Carnegie Mellon University Research Showcase @ CMU

Robotics Institute

School of Computer Science

2006

Opportunistic Use of Vision to Push Back the Path-Planning Horizon

Bart Nabbe Intel Research

Derek Hoiem Carnegie Mellon University

Alexei A. Efros Carnegie Mellon University

Martial Hebert Carnegie Mellon University

Follow this and additional works at: http://repository.cmu.edu/robotics



Part of the Robotics Commons

Published In

Proceedings: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS '06), 2388-2393.

This Conference Proceeding is brought to you for free and open access by the School of Computer Science at Research Showcase @ CMU. It has been accepted for inclusion in Robotics Institute by an authorized administrator of Research Showcase @ CMU. For more information, please contact research-showcase@andrew.cmu.edu.

Opportunistic Use of Vision to Push Back the Path-Planning Horizon

Bart Nabbe, Derek Hoiem, Alexei A.A. Efros and Martial Hebert

Intel Research Pittsburgh 4720 Forbes Avenue, Suite 410 Pittsburgh, PA, 15213, USA bart.c.nabbe@intel.com The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA, 15213, USA
{dhoiem,efros,hebert}@cs.cmu.edu

Abstract—Mobile robots need maps or other forms of geometric information about the environment to navigate. The mobility sensors (LADAR, stereo, etc.) on these robotic vehicles can however populate these maps only up to a distance of a few tens of meters. A navigation system has no knowledge about the world beyond this sensing horizon. As a result, path planners that rely only on this knowledge are unable to anticipate obstacles sufficiently early and have no choice but to resort to an inefficient local obstacle avoidance behavior.

However, recent developments in the computer vision community allows us to collect geometric information about the environment far beyond this sensing horizon. The coarse 3D geometric estimation that can be recovered is derived from an appearance-based model. That uses a multiple-hypothesis framework to robustly estimate scene structure from a single image and estimating confidences for each geometric label. This 3D geometric estimation is used with a previously presented navigation strategy that reasons about sensor constraints and plans for measurements while navigating towards the goal.

The validity of the sensing method and navigation strategy is supported by results from simulations as well as field experiments with a real robotic platform. These results also show that significant reduction in path length can be achieved by using this framework.

I. INTRODUCTION

Mobile autonomous systems use some form of geometric representation of the environment to plan a path to their destination. This information is often represented as a 2.5D map of the environment. When such a mobile robot is to navigate in an unstructured outdoor environment, it can not rely on the availability of maps. The system will need to autonomously assess the drivability of the terrain.

Typical mobility sensors, such as Laser Radar (LADAR) or passive stereo vision are used to build the required geometric representation of the terrain; however, these systems will only acquire data up to a few tens of meters to 100 meters. As a result, the planner has no knowledge about what to encounter beyond the sensed perimeter and is thus unable to anticipate obstacles sufficiently early and has no choice but to plan paths close to obstacle boundaries. This effect is known as the myopic planning effect (Figure 1). We have shown in previous work [14] that we can alleviate this problem by using midrange sensing.

Unfortunately, these sensing methods come with their own constraints. These may include increased computation time,

which prevents continuous acquisition of data; narrow field of view; sparse data; constraints on data acquisition procedure (for example, structure from motion (SFM) techniques require that video data be acquired over a sufficiently long traverse in order to establish a long enough baseline [1], [2]).

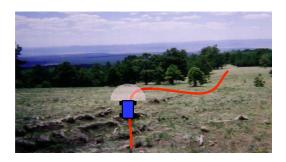




Fig. 1. Top: poor performance due to lack of sensor planning and midrange sensing, the robot has to rely on mobility sensor data only. Bottom: the pinch point is detected earlier, which allows the planner to adjust the path accordingly.

Current developments in the computer vision community [3], [4] has given us another method of obtaining midrange sensing data. This technique allows us to estimate the geometric structure of the environment from a single image. In the remainder of this paper, we will show how this novel reconstruction method can be used as an acquisition technique for mid-range data. We have extended our previous work in the area of mid-range sensing and path planning to leverage this new reconstruction technique to plan paths beyond the typical sensing horizon.

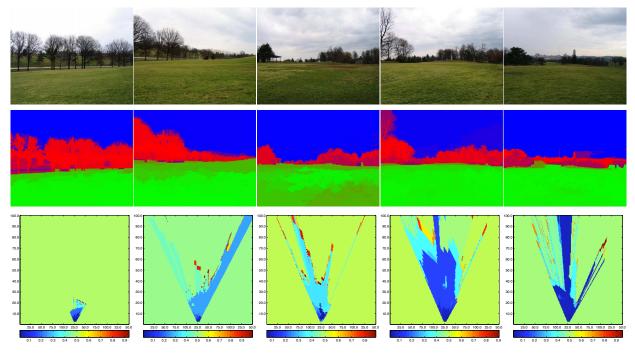


Fig. 2. Geometric labelling and reconstruction for example unstructured environments. Top: original image, center row, the labelled classes "ground" (green), "vertical" (red) and "sky" (blue) a darker shade represents a lower confidence, distances are in meters. The bottom row shows the corresponding navigation map with 0 denoting traversable and 1 representing an obstacle.

II. MID-RANGE SENSING FROM SINGLE-VIEW GEOMETRY

In previous work [3], we demonstrated the ability to estimate 3D scene geometry from a single outdoor image. Although it is generally mathematically impossible to recover 3D information from a single 2D projection, we are able to succeed by taking advantage of the natural structure in our world. Most scenes can be characterized simply by the ground plane, the things that stick out of the ground (usually at right angles due to gravity), and the sky. After *learning* the structure of our world through appearance-based models of geometry, we are able to estimate the coarse geometry of the scene given a single input image.

Our models are learned from a diverse set of images, obtained from Google image search, that are representative of outdoor scenes. We learn how to segment the image into geometric classes ("ground", "vertical", and "sky") and train classifiers to recognize each of these classes given a segmented image. After performing an oversegmentation [5] into superpixels, we group the superpixels according to the probability that pairs of superpixels belong to the same geometric class. We then estimate the likelihood that each segment is "good" (i.e., that it is entirely composed of *one* geometric label) and the likelihood of each possible geometric label, given that the segment is good. These likelihoods are learned by boosted decision trees [6] using a logistic regression version of Adaboost [7] that provides probabilistic estimates. The robustness of our algorithm results from using a large and varied set of image cues, including color, texture, location, perspective cues, and from combining estimates from multiple segmentations to get the final estimate of geometry.

On a test set of 250 images, we obtain an average pixel-wise accuracy of 86% when labeling the image into "ground", "vertical", and "sky". More results and details on the method are presented in [3]. Figure 2 shows an example of the geometric information provided by our algorithm, with probabilistic confidences displayed in the RGB channels of the output image. MATLAB code for this algorithm can be downloaded from:

http://www.cs.cmu.edu/~dhoiem/projects/popup

By making the assumption that the terrain local to the robot is planar we can use the recovered geometric information to produce a coarse estimate of the geometry of the environment. This process is fairly straightforward if the intrinsic and extrinsic parameters (with respect to the robot coordinate system) are known. With the camera parameters known we can reproject the geometric labels back into a navigation map. For all ground labels with a confidence greater than 0.5 a ray of free space is traced from the camera origin until the first vertical geometric label is encountered. The remainder of the cells that belong to this vertical segment are projected onto a vertical surface which gives us the height of the obstacle. No assumptions are made about cells that are occluded by an obstacle.

For each location in the map we can find the geometric label in the image such that we can mark cells traversable if they belong to the ground class and obstacles are constructed at the boundary of the ground/vertical intersection by projecting the remainder of the vertical segment on the vertical axis.

III. PLANNING VIEW-POINTS

Since the camera has a limited and possibly obstructed field of view it is necessary to reason about from which position and in which direction we should take our next measurement. In addition we would like to limit ourselves to only those measurements that are strictly necessary for navigation. We would therefore like to minimize the required amount of measurements since it is fairly expensive to compute geometric information from images (approximately a minute).

Let us formulate this notion of "useful measurement". Consider a candidate observation position \mathbf{u} in our map. At this position, we can point our sensor at an angle θ . Given the field of view and minimum and maximum sensing distance of the sensor, we can define a sensor footprint. In addition we define \mathcal{L} as the set of all possible binary navigation maps. Each labelling of a binary navigation map $(L_i \in \mathcal{L})$ has an

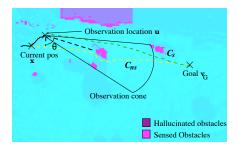


Fig. 3. Illustration of the path computed with (C_s) and without sensing (C_{ns}) in an example world configuration.

associated probability $P(L_j)$. Given such a labelling and the map cells the camera can see, one can compute the cost of getting from the current robot location $\mathbf x$ to the goal $\mathbf v_G$ by not making the observation at $\mathbf u$: $C_{ns}(L_j,\mathbf x,\mathbf v_G)$, and the cost of getting to the goal making the observation: $C_s(L_j,\mathbf x,\mathbf u,\theta,\mathbf v_G)$ (Figure 3). Intuitively, $C_{ns}(L_j,\mathbf x,\mathbf v_G)$ is the cost of the path executed if no sensors other than the short-range mobility sensors are used and the map is in configuration (L_j) . The utility ψ for visiting $\mathbf u$ while navigating to the goal is defined as:

$$\psi_{\mathbf{u}}(\mathbf{x}, \mathbf{v}_G) = 1 - \max_{\theta} \left(\sum_{j} \frac{P(L_j)(C_{ns}(L_j, \mathbf{x}, \mathbf{v}_G) - C_s(L_j, \mathbf{x}, \mathbf{u}, \theta, \mathbf{v}_G))}{C_{ns}(L_j, \mathbf{x}, \mathbf{v}_G)} \right)$$
(1)

In principle, the utility can be computed everywhere in the map and the resulting utility map can be used by the planner to select the most favorable sensing position. In principle, the utility $\psi_{\bf u}$ can be computed at every ${\bf u}$ by simulating the two paths of cost C_{ns} and C_s . This "forward simulation" [8]–[10] is at the heart of our approach.

Both C_s and C_{ns} can be evaluated by running the planner on different "virtual" configurations of the world map corresponding to different configurations L_j . In reality, however, this would require the enumeration of *all* possible configurations of the world which is clearly a combinatorially large set. Furthermore, the sum above must be evaluated, in principle, not only for all cell locations, but also for all possible sensor

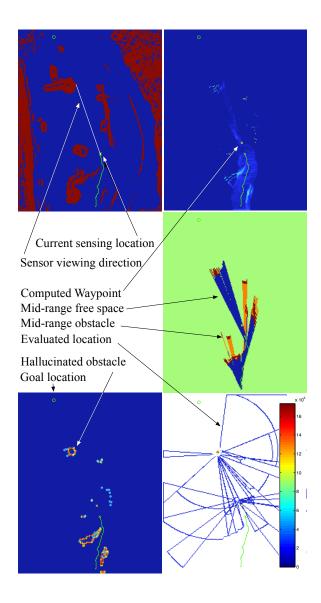


Fig. 4. A snapshot of an example run that shows the selection of observation locations. Clockwise from the top left is depicted: Terrain map (ground truth), Observed terrain map, Mid-range pencils of rays, binary hallucinated obstacle map and the utility map.

orientations. Therefore, the computation of the optimal utility function defined above is not tractable so we use a common approximation method that instead of averaging over all possible environments uses the Maximum A Posteriori (MAP) estimate or most probable world [11]–[14].

This most likely world or "Hallucinated world" gives us the locations in which obstacles are anticipated. If observing these anticipated obstacles yields a positive utility (Eq.1), the camera can be directed to confirm the existence of these obstructions [14].

We can use this utility for gathering geometric information from images to find view-points along our route to the goal. The complete navigation strategy is composed of the following steps: 1) capture an image and compute geometric information, 2) infer obstacles, compute utility of an extra measurement, 3) hand-off the location of maximum utility as a way-point/view-point to the vehicle planner, 4) repeat (Figure 4). More details of the view planning algorithm can be found in [15].

IV. EXPERIMENTAL EVALUATION

Several experiments were carried out to validate the proposed algorithm. These include results from field tests with an autonomous vehicle, as well as results from a much larger set of simulated experiments.

In the example field experiment, we used an ATRV-JR mobile robot equipped with a firewire camera and SICK laser range finder to navigate a simple course on campus (Figure 5). The robot was to navigate to a location behind the shown obstacle. The navigation map used on the robot was a

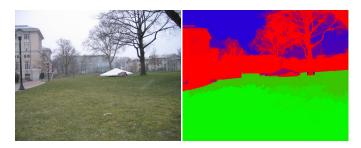


Fig. 5. Our experimental plaform is based on an ATRV-JR equiped with a firewire camera and SICK laser range finder. The vehicle planner [16] and laser processing run on the robot. The image processing is done off-line and results are send back to the robot.

150x150m grid with a 0.5m resolution. The camera resolution was 1024x768 which covered a field of view of about 50°. The Laser data was used to compute mobility terrain data and was artificially limited to 2m. For logistical reasons only, the experiment was setup with a scaled down sensing range and environment size. Localization was provided by an on-board GPS and laser ring gyro. The vehicle planner [16] and laser processing run on the robot. The image processing is done off-line and the robot stays in place until results are received send back.

In Figure 6 and 7 the results of this example run are shown. The displays in Figure 6 shows the initial measurement and corresponding geometry computed by the robot. Interesting to note is that it finds a clear path to the immediate left of the obstacle. The robot navigates to the goal through this corridor while planning for one more observation location closer to the obstacle (Figure 7).

The field experiments are encouraging, but it is difficult to evaluate performance over a large set of conditions (environment configuration, sensor characteristics, start and goal point locations) that are statistically meaningful. This is particularly hard given the wide range of variation in possible terrain configurations observed in the real world, which is very difficult to generate in physical experiments. In order to analyze more closely the performance of the algorithm, we conducted controlled experiments for a substantial number of runs using



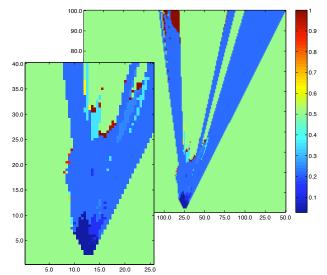


Fig. 6. Example range measurement from a single image. Top left: the original image, top right: the geometric labels. The navigation map is displayed in the lower panel. As before, the traversability ranges from 0 for traversable to 1 for an obstacle. Distance is in meters

a simulator on data from the U.S. Geological Survey with a simple simulated sensor model and the robot modeled by a single point.

We have executed 500 trial runs with randomly chosen start and goal positions and analyzed the data in the following way: First, we compare the lengths of the paths generated by using the mid-range sensor planning method with the paths executed by using mobility sensing only. This is shown in Figure 8, in which the runs are sorted in the order of increasing gain for the sensor-based planning method. This first type of analysis is intended to evaluate the amount gained from using sensor planning. It is important to note that the gain can vary dramatically depending on the start and goal points. Intuitively, little gain can be expected if the area between the start and goal points is completely unobstructed, in which case *any* planning strategy would perform well.

Second, it is important to compare the paths obtained with our sensor planning heuristics with the paths generated by using a mid-range sensor that senses all the time in every direction, since our claim is that the algorithm generates a "good" selection of when/where to sense during motion of the robot. If our planner were to generate paths that have substantially higher cost than those generated by using the "sense all the time/everywhere" strategy, it would indicate that

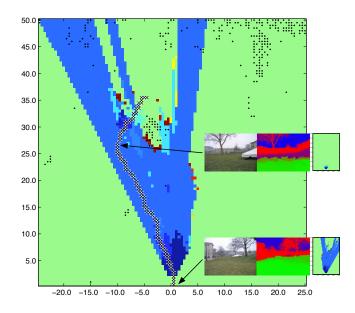


Fig. 7. The final map (in meters) of the example experiment in which the mobile robot navigated to a location out of initial view. The robot took two mid-range sensor measurements at the positions annotated with the recovered geometric information. The black dots display sensor returns from the SICK laser range finder plotted for ground truth comparison.

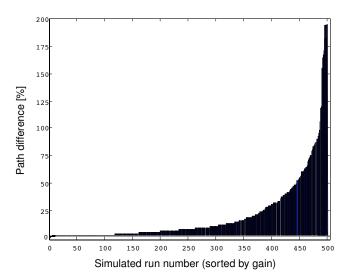


Fig. 8. The Mid-range planning sensing algorithm compared to the standard navigation approach. A positive percentage shows how much shorter the path from the mid-range sense/planning is over the mobility only approach. The average gain in path length is +18%. However, more interesting is that +76% of the runs exhibit positive gain (up to almost +200% gain) and that of those that have negative gain, the maximum loss is -6%. This is due to the overhead introduced by the explore behavior that tries to find better paths. For 22% of the runs there was no gain or loss.

our heuristics can be improved. This part of the analysis is summarized in Figure 9, in which the lengths of the paths generated from our planning approach and from the "sense all the time/everywhere" approach are plotted (with (+)) as a scatter plot. The plot indicates that the lengths are similar, i.e., the values are scattered near the diagonal (the correlation

coefficient is $\rho=0.99$). This result verifies empirically our hypothesis that continuous sensing of the environment is not necessary, provided that suitable heuristics are used for computing when and where to sense.

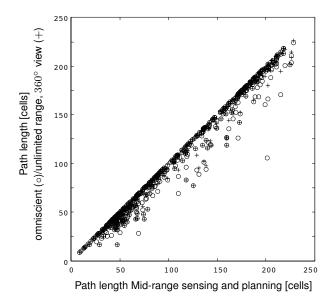


Fig. 9. The Mid-range planning sensing algorithm compared in a scatter plot to a hypothetical continuous sensing method with an unlimited range, 360° view range finder (no viewpoint planning (+)). and also compared to the true shortest path (o). The vertical axis is the path length from the omniscient planner and also the planner using continuous sensing and planning. The horizontal axis is the path length from our Mid-range sensing/planning method. Results show that the performance of our mid-range sensing approach are almost as close to the optimal as they were to the continuous sensing planner.

Finally, a third type of analysis uses the paths generated by an omniscient planner, which has complete knowledge of the entire map prior to execution, as the baseline for comparison. Such an omniscient planner generates the shortest paths that can be generated by any planner given the environment and the selections of start and goal locations. As such, it is a useful baseline to quantify the degradation of performance due to limited sensor horizon. This analysis is shown in Figure 9 as well, in which the paths generated by using our mid-sensor planning strategy and the paths executed by omniscient planner (marked with o) are plotted again as a scatter plot. The graph shows that the paths generated with mid-range sensor planning are almost as close to the optimal from the omniscient planner as they were to the continuous sensing planner. Specifically, the correlation coefficient is $\rho = 0.98$ between the paths planned by the omniscient planner and our mid-range sensor planning approach.

V. DISCUSSION

Failures in the single-view geometry estimation can arise from difficulties in segmentation or shadowing. For instance, in Figure 6, the sidewalk is thought to be an obstacle because the segmentation algorithm often groups it with the building (which is more similar in color and texture than the grass).

The geometry estimation may be improved by training from images taken by the robotic system used for testing in scenes that the robot is likely to encounter, rather than using general user photographs for training as was done here.

Additional problems can arise from the assumption that the environment close to the robot is planar. This can however be overcome by using the data from a scanning laser range finder to estimate the orientation of the ground plane. In addition other cues from the image, specifically the separation between ground and vertical classes, can be used to estimate the horizon and register it with the extrinsic camera parameters.

Overall, our experiments have shown that even error prone unreliable data from monocular vision can aid significantly to improve long range navigation performance by allowing to plan a path past the traditional sensing horizon.

Acknowledgments

This work was supported in part by the Robotics Consortium sponsored by the U.S. Army Research Laboratory under the Collaborative Technlogy Alliance Program, Cooperative Agreement DAAD-19-01-2-0012. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government.

REFERENCES

 K. Konolige and M. Agrawal, "Outdoor mapping and navigation using stereo vision," Submitted to International Symposium on Eperimental Robotics, 2006.

- [2] D. Nister, "An efficient solution to the five-point relative pose problem," in Conference on Computer Vision and Pattern Recognition, June 2003.
- [3] D. Hoiem, A. A. Efros, and M. Hebert, "Geometric context from a single image," in *International Conference of Computer Vision (ICCV)*. IEEE, October 2005.
- [4] A. Saxena, S. H. Chung, and A. Ng, "Learning depth from single monocular images," in *Advances in Neural Information Processing Systems* 18, Y. Weiss, B. Schölkopf, and J. Platt, Eds. Cambridge, MA: MIT Press, 2006, pp. 1163–1170.
- [5] P. Felzenszwalb and D. Huttenlocher, "Efficient graph-based image segmentation," *International Journal of Computer Vision*, vol. 59, no. 2, 2004.
- [6] J. Friedman, T. Hastie, and R. Tibshirani, "Additive logistic regression: a statistical view of boosting," *Annals of Statistics*, vol. 28, no. 2, 2000.
- [7] M. Collins, R. Schapire, and Y. Singer, "Logistic regression, adaboost and bregman distances," *Machine Learning*, vol. 48, no. 1-3, 2002.
- [8] A. Kelly and A. T. Stentz, "Rough terrain autonomous mobility part 2: An active vision, predictive control," *Autonomous Robots*, vol. 5, pp. 163 – 198, May 1998.
- [9] J. M. Esposito and V. Kumar, "Closed loop motion plans for mobile robots," in *International Conference on Robotics and Automation*, IEEE, Ed., April 2000, pp. 2777–2782.
- [10] S. LaValle, "A game-theoretic framework for robot motion planning," Ph.D. dissertation, University of Illinois, Urbana, IL, 1995.
- [11] J. Besag, "On the statistical analysis of dirty pictures," *Journal of Royal Statistical Soc.*, vol. B-48, pp. 259–302, 1986.
- [12] S. Z. Li, Markov Random Field Modeling in Image Analysis. Tokyo: Springer-Verlag, 2001.
- [13] C. S. Won and H. Derin, "Unsupervised segmentation of noisy and textured images using markov random fields," CVGIP, vol. 54, pp. 308– 328, 1992.
- [14] B. Nabbe, S. Kumar, and M. Hebert, "Path planning with hallucinated worlds," in *Proceedings: IEEE/RSJ International Conference on Intelli*gent Robots and Systems. IEEE, October 2004.
- [15] B. Nabbe, "Extending the path-planning horizon," Ph.D. dissertation, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, July 2005.
- [16] D. Ferguson and A. Stentz, "Field D*: An Interpolation-based Path Planner and Replanner," in *Proceedings of the International Symposium* on Robotics Research (ISRR), 2005.