Submodular Surrogates for Value of Information

Yuxin Chen, Shervin Javdani, Amin Karbasi, J. Andrew Bagnell, Siddhartha Srinivasa, Andreas Krause

ETH zürich
Carnegie Mellon University
Yale University

Motivating Application

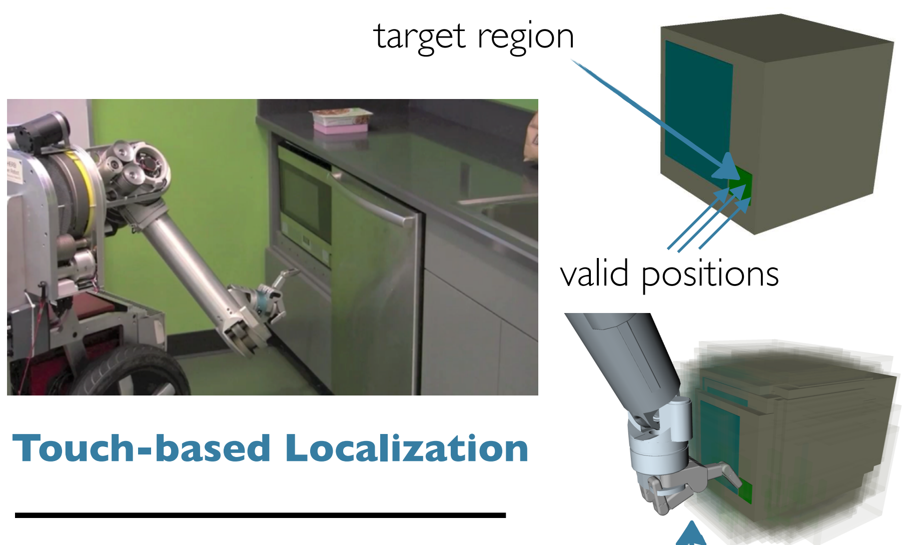

Touch-based Localization

Movie Recommendation

Comparison-based Search

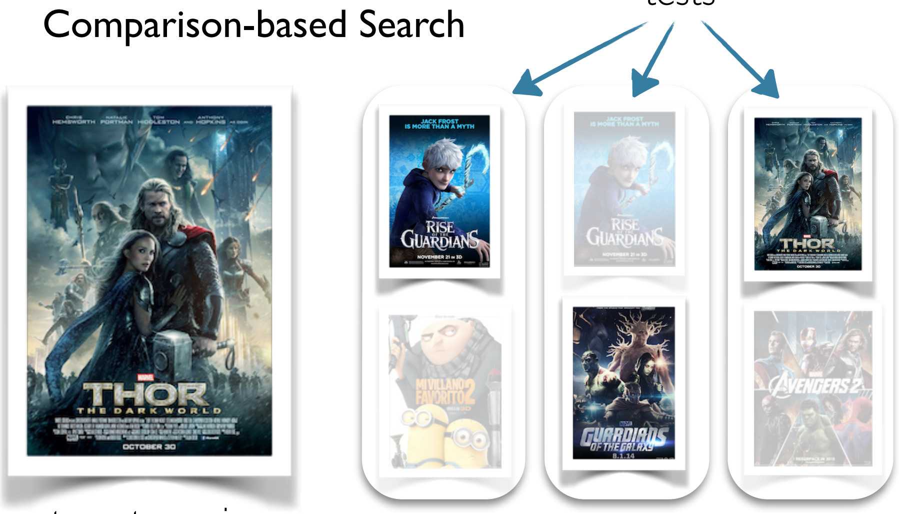

target movie

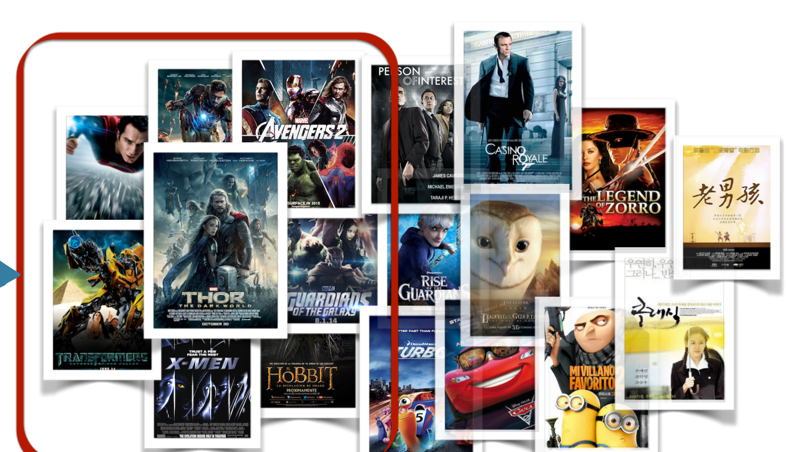

Movies of the same Genre

The Nonmyopic VoI Problem

Decisions: $D = \{d_1, d_2, \dots, d_m\}$, Unknown hidden RV: Y

Utility of making a decision $d \in D$ for y is $u(y, d)$

Test $t \in T = \{1, \dots, n\}$ are correlated with Y; has cost $c(t)$

After performing a set of tests A, and observe outcomes $\mathbf{x}_A$, we define the ***Value of Information*** as:

$$\mathrm{VoI}(\mathbf{x}_{\mathcal{A}}) = \max_{d \in \mathcal{D}} \mathbb{E}_y[u(y,d) \mid \mathbf{x}_{\mathcal{A}}]$$

The ***regret*** of a decision:

$$R(d \mid \mathbf{x}_{\mathcal{A}}) = \max_{\mathbf{x}_{\mathcal{T}}: \mathbb{P}[\mathbf{x}_{\mathcal{T}} \mid \mathbf{x}_{\mathcal{A}}] > 0} \underbrace{[\mathrm{VoI}(\mathbf{x}_{\mathcal{T}})}_{} - \underbrace{\mathbb{E}_y[u(y,d) \mid \mathbf{x}_{\mathcal{A}}]]}_{}$$

value of the most informed decision after observing all test outcomes — expected value of d after observing $\mathbf{x}_A$

We seek a min-cost policy π^*, which suffers regret of at most ε (comparing with hindsight optimal):

$$\pi^* \in \arg\min_{\pi} \mathrm{cost}(\pi), \text{ s.t.}$$
$$\forall \mathbf{x}_{\mathcal{T}}\ \exists d: R(d \mid \mathcal{S}(\pi, \mathbf{x}_{\mathcal{T}})) \le \varepsilon \text{ whenever } \mathbb{P}[\mathbf{x}_{\mathcal{T}}] > 0.$$

An Alternative View

Hypothesis h: a test outcome vector $\mathbf{x}_T$

Decision Region $\mathcal{R}_d$: the set of hypotheses h, for which decision d is an ε-optimal action:

We seek a min-cost policy π^*, such that once terminated, all remaining hypotheses are within one region.

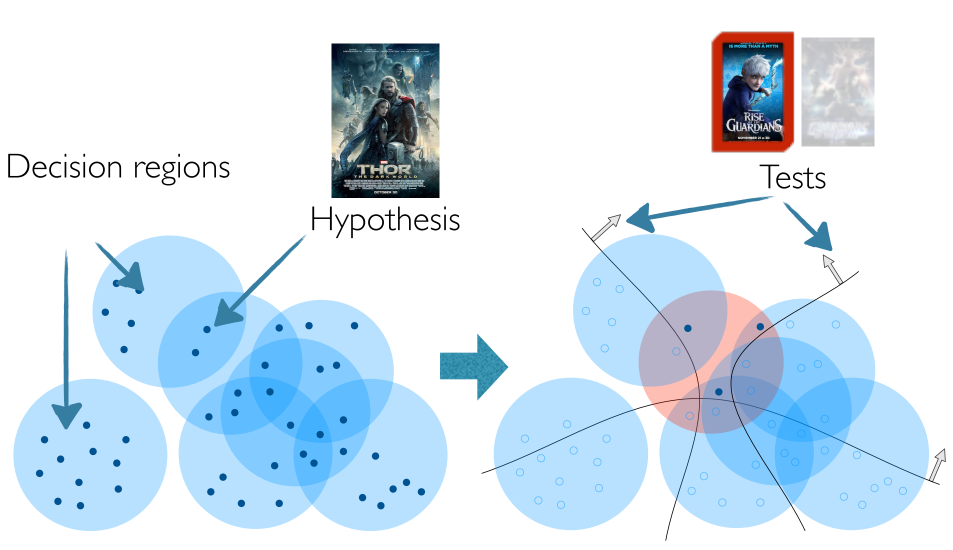

The Decision Region Determination (DRD) Problem

Special Cases

The Optimal Decision Tree Problem

Each decision region corresponds to a single h

Example: 1-dimensional threshold function

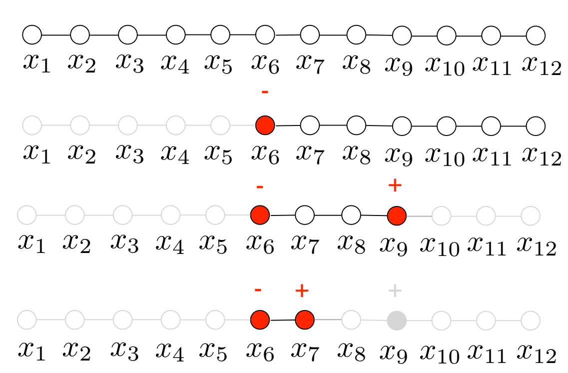

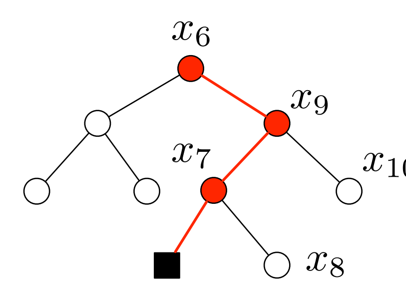

Generalized binary search

[Garey & Graham, 1974; Dasgupta, 2004; Golovin & Krause, 2010; …]

Greedy is near-optimal: $\mathrm{cost}(\pi) = (\ln(1/p_{\min}) + 1) \cdot \mathrm{cost}(\pi^*)$

The Equivalence Class Determination Problem

Decision regions do not overlap

Example: Medical Diagnostics

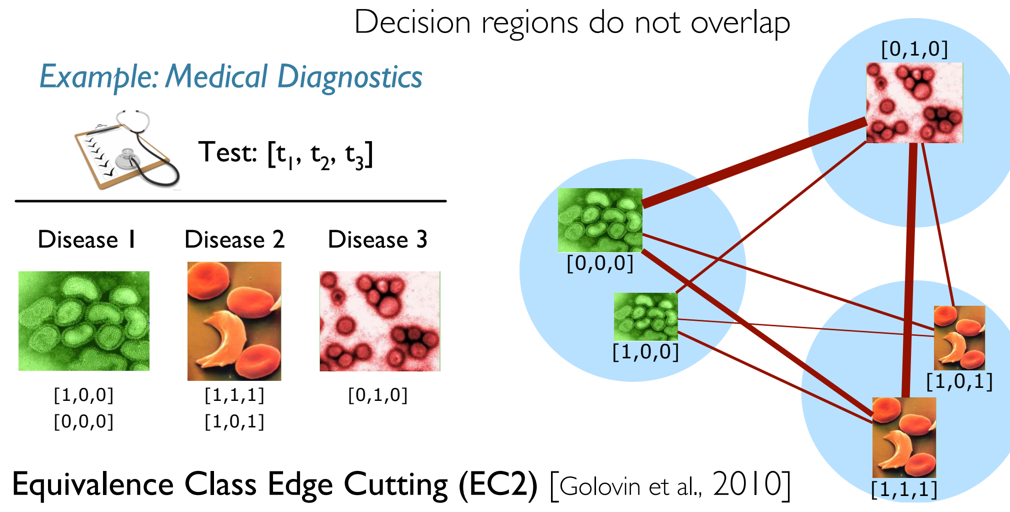

Equivalence Class Edge Cutting (EC2) [Golovin et al., 2010]

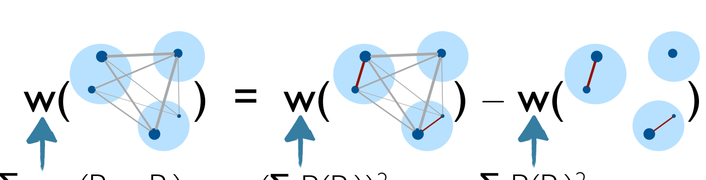

$\Sigma_{i \neq j} w(R_i \times R_j)$ — $(\Sigma_i P(R_i))^2$ — $\Sigma_i P(R_i)^2$

Greedy is near-optimal: $\mathrm{cost}(\pi) = (2\ln(1/p_{\min}) + 1) \cdot \mathrm{cost}(\pi^*)$

Essentially, the EC2 objective is efficiently computed as **elementary symmetric polynomials**. We can compute the sum of all edge weights, and then subtract those that share a region.

HEC: The Hyperedge Cutting Algorithm

Hyperedge ⇔ a (multi-) set of hypotheses that do not share a region

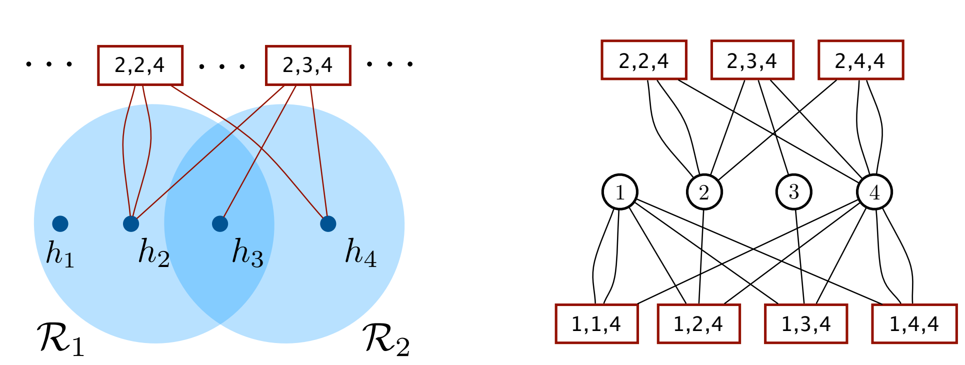

Decision regions and hypotheses — The hypergraph of cardinality k=3

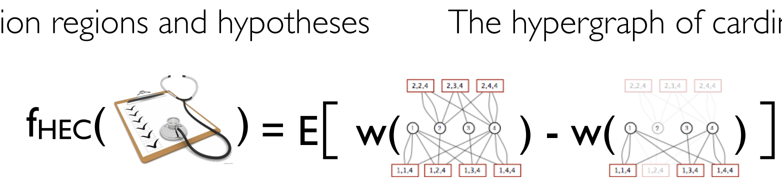

$\mathrm{cost}(\pi) = (k\ln(1/p_{\min}) + 1) \cdot \mathrm{cost}(\pi^*)$

Constructing the hypergraph is **NP-hard** [Javdani et al. 2014]

Adaptive Submodularity [Golovin & Krause, 2010]

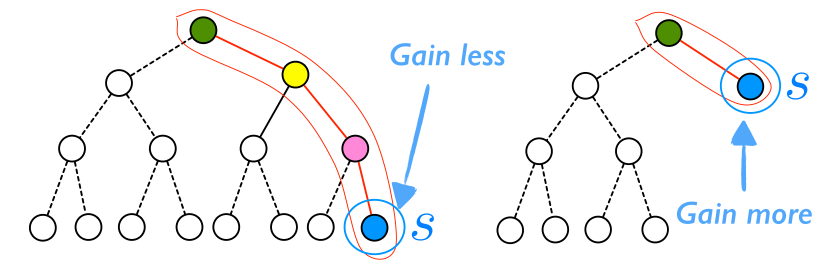

Playing an action earlier only increases its **expected** marginal benefit.

The "Or" Decomposition of DRD

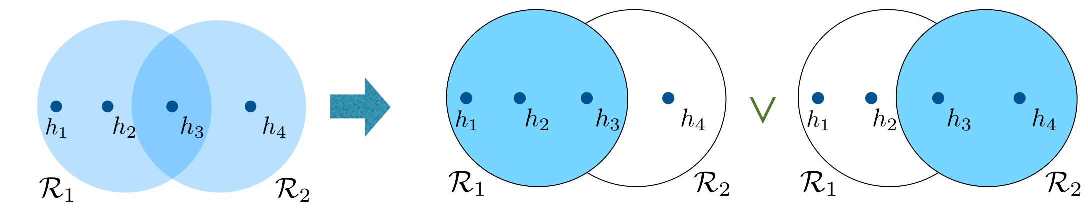

DiRECt: The Noisy-Or Construction

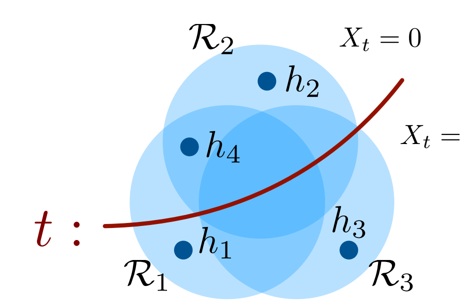

The DRD Problem

Main idea: To jointly optimize multiple EC2 objective functions corresponding to the decomposed graphs.

Cutting all edges in (at least) one of the graphs is both **sufficient and necessary** to identify the true region!

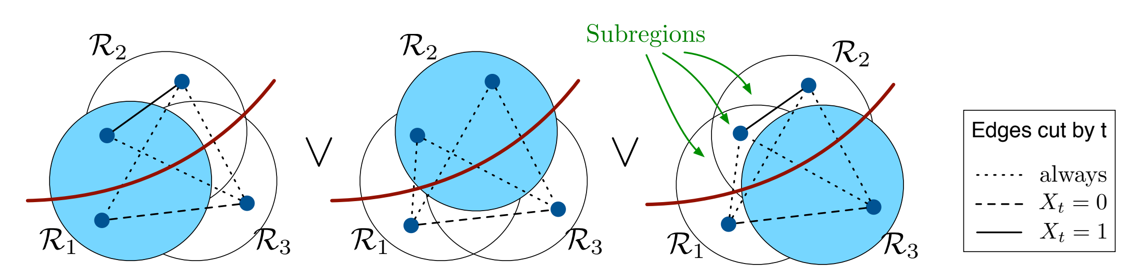

The **D**ecision **R**egion **E**dge **Cut**ting (DiRECt) Objective

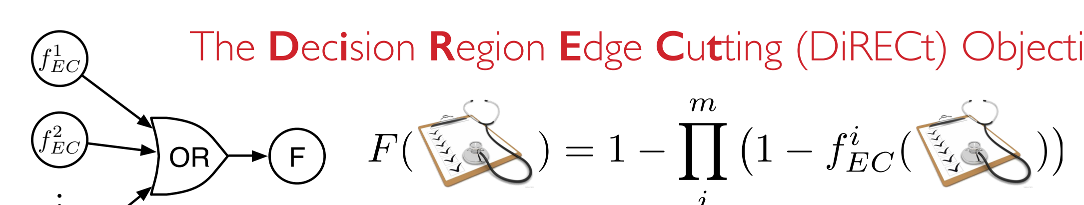

Lemma: Noisy-Or of f_{EC}'s preserves adaptive submodularity.

number of decomposed graphs

Greedy is near-optimal: $\mathrm{cost}(\pi) = (2m\ln(1/p_{\min}) + 1) \cdot \mathrm{cost}(\pi^*)$

Improving the Bound via Graph Coloring

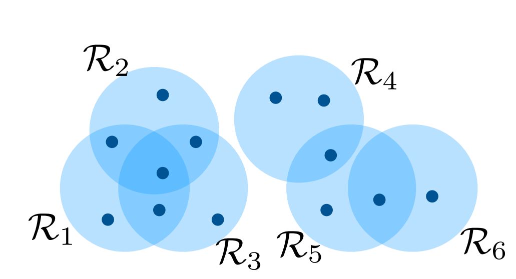

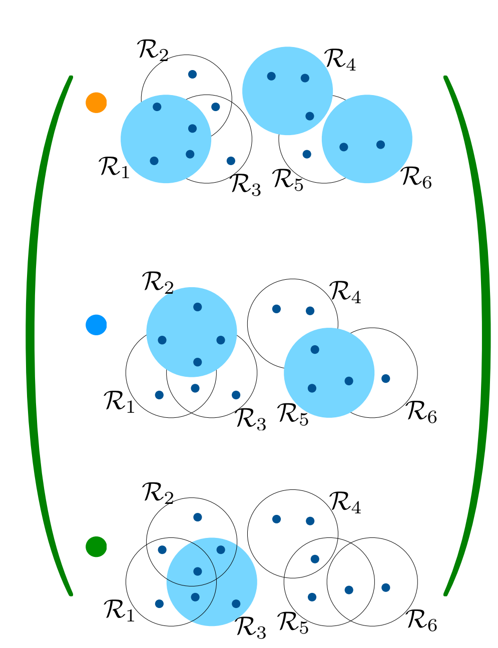

One can partition the regions via **Graph Coloring**

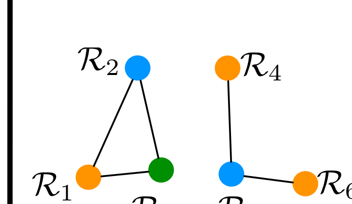

# of colors is bounded by the max degree of the color graph

$\mathrm{cost}(\pi) = (2 \times (\#\mathrm{Colors}) \times \ln(1/p_{\min}) + 1)\, \mathrm{cost}(\pi^*)$

Experimental Results

Candidate Algorithms

- GBS: Generalized Binary Search
- GBS-DRD: GBS with DRD stopping criteria
- EC2: Equivalence Class Edge Cutting
- EC2-DRD: EC2 with DRD stopping criteria
- VoI: Myopic Value of Information
- HEC: Hyperedge Cutting
- DiRECt: Decision Region Edge Cutting

Approximate comparison-based Search

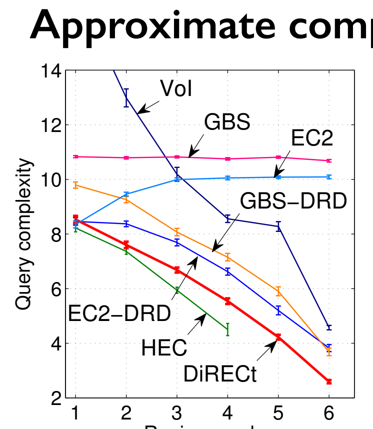

MovieLens

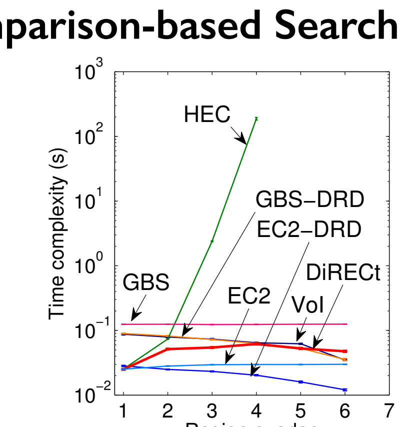

MovieLens - time

Active Touch-based localization

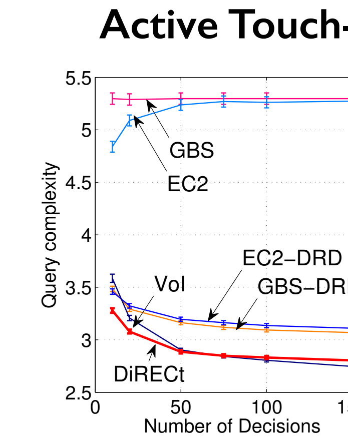

Robot - varying |D|

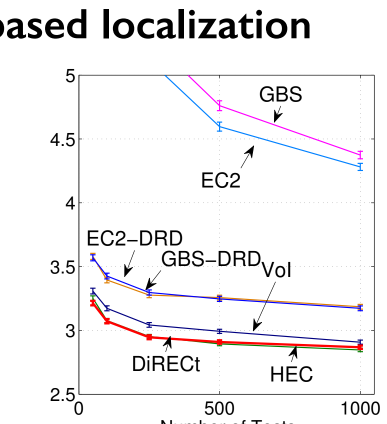

Robot - varying |T|

Adaptive management for biodiversity conservation

Hypotheses	Strategies							
		1	2	3	4	5	6	7
Description	Weight (%)	Wait	Kill Flies	Swap eggs	Restore mea	April DD and	No salvage	No disturban
Too young	9.4	0.586	0.491	0.581	0.735	0.663	0.539	0.402
Black flies	29.1	0.021	0.425	0.242	0.373	0.253	0.589	0.139
Social condit	11.9	0.093	0.145	0.22	0.429	0.321	0.218	0.485
Nutrient limi	22.8	0.036	0.081	0.254	0.992	0.863	0.128	0.166
Nutrient limi	5.9	0.093	0.119	0.26	0.466	0.405	0.185	0.223
Nutrient limi	6.6	0.036	0.077	0.243	0.792	0.703	0.172	0.179
Egg salvage	4.4	0.147	0.622	0.662	0.436	0.291	0.354	0.158
Disturbance	10	0.12	0.393	0.74	0.363	0.216	0.168	0.256
Expected value		0.106	0.284	0.343	0.59	0.475	0.331	0.231

EMPCranes: the hypotheses-decision utility table

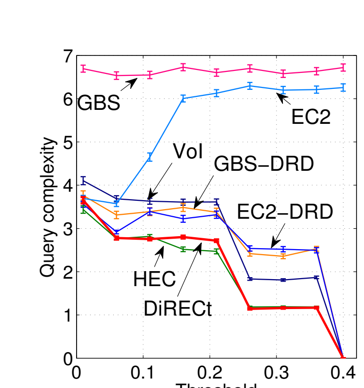

EMPCranes

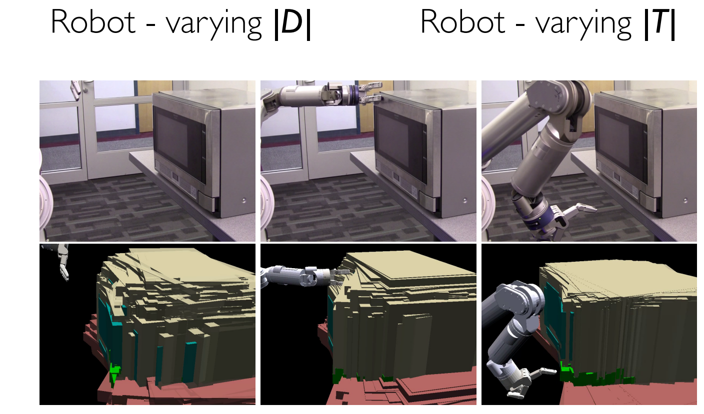

Experimental setup for touch-based localization