Improving Heuristics of Optimal Perception Planning using Visibility Maps

Tiago Pereira*†‡
*Faculty of Engineering, University of Porto
Porto, Portugal
{tiago.raul, amoreira}@fe.up.pt

António Moreira*†
†INESC-TEC
INESC Technology and Science
Porto, Portugal

Manuela Veloso‡
‡Carnegie Mellon University
Pittsburgh, USA
tpereira@cmu.edu, mmv@cs.cmu.edu

***Abstract*—In this paper we consider the problem of motion planning for perception of a target position. A robot has to move to a position from where it can sense the target, while minimizing both motion and perception costs. The problem of finding paths for robots executing perception tasks can be solved optimally using informed search. In perception path planning, the solution for the perception task considering a straight line without obstacles is used as heuristic. In this work, we propose a heuristic that can improve the search efficiency. In order to improve the node expansion using a more informed search, we use the robot Approximate Visibility Map (A-VM), which is used as a representation of the observability capability of a robot in a given environment. We show how the critical points used in A-VM provide information on the geometry of the environment, which can be used to improve the heuristic, increasing the search efficiency. The critical points allow a better estimation of the minimum motion and perception cost for targets in non-traversable regions that can only be sensed from further away. Finally, we show the contributed heuristic dominates the common heuristic (based on the euclidian distance), and present the results of the performance increase in terms of node expansion.**

***Keywords*-perception planning; visibility maps; improved heuristics;**

I. Introduction

In this work we deal with motion planning for perception tasks, where both the motion and sensing costs have to be considered in order to find an optimal path.

As we show in Figure 1, a path has motion cost $cost_m$, proportional to the distance traveled, and a perception cost $cost_p$, which is a function of the perceiving distance between the goal position and the target. Here we assume the sensing cost is a function of the minimum distance between the path and the target. As a result, the final position, i.e., the goal position, has the minimum distance to the target. The path selection changes depending on the relative costs of motion and perception, in order to minimize the overall cost.

Using an informed search algorithm, such as PA* [1], it is possible to explore the space and find the optimal path while reducing the number of expanded nodes compared to breadth-first search. However, the basic algorithm searches the environment without considering any information from the world. As in the common motion planning problem with A*, the heuristic in a certain node is determined using the solution for the straight line perception task between the node and the target, without considering obstacles.

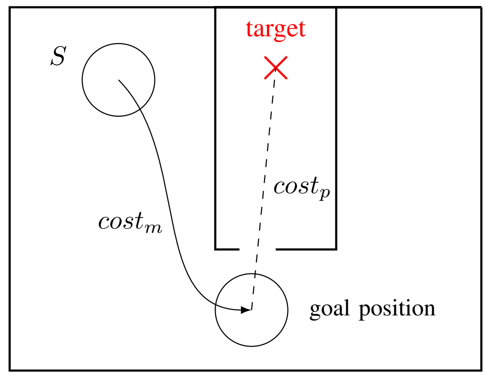

Figure 1. The cost of a path is given by the sum of the motion cost and the perception cost; an informed heuristic search can find the optimal path; in this example, the target is inside an unreachable region, so there is always a minimum perception distance independent of the optimal path.

Here we contribute a new heuristic that can reduce the number of expanded nodes when the perception target location lies in non-traversable regions, where there is a minimum perception distance. For that purpose, we use the Approximate Visibility Map, a transformation on the original map that can be used to solve the observability problem, i.e., determining which regions are visible by a circular robot from any point reachable from the initial robot position. The parameters of this transformation are the robot size and maximum sensing range.

In the Approximate Visibility Map critical points are used in order to estimate the visibility inside non-traversable regions, while reducing the computation time compared to the brute-force algorithm. In this work we prove the critical points, by definition the points in the configuration space that generate the frontiers of motion reachability, are the closest to any target position inside unreachable regions. Therefore, we use the distance from the target to each critical point as an estimate of the minimum perception distance. The distance to the critical point can be used to create an admissible heuristic that outperforms the heuristic in PA*, which is only a function of the distance between the current node and the target. The critical point can also be used to have an estimate of the minimum motion distance.

We present related work, then describe the new heuristic used in the search algorithm of motion planning for perception tasks, proving it is admissible and dominant over the basic heuristic, yielding a faster convergence to the optimal path. We show some results on the increased efficiency in terms of node expansion, and then present our conclusions and future work directions.

II. Related Work

Many robotic applications consider perception separately from planning, with both being computed interleaved [2]. It has been used for tasks as varied as SLAM [3], robot localization [4], in exploration to guide robots to unexplored regions [5], and for object recognition [6].

However, perception got recently a more active role in planning. An example is object detection, where the next moves of the robot should be planned to maximize the likelihood of correct object detection and classification [7], [8]. In [9], probabilistic active perception is planned for realistic environments, with arbitrary object positions.

Another class of problems for perception planning is the inspection problem. In order to determine a path that can sense multiple targets, a neural network approach was used to solve the NP-hard Watchman Routing Problem. In order to do so efficiently, a fast method was proposed to answer visibility queries [10]. In that work, the solution is not optimal due to the existence of multiple targets.

PA*, a heuristic search, was proposed to solve the motion planning problem for perception of a target position in 2D gridmaps, given motion and perception costs [1].

It was also proposed in the past that robots maintain reachabilities and visibilities information, both of a robot and a human partner in a shared workspace. However, it uses non-mobile robotic platforms [11]. Visibility graphs are considered in [12], but the focus is on generating points for a patrolling motion plan. Moreover, it assumes vectorial obstacles, so visibility can easily be calculated using ray casting at the extremes of lines.

Morphological operations have also been used in robotics to determine the actuation space of a robot [13], which is later used to coordinate multi-robot teams. In our work we use a similar technique, but we can determine visibility for a sensing range bigger than robot size. We use the idea of critical points to increase the efficiency of computation, and also as a tool to improve the heuristic of PA*.

III. Background

In this section we will summarize our previous technique on informed heuristic search for optimal perception planning. We will also show how to use morphological operations in order to build approximate visibility maps. The visibility maps can then be used to obtain information that, when considered in the heuristic of perception planning, can improve the search efficiency.

A. Optimal Perception Planning

PA* is a heuristic search for motion planning that returns the optimal path to perceive a target, considering both motion and perception cost [1]. Given a path ρ, the overall cost is given by

$$\text{cost}(\rho) = \text{cost}_m(\rho) + \lambda \text{cost}_p(\rho, T) \tag{1}$$

where $\text{cost}_m(\rho)$ is proportional to the path size, and $\text{cost}_p(\rho, T)$ is a function of the minimum distance between the path and target T. We assume the perception cost is function of the minimum sensing distance from the path. The parameter λ is a weight parameter that trades-off motion and perception cost in the overall cost function.

As in the A* algorithm, the total cost estimate is given the sum of $g(S, n)$, the path distance from the starting position S to the current node n, and $h(n, T)$, a heuristic of both the motion and perception costs from n to T.

$$f(n) = g(S, n) + h(n, T) \tag{2}$$

If the heuristic used is admissible, i.e., always less or equal than the true value, then the path returned is guaranteed to be optimal. Therefore, the choice for the heuristic is based on the euclidean distance between the current node and the target, without considering any obstacles.

$$h(n, T) = \min_q \Big(||n - q|| + \lambda c_p(||q - T||) \Big) \tag{3}$$

We assume that from position n the robot can still approach the target by moving to other location q, from where it senses the target. There is a trade-off between the possible increase of motion cost, and the decrease of perception cost. We take the distance between points n and q, $||n - q||$, as the approaching cost.

We proved that if c_p is an increasing function with a maximum sensing range, then the optimal point q^* lies in the straight line between n and T, as shown in Figure 2. The perception cost function c_p can be any monotonically increasing function, allowing flexibility to represent the cost of multiple perception models.

Figure 2. Given a robot at position n and a perception target T in a scenario without obstacles, the optimal sensing goal position lies in the straight line connecting those two points.

For any specific perception cost, it is possible to find the optimal sensing position q^* as a function of the distance $||n - T||$. And with q^* known before-hand, the heuristic h is only a function of n and T, and easy to use during search.

We give two examples for the function c_p, where the perception cost is either a linear or quadratic function of the distance to the target. In our model we assume circular omnidirectional sensing, with a limited range r_p. The heuristic $h(n, T)$ can then be determined easily for each specific perception cost functions.

We define the optimal sensing distance as $d^* = ||q^* - T||$.

In the linear case, where $c_p(||q - T||) = ||q - T||$, there are two cases for the optimal distance d^*:

- $\lambda < 1$: Cost of motion is greater than cost of sensing, so robot minimizes motion by sensing form as far apart as possible (limited by maximum sensing range r_p);
- $\lambda \geq 1$: Cost of sensing is greater than cost of motion, so robot moves as close to the target as possible.

The optimal sensing distance d^* for linear perception is:

$$d^* = \begin{cases} r_p & \lambda < 1 \\ 0 & \lambda \geq 1 \end{cases} \tag{4}$$

With a quadratic sensing cost, $c_p(||q - T||) = ||q - T||^2$, we can solve for q^* using equation 3. Ideally the robot would move to a fixed distance $1/(2\lambda)$ of the target to sense it optimally. The sensing distance d^* also depends on r_p:

$$d^* = \begin{cases} \frac{1}{2\lambda} & 1/(2\lambda) \leq r_p \\ r_p & 1/(2\lambda) > r_p \end{cases} \tag{5}$$

This optimal distance d^* can now be used in the heuristic:

$$h(n, T) = \begin{cases} (||n - T|| - d^*) + \lambda c_p(d^*) & ||n - T|| \geq d^* \\ \lambda c_p(||n - T||) & ||n - T|| < d^* \end{cases} \tag{6}$$

B. Approximate Visibility Maps

We assume robots have a circular shape, and a maximum sensing range, r_p. The goal of visibility maps is to efficiently determine the observability of a robot in a certain environment, i.e., determine what regions can be sensed from any point that is reachable from the initial robot position. The algorithm is a function of robot size and sensing range. We show in Figure 3 a simulated environment with obstacles, and the Approximate Visibility Map (A-VM).

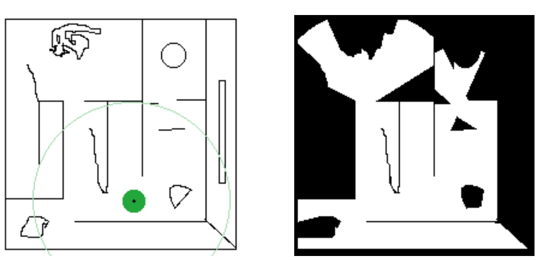

(a) Map (b) A-VM

Figure 3. Given a black and white gridmap, an omnidirectional circular robot (green), and a sensing range (green circumference), the A-VM determines what can be sensed from reachable positions.

We use morphological operations, which can be applied on images using a structuring element with a given shape. Here the structuring element R is a circle representing a circular robot. The domain is a grid of positions G. The input is a black and white binary image representing the map, with M being the set containing the positions that correspond to obstacles. The morphological operation dilation on the obstacle set M by R is

$$M \oplus R = \bigcup_{r \in R} M_r \tag{7}$$

where $M_r = \{z \in G \mid z = m + r, m \in M\}$, i.e., the translation of M by r over the grid G.

When applying the dilation operation to black points in the image, the algorithm inflates the obstacles of the map by the robot size, achieving the configuration space.

$$C^{free} = \{z \in G \mid z \notin M \oplus R\} \tag{8}$$

Given the free configuration space C^{free}, it is possible to find the points that are reachable from the initial robot position S.

$$Reach(S) = \{z \in C^{free} \mid z \text{ connected to } S\} \tag{9}$$

The partial morphological closing applies the second morphological operation only to the reachable set, $Reach(S)$, instead of C^{free}. Morphological closing is the combination of a dilation operation followed by an erosion. Dilation and erosion are dual operations. In order to apply the closing operation to the obstacles, we apply an erosion to the dilated obstacles. But being dual operations, the partial morphological closing of obstacles is equivalent to the dilation of the reachable space:

$$A(S) = Reach(S) \oplus R \tag{10}$$

The actuation space, $A(S)$, can be seen as a first approximation of the visibility map, if the maximum sensing range considered is less than the robot size (Figure 4).

From $A(S)$, it is possible to define the unreachable regions, i.e., regions that are not reachable to the robot body, and thus cannot be actuated.

$$U(S) = \{z \in G \mid z \notin A(S) \wedge z \notin M\} \tag{11}$$

$U(S)$ is then divided in a set of different disconnected components $U^l(S)$. The separation of the unreachable regions of the actuation space in disconnected parts is useful, allowing to determine visibility independently (Figure 4). Each region $U^l(S)$ has their unique openings to the actuation space, from where visibility inside $U^l(S)$ is possible. These openings are the frontiers, defined as the points of the unreachable space that connect with $A(S)$:

$$F^l(S) = \{z \in U^l(S) \mid \exists z' : z' \text{ is adjacent to } z \wedge z' \in A(S)\} \tag{12}$$

$F^l(S)$ might be composed of multiple disconnected frontiers, so $F^{li}(S)$ represents the multiple components of the frontier of region $U^l(S)$.

When determining visibility for sensing range greater than robot size, it is necessary to find points that have line of sight inside of $U^l(S)$ through some $F^{li}(S)$. There are multiple candidate points, and all of them have to be in $Reach(S)$, the feasible positions for the robot center position. All those points should be considered for the true visibility map.

Because the brute-force solution is computational expensive, we proposed an alternative, where the visibility inside unreachable regions is considered only from one point of the reachable space for each frontier $F^{li}(S)$ (see Figure 4).

As only one point is being used, the final visibility map is an approximation of the ground-truth. In order to obtain a better approximation, the point chosen has to maximize the expected visibility inside the unreachable region. That is accomplished by choosing a point close to the frontier, as being closer is equivalent to having a deeper and wider expected visibility inside $U^l(S)$, maximizing the expected visibility area. In order to find this critical point, $c^*_{li}(S)$, the sum of the distance to all frontier points is minimized:

$$c^*_{li}(S) = \underset{z \in Reach(S)}{\operatorname{argmin}} \sum_{\zeta \in F^{li}(S)} \|z - \zeta\|^2 \tag{13}$$

Ray casting is used to test visibility inside unreachable regions from the critical point, in order to consider occlusions, resulting in the Approximate Visibility Map.

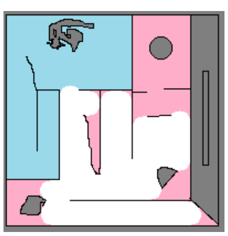

(a) Unreachable Regions

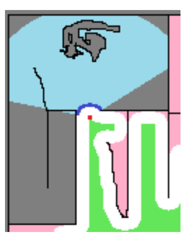

(b) Critical Point

Figure 4. In (a) we show $A(S)$ in white, in pink multiple disconnected unreachable regions, and in blue an example of one disconnected unreachable region $U^l(S)$; in (b) we highlight one disconnected region, showing in dark blue the frontier points $F^{li}(S)$, and in red the critical point $c^*_{li}(S)$; the points from the reachable set are shown in green; light blue represents the expected visibility from the critical point, through the frontier, into the unreachable region, before accounting for occlusions.

IV. Perception Planning with A-VM

In this section we show that with an initial fixed cost of building the Approximate Visibility Map (A-VM), it is possible to use the critical points from the A-VM to improve the search heuristic of PA*.

The visibility map gives information on the feasibility of perception, while not giving any information about the positions from where points can be perceived. Nevertheless, the transformation provides structured information about the environment, and it is possible to separate grid points into three categories:

1) **Reachable Space**: points that can be reached by the robot center, $Reach(S)$;
2) **Actuation Space**: points that can be "touched" by the robot body, $A(S)$;
3) **Unreachable Regions**: points the robot cannot cover with its body and motion only, because they lie in positions not traversable by the robot, $U(S)$.

Only points in the unreachable regions have some information about a possible position from where they can be sensed, because they have associated a critical point. The distance between a target in region $U^l(S)$ and a critical point $c^*_{li}(S)$ can be used as a better estimate of the perception distance in the heuristic for perception planning. Therefore, we will focus our discussion only to points that belong to the *Unreachable Regions*.

Considering the base heuristic of PA*, independently of the perception cost function, we know it is associated with the cost of moving to a better sensing position and perceiving the target from there. We also know that the heuristic does not consider obstacles, and the best sensing position lies in the straight line between the current node n and the target T. We assume we can solve the heuristic minimization problem (equation 3) for a specific cost function c_p, and find the optimal sensing distance d^*. Assuming $0 \leq d^* \leq ||n - T||$, the heuristic can be given as

$$h(n, T) = ||n - T|| - d^* + \lambda c_p(d^*) \tag{14}$$

The visibility map gives information about the minimum sensing distance from any point in the reachable space to a point in the unreachable region, which can be used in the heuristic instead of the optimal sensing distance d^* given by the straight line solution. When using the Approximate Visibility Map, and being the distance from $T \in U^l(S)$ to critical point $d^c_{li}(T) = ||T - c^*_{li}(S)||$, the heuristic becomes:

$$h^1(n, T) = ||n - T|| - d^c_l(T) + \lambda c_p(d^c_l(T)) \tag{15}$$

where $d^c_l(T) = \min_i d^c_{li}(T) \geq d^*$. Here we are not considering the possibility that d^c_l is bigger than r_p. In that case, the point would not be visible. If $||n - T|| < d^c_l(T)$, then $h^1(n, T) = \lambda c_p(d^c_l(T))$. In order to use this heuristic as admissible and guarantee an optimal path, we only need to prove the distance from T to any other point in the reachable space is bigger than $d^c_l(T)$.

Theorem 1. *Distance of points inside unreachable regions to the critical point is minimal in comparison to distance to any other point in the Reachable Space.*

Proof: We assume only one critical point and frontier, for sake of simplicity. As shown in Figure 5, we consider only the frontier extremes, the two obstacles at points $O_1(0,-\zeta)$ and $O_2(0,\zeta)$, with $\zeta < R$, being R the robot radius. The frontier is between those two obstacles. If the robot starts at some point with $x > 0$, then the unreachable region consists of points with $x \leq 0$. Following this description, the critical point results as the point that is at R distance from both obstacles, $(\sqrt{R^2-\zeta^2}, 0)$. For any point (a,b), with $a < 0$, the distance to the critical point has to be the minimum distance between (a,b) and any point in the reachable space, (α', β'), with $\alpha' > 0$. As we can see in Figure 5, for any point (α', β') there is a point (α, β) in the border of reachability that has lower distance to (a,b). And the distance between (a,b) and (α,β) is given by

$$\begin{aligned} d^2 &= (\gamma + R\cos\theta)^2 + (R\sin\theta)^2 \\ &= \gamma^2 + R^2\cos^2\theta + 2\gamma R\cos\theta + R^2\sin^2\theta \\ &= \gamma^2 + R^2 + 2\gamma R\cos\theta \end{aligned} \quad (16)$$

As we can see from the equation, the distance is minimized increasing the angle θ, and the angle θ is maximized at the critical point. Thus, we prove the distance to any unreachable point is minimized by the critical point.

■

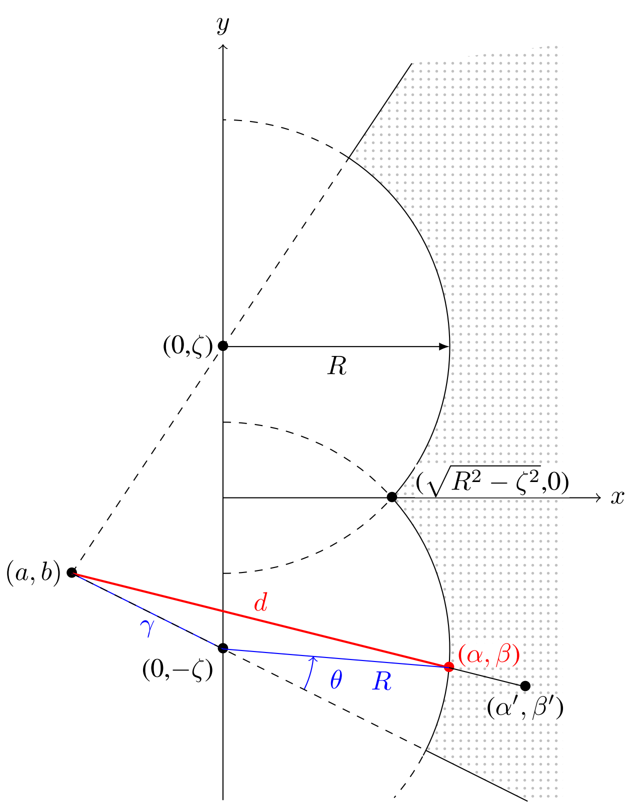

Figure 5. Given two obstacles at positions $(0,\zeta)$ and $(0,-\zeta)$, the set of points in the reachable space that can sense the point (a,b) is represented with the filled region; the critical point $(\sqrt{R^2-\zeta^2}, 0)$ is the point with the minimum distance to any (a,b) in the unreachable region.

However, it is still possible to improve the proposed heuristic. Instead of using the critical point to have only a lower bound estimate on the perception distance, we can use it to estimate a lower bound for motion cost as well.

We assume the optimal sensing distance d^* is lower than the distance to the critical point, d^c_{li}, thus, from any position with line of sight to the target, the robot will move to a point as close as possible, i.e., a border of the reachable space. Therefore, we know that the minimum motion cost will be the distance between the current node n and the closest point in the border of the Reachable Space. From Figure 5, we can see that in the worst case scenario, the distance between the critical point and any other point of the Reachable Space border, with line of sight to the target, is $2R$. Therefore, we can have a new estimate for an admissible heuristic:

$$h^2(n,T) = \min_i \left(||n - c^*_{li}(S)|| - 2R + \lambda c_p(d^c_{li}(T)) \right) \quad (17)$$

In case $d^* > d^c_{li}(T)$, we can update the heuristic to consider $\delta = \max(d^* - d^c_{li}(T), 0)$.

$$h^2(n,T) = \min_i \left(||n - c^*_{li}(S)|| - 2R - \delta + \lambda c_p(d^c_{li}(T) + \delta) \right) \quad (18)$$

Finally, the first heuristic h^1 might be a better estimate in cases there is line of sight between n and T, so in order to always use the best heuristic, we choose the one closest to the true value, considering they are both admissible.

$$h^{AVM}(n,T) = \max(h^1(n,T), h^2(n,T)) \quad (19)$$

Theorem 2. *Heuristic using Approximate Visibility Map dominates original heuristic in PA*.*

Proof: The original heuristic $h(n,T)$ in PA* is always less than the real cost, because it uses the optimal solution for the euclidean distance without any obstacles, assuming optimal motion and perception distances. The first heuristic using the A-VM, $h^1(n,T)$, uses the same straight line assumption, but replaces the perception and motion costs by better estimates. Thus the estimate $h^1(n,T)$ is always greater or equal than $h(n,T)$, because we proved the sensing distance to the critical point is the minimal perception distance. Therefore, $h(n,T) \leq h^{AVM}(n,T)$. And because $h^{AVM}(n,T)$ is admissible, it is also dominant over $h(n,T)$.

■

V. RESULTS

We tested in a simulated map the performance of our proposed algorithm with a dominant heuristic (PA-VM), against the original PA* heuristic. We compared the number of expanded nodes, presenting the results in Figure 6. We analyzed the change in efficiency with changes of the weight parameter λ, with 7 values ranging between 0.008 and 125. We set 8 different initial robot positions distributed uniformly in the reachable space, and 25 different target positions, also uniformly distributed, resulting in 1400 different search instances over the range of λ. Because the only difference is for points in $U(S)$, we only tested targets in the unreachable regions. We show in Figure 6 that for λ greater than one (small optimal sensing distances d^*), there is a great improvement in the average number of nodes expanded, with our method expanding only 35% of nodes expanded by PA*. The results depend highly on the environment topology, having a high variance. Depending on the target position, the node expansion percentage can change from almost 0 to 90%, for large λ. Even for low λ, where both heuristics have similar average results, there were some instances with a gain of 50% in the number of nodes expanded.

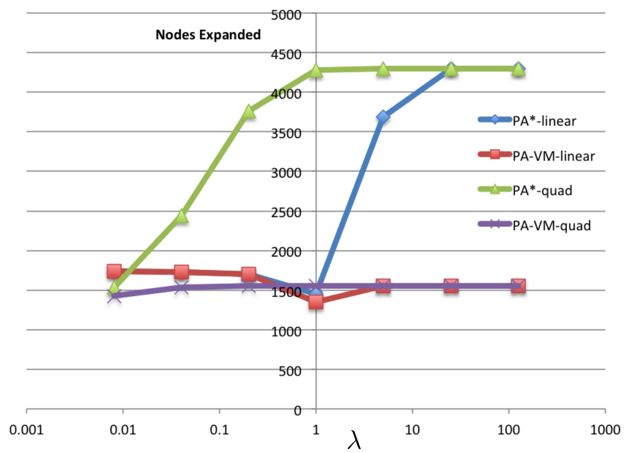

Figure 6. Comparison of average number of nodes expanded by PA* and PA* with Visibility Maps (PA-VM) as a function of λ, for both linear and quadratic perception cost function.

VI. Conclusions

We reviewed both PA* and Approximate Visibility Map algorithms. The first is an informed search method to find optimal paths for perception tasks. The latter is a map transformation that represents the observable regions in a 2D environment by a given robot. We showed how the added information about the structure of environment can be used to improve the heuristic, resulting in a reduced search, expanding less nodes. We used the critical points from the Approximate Visibility Map to create better estimates of the motion and perception costs, while proving they can be used as an admissible and dominant heuristic compared to the one proposed for PA*. In this work we considered circular robots, with motion and perception cost functions that are rotation invariant. In the future we would like to consider more complex motion and perception cost functions for any-shape non-holonomic robots.

Acknowledgment

This work is financed by the North Portugal Regional Operational Programme (NORTE 2020), under the PORTUGAL 2020 Partnership Agreement, and through the European Regional Development Fund (ERDF), within project NORTE-01-0145-FEDER-000020, and under Carnegie Mellon Portugal Program grant SFRH/BD/52158/2013.

References

[1] T. Pereira, A. P. Moreira, and M. Veloso, “PA*: Optimal path planning for perception tasks,” *Submitted to International Conference*, 2016.

[2] R. B. Rusu, I. A. Şucan, B. Gerkey, S. Chitta, M. Beetz, and L. E. Kavraki, “Real-time perception-guided motion planning for a personal robot,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2009.

[3] L. Carlone, M. Ng, J. Du, B. Bona, and M. Indri, “Rao-blackwellized particle filters multi robot slam with unknown initial correspondences and limited communication,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2010.

[4] J. Biswas and M. M. Veloso, “Localization and navigation of the cobots over long-term deployments,” *The International Journal of Robotics Research*, vol. 32, no. 14, 2013.

[5] A. Howard, L. Parker, and G. Sukhatme, “Experiments with a large heterogeneous mobile robot team: Exploration, mapping, deployment and detection,” *The International Journal of Robotics Research*, vol. 25, no. 5-6, pp. 431–447, 2006.

[6] S. Ekvall, P. Jensfelt, and D. Kragic, “Integrating active mobile robot object recognition and slam in natural environments,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2006.

[7] C. Potthast and G. S. Sukhatme, “A probabilistic framework for next best view estimation in a cluttered environment,” *Journal of Visual Communication and Image Representation*, vol. 25, no. 1, pp. 148–164, 2014.

[8] J. Velez, G. Hemann, A. Huang, I. Posner, and N. Roy, in *International Conference on Automated Planning and Scheduling (ICAPS)*, 2011.

[9] R. Eidenberger and J. Scharinger, “Active perception and scene modeling by planning with probabilistic 6D object poses,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2010.

[10] J. Faigl, “Approximate solution of the multiple watchman routes problem with restricted visibility range.” *IEEE Transactions on Neural Networks*, vol. 21, no. 10, 2010.

[11] A. K. Pandey and R. Alami, “Mightability maps: A perceptual level decisional framework for co-operative and competitive human-robot interaction,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2010.

[12] S. M. LaValle, *Planning algorithms*. Cambridge university press, 2006.

[13] T. Pereira, M. Veloso, and A. Moreira, “Multi-robot planning using robot-dependent reachability maps,” in *Robot 2015: Second Iberian Robotics Conference*, 2015.