Uncertainty-based mapping and planning strategies for safe navigation of robots with stereo vision

Martim Brandao*, Kenji Hashimoto† and Atsuo Takanishi‡§

*Grad. School of Adv. Science and Engineering
Waseda University, Tokyo, Japan
Email: contact@takanishi.mech.waseda.ac.jp

†Faculty of Science and Engineering
Waseda University, Tokyo, Japan

‡Department of Modern Mechanical Engineering
Waseda University, Tokyo, Japan

§Humanoid Robotics Institute (HRI)
Waseda University, Tokyo, Japan

***Abstract*—We describe our recent developments in probabilistic modeling of 3D reconstruction with stereo vision, applied to planning strategies for locomotion and gaze. We first overview the use of probabilistic occupancy grids for 3D reconstruction, and the sensor models of stereo best suited to the problem. These grids are then used for robot navigation, which is tackled at two levels: 1) At the locomotion level, trajectories are computed from the grid using an A* search algorithm that minimizes the total probability of occupancy over the trajectory. 2) At the grid level, we propose two task-relevant active strategies which redirect the sensor to "maximum visible entropy" and "maximum visible occupancy" points along the planned locomotion trajectories. Steps 1) and 2) are executed alternately until the locomotion trajectory converges to a high certainty, safe solution.**

Results of the proposed gaze and locomotion planning strategies were obtained on simulated scenarios and a real robot. Estimates of the uncertainty that occupancy grids are subjected to in real outdoor scenarios were computed for different stereo sensor models. These estimates were used in active gaze simulations for an extensive comparison of gaze strategies across 400 randomly generated environments. The results show that careful modeling of stereo vision sensor uncertainty and the proposed task-relevant planning strategies lead to more complete and consequently collision-free reconstructions of the environment along planned robot trajectories.

I. Introduction

In order to navigate environments safely, robots require accurate models of their sensors and reliable strategies for planning locomotion and gaze trajectories. Gaze, or sensor orientation, can be especially useful when the robot's sensors have a narrow field of view such as stereo sensors (i.e. two cameras). In such cases, only a part of the space around the robot can be sensed at each time and so measures must be taken in order to reduce the uncertainty of the environment and probability of collision. Researchers have, however, focused mainly on active gaze strategies for the localization [1], [2] and exploration [3], [4], [5], [6] tasks, which could lead to poor safety of planned trajectories. Robots with complex dynamics such as humanoids robots require special care during planning since they can be slower to respond to a sudden obstacle and unstable if collision occurs. We argue that if safety is to be prioritized, then gaze should be guided to points along the current trajectory plan until the probability of collision is negligible [7]. Other recent work focusing on active gaze for obstacle avoidance include autonomous policy learning methods [8] and utility theory-based planning [9]. However, either sparse environment representations or limited sensor models were used, which can also lead to unsafe trajectories being generated.

*Corresponding author

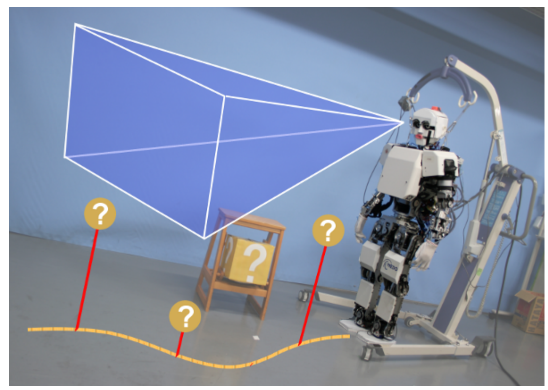

Fig. 1. The KOBIAN humanoid robot faces an obstacle which is not sensed due to narrow field of view. How should the uncertainty of the current measurement be modeled? When walking up to its goal, where should the robot gaze to?

Both reliable environment representations and sensor models are required so that uncertainty of the environment can be estimated and exploited for guiding gaze. In this paper we overview the use of probabilistic occupancy grids for mapping (i.e. 3D reconstruction) of the environment, as well as the sensor models of stereo that can be used with such reconstruction approach. This is discussed in Section II. Occupancy grids [10] are not only useful for robot trajectory planning, but also provide an environment uncertainty measure in a straightforward way. We then introduce locomotion planning and active gaze strategies that focus on safety of the robot by minimizing collision probability (Section III). The methods were evaluated on both real data and simulated environments. Experiments with occupancy grids are discussed in Section IV, while active gaze experiments are reported in Section V and VI.

II. Uncertainty-based mapping with stereo and grids

Consider a three-dimensional grid of cells which can be in one of two states: occupied O or free $\overline{O}$. The objective of an occupancy grid algorithm is to compute or update the probabilities $p(O_i|e_{0...t}, x_{0...t})$ for each cell i, at each time instant t, given measurements $e_{0...t}$ and sensor locations $x_{0...t}$ until time t. This is implemented as a Bayes filter at each cell, which updates occupancy probabilities every time a new measurement is taken [11].

Uncertainty of the binary random variables O_i can be measured using Shannon entropy:

$$h_i = -p_i log(p_i) - q_i log(q_i) \tag{1}$$

where $p_i = p(O_i|e_{0...t}, x_{0...t})$ is the probability of occupation of that cell and $q_i = 1 - p_i$.

A measurement e_t can consist, for example, of a set of rays from a laser rangefinder. Each ray is in turn associated to a distance of the nearest object. In the case of stereo vision e_t corresponds also to a set of rays, but each ray is associated to a pixel in one image and a cost function of distance. Probability of occupancy at each distance can be computed directly from these cost functions, as described next.

A. Occupancy grid formulation

We recently proposed a novel occupancy grid framework which integrates all information returned by stereo vision measurements into occupancy grids [12]. Briefly, we proposed to compute occupancy of a cell i as

$$p(O_i|E) = p(O_i|EV_i)p(V_i|E) + p(O_i|E\overline{V}_i)(1 - p(V_i|E)), \tag{2}$$

where E represents a cost-function of distance taken along the ray which intersects cell i. $V_i = \overline{O}_{i-1}...\overline{O}_2, \overline{O}_1$ represents visibility of cell i. Under the assumption that O_i and the measurement E are conditionally independent on invisible cells $\overline{V}_i$, we have $p(O_i|E\overline{V}_i) = p(O_i|\overline{V}_i)$. For sake of readability and compactness, the equations shown here are for a one-dimensional grid aligned with the sensor - correspondent to the intersection of a camera ray with the three-dimensional grid. Bresenham's line algorithm [13] in 3D can be used to efficiently compute the set of discrete cells along the ray from cell i to the sensor origin. In the original publication we also showed that

$$p(V_i|E) = \prod_{j=0...i-1} p(\overline{O}_j|EV_j), \tag{3}$$

where $p(\overline{O}_j|EV_j) = 1 - p(E|O_iV_i)$ holds under certain assumptions. Equation (2) is then reduced to the computation of the following models (and their complements):

- $p(E|O_iV_i)$. The probability of measuring costs E given cell i is both occupied and visible.
- $p(O_i|V_i)$. A prior on the environment geometry. We empirically set the prior equal to 0.5 in our experiments so that occupied and free cells are equally probable.

B. Stereo vision

Consider two images I_1 and I_2, aligned along the x axis. In stereo vision, the cost-curve $E(d)$ of assigning $I_2(x, y)$ to $I_1(x+d, y)$ is computed at each pixel (x, y) and d is called the disparity. In perfect conditions (i.e. no image noise, occlusion, discontinuity or sampling problems), the cost-curve $E(d)$ has its minimum d_{min} at true disparity d^*.

The conditional probability $p(E(d)|d = d^*)$, or alternatively $p(E(d)|O_iV_i)$ in grid terms, is called the direct sensor model and can be formulated in either a winner-take-all or whole-cost-curve way.

1) Whole-cost-curve (WCC) model: The conditional probability function of measuring $E(d)$ at true disparity d^* can be defined assuming a normal distribution of costs [14],

$$p(E(d)|d = d^*) \propto e^{-\frac{E(d)}{2\sigma_{px}^2}}, \tag{4}$$

where σ_{px}^2 represents the variance of pixel intensity noise. Although out of scope of this paper, different confidence measures also exist to compute $p(E(d)|d = d^*)$. For a thorough review please refer to [15]. In that review, (4) ranks within the highest confidence measures considering the whole cost-curve.

2) Winner-take-all (WTA) model: The WTA model is arguably the most popular one in stereo, although it originates from laser rangefinder sensors that measure distances to targets instead of costs for each possible distance. This model thus depends only on the least-cost disparity d_{min}. Depending on the literature, the shape of $E(d)$ around d_{min} (i.e. the curvature of the minimum) is used as a measure of uncertainty. This approximate model is given by

$$p(E(d)|d = d^*) \sim \mathcal{N}(d_{min}, \sigma_d^2), \tag{5}$$

where σ_d^2 represents the variance of a disparity measurement, which is either fixed or could be a function of the curvature at $E(d_{min})$.

III. Safety-centered planning with uncertainty

A. Robot trajectory planning

The trajectory of a robot from its initial state to a target state can be computed from the occupancy grid such as to minimize occupancy probability (i.e. collision probability) along this path. Such an approach focuses on safety of the robot. To reduce the complexity of the problem, a 3D grid can be projected into a 2D top-view map where each cell's value corresponds to the maximum probability of occupation along the vertical axis. In this work we opted for an A* approach [16] to the search problem. We chose this method for its simplicity, although other more efficient approaches could be used. We use a set of predefined robot motions to build the search graph, adjusted to the motion capabilities and limitations of our robot such as maximum turning angle. The cost

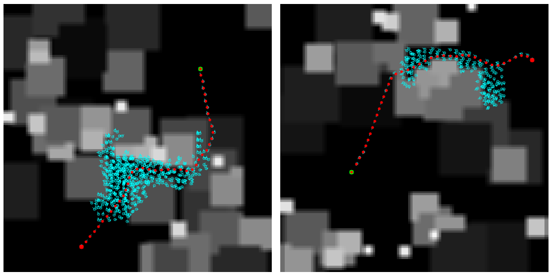

Fig. 2. Two examples of simulated scenarios with regions of different occupancy probability. The brighter the pixel the higher the occupancy probability. Trajectory nodes explored (closed list of the A* algorithm) are in cyan and final solution in red. With this approach we look for minimum cost trajectories preferring regions with low occupation probability.

associated to a certain motion is set as to grow exponentially with occupancy probability of the cells it traverses. The Euler distance multiplied by the minimum motion cost (occupancy $P = 0$) was used as a heuristic for the cost to the goal. The result is exemplified in Figure 2.

For planning purposes it is useful to consider the robot as a point in the grid and hence obstacles are dilated according to the robot's dimensions. Here we keep the grid probabilistic, without classifying cells into occupied or free. Therefore the grid can be dilated by taking for each grid cell the maximum of occupancy probability in the robot's area around that cell.

B. Safety-centered active gaze: next-best-view strategies

The occupancy grid framework can also be used to guide an active policy for reducing uncertainty of the robot trajectory. We propose to use next-best-view strategies to lower the chance of collision between robot and environment, thus increasing the robot's safety. The purpose of these strategies is to

$$minimize\left(1-\prod_{k\in K}p(\overline{O}_k|e_{0\ldots t+\Delta t},x_{0\ldots t+\Delta t})\right), \quad (6)$$

where K is the set of cells intersected by the robot trajectory and Δt represents the time it takes for the robot to execute the gaze command. Robots with controllable gaze direction can thus actively use it such that the expected value of collision probability is lowered. Empirically, occluded regions should not be gazed at because they will not lead to new measurements on the occluded area. Priority should go to visible cells of higher uncertainty or occupancy probability, so that this information can be confirmed or denied: and the robot trajectory plan consequently adjusted. In this paper we will focus on the following greedy formulations of the problem.

1) Maximum visible entropy: Gazing point will be cell g, such that

$$g = \underset{k\in K}{argmax}\ p(V_k|e_{0\ldots t+1},x_{0\ldots t})h_k. \quad (7)$$

2) Maximum visible occupancy: Gazing point will be cell g, such that

$$g = \underset{k\in K}{argmax}\ p(O_kV_k|e_{0\ldots t},x_{0\ldots t}). \quad (8)$$

Both formulations are greedy in the sense that they attempt to minimize uncertainty only of the points along the current trajectory which are most likely to contribute to a safer trajectory after measurement and re-planning. These are thus purely exploitation strategies, where in turn exploration is guided by robot trajectory planning itself. When a new plan leads to unexplored cells being traversed, their uncertainty or occupancy probability will be high and thus gazed at by the greedy gaze strategies. The process can go on until no new gaze points are generated (local minimum), or until a certain trajectory safety threshold is reached.

In the experimental section we will also evaluate simplified versions of (7) and (8), where visibility is not considered (i.e. all cells are considered visible). We call these two strategies "maximum entropy" and "maximum occupancy" respectively. The former was originally introduced by us in [7].

IV. Occupancy grid experiment: the WTA and WCC models

We first report on an experimental evaluation of the occupancy grid method. For this evaluation we used the publicly available KITTI dataset of outdoor stereo images [17]. The real-world, noisy stereo measurements and localization data in this dataset provide a challenging scenario for 3D reconstruction.

To measure the performance of our occupancy grid method we used the precision and recall ratios of the grid's cells. Precision is defined as $\frac{tp}{tp+fp}$, where tp (true positives) refers to the number of cells correctly classified as occupied (i.e. occupancy $P > 0.5$) and fp (false positives) refers to the number of cells incorrectly classified as occupied. On the other hand, recall is defined as $\frac{tp}{n}$, where n refers to the total number of occupied cells on ground-truth data.

We ran the probabilistic occupancy grid algorithm on the KITTI residential area dataset "2011_09_26_drive_0079" using a grid of cubic cells with dimensions 0.20x0.20x0.20 meters. In Figure 3 we show the reconstruction performance along time when both the WTA and WCC stereo models were used. The data shown were obtained using a Sum of Squared Differences (SSD) with 9x9 window size as the cost function $E(d)$. The whole-cost-curve model of stereo achieved higher performance at the cost of lower recall. We also point to the ascending precision which is observed as more stereo measurements are taken, which shows that the occupancy grid method can correctly fuse the noisy uncertain measurements into a coherent and precise representation. Other promising results were also recently reported in [12].

The initial precision and recall ratios will be used in our active gaze simulation experiment in Section VI to model the noise our method is subjected to on real environments.

V. Active gaze experiment 1 - proof of concept on robotic platform

As we discussed in Section III-B it is important to actively control gaze (i.e. sensor) orientation in order to decrease uncertainty on an occupancy grid, especially so along the

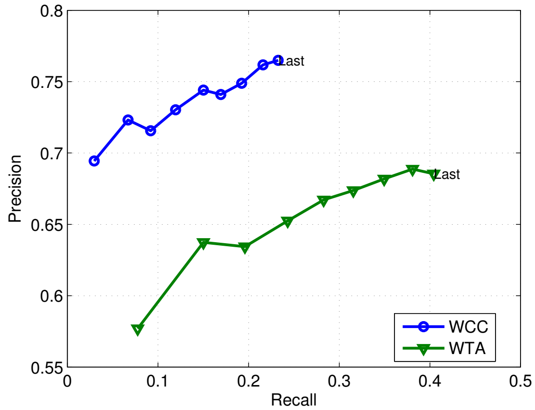

Fig. 3. Precision and recall of the obtained occupancy grid along time, for both a WTA and WCC stereo model. Each marker represents the update of the grid after a new stereo measurement. The last measurement is indicated with the word ”Last”.

planned robot trajectory. Our first active gaze experiment was originally presented in [7] and involves testing the introduced algorithms on a real robotic platform. The environment was carefully designed such that several gaze actions would be required in order to reconstruct all obstacles present between the robot and target. A naive gaze strategy such as target fixation would lead to certain obstacles being unnoticed and to a collision, were the planned trajectory to be executed.

A. Experimental setup

We tested the occupancy grid method and the ”maximum entropy” gaze strategy on the biped humanoid robot platform KOBIAN [18]. KOBIAN is 1.4m tall, weighs 62kg and has a total 48 DoF. The vision system uses two CMOS cameras working at a 30Hz acquisition rate. Camera images were used at a 320x240 pixel resolution.

In this experiment the dimensions of the occupancy grid's cells were set according to the physical dimensions of the robot. Having in mind the average step size of the robot, we used cubic cells of 0.15x0.15x0.15 meters. Occupied cells were dilated taking into account the robot's dimensions (approximated as a 2D square of 0.60 meters each side) so that the robot can be represented as a single cell in the grid. The WTA stereo model was chosen for slightly faster computation. For a comparison of computation times please refer to [12].

The tested scenario is as follows: the robot stands in a room looking forward, having a target where it has to walk to, fixed in the world (3m ahead, 2m to the left). Between the robot and the target, some common obstacles such as chairs were placed. The proposed algorithm is started once the robot is on the floor, successfully generating gaze targets along the trajectory. All was generated online and automatically without human intervention. The duration of the experiment was approximately 1 minute.

B. Results

Figure 4 shows the camera image sequence and occupancy grid results for this experiment. Even though certain obstacles were not initially visible, the robot successfully managed to find a safe (i.e. collision free) trajectory after 3 gaze actions. As seen in the figure, gaze targets were lower than the starting one. This is due to the narrow field of view of the stereo sensor that leads to regions closer to the floor not being sensed in the beginning (the robot looks straight ahead). A constant gaze at the target, however, which is a common strategy in robot navigation through visual servoing [19], would lead to unnoticed obstacles being ignored and a thus a high chance of collision, as we reported in [7].

VI. Active gaze experiment 2 - extensive evaluation in simulation

On a second and new experiment we developed a simulation system so that an extensive and significant evaluation of the active gaze strategies could be taken across a large number of random environments.

A. Experimental setup

Our simulation system consists of a computer program that simulates noisy sensor measurements and perfect gaze actions that instantly update the sensor orientation. The following functionalities were implemented.

1) Randomly simulated environments: A total of 400 different random environments was generated for evaluation. Robot starting and target points were set constant and 15 meters apart. At each environment, 100 random squared obstacles are generated with the constraint of minimum distance of 1 meter to both the robot starting and target points. Obstacles have a random height of up to 0.75m and random side of up to 3 meters. Two of the generated environments are shown in Figure 5 with color-encoded obstacle height.

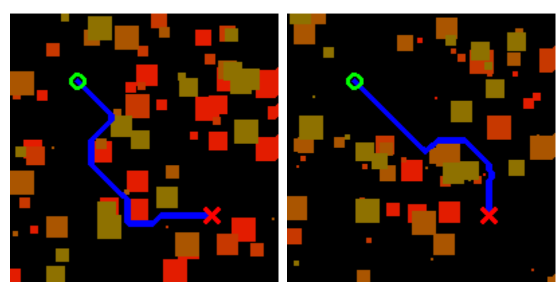

Fig. 5. Top-view of the random environments number 0 and 1, out of the set of 400 used for evaluation. Squares represent obstacles that are higher as their color gets closer to red (red is robot height, green is 0 height). Trajectory solution is shown in blue, the green circle represents the robot and the cross represents the target.

2) Gaze strategies: We implemented a total of 6 gaze strategies, which were run on the same scenarios. Besides the strategies defined in equations (7) (8) and their unconstrained visibility versions, we implemented a ”random gaze” strategy and a ”fixate target” strategy. Random gaze generates a random

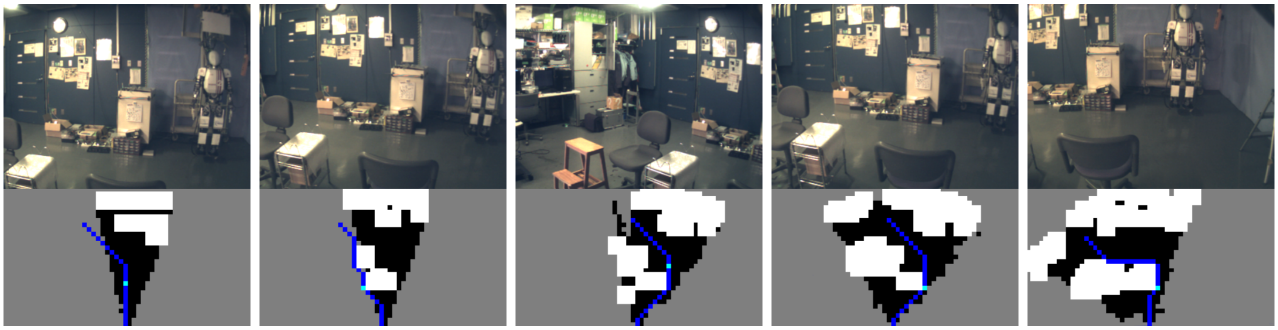

Fig. 4. Results of the "maximum entropy" gaze strategy. Top: Right camera image; Bottom: 2) Occupancy grid (projected to 2D and dilated to robot dimensions). Egocentric representation: vertical direction in the image corresponds to current sensor direction. Black is probability of occupation P=0 and white P=1, generated trajectory blue, and gaze target cyan point. From left to right: frames 112, 208, 310, 417, 523.

sensor orientation within physical limits at each instant of time, while the target strategy simply keeps the robot gaze at the target throughout the whole experiment.

3) Measurement model: We implemented a measurement model that updates the occupancy grid after each (simulated) gaze command is executed. The grid is updated at all cells lying inside the sensor field of view (FOV). In these experiments, we selected a narrow FOV of 60 and 40 degrees in the horizontal and vertical axis respectively. At each cell i, the sensor measures either $p(O_i|E) = 0.4$ or $p(O_i|E) = 0.6$ depending on whether a "free" or "occupied" contribution is measured, respectively. The sensor is assumed to provide noisy measurements that lead to a false-negative ratio fn and false-positive ratio fp of cell classification. In other words, occupied cells receive an occupancy probability $p(O_i|E) = 0.4$ with fn chance, while free cells receive $p(O_i|E) = 0.6$ with fp chance. Visible cells will get a correct measurement with $(1-fp)$ and $(1-fn)$ rates. Visibility is taken from the ground-truth map (i.e. the randomly generated map) by line drawing using Bresenham's line algorithm. We selected $fp = 0.01$ and $fn = 0.77$, which correspond to the values of precision and recall obtained on the real outdoors dataset, reported in Section IV.

B. Results

Each gaze strategy affects the safety of the trajectory differently along time. For each strategy we computed the safety $s(t)$ of obtained trajectories as the average probability $p(\overline{O}_k|e_{0\ldots t}, x_{0\ldots t})$ for cells k along that trajectory. In Figure 6 we show the average safety $\overline{s}(t)$ of generated trajectories, averaged over all random maps for each gaze strategy. We also show the actual probability of collision over time, obtained by counting at each instant of time the number of maps where the trajectory collides with obstacles in the ground-truth map. The data clearly show that on average the active gaze strategies here discussed lead to robot trajectories that are safer than a random gaze strategy at all times, and safer than a "gaze at target" strategy after less than 5 gaze actions. The figure also shows that considering visibility of cells during gaze selection slightly improves $s(t)$.

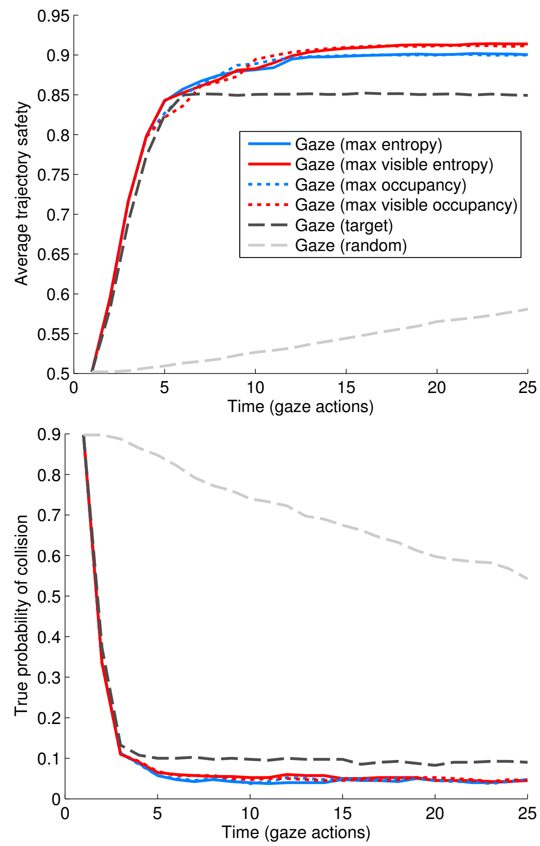

Fig. 6. Comparison of the several active gaze strategies. Top: Average safety $\overline{s}(t)$ of the generated trajectories along time. Maximum safety is $\overline{s}(t) = 1$. Bottom: True probability of collision of the trajectories, computed from ground-truth. Values are averaged across the whole set of random environments at each instant of time.

We finally show, in Figure 7, results of the simulated grid algorithm and planned trajectory for the "maximum visible occupancy" strategy on an example random environment. In this experiment, the average safety $s(t)$ of the generated trajectory goes from 0.5 at the initial condition to 0.63, 0.77, 0.82, 0.88 and finally 0.93 at 5th gaze action. Similarly to the proof of concept experiment in Section V, gazing actions

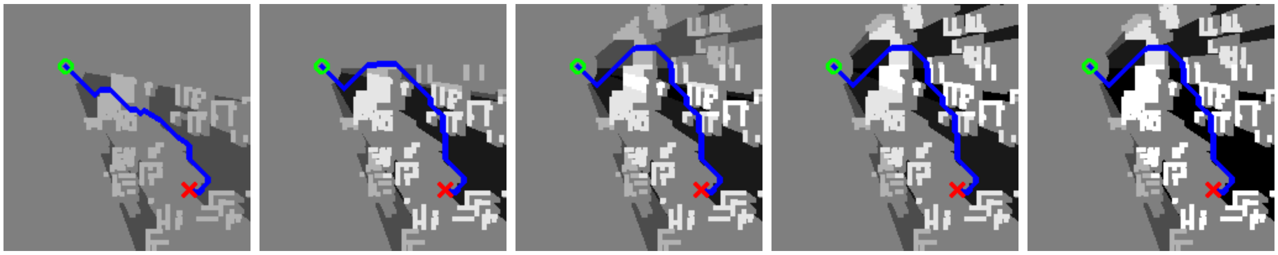

Fig. 7. Grid and robot trajectories at each instant of time, when the "maximum visible occupancy" gaze strategy is used. Random environment number 0. The brighter the pixel the higher the occupancy probability. Trajectory solution is shown in blue, the green circle represents the robot and the cross represents the target.

successively lead to new obstacles being found and trajectories re-planned until a safe path is obtained. A safe collision-free path is obtained by the end of the 3rd gaze action.

VII. Conclusion

In this paper we introduced a set of methods that deal with stereo sensor uncertainty during robot navigation, with a special focus on robot trajectory safety. We proposed methods for stereo mapping (winner-take-all and whole-cost-curve occupancy grids) and methods for planning gaze and locomotion.

We obtained estimates of the uncertainty that occupancy grids are subjected to in real outdoor scenarios and showed that WCC stereo sensor models can lead to higher precision maps than WTA, at the cost of lower recall. We empirically showed, on a real robot, that gazing at maximum entropy points along the planned robot trajectories is an efficient way to exploit occupancy grids to increase safety and confidence on a trajectory. We also validated this observation with extensive simulation on random environments. Several gaze strategies were compared, of which the "maximum visible entropy" and "maximum visible occupancy" scored best on average. With these strategies, highest safety trajectories were obtained in less than 5 gazing actions for a narrow sensor of 60 by 40 degrees field-of-view.

Acknowledgment

This study was conducted as part of the Research Institute for Science and Engineering, Waseda University, and Humanoid Robotics Institute, Waseda University. It was also supported in part by JSPS KAKENHI (Grant Number: 24360099 and 25220005), and Strategic Young Researcher Overseas Visits Program for Accelerating Brain Circulation, JSPS, Japan.

References

[1] A. J. Davison and D. W. Murray, "Mobile robot localisation using active vision," in *Proceedings of the 5th European Conference on Computer Vision - Volume II*, ser. ECCV '98. London, UK: Springer-Verlag, 1998, pp. 809–825.

[2] G. Lidoris, K. Kuhnlenz, D. Wollherr, and M. Buss, "Information-based gaze direction planning algorithm for slam," in *2006 6th IEEE-RAS International Conference on Humanoid Robots*, 2006, pp. 302–307.

[3] R. Sim and N. Roy, "Global a-optimal robot exploration in slam," in *2005 IEEE International Conference on Robotics and Automation*, 2005, pp. 661–666.

[4] M. Strand and R. Dillmann, "Using an attributed 2d-grid for next-best-view planning on 3d environment data for an autonomous robot," in *International Conference on Information and Automation*, 2008, pp. 314–319.

[5] B. Yamauchi, "A frontier-based approach for autonomous exploration," in *1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation*, 1997, pp. 146–151.

[6] R. Shade and P. Newman, "Choosing where to go: Complete 3D exploration with stereo," *2011 IEEE International Conference on Robotics and Automation*, pp. 2806–2811, May 2011.

[7] M. Brandao, R. Ferreira, K. Hashimoto, J. Santos-Victor, and A. Takanishi, "Active gaze strategy for reducing map uncertainty along a path," in *3rd IFToMM International Symposium on Robotics and Mechatronics*. IFToMM, 2013, pp. 455–466.

[8] M. Suzuki, T. Gritti, and D. Floreano, "Active vision for goal-oriented humanoid robot walking," in *Creating Brain-Like Intelligence*, ser. Lecture Notes in Computer Science, B. Sendhoff, E. Krner, O. Sporns, H. Ritter, and K. Doya, Eds. Springer Berlin Heidelberg, 2009, vol. 5436, pp. 303–313.

[9] J. Seara and G. Schmidt, "Intelligent gaze control for vision-guided humanoid walking: methodological aspects," *Robotics and Autonomous Systems*, vol. 48, no. 4, pp. 231 – 248, 2004.

[10] A. Elfes, "Sonar-based real-world mapping and navigation," vol. 3, no. 3, pp. 249–265, 1987.

[11] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*. The MIT Press, 2005.

[12] M. Brandao, R. Ferreira, K. Hashimoto, J. Santos-Victor, and A. Takanishi, "Integrating the whole cost-curve of stereo into occupancy grids," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 4681–4686.

[13] J. E. Bresenham, "Algorithm for computer control of a digital plotter," *IBM Systems Journal*, vol. 4, no. 1, pp. 25–30, 1965.

[14] L. Matthies and M. Okutomi, "A Bayesian foundation for active stereo vision," *Proc. SPIE Sensor Fusion II: Human and Machine Strategies*, pp. 1–13, 1989.

[15] X. Hu and P. Mordohai, "A Quantitative Evaluation of Confidence Measures for Stereo Vision," vol. 34, no. 11, pp. 2121–2133, 2012.

[16] P. Hart, N. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," *IEEE Transactions on Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.

[17] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2012, pp. 3354–3361.

[18] N. Endo, S. Momoki, M. Zecca, M. Saito, Y. Mizoguchi, K. Itoh, and A. Takanishi, "Development of whole-body emotion expression humanoid robot," in *2008 IEEE International Conference on Robotics and Automation*, 2008, pp. 2140–2145.

[19] M. Brandao, L. Jamone, P. Kryczka, N. Endo, K. Hashimoto, and A. Takanishi, "Reaching for the unreachable: integration of locomotion and whole-body movements for extended visually guided reaching," in *13th IEEE-RAS International Conference on Humanoid Robots*. IEEE, 2013, pp. 28–33.