

Active Exploration for Robust Object Detection

Javier Velez[†] Garrett Hemann[†] Albert S. Huang[†] Ingmar Posner[‡] Nicholas Roy[†]

[†]MIT Computer Science and
Artificial Intelligence Laboratory
Cambridge, MA, USA

[‡]Mobile Robotics Group
Dept. of Engineering Science, Oxford University
Oxford, UK

Abstract

Today, mobile robots are increasingly expected to operate in ever more complex and dynamic environments. In order to carry out many of the higher-level tasks envisioned a semantic understanding of a workspace is pivotal. Here our field has benefited significantly from successes in machine learning and vision: applications in robotics of off-the-shelf object detectors are plentiful. This paper outlines an online, any-time planning framework enabling the active exploration of such detections. Our approach exploits the ability to move to different vantage points and implicitly weighs the benefits of gaining more certainty about the existence of an object against the physical cost of the exploration required. The result is a robot which plans trajectories specifically to decrease the entropy of putative detections. Our system is demonstrated to significantly improve detection performance and trajectory length in simulated and real robot experiments.

INTRODUCTION

Years of steady progress in robotic mapping and navigation techniques have made it possible for robots to construct accurate traversability maps of relatively complex environments and to robustly navigate within them (see, for example Newman et al., 2009). Such maps generally represent the world as regions which are traversable by a robot and regions which are not (see Fig. 1(a)) and are ideal for low-level navigation tasks such as moving through the environment without collisions. However, mobile robots are increasingly tasked to perform high-level requests in complex and dynamic environments. Sophisticated interactions between an agent and its workspace require the addition of semantic information into traditional environment representations, such as information about the location and identity of objects in the workspace (see Fig. 1). For example, a cleaning robot knows that dirty dishes are usually on top of tables and benefits from knowing where the tables are in the environment.

This work was performed under NSF IIS Grant #0546467 and ONR MURI N1141207-236214 and their support is gratefully acknowledged.

Recently, advances in vision- and laser-based object recognition have been leveraged to enrich maps with higher-order, semantic information (e.g., Posner, Cummins, and Newman, 2009; Douillard, Fox, and Ramos, 2008; Mozos, Stachniss, and Burgard, 2005; Anguelov et al., 2004). A straightforward approach to adding semantic information is to accept the results of a standard object detector framework *prima facie*, irrespective of sensor noise. A consequence of directly using the results of an object detector is that the quality of the map built strongly depends on the shortcomings of the object detector. Vision-based object detection, for example, is often-times plagued by significant performance degradation caused by a variety of factors including a change of aspect compared to that encountered in the training data and, of course, occlusion (e.g., Coates and Ng, 2010; Mittal and Davis, 2008). Both of these factors can be addressed by choosing the location of the camera carefully before acquiring an image and performing object detection.

Rarely, however, are the capabilities of the mobile robot building the map exploited to improve the robustness of the detection process by specifically counteracting known detector issues. Rather than placing the burden of providing perfect detections with the detector itself, the robot can *act* to improve its perception.

In this work, we explore how a robot’s ability to selectively gather additional information about a possible object by moving to a specified location — a key advantage over more conventional, static deployment scenarios for object detectors — can improve the precision of object detections included in the map. In particular, we propose a planning framework that reasons about taking detours from the shortest path to a goal in order to explore potential object detections based on noisy observations from an off-the-shelf object detector. At each step, the agent weighs the potential benefit of increasing its confidence about a potential object against the cost of taking a detour to a more suitable vantage point.

We make two primary contributions in this paper. Firstly, we explicitly incorporate the costs of motion when planning sensing detours. Secondly, we give an approximate sensor model that captures correlations between subsequent observations (objects). Previous work has largely ignored motion costs and has typically assumed observations are conditionally independent given the sensor position. Inspired by recent progress in forward search for planning under uncertainty, we




Figure 1: A traditional environment map (a) with regions that are traversable (cyan) and not (red) useful for navigation, (b) the traversable map augmented with semantic information about the identity and location of objects in the world allowing for richer interactions between agent and workspace.

show that motion costs can be incorporated and measurement correlations modeled, allowing us to efficiently find robust observation plans.¹

RELATED WORK

Planning trajectories for a mobile sensor has been explored in various domains. The approach presented here is inspired by forward search strategies for solving partially observable Markov decision processes (Ross et al., 2008; Prentice and Roy, 2009) but incorporates a more complex model that approximates the correlations between observations.

The controls community and sensor placement community formulate the problem as minimizing a norm of the posterior belief, such as entropy, without regard to motion cost. As a consequence, greedy strategies that choose the next-most valuable measurement can be shown to be boundedly close to the optimal, and the challenge is to generate a model that predicts this next-best measurement (Guestrin, Krause, and Singh, 2005; Krause et al., 2008). Formulating the information-theoretic problem as a decision-theoretic POMDP Sridharan, Wyatt, and Dearden (2008) showed that true multi-step policies did improve the performance of a computer vision system in terms of processing time. However, the costs of the actions are independent (or negligible), leading to a submodular objective function and limited improvement over greedy strategies.

A few relevant pieces of work from the active vision domain include Arbel and Ferrie (1999) and more recently Laporte and Arbel (2006) who use a Bayesian approach to model detections that is related to ours, but only search for the next-best viewpoint, rather than computing a full plan. The work of Deinzer, Denzler, and Niemann (2003) is perhaps most similar to ours in that the viewpoint selection problem is framed using reinforcement learning, but again the authors “neglect costs for camera movement” and identify the absence of costs as a limitation of their work.

The contribution of our work over the existing work is primarily to describe a planning model that incorporates both action costs and detection errors, and specifically to give an approximate observation model that captures the correlations

¹This paper presents results that originally appeared at ICAPS 2011 (Velez et al., 2011).




Figure 2: A conceptual illustration of (a) the robot at view-point x while following the original trajectory (bold line) towards the goal (red star), (b) the perception field for a particular object detector centered around an object hypothesis and alternative path (bold dash-dotted line) taken by a curious robot to investigate the object of interest. Lighter cell shading values indicate a lower conditional entropy and therefore desirable vantage points.

between successive measurements that still allows forward-search planning to operate, leading to an efficient multi-step search to improve object detection.

PROBLEM FORMULATION

Consider an agent navigating through an environment with potential objects of interest at unknown locations. The agent has a goal destination but is allowed to inspect the locations of said objects of interest; for example, a rescue robot looking for people in a first-responder scenario. Traditionally, an object detector is used at waypoints along the way and an object is either accepted into the map or rejected based upon a simple detector threshold.

However, the lack of introspection of this approach regarding both the confidence of the object detector and the quality of the data gathered can lead to an unnecessary acceptance of spurious detections; what looked like a ramp from a particular viewpoint may in fact have been a trick of the light. Most systems simply discard lower confidence detections, and have no way to improve the estimate with further, targeted measurements. In contrast, we would like the robot to modify its motion to both minimize total travel cost and the cost of errors when deciding whether or not to add newly observed objects to the map.

Let us represent the robot as a point $x \in \mathbb{R}^2$; without loss of generality, we can express a robot trajectory as a set of waypoints $\mathbf{x}^{1:K}$. As the robot moves, it receives output from its object detector that gives rise to a belief over whether a detected object truly exists at the location indicated. Let the presence of an object at some location (x, y) be captured by the random variable $Y \in \{\text{object}, \text{no-object}\}^2$. Let us also define a *decision* action $a \in \{\text{accept}, \text{reject}\}$, where the detected object is either accepted into the map (the detection is determined to correspond to a real object) or rejected (the detection is determined to be spurious). Additionally, we have an ex-

²We assume the robot knows its own location, and has a sufficiently well-calibrated camera to assign the location of the object. In this work, the uncertainty is whether or not an object is present at the given location (x, y) .

explicit cost $\xi_{dec} : \{\text{accept, reject}\} \times \{\text{object, no-object}\} \mapsto \mathbb{R}$ for a correct or incorrect accept or reject decision. We cannot know the true costs of the decisions because we ultimately do not know the true state of objects in the environment. But, we can use a probabilistic sensor model for object detections to minimize the expected cost of individual decision actions ξ_{dec} given the prior over objects. We therefore formulate the planning problem as choosing a sequence of waypoints to minimize the total travel cost along the trajectory and the expected costs of the decision actions at the end of the trajectory.

OBJECT DETECTION SENSOR MODEL

In order to compute the expected cost of decision actions, we must estimate the probability of objects existing in the world, and therefore require a probabilistic model of the object detector. The key idea is that we model the object detector as a spatially varying process; around each potential object, we characterize every location with respect to how likely it is to give rise to useful information.

A measurement, z , at a particular viewpoint consists of the output of the object detector, assumed to be a real number indicating the confidence of the detector that an object exists. The distribution over the range of confidence measurements is captured by the random variable Z defined over a range \mathcal{Z} of discretized states (bins). At every location \mathbf{x} the posterior distribution over Y can be expressed as

$$p(y|z, \mathbf{x}) = \frac{p(z|y, \mathbf{x})p(y)}{\sum_{y'} p(z|y', \mathbf{x})p(y')}, \quad (1)$$

where $p(z|y, \mathbf{x})$ denotes the likelihood of observing a particular detector confidence at \mathbf{x} given the true state of the object. This likelihood can be obtained empirically.

When observations originate from a physical device such as a camera, a straightforward approach is to treat observations as conditionally independent given the state of the robot (see Fig. 3(a)). This model is easy to work with, but assumes that measurements vary only as a result of sensor noise and the object state. In practice, measurements vary as a function of many different factors such as scene geometry, lighting, occlusion, etc. The approach in Fig. 3(a) simply ignores these factors. If all of these (often unobservable) sources of variation were correctly modeled and estimated in an environmental variable Ψ (Fig. 3(b)), then the conditional independence would hold, but constructing and maintaining a model sufficient to capture the image generation process is an intractable computational and modeling burden.

To correct our observation model without an explicit model Ψ , we can maintain a history of observation viewpoints. As more viewpoints are visited, knowledge regarding Y and future observations can be integrated recursively. As shown in Fig. 3(c), we remove Ψ and add a dependency between previous viewpoints and the current observation z^K . Let \mathcal{T}^K denote a trajectory of K locations traversed in sequence. At each location a measurement is obtained, giving a possible detection and the corresponding confidence. The trajectory is thus described by a set $\mathcal{T}^K = \{\{\mathbf{x}^1, z^1\}, \{\mathbf{x}^2, z^2\}, \dots, \{\mathbf{x}^K, z^K\}\}$ of K location-observation pairs. Knowledge gained at each step along the

trajectory can be integrated into the posterior distribution over Y such that

$$p(y|\mathcal{T}^K) = \frac{p(z^K|y, \mathbf{x}^K, \mathcal{T}^{K-1})p(y|\mathcal{T}^{K-1})}{p(z^K|\mathbf{x}^K, \mathcal{T}^{K-1})}, \quad (2)$$

where z^K is the K^{th} observation, which depends not only on the current viewpoint but also on the history of measurements \mathcal{T}^{K-1} . In principle, K can be arbitrarily long, so the primary challenge is to develop an efficient way of conditioning our observation model on previous viewpoints.

To overcome this difficulty, we approximate the real process of object detection with a simplistic model of how the images are correlated. We replace the correlating influence of environment Ψ with a convex combination of a fully uncorrelated and a fully correlated model such that the new posterior belief over the state of the world is computed as

$$p(y|\mathcal{T}^K) = \left((1-\alpha) \frac{p(z^K|y, \mathbf{x}^K)}{p(z^K|\mathbf{x}^K)} + \alpha \right) p(y|\mathcal{T}^{K-1}) \quad (3)$$

This captures the intuition that repeated observations from the same viewpoint add little to the robot's knowledge about the state of the world. Observations from further afield, however, become increasingly independent; Ψ has less of a correlating effect. The mixing parameter, α , can be chosen such that no information is gained by taking additional measurements at the same location and the information content of observations increases linearly with distance from previous ones (see Velez et al., 2011).

Perception Field

Using the observation model and Equ. 3, it is possible to evaluate how much additional information a future viewpoint \mathbf{x}^K provides on the object state Y . Given \mathbf{x}^K and the trajectory \mathcal{T}^{K-1} visited thus far, the expected reduction in uncertainty is captured by the mutual information between Y and the observation Z received at \mathbf{x}^K :

$$I(Y, Z; \mathbf{x}^K, \mathcal{T}^{K-1}) = H(Y; \mathcal{T}^{K-1}) - H(Y|Z; \mathbf{x}^K, \mathcal{T}^{K-1}), \quad (4)$$

Of these two terms, the first $H(Y; \mathcal{T}^{K-1})$ term is independent of \mathbf{x}^K and can be ignored. The second term, the conditional entropy, can be readily evaluated using empirically determined quantities. In particular, we use the conditional entropy evaluated at all viewpoints in the robot's surrounding workspace to form the *perception field* (Velez et al., 2011) for a particular object hypothesis (see Fig. 2(b)). This field denotes locations with high expected information gain. Due to the correlations between individual observations made over a trajectory of viewpoints, the perception field changes as new observations are added. Note that if the robot's only goal was to determine Y with the greatest confidence, it would repeatedly visit locations with least conditional entropy, as indicated by the perception field.

PLANNING DETOURS

We now describe a planning algorithm that trades off the benefit of gaining additional information about an object hypothesis against the operational cost of obtaining this information.

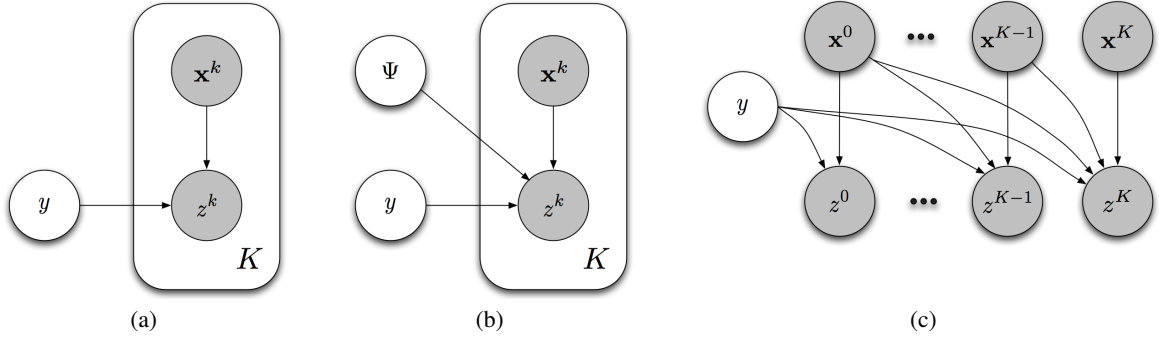


Figure 3: Different graphical models representing the observation function. (a) A naive Bayes approximation, that assumes that every observation is conditionally independent given knowledge of the object. (b) The true model that assumes that observations are independent given knowledge of the environment and the object. (c) The model employed here, in which the correlations are approximated by way of a mixture model parameterized by α as per Equ. 3.

When an object is first detected, a new path to the original goal is planned based on the total cost function, which includes both the motion cost along the path and the value of measurements from locations along the path, expressed as a reduction in the expected cost of decision actions. The cost function consists of two terms: the motion cost $c_{mot}(\mathbf{x}^{1:K})$ and the decision cost $c_{dec}(\mathbf{x}^{1:K}, a)$, such that the optimal plan π^* is given by

$$\pi^* = \arg \min_{\mathbf{x}^{1:K}, a} (c_{mot}(\mathbf{x}^{1:K}) + c_{dec}(\mathbf{x}^{1:K}, a)), \quad (5)$$

$$c_{dec}(\mathbf{x}^{1:K}, a) = E_{y|\mathbf{x}^{1:K}}[\xi_{dec}(a, y)], \quad (6)$$

where $E_{y|\mathbf{x}^{1:K}}[\cdot]$ denotes the expectation with respect to the robot's knowledge regarding the object, after having executed path $\mathbf{x}^{1:K}$. We choose $c_{mot}(\mathbf{x}^{1:K})$ to be a standard motion cost function proportional to the path length.

The planning process therefore proceeds by searching over sequences of $\mathbf{x}^{1:K}$, evaluating paths by computing expectations with respect to both the observation sequences and the object state. The paths with the lowest decision cost will tend to be those leading to the lowest posterior entropy, avoiding the large penalty for false positives or negatives.

Multi-step Planning A naive approach to searching over trajectories scales exponentially with the planning horizon K and rapidly becomes intractable as more observations are considered. We therefore adopt a roadmap scheme in which a fixed number of locations are sampled every time a new viewpoint is to be added to the current trajectory. A graph is built between the sampled poses, with straight-line edges between samples. The perception field is used to bias sampling towards locations with high expected information gain.

The planning approach described so far can be extended to planning in an environment with M object hypotheses by considering a modified cost function which simply adds the cost for each object and treating the existence of each object independently such that individual perception fields add at a particular location. See Velez et al. (2011) for more details.

EXPERIMENTS

We tested our approach on both a simulated robot with an empirically derived model of an object detector and on an actual



Figure 4: Robotic wheelchair platform

autonomous wheelchair (Fig. 4) using a vision-based object detector (Felzenszwalb, McAllester, and Ramanan, 2008). In both the simulation and physical experiments, the robot was tasked with reaching a manually specified destination. The robot was rewarded for correctly detecting and reporting the location of doors in the environment, penalized for false alarms, and incurred a cost proportional to the length of its total trajectory. We chose doors as objects of interest due to their abundance in indoor environments and their utility to mobile robots – identifying doorways is often a component of higher level tasks.

Our autonomous wheelchair is equipped with onboard laser range scanners, primarily used for obstacle sensing and navigation, and a Point Grey Bumblebee2 color stereo camera. The simulation environment is based on empirically constructed models of the physical robot and object detector. We set $\xi_{dec}(\text{reject}, \cdot)$ to zero, indicating no penalty for missed objects.

Learned Perception Field

Fig. 5 shows the perception field for the detector model learned from 3400 training samples, with each cell indicating the conditional entropy of the posterior distribution over

Average	Random $_{\beta=0.8}$	Random $_{\beta=0.6}$	Greedy $_{\beta=0.8}$	Greedy $_{\beta=0.6}$	Planned	RTBSS
Precision	0.27 \pm 0.03	0.26 \pm 0.04	0.31 \pm 0.06	0.60 \pm 0.07	0.75 \pm 0.06	0.45 \pm 0.06
Recall	0.72 \pm 0.06	0.60 \pm 0.07	0.44 \pm 0.07	0.62 \pm 0.07	0.80 \pm 0.06	0.58 \pm 0.07
Path Length (m)	62.63 \pm 0.02	62.03 \pm 0.67	67.08 \pm 2.23	41.95 \pm 0.88	54.98 \pm 3.04	47.57 \pm 0.19
Total Trials	50	50	50	50	50	50

Table 1: Simulation performance on single door scenario, with standard error values.

Average	Random $_{\beta=0.8}$	Random $_{\beta=0.6}$	Greedy $_{\beta=0.8}$	Greedy $_{\beta=0.6}$	Planned	RTBSS
Precision	0.64 \pm 0.03	0.67 \pm 0.03	0.64 \pm 0.03	0.54 \pm 0.03	0.53 \pm 0.05	0.70 \pm 0.03
Recall	0.64 \pm 0.04	0.69 \pm 0.03	0.63 \pm 0.02	0.57 \pm 0.03	0.76 \pm 0.03	0.66 \pm 0.03
Path Length (m)	199.62 \pm 11.24	161.36 \pm 6.13	153.32 \pm 4.37	121.35 \pm 1.32	138.21 \pm 7.12	160.74 \pm 6.08
Total Trials	50	50	50	50	50	50

Table 2: Simulation performance on multiple door scenario, with standard error values.




Figure 5: Learned perception field for a possible door. The unknown object is centered at the origin (blue). Brighter regions correspond to viewpoints more likely to result in higher confidence posterior beliefs.

the true state of an object given an observation from this viewpoint, $p(y|z, \mathbf{x})$ and a uniform initial belief. Intuitively, this suggests that the viewpoint from which a single observation is most likely to result in a low-entropy posterior belief is approximately 9 m directly in front of the door. Viewpoints too close to or far from the door, and those from oblique angles are less likely to result in high-confidence posterior beliefs.

Simulation Results

We first assessed our planning approach using the learned model in a simulated environment. Our simulation environment consisted of a robot navigating through an occupancy map, with object detections triggered according to the learned object detection model and correlation model α . We also simulated false positives by placing non-door objects that probabilistically triggered object detections using the learned model for false-alarms. The processing delay incurred by the actual object detector was also simulated (the object detector requires approximately 2 seconds to process a spatially decimated 512x384 pixel image).

Comparison Algorithms

For the simulation trials we compared our algorithm against three other algorithms. The Random $_{\beta}$ algorithm repeatedly obtained observations from randomly selected viewpoints near detected objects until the belief of each object exceeded a threshold β , and then continued on to the original destination. The Greedy $_{\beta}$ algorithm selected the *best* viewpoint according to our perception field for each potential object until




Figure 6: Simulation scenario with multiple doors (cyan triangles) and non-doors (black bars). The robot starts at the bottom-left (S) with a goal near the bottom-right (G). The robot's object detector responds to both doors and non-doors. Top: example trajectory (purple) executed using random viewpoint selection. Bottom: example trajectory (purple) executed using planned viewpoints.

the belief of each object exceeded a threshold β . Lastly, we compared our algorithm against the RTBSS online POMDP algorithm (Paquet, Tobin, 2005) with a maximum depth of 2.

Single Door Simulation

First, we tested our planning algorithm on a small simulation environment with one true door and two non-doors. Table 1 shows the results of 50 trials. Overall, explicitly planning viewpoints resulted in significantly higher performance. The planned viewpoints algorithm performed better than RTBSS in terms of precision and recall, most likely because our algorithm sampled continuous-space viewpoints and the RTBSS algorithm had a fixed discrete representation, while RTBSS paths were shorter.

Multi Door Simulation

Next, we evaluated our algorithm in a larger, more complex scenario containing four doors and six non-door objects. Figure 6 shows the multiple door simulation environment and ex-




Figure 7: Trajectory executed on the actual robot wheelchair using planned viewpoints from 'S' to 'G' where the robot discovers one true door (cyan triangle). Near the goal, it detects two more possible doors (red dots), detours to inspect them, and (correctly) decides that they are not doors.

Average	Greedy $_{\beta=0.8}$	Planned
Precision	0.53 \pm 0.14	0.7 \pm 0.15
Recall	0.60 \pm 0.14	0.7 \pm 0.15
Path Length (m)	153.86 \pm 33.34	91.68 \pm 15.56
Total Trials	10	10

Table 3: Results of real-world trials using robot wheelchair.

ample trajectories planned and executed by the planned and Random $_{\beta}$ algorithms.

Table 2 shows the simulation results for the multi-door scenario. Our planned viewpoints algorithm resulted in the second shortest paths after Greedy $_{\beta=0.6}$ but with superior detection performance. Planned viewpoints also resulted in significantly shorter paths than RTBSS given the same operating point on the ROC curve.

Physical Wheelchair Trials

We conducted a small experiment comparing our planned viewpoints algorithm and the Greedy $_{\beta=0.8}$ on a robot wheelchair platform. The robot was given a goal position such that a nominal trajectory would bring it past one true door, and near several windows that trigger object detections.

Figure 7 illustrates the trajectory executed during a single trial of the planned viewpoints algorithm, and Table 3 summarizes the results of all trials. Our planned viewpoints algorithm resulted in significantly shorter trajectories while maintaining comparable precision and recall. For doors detected with substantial uncertainty, our algorithm planned more advantageous viewpoints to increase its confidence and ignored far away detections because of high motion cost.

CONCLUSIONS AND FUTURE WORK

Previous work in planned sensing has largely ignored motion costs of planned trajectories and used simplified sensor models with strong independence assumptions. In this paper, we presented a sensor model that approximates the correlation in observations made from similar vantage points, and an efficient planning algorithm that balances moving to highly informative vantage points with the motion cost of taking de-

tours. The result is a robot which plans trajectories specifically to decrease the entropy of putative detections. The performance of our algorithm could be further improved by future work in both the sensor model and planning technique.

References

- Anguelov, D.; Koller, D.; Parker, E.; and Thrun, S. 2004. Detecting and modeling doors with mobile robots. In *Proc. ICRA*.
- Arbel, T., and Ferrie, F. P. 1999. Viewpoint selection by navigation through entropy maps. In *Proc. ICCV*.
- Coates, A., and Ng, A. Y. 2010. Multi-camera object detection for robotics. In *Proc. ICRA*.
- Deinzer, F.; Denzler, J.; and Niemann, H. 2003. Viewpoint selection - planning optimal sequences of views for object recognition. In *Proc. ICCV*. Springer.
- Douillard, B.; Fox, D.; and Ramos, F. 2008. Laser and Vision Based Outdoor Object Mapping. In *Proc. RSS*.
- Felzenszwalb, P.; McAllester, D.; and Ramanan, D. 2008. A discriminatively trained, multiscale, deformable part model. In *Proc. CVPR*.
- Guestrin, C.; Krause, A.; and Singh, A. 2005. Near-optimal sensor placements in gaussian processes. In *Proc. ICML*.
- Krause, A.; Leskovec, J.; Guestrin, C.; VanBriesen, J.; and Faloutsos, C. 2008. Efficient sensor placement optimization for securing large water distribution networks. *Journal of Water Resources Planning and Management* 134:516.
- Laporte, C., and Arbel, T. 2006. Efficient discriminant viewpoint selection for active bayesian recognition. *International Journal of Computer Vision* 68(3):267–287.
- Mittal, A., and Davis, L. 2008. A general method for sensor planning in multi-sensor systems: Extension to random occlusion. *International Journal of Computer Vision* 76:31–52.
- Mozos, O. M.; Stachniss, C.; and Burgard, W. 2005. Supervised Learning of Places from Range Data using Adaboost. In *Proc. ICRA*.
- Newman, P.; Sibley, G.; Smith, M.; Cummins, M.; Harrison, A.; Mei, C.; Posner, I.; Shade, R.; Schroeter, D.; Murphy, L.; Churchill, W.; Cole, D.; and Reid, I. 2009. Navigating, recognising and describing urban spaces with vision and laser. *International Journal of Robotics Research*.
- Paquet, S. and Tobin, L. 2005 Real-Time Decision Making for Large POMDPs. In *Proc. Canadian Conference on Artificial Intelligence*.
- Posner, I.; Cummins, M.; and Newman, P. 2009. A generative framework for fast urban labeling using spatial and temporal context. *Autonomous Robots*.
- Prentice, S., and Roy, N. 2009. The belief roadmap: Efficient planning in belief space by factoring the covariance. *International Journal of Robotics Research* 8(11-12):1448–1465.
- Ross, S.; Pineau, J.; Paquet, S.; and Chaib-draa, B. 2008. Online planning algorithms for POMDPs. *Journal of Artificial Intelligence Research* 32(1):663–704.
- Sridharan, M.; Wyatt, J.; and Dearden, R. 2008. HiPPo: Hierarchical POMDPs for Planning Information Processing and Sensing Actions on a Robot. In *Proc. ICAPS*.
- Velez, J.; Hemann, G.; Huang, A.; Posner, I.; and Roy, N. 2011. Planning to Perceive: Exploiting Mobility for Robust Object Detection. In *Proc. ICAPS*.