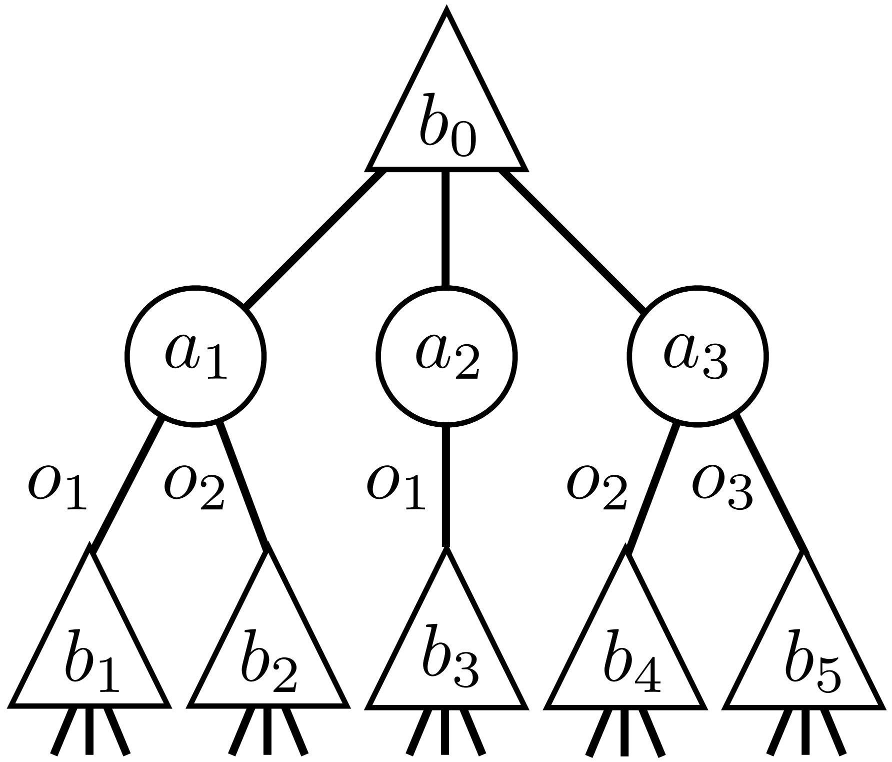

Online Planning

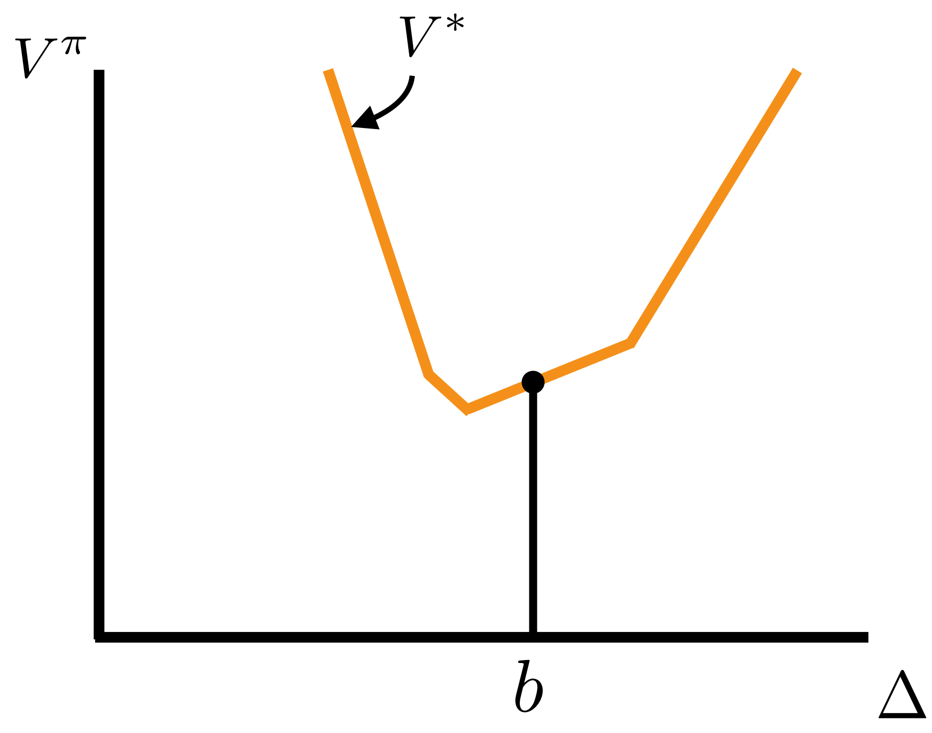

Offline Planning

Point-Based Methods

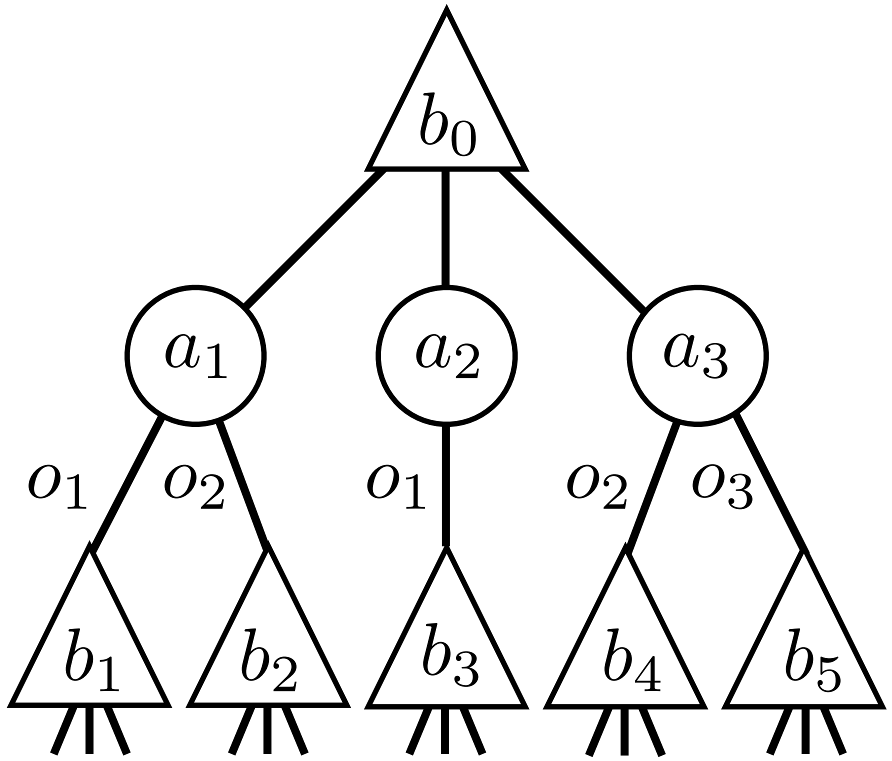

Online Planning

Heuristics / Bounds

Combine online and offline planning.

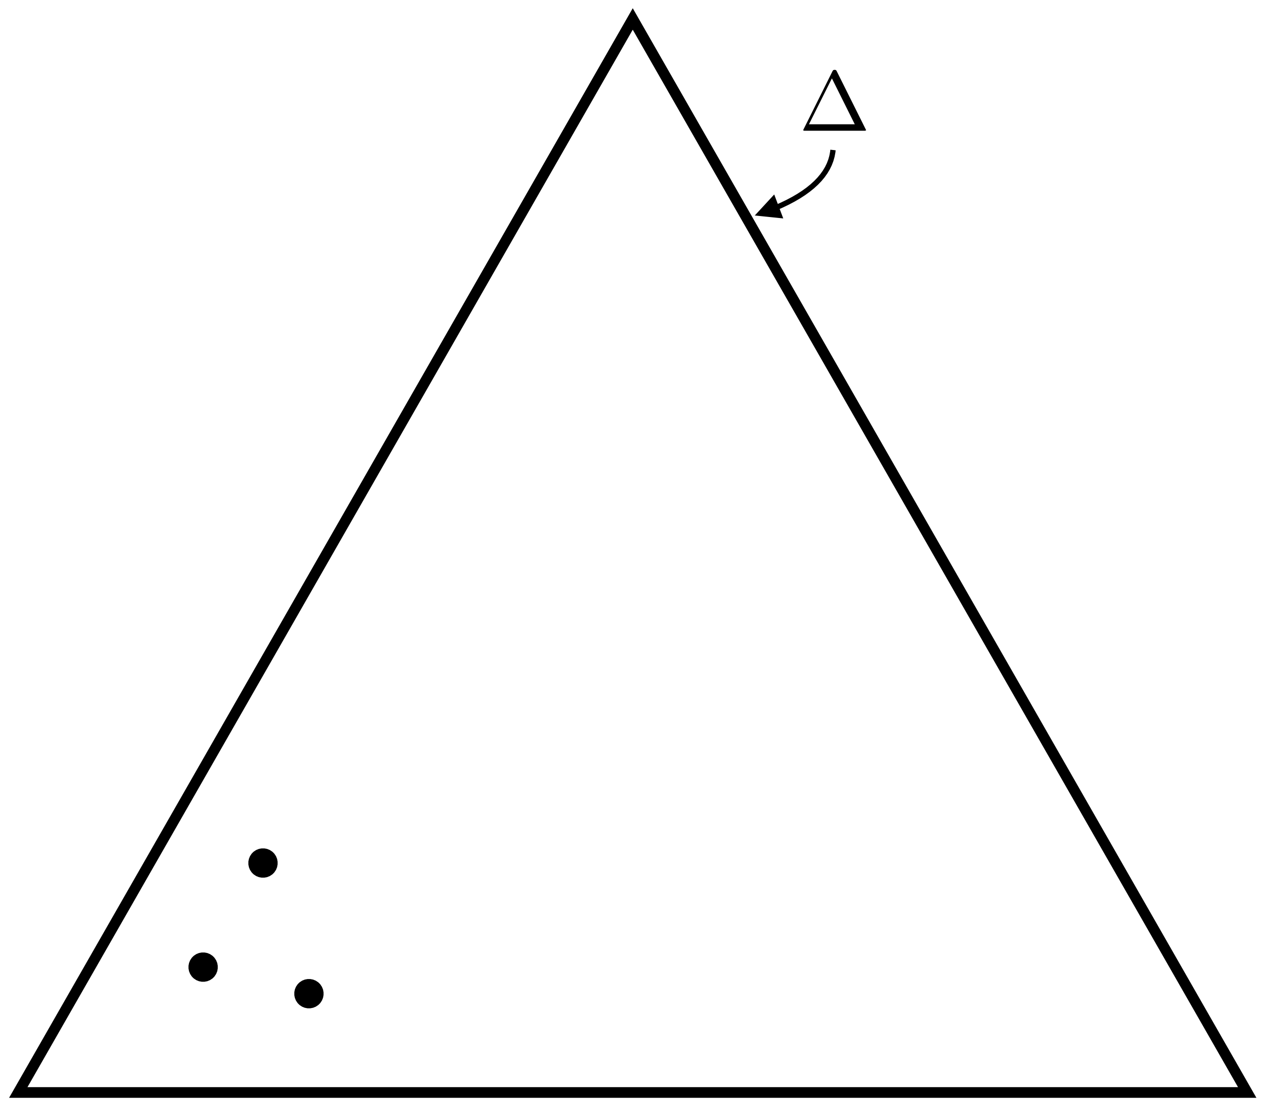

The post-contact belief space is small

M. Koval, N. Pollard, and S. Srinivasa. Pre- and post-contact policy decomposition for planar contact manipulation under uncertainty. IJRR, 2016.
M. Koval, N. Pollard, and S. Srinivasa. Pre- and post-contact policy decomposition for planar contact manipulation under uncertainty. RSS, 2014.

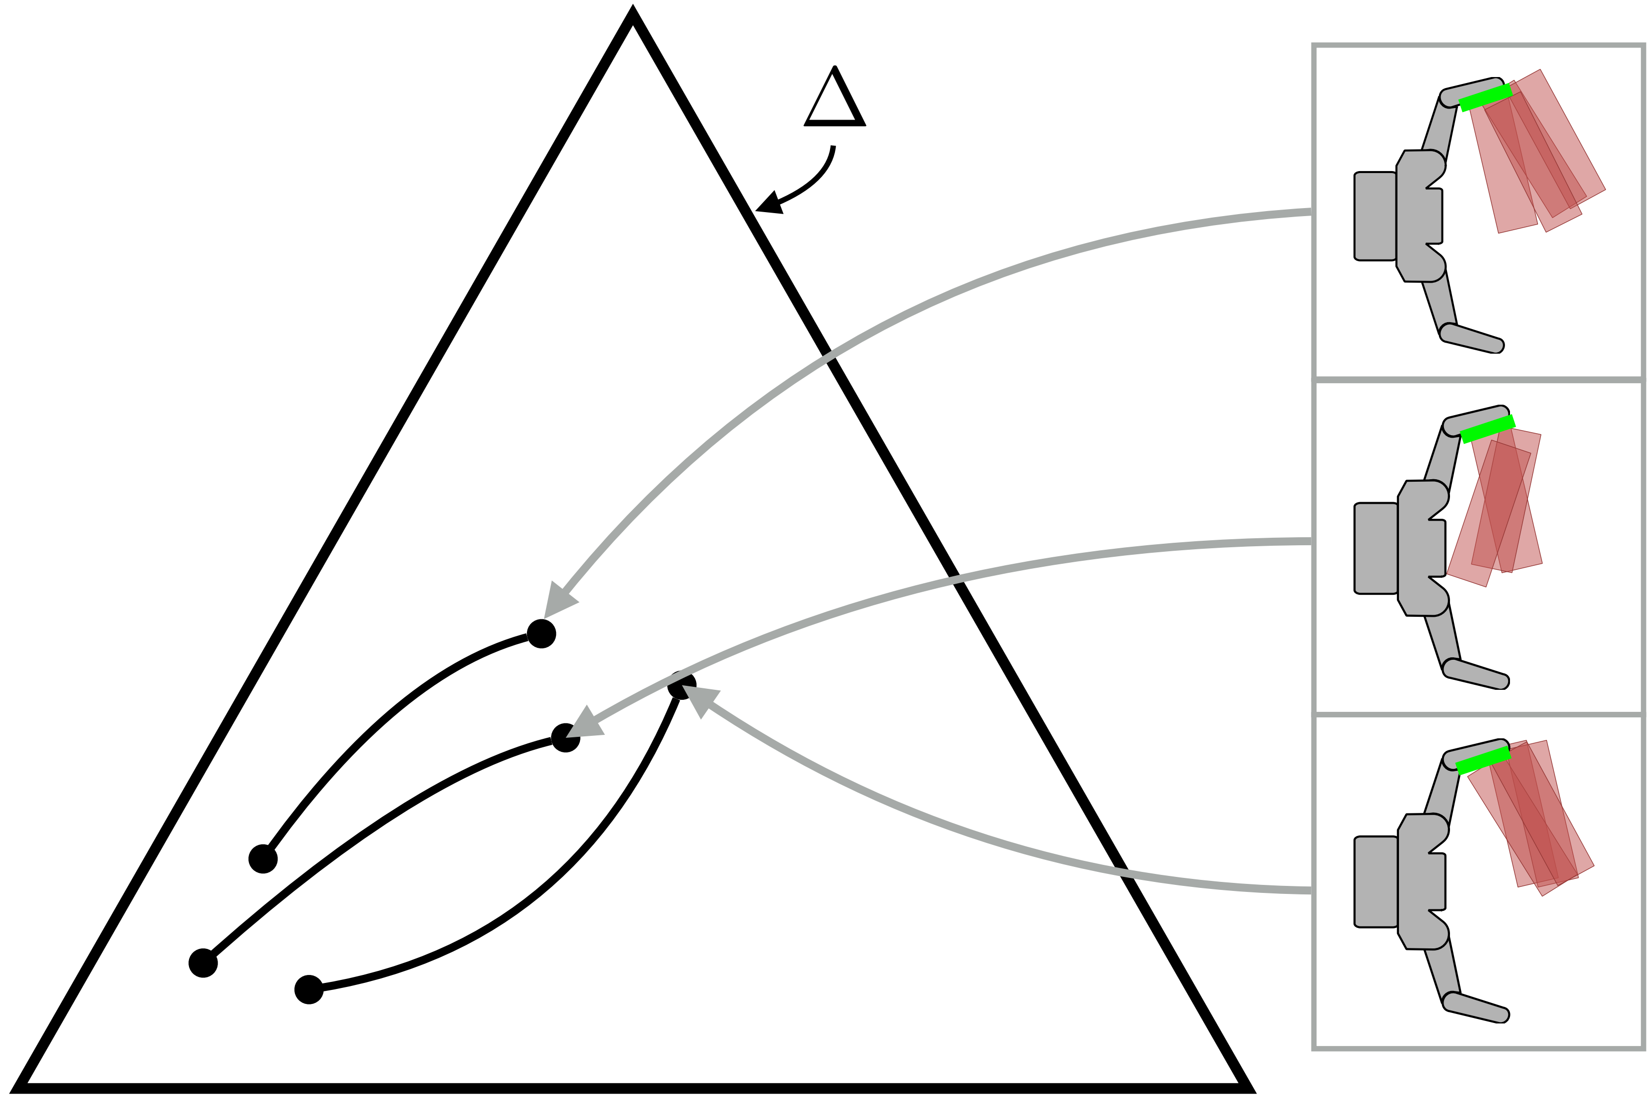

The post-contact belief space is small

M. Koval, N. Pollard, and S. Srinivasa. Pre- and post-contact policy decomposition for planar contact manipulation under uncertainty. IJRR, 2016.
M. Koval, N. Pollard, and S. Srinivasa. Pre- and post-contact policy decomposition for planar contact manipulation under uncertainty. RSS, 2014.

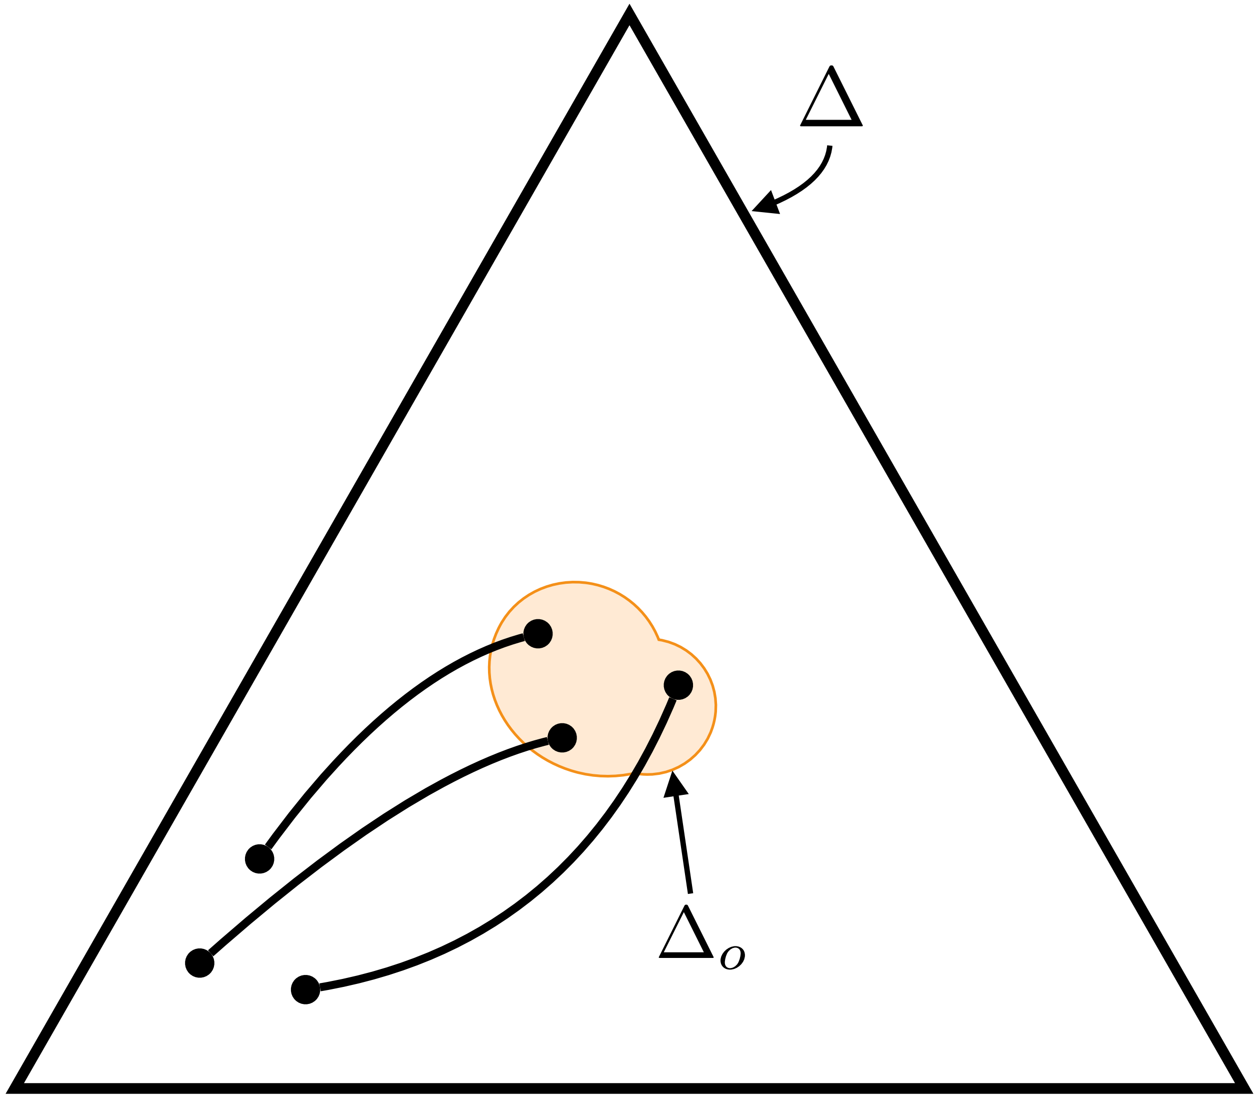

The post-contact belief space is small

M. Koval, N. Pollard, and S. Srinivasa. Pre- and post-contact policy decomposition for planar contact manipulation under uncertainty. IJRR, 2016.
M. Koval, N. Pollard, and S. Srinivasa. Pre- and post-contact policy decomposition for planar contact manipulation under uncertainty. RSS, 2014.

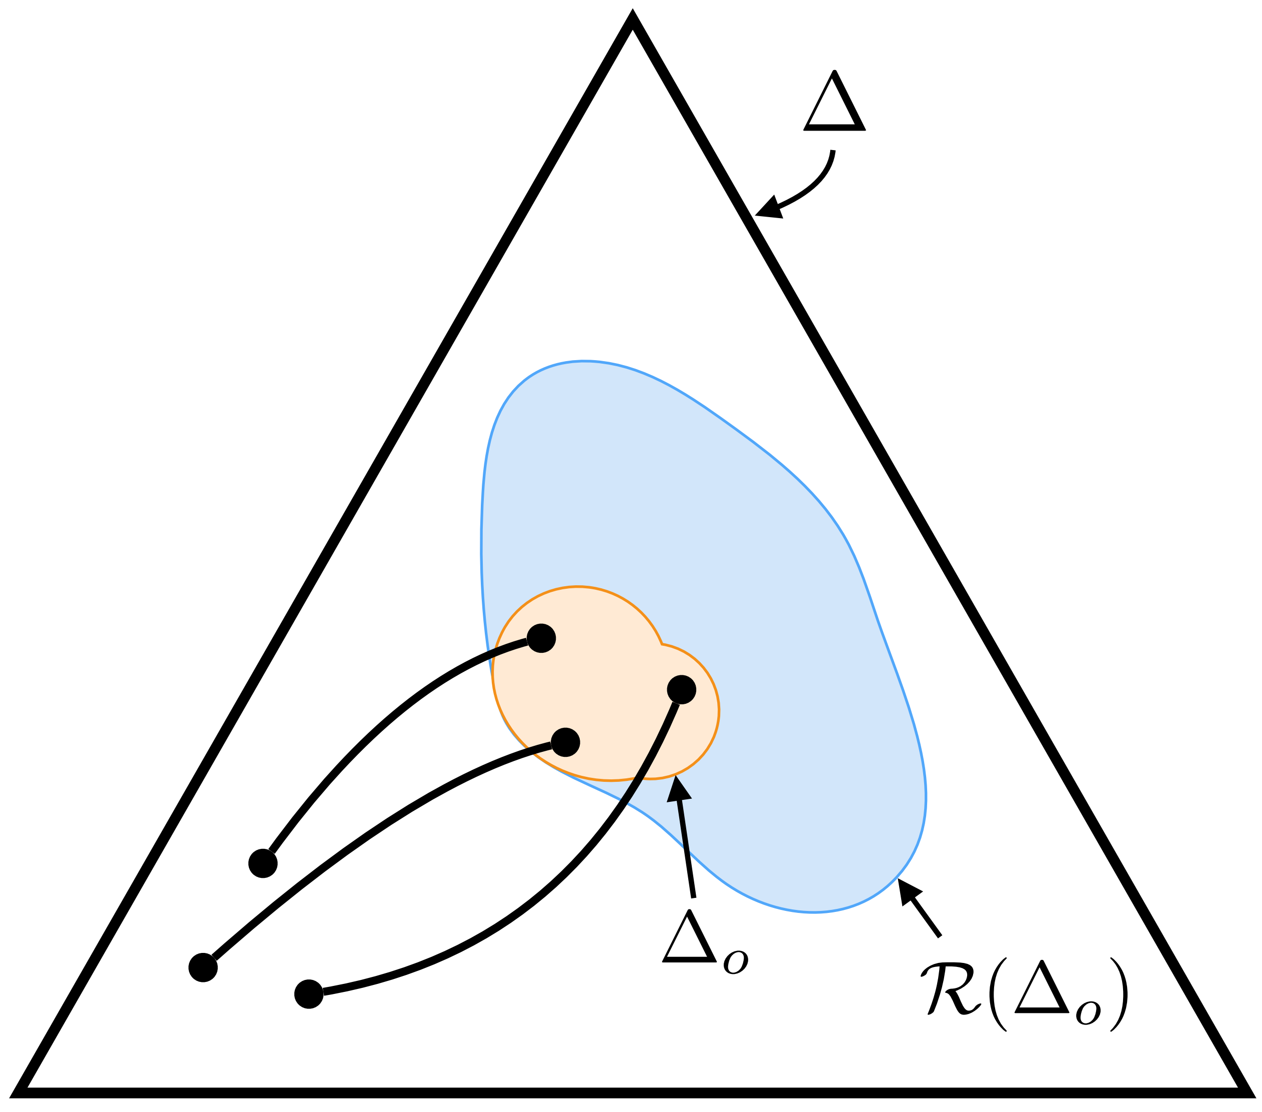

The post-contact belief space is small

M. Koval, N. Pollard, and S. Srinivasa. Pre- and post-contact policy decomposition for planar contact manipulation under uncertainty. IJRR, 2016.
M. Koval, N. Pollard, and S. Srinivasa. Pre- and post-contact policy decomposition for planar contact manipulation under uncertainty. RSS, 2014.

pre-contact policy
computed online
move-until-touch
once per problem

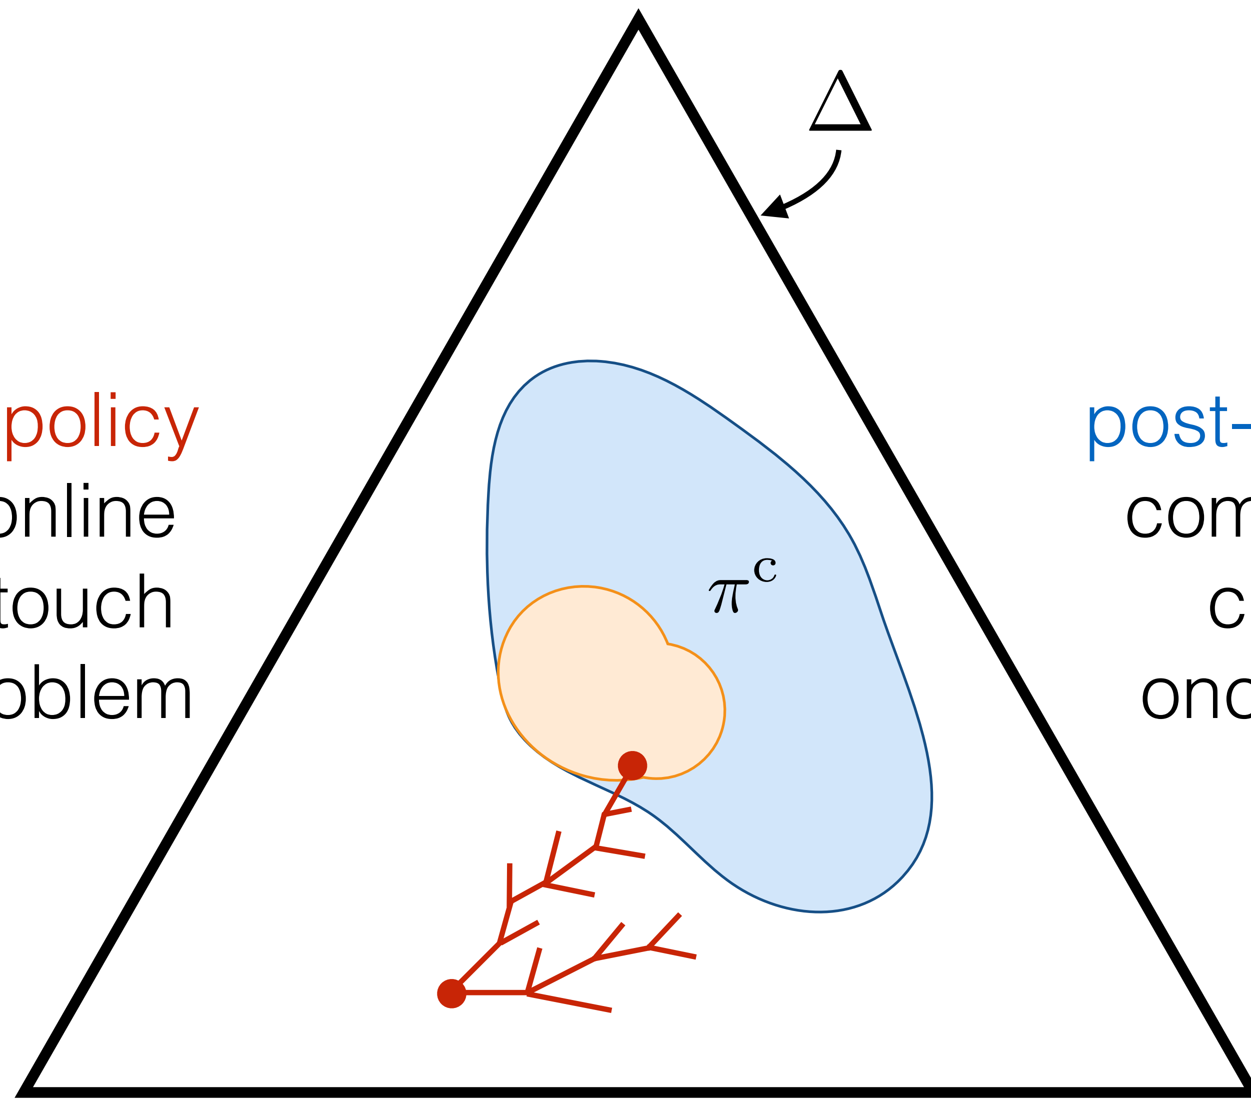

post-contact policy
computed offline
closed-loop
once per object

Decompose into pre- and post-contact policies

M. Koval, N. Pollard, and S. Srinivasa. Pre- and post-contact policy decomposition for planar contact manipulation under uncertainty. IJRR, 2016.
M. Koval, N. Pollard, and S. Srinivasa. Pre- and post-contact policy decomposition for planar contact manipulation under uncertainty. RSS, 2014.

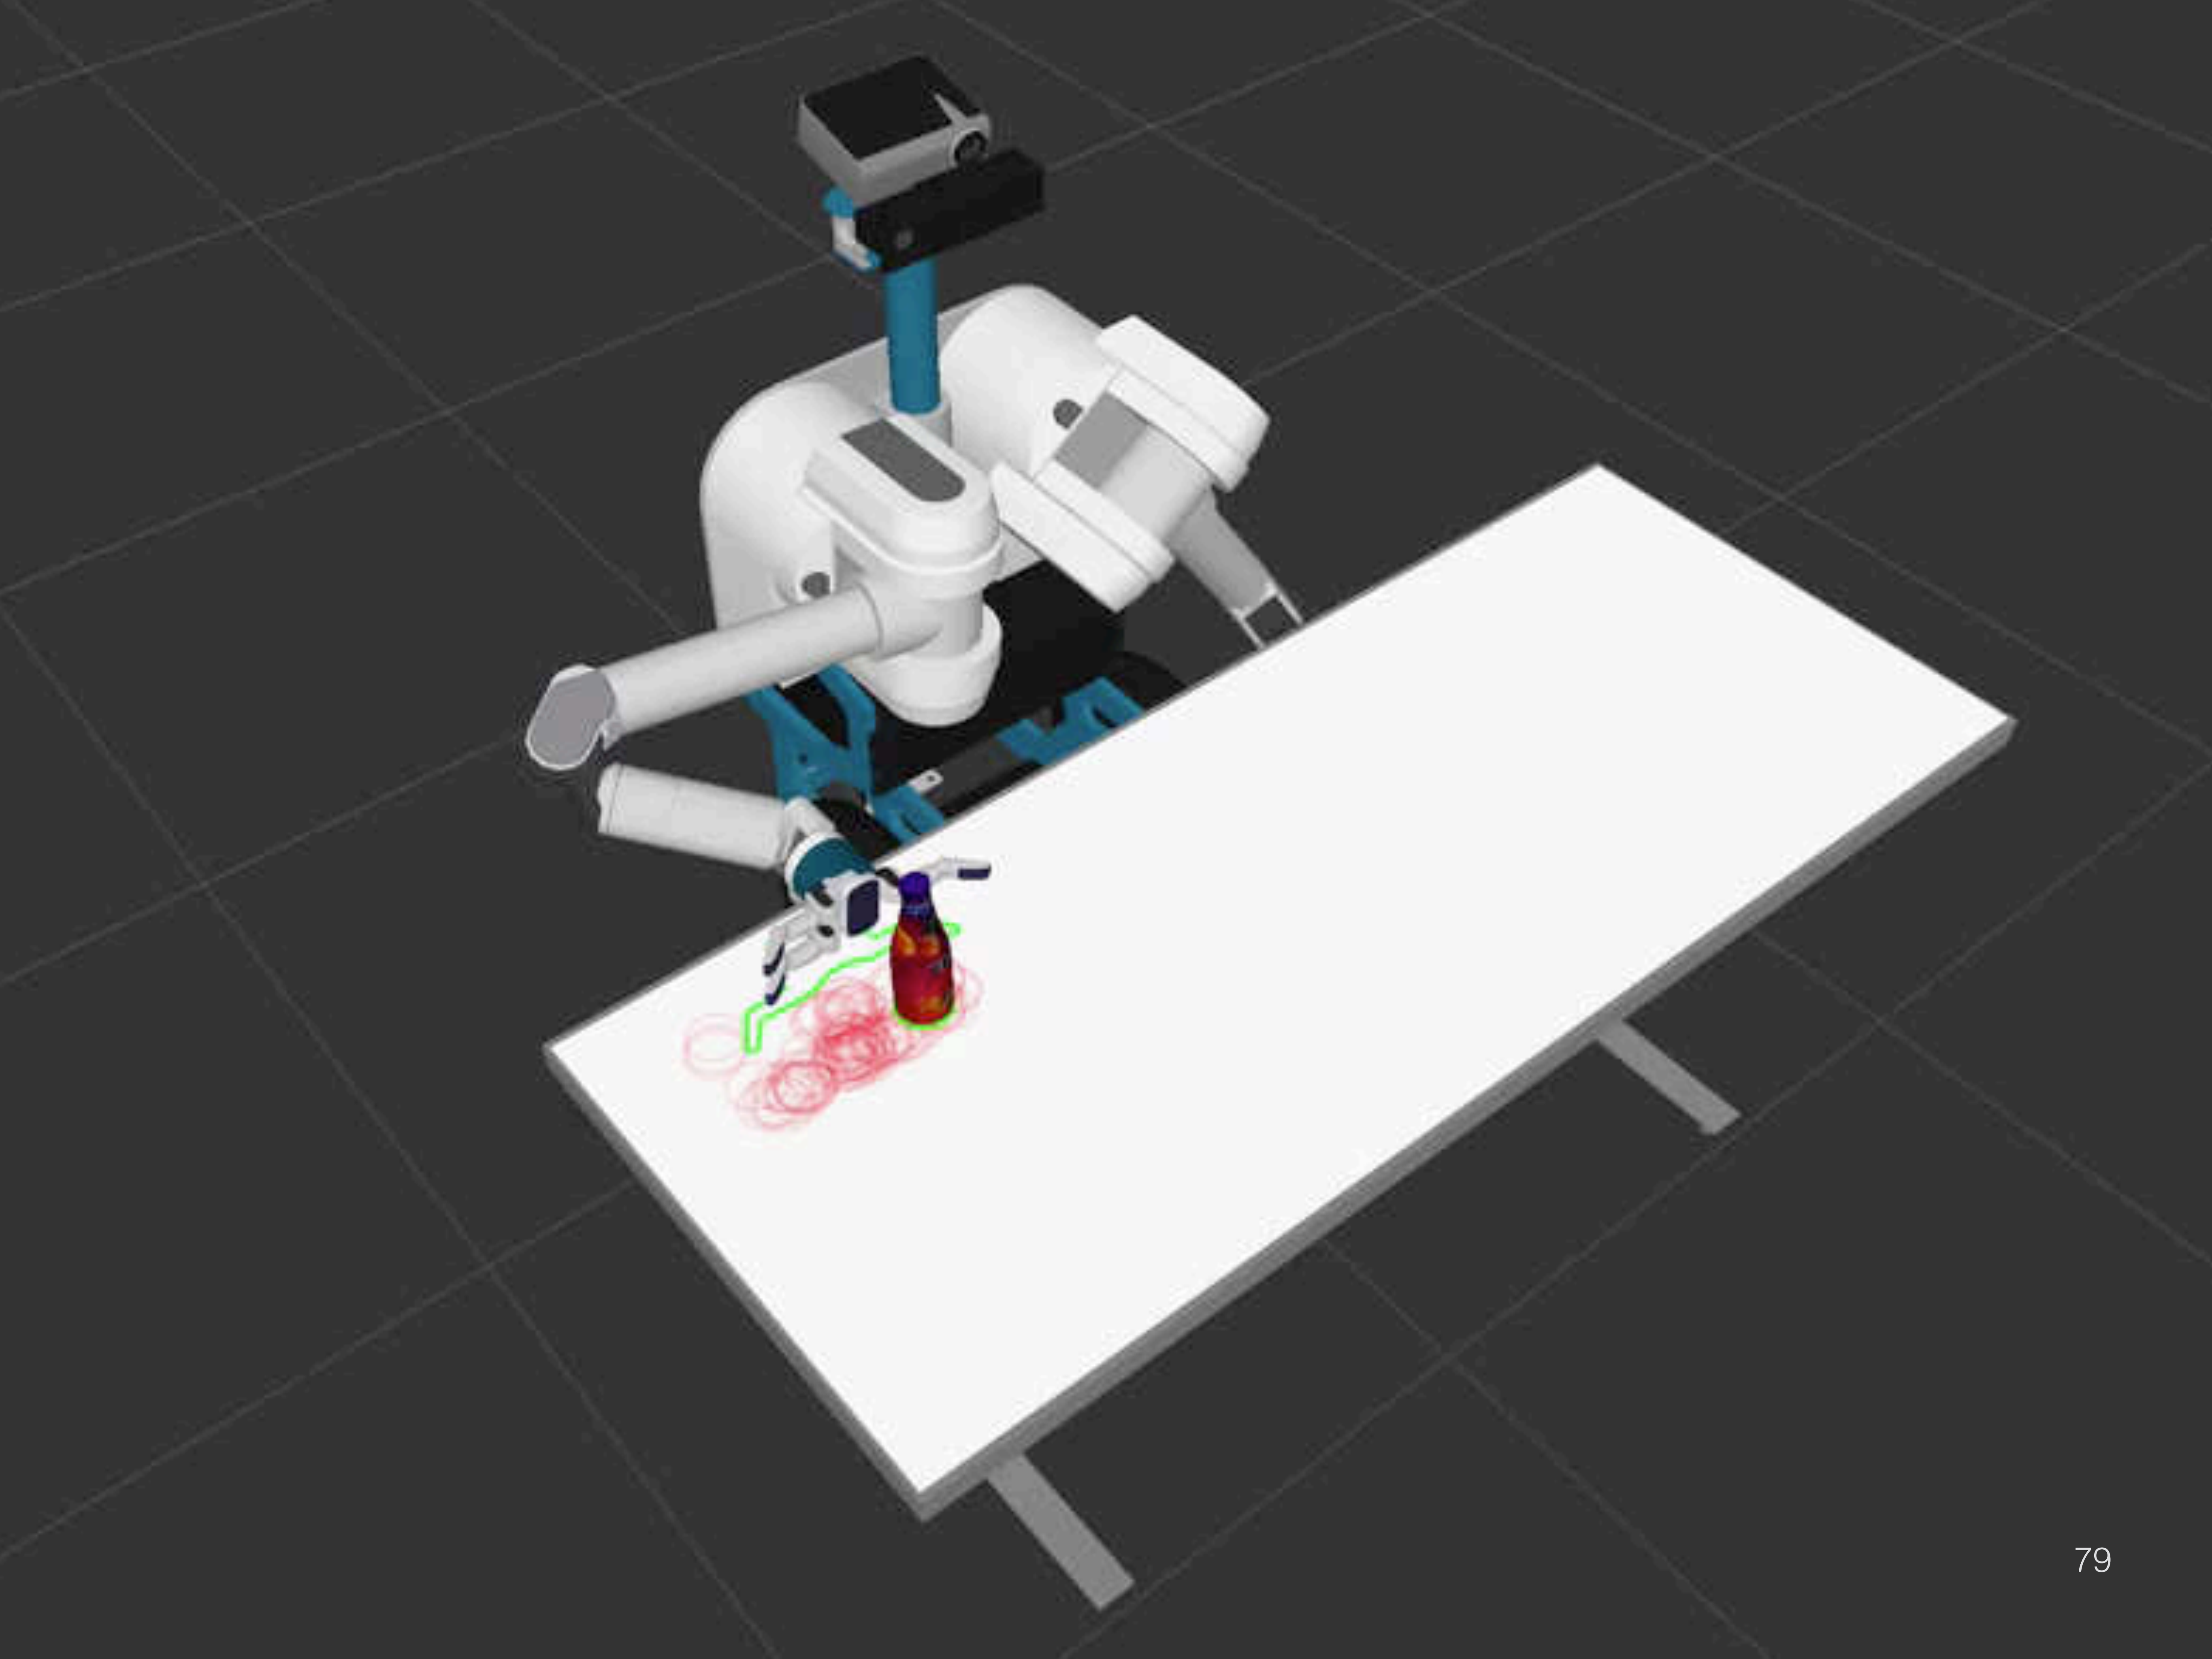

Belief 3

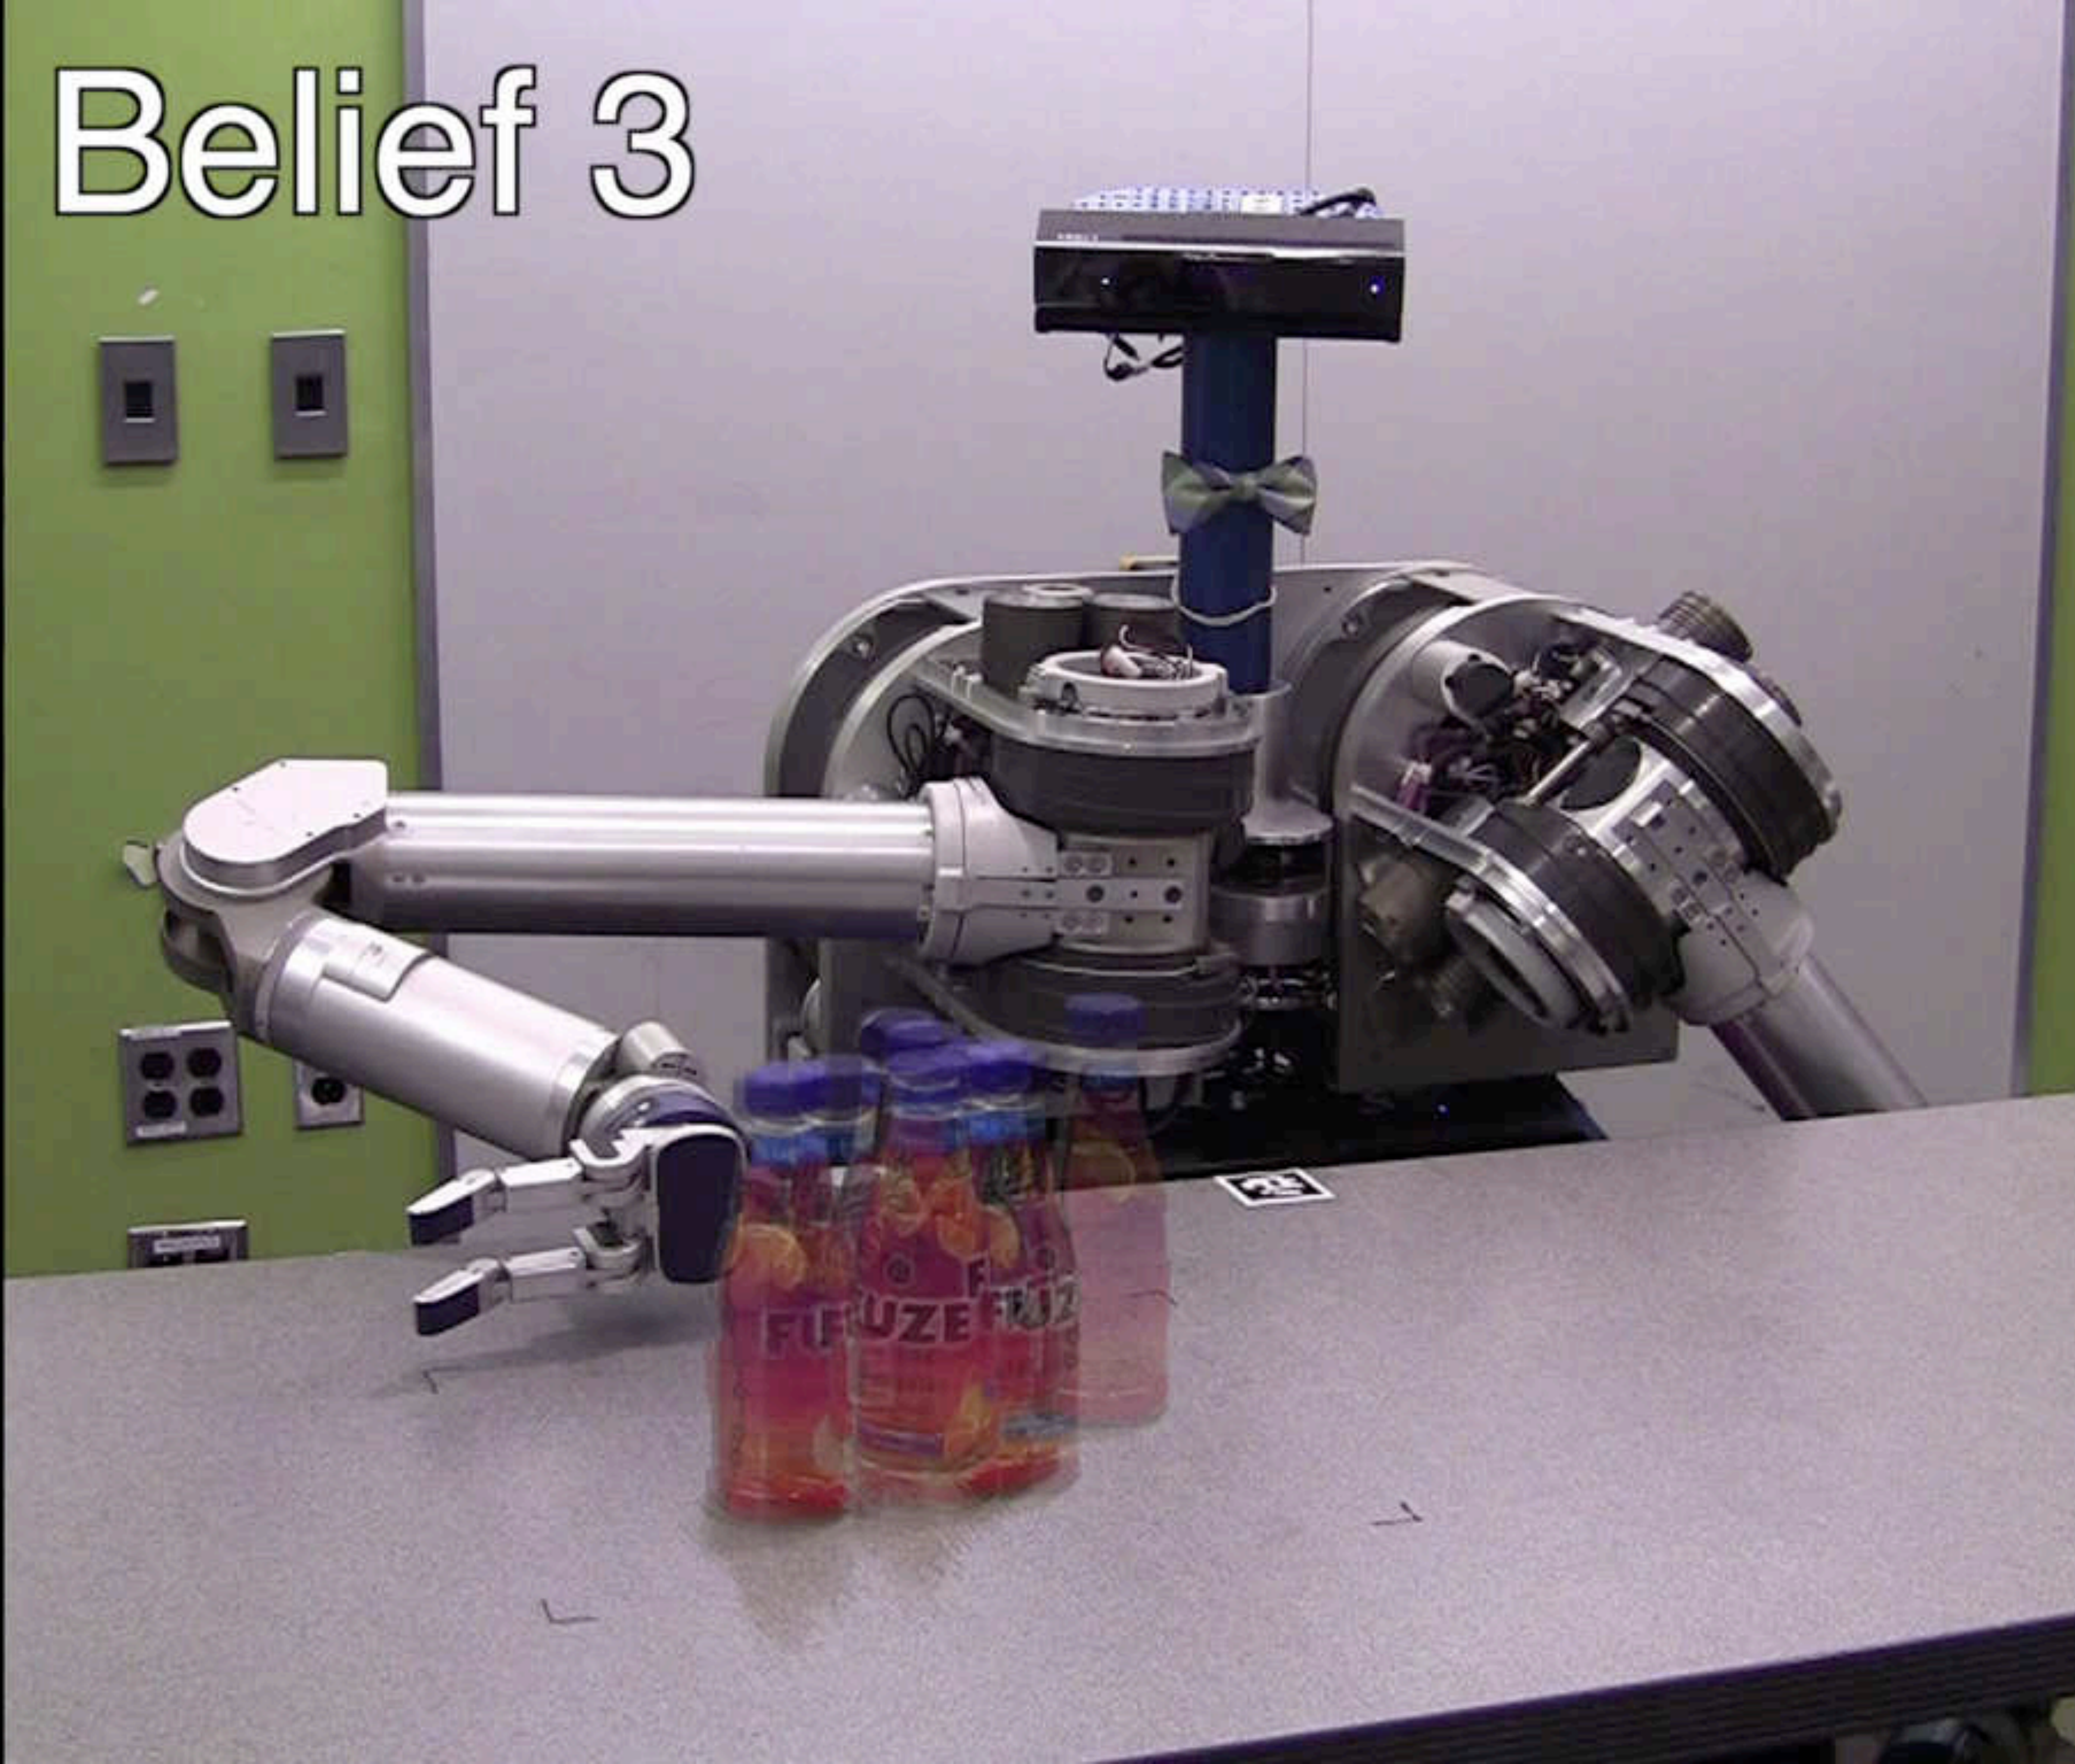

Physics-Based Manipulation under Uncertainty

Michael Koval
mkoval@cs.cmu.edu

February 9, 2016
Hands: Design and Control for Dexterous Manipulation