

Path Diversity Is Only Part of the Problem

Ross A. Knepper and Matthew T. Mason

Abstract—The goal of motion planning is to find a feasible path that connects two positions and is free from collision with obstacles. *Path sets* are a robust approach to this problem in the face of real-world complexity and uncertainty. A path set is a collection of feasible paths and their corresponding control sequences. A path-set-based planner navigates by repeatedly testing each of these robot-fixed paths for collision with obstacles. A heuristic function selects which of the surviving paths to follow next. At each step, the robot follows a small piece of each path selected while simultaneously planning the subsequent trajectory. A path set possesses high *path diversity* if it performs well at obstacle-avoidance and goal-seeking behaviors. Previous work in path diversity has tacitly assumed that a correlation exists between this dynamic planning problem and a simpler, static path diversity problem: a robot placed randomly into an obstacle field evaluates its path set for collision a single time before following the chosen path in entirety. Although these problems might intuitively appear to be linked, this paper shows that static and dynamic path diversity are two distinct properties. After empirically demonstrating this fact, we discuss some of the factors that differentiate the two problems.

I. INTRODUCTION

A. Path Sets

A path set consists of a fixed set of control sequences, each paired with its corresponding trajectory. These controls are often selected, and their responses precomputed, off-line using a high-fidelity vehicle model. Since the paths are represented by actions, they are inherently constructed in the vehicle's coordinate frame and originate from the robot's current position (i.e. they are robot-fixed). The salient feature of a path set considered in this paper is the spatial separation, or *path diversity*, of its constituent paths. This paper examines the relationship between path diversity and planner success rate.

Path sets are the basis of an approach to planning and control suitable for use in real-world planning problems. As an alternative to planning in the continuum of motions, a path set considers a discrete set of control inputs, picking the best one according to its predicted trajectory. The planner selects one path for execution each cycle, but it may switch paths with each consecutive cycle, at a rate of roughly 10 Hz.

Fig. 1 depicts an example path set within a simulated planning context. We employ path sets in a hierarchical planning approach with a division of labor between local and global planners. The global planner is responsible for discovering the connectivity of freespace and guiding the robot in the direction of the goal. For scalability, the global level abstracts vehicle state into a square grid with 8-way




Fig. 1. Path sets used for planning. At center, a path set in black is shown fixed to the robot (red disc), which performs interleaved planning and execution. The robot originated at the lower right. Every cycle the planner evaluates each node in the path tree with respect to the global guidance navigation function—depicted here as shades of gray in the background. A heuristic function selects the best path, resulting in the feasible path in yellow and the infeasible, grid-based path to the goal at left, in red. The robot takes a step along the best path and then repeats the cycle.

connectivity. Consequently, only the local planner is aware of nonholonomic constraints that restrict the robot's mobility in the world. A heuristic function mediates between the hierarchy levels by selecting the best available combination of local and global paths. In such a hierarchical planning scenario, the global planner replans infrequently after the initial plan is generated, while the local planner replans regularly in order to perform the dual functions of path-finding and control.

While the replan rate may be viewed as a tunable parameter, there are two significant planning paradigms where replan rate is concerned. In the first, which we call the dynamic paradigm, the robot traverses much less than the length of a local path between replan cycles. As an alternative, the static paradigm generates only a single plan before traversing an entire trajectory. In this paper, we show that this choice of planning paradigm has a significant impact on the relationship between path diversity and planner success rate. This study is motivated by the desire to understand the fundamental behavior of real-world motion planners. This desire, in turn, is driven by the need to offer insight and support to the designers of planning systems for field robotics applications.

B. Path Diversity Paradigms

This paper examines the concept of path diversity in two different planning paradigms: static and dynamic. In the

R.A. Knepper and M.T. Mason are with the Robotics Institute, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA, USA
rak@ri.cmu.edu, mason@cs.cmu.edu




Fig. 2. Path sets: (a) Full 2,401-path data set; (b) path set generated with the Green-Kelly approximate-area metric; (c) path set generated with the Hausdorff metric; (d) path set generated using a mutual-collision metric, which estimates the probability of two paths colliding with the same obstacle.

static case, we place a stationary robot arbitrarily in a space cluttered with obstacles and test whether at least one of the paths in the path set is collision-free. In this simple paradigm, a single planning step precedes execution of the full plan. This formulation of the static paradigm is designed to obey the assumptions made by many other works on path diversity [9], [3], [6].

In the dynamic case, the robot simultaneously drives the beginning of each path plan while computing the subsequent replan step. In most implementations of this planning approach, new plans, which look ahead five seconds or more, are generated at about 10 Hz. Consequently, only a small fraction of each plan is ever executed on the robot. Although the remainder of the path is never executed, it still serves two important functions. First, the planner must look ahead beyond its minimum stopping distance in order to guarantee safety. Second, the remaining path segment approximates what future invocations of the planner might choose to do, although there is no guarantee that the next cycle's path set will contain the remaining path section. This property causes complex emergent behavior to dominate performance and confound simple theoretical analysis. Thus, we are led to the simulation-based analysis which makes up the bulk of this work.

We are not aware of any previous work in deterministically generating dynamic path sets that optimize performance, although our own earlier work [11] adopts a probabilistic approach to this task. Our goal in this work is to discover the fundamental principles governing the performance achievable under the constraints imposed by the local planning problem.

C. Prior Work in Path Diversity

The concept of path diversity has only recently been appreciated as an aspect of the path set planning problem that can make or break the planner on challenging problems. But what is path diversity?

Green and Kelly [9] define *relative completeness*, as “the prior probability, before the environment is specified, of pro-

ducing a solution path in a bounded amount of computation.” The authors go on to introduce an approximate-area metric between paths, and they show that relative completeness is related to the dispersion of paths. Additionally, they provide a greedy algorithm for selecting a diverse path subset from a large, densely-sampled path set. We will refer to these as the Green-Kelly metric and Green-Kelly algorithm.

Branicky et al. [3] introduce the term “path diversity,” defining it as “the probability of the survival of paths, averaged over all possible obstacle environments.” In this context, a path set survives when at least one of its member paths is free of collision. The essence of the algorithms presented in that paper is to minimize mutual overlap between paths, recognizing that points of overlap represent vulnerabilities where a single obstacle could block several paths.

A variety of applications from earlier works ([4], [5], [8], [10], [16]) face the path diversity problem in the dynamic paradigm, but none explicitly address it. All of these works use a path set composed of constant-curvature arcs, either in the continuum or discretely sampled. The intuition supporting arc-based path sets appears to be that an arbitrary path can be approximated to desired precision by a curve of piecewise-constant curvature. While the basis of this argument is correct, the planner is not made aware of all these possibilities. Rather, the planner believes at each cycle that it must commit to follow a circular arc out to the path set horizon. Consequently, the robot saddled with an arc-based planner cannot plan to follow an S-curve, a J-curve, or any other non-arc trajectory.

The ego-graph [12] and lattice [14] planners take on the problem of path diversity to a greater degree by incorporating a set of paths which is diverse in shape space for the purpose of achieving arbitrary configurations. Although path diversity is not of explicit concern in these two papers, it is an ancillary benefit in both approaches.

Finally, many of the finalists in the 2007 DARPA Urban Challenge utilized path set variants in planning. Several competitors relied on constant-curvature arcs [2], [15], [18]. Team MIT generated non-fixed random path sets using Rapidly-exploring Random Trees [13], while VictorTango pre-computed a path set in the form of an ego-graph designed for on-road operations [1].

Tartan Racing team's Boss delivered a clear Urban Challenge victory utilizing both the static and dynamic path set approaches. For path following, they dynamically generated a path set adapted for lane following by generating trajectories constrained to end parallel to the direction of motion [7], resembling the approach of 2005 DARPA Grand Challenge winner, Stanley [17]. As in the dynamic path set paradigm, Boss followed only the first part of each control before planning a new trajectory. In unconstrained off-road areas, Boss employed a lattice planner, in which a fixed path set is tessellated through space. After planning a motion through the lattice, Boss tried to follow the path without replanning, as in the static path set paradigm.

II. EXPERIMENTAL SETUP

This paper discusses a series of planning experiments that we conducted in simulation. We find that it is necessary to perform planner experimentation primarily in the simulation domain in order to perform the number of trials necessary to obtain statistically significant results. This work builds on our previous experimental contribution [11].

As before, we based our experiments on the Nomad Scout, a differential drive cylindrical robot of radius 0.206 m. After developing a high-fidelity dynamic vehicle model, we artificially limited the robot’s action set through the choice of path set. All candidate paths were selected from a “full” path set of size 2,401 (Fig. 2), based on a car-like steering model. Actions in the path set were limited to a linear velocity of 0.3 m/s and angular velocities with magnitude not more than 0.63 rad/s.

All experiments take place in a continuous, bounded 2D world measuring 20 m on a side. For planning and mapping purposes, the world is discretized into square 0.1 m cells. We generated a large set of planning problems by producing maps with different obstacle configurations. In each problem, the locations of the cell-sized obstacles were sampled from the uniform distribution with 2.8% coverage. Following obstacle sampling, we generated an approximate C-space expansion in the map grid based on the robot’s radius, coloring these cells yellow on the map (See Fig. 1). By remaining in the gray cells, the planner can ensure robot safety. The gray cells depict a navigation function, which we generated using the brushfire algorithm in an 8-connected grid, originating from the goal location. We use the navigation function to provide global guidance to the robot. Each problem also contains a unique start/goal pair, whose positions are sampled uniformly from the continuous world under the constraint that the distance between start and goal equals 10.5 m.

At the start of dynamic planning problems, the robot is oriented along the navigation function’s gradient. The dynamic planner then proceeds to replan at 5 Hz while the robot follows the most recently chosen action from the path set at each stage. The planner returns success when the robot reaches the goal position. If the heuristic function returns a better score at the robot’s current position than at any path set leaf node, then the planner commands zero velocity. Only when the robot comes to a complete stop without finding a way to progress does the planner return failure.

III. STATIC PLANNING PARADIGM

Most of the recent work focusing on the design of diverse path sets has concentrated on the static case, which essentially amounts to the following question: if the robot “parachutes” into a workspace cluttered with randomly-arranged obstacles, what is the probability that at least one of the paths in its path set will not collide with an obstacle? Many motion planning problems are handled by planning for a period of time and then trying to follow a single planned trajectory from start to goal.

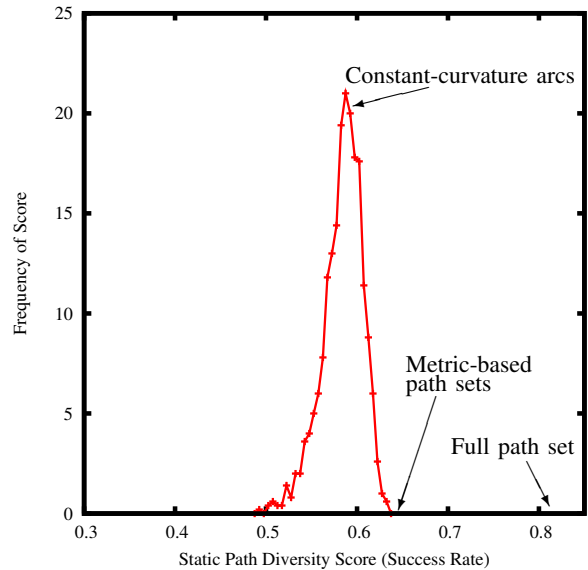


Fig. 3. Histogram of static path diversity. This plot shows performance of randomly-generated path sets compared to several deterministic path sets in the static parachuting robot experiment. The red curve represents a histogram of random path set performance. Other path sets of interest are picked out along the curve. The greedy “metric-based path sets,” including Green-Kelly, Hausdorff, and mutual collision, all have virtually the same performance in this test.

The key to achieving static path diversity is reasonably well understood. A diverse sample of paths is able to forecast all possible future motions to some precision. The argument is based on the locality assumption: without regard to the particular configuration of obstacles, a given path being collision-free implies that paths within its neighborhood in path space are also likely to be collision-free. If one chooses a set of paths that minimizes dispersion—the distance between any arbitrary solution path and the closest member of the set, using some metric on paths—then one will maximize the probability of at least one path surviving.

In our first experiment, we were interested in how several low-dispersion path sets constructed using the Green-Kelly algorithm performed in comparison to a variety of randomly-generated, fixed path sets. We measured the performance of these and several other path sets in the parachuting robot scenario. The goal of these experiments was to approximately quantify the range in path diversity across the space of path sets. These simulations were conducted with the framework discussed in Section II.

To generate these random path sets, we sampled 1,000 symmetric path sets of size twenty-four from the full set of 2,401-paths. For each path set, the simulator repeated the following process 10,000 times. First, a random world was selected as described previously, and the vehicle was placed in a random configuration. Second, with the robot in this configuration, the path set was expanded and tested for collision with obstacles. We recorded statistics at each trial for the surviving number of actions and leaves in the path set tree—representing the first and last levels in the tree, respectively. Some configurations were inherently

infeasible, such as when the robot landed on top of an obstacle. Therefore, we limited our analysis to those cases in which at least one action survived. From that set of trials, we reported the fraction of test cases in which at least one leaf node in the tree survived. We defined this fraction (or survival rate) as a static path diversity score.

Fig. 3 shows a histogram of this performance measure on each of the 1,005 path sets we tested. A set of 24 constant-curvature arcs put in a mediocre performance compared to the random path sets. Path sets generated using the Green-Kelly algorithm with various metrics (Fig. 2) performed just beyond the best randomly-generated path set, suggesting that despite the suboptimal greedy algorithm and imperfect metrics, these path sets may still approach the theoretical maximum static path diversity for path sets of size twenty-four. At far right in Fig. 3 is the full path set. It naturally follows that this path set would outperform the rest in this static experiment, since it contains 100 times as many paths as the others, each offering a chance of survival.

These results are consistent with both intuition and experience. As we have seen, path sets have generally been designed using a static paradigm assumption, even though they are often pressed into service in the dynamic paradigm. Thus, we must address the question of how well static path set performance translates into the dynamic replanning case.

IV. DYNAMIC PLANNING PARADIGM

There are applications, especially in field robotics, where it becomes necessary for the planner to deal with dynamic situations; as the robot moves through its environment, perception data changes continually due to noise, occlusion, and range limitations (in addition to true dynamic obstacles). Researchers have developed sophisticated techniques for handling evolving cost maps, but one very simple technique is often employed to tackle these problems: by throwing out the rest of the old path in favor of generating a new plan from scratch each cycle, the effects of stale plans are elegantly avoided.

One can argue the merits of the dynamic planning paradigm, but the recent DARPA Grand Challenges offer ample evidence for the efficacy of this approach. Despite considerable exposure in these and other robotic vehicles, it was not widely appreciated, prior to this work, how different the static and dynamic planning cases actually are. After examining the spectrum of dynamic planning performance, we present evidence for the high degree of the distinction between the two below in Section V.

We have previously reported on the results of our experiments examining the dynamic replanning case [11]. We repeated those experiments with slightly different parameters to match those used in the static experiments above. In this trial, we randomly generated a set of 1000 symmetric path sets, again comprising twenty-four paths each. We tested each of these random path sets, along with the special path sets mentioned in Section III above, on a group of 1000 planning problems of the type defined in Section II.




Fig. 4. Histogram of dynamic path diversity. 1000 random path sets were each tested on 1000 dynamic planning problems and scored based on their success rate and completion time. The curve represents a histogram of the scores of these path sets. The scores of several special path sets are also shown.

We report the results of these tests in Fig. 4. Our dynamic path diversity score is based on success rate, which in this experiment is defined as the fraction of test runs in which the robot successfully navigates to the goal position via dynamic replanning. As with the static paradigm, a histogram depicts the performance range of the random path sets. The set of constant-curvature arcs, which last time produced average performance, this time delivered a result considerably below every single other path set under test. The metric-based Green-Kelly path sets still performed near the top, but this time they were outraced by a handful of random path sets. Finally, we have a counterintuitive result: the full path set performed worse than nearly half of the random path sets. This result occurred despite the fact that the simulation was permitted to “cheat” and pretend that all 2,401 path sets could be computed within the 0.2 second deadline.

One possible explanation for this phenomenon goes as follows. Although the full path set approximates all feasible vehicle trajectories, the horizon of the path set is still only six seconds in the future. Beyond this horizon, the same infeasible 8-connected grid supplies guidance. Thus, the full path set exhibits the same greedy behavior as the other path sets. In fact, it can exercise greater control in driving near obstacles. A greedy planner naturally attempts to graze every obstacle (and follow narrow passages) where doing so minimizes the heuristic solution cost. Since our simulator is not subject to error from noise, path-following is not of concern. However, driving close to obstacles has an adverse effect not considered by the planner. Doing so drastically reduces the set of future collision-free paths. The lesson from this example is that a planner should attempt to maximize future flexibility during each plan step.

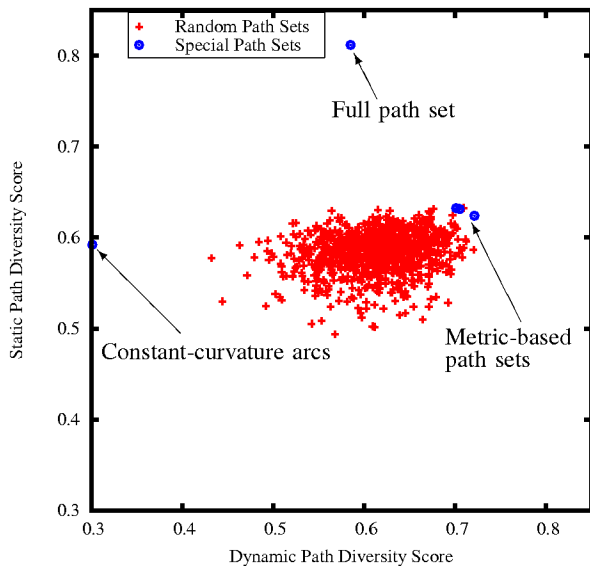


Fig. 5. Annotated scatter plot of path diversity correlation. Each path set’s scores are plotted for the static (vertical axis) and dynamic (horizontal axis) cases. A path set that performs well in both would be at the top-right, while one which performs poorly overall would appear at the bottom-left of the image. The correlation coefficient between these two distributions is 0.2175, reflecting only a very mild connection between performance in these two distinct problems.

V. STATIC AND DYNAMIC PATH DIVERSITY ARE DIFFERENT PROBLEMS

A. Experimental Results

In earlier sections, we showed experimental results in which path sets sampled at random from a uniform distribution were tested in a static “parachuting” context as well as a dynamic navigation context. In the next experiment, we tried to identify a correlation in performance between the two cases. We ran both the static and dynamic paradigm tests on the same group of 1,000 path sets (Fig. 5). In this figure, we see scores centered on a mean of (0.613, 0.584)—the peak of an approximate 2D Gaussian distribution. This distribution is striking because these two scores appear to be independent of each other. The correlation value between the static and dynamic distributions is only 0.2175 on a scale from 0 (no correlation) to 1 (linear relationship). This result suggests that while there is a mild correlation between performance in the two cases, one should not select path sets for use in the dynamic paradigm by static paradigm means.

B. Obstacle Distribution

At this point, we start to speculate about the reasons why the static and dynamic path diversity problems should be so distinct from each other. Intuition tells us that shape diversity, which proves effective at avoiding obstacles statically, should also be a factor in dynamic obstacle avoidance since the robot constantly discovers new obstacles while navigating. However, due to the small elapsed time between replan steps, only the edges of the path tree encounter novel terrain. At any given replan step, the majority of the tree covers ground which has been seen and reacted to at least once already.

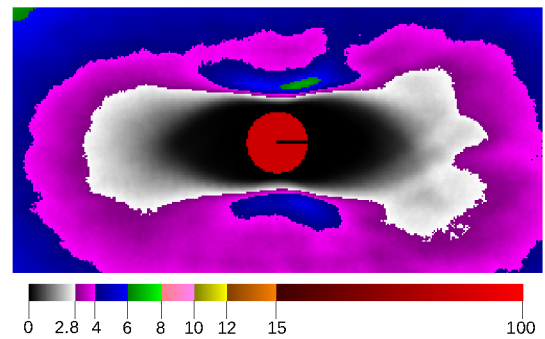


Fig. 6. Distribution of obstacles during navigation from the perspective of the robot at center. This distribution was generated by planning with a path set (Green-Kelly) across a range of over 300 worlds. During each planning cycle in which the planner successfully finds a next action, the position of local obstacles in the vehicle frame is averaged into the obstacle distribution map. Obstacles are randomly distributed within each test world at a 2.8% density, but they are clearly not uniformly distributed in the robot’s frame of reference. All points in this image below this nominal density are shown in shades of gray.

Although the distribution of obstacles in the world is uniform and fixed at 2.8% coverage, the obstacle distribution in the robot’s frame of reference—where the local planner evaluates path sets—is not at all uniform. The act of obstacle avoidance alters this distribution, as shown in Fig. 6. In particular, we can see that obstacles seldom occur immediately in front of the vehicle, where it is committed to drive. This phenomenon occurs because the planner is actively steering the robot away from obstacles in its path. By contrast, the unreachable areas beside the vehicle possess an increased likelihood of containing an obstacle.

To date, all work we are familiar with in generating diverse path sets has assumed that such diversity is equally important throughout the reachable workspace. From this result, it seems likely that more effective path sets could be designed by incorporating the expected distribution of obstacles in the body frame of the robot.

C. Recurrence Property

Next, we consider an aspect of the temporal interaction between consecutive replan cycles. A desirable path set attribute is *recurrence*, which states that for a solution path chosen during one cycle, the remaining unexecuted portion will exist as an option during the following cycles. The primary example is the set of constant-curvature arcs. This fact may help further explain arcs’ popularity in local planning over the years because it guarantees that once found, a solution will not disappear in future iterations.

While intuition suggests that recurrence should lead to good performance in replanning, the arcs path set actually stands as a counterexample to this notion. To test the proposition further, we performed another experiment. We tested each of the same 1000 random path sets from Section IV with one new feature: during each replan cycle, we include the continuation (or remaining unexecuted portion) of the previous cycle’s best path in the new path set. So, while we are no longer strictly employing fixed path sets, the

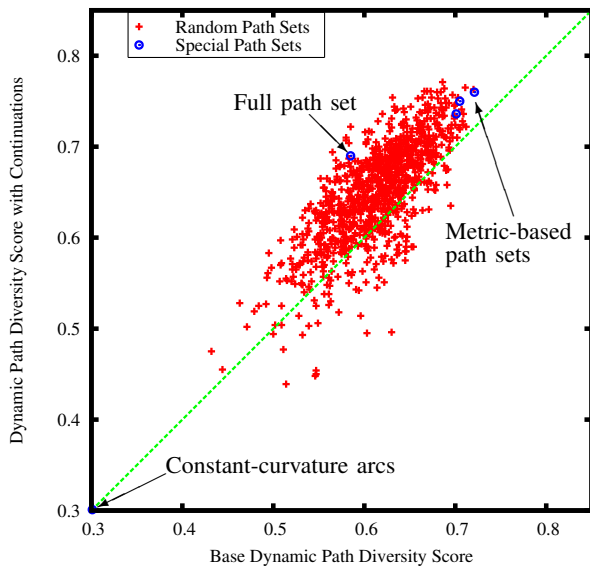


Fig. 7. The effect of continuations on navigation performance. To better understand the effect of recurrence on dynamic planning performance, each path set tested in Section IV was retested with an extra path: the untraversed remainder of the previously-commanded path. Path sets which degrade in performance appear below the diagonal line, while those path sets that improve are shown above the line. About 15% of path sets are adversely affected by the additional continuation option.

recurrence property is guaranteed during each replan cycle.

To interpret the results, we plotted in Fig. 7 a comparison of performance with and without continuations for each path set. About 90% of path sets improved their performance with continuations. This result is of interest to us because continuations partly bridge the gap between the static and dynamic paradigms. However, fully understanding the implications of this result remains a subject for future work.

VI. DISCUSSION AND FUTURE WORK

In this paper, we probed the difference in planning performance between the static and dynamic planning paradigms as influenced by path diversity. We demonstrated through simulation studies that a random path set's static and dynamic diversity are weakly correlated. While it is clear that a substantial performance boost may be realized by careful path set selection, one must also use care in ensuring that the basis by which a path set is selected ensures applicability in its planning context. These foundational principles affect the performance of many robotic planning systems, and robot designers would benefit from considering them as part of the design process.

In the future, we hope to develop a more sophisticated model of dynamic path diversity that is capable of predicting planner performance based on a variety of path set properties. In the longer term, we hope to use this model to devise algorithms capable of generating high-performance path sets that are more robust to the perils of real-world navigation.

VII. ACKNOWLEDGMENTS

This work is sponsored by the Defense Advanced Research Projects Agency. This work does not necessarily

reflect the position or the policy of the Government. No official endorsement should be inferred.

We would like to thank Alonzo Kelly and Michael Erdmann for their continuing contributions to this work.

REFERENCES

- [1] A. Bacha, C. Bauman, R. Faruque, M. Fleming, C. Terwelp, C. Reinholtz, D. Hong, A. Wicks, T. Alberi, D. Anderson, S. Cacciola, P. Currier, A. Dalton, J. Farmer, J. Hurdus, S. Kimmel, P. King, A. Taylor, D. Van Covern, and M. Webster. Odin: Team VictorTango's entry in the DARPA urban challenge. *Journal of Field Robotics*, 25(8):467–492, 2008.
- [2] J. Bohren, T. Foote, J. Keller, A. Kushleyev, D. Lee, A. Stewart, P. Vernaza, J. Derenick, J. Spletzer, and B. Satterfield. Little ben: The ben franklin racing team's entry in the 2007 DARPA urban challenge. *Journal of Field Robotics*, 25(9):598–614, 2008.
- [3] M.S. Branicky, R.A. Knepper, and J.J. Kuffner. Path and trajectory diversity: Theory and algorithms. In *Proc. Intl. Conf. Robotics and Automation*, Pasadena, CA, 2008.
- [4] J. Buhmann, W. Burgard, A.B. Cremers, D. Fox, T. Hofmann, F. Schneider, J. Strikos, and S. Thrun. The mobile robot RHINO. *AI Magazine*, 16(1), 1995.
- [5] M. Daily, J. Harris, D. Keirse, D. Olin, D. Payton, K. Reiser, J. Rosenblatt, D. Tseng, and V. Wong. Autonomous cross-country navigation with the ALV. In *Proc. of IEEE Intl. Conf. Robotics and Automation*, pages 718–726, Philadelphia, PA, 1988.
- [6] L.H. Erickson and S.M. LaValle. Survivability: Measuring and ensuring path diversity. In *Proc. IEEE International Conference on Robotics and Automation*, Kobe, Japan, May 2009.
- [7] D. Ferguson, T. Howard, and M. Likhachev. Motion planning in urban environments. *Journal of Field Robotics*, 25(11/12):939–960, 2008.
- [8] D. Fox, W. Burgard, and S. Thrun. The dynamic window approach to collision avoidance. *IEEE Robotics and Automation*, 4(1), 1997.
- [9] C. Green and A. Kelly. Toward optimal sampling in the space of paths. In *13th International Symposium of Robotics Research*, Hiroshima, Japan, November 2007.
- [10] A. Kelly and A. Stentz. Rough terrain autonomous mobility - part 2: An active vision, predictive control. *Autonomous Robots*, 5(2):163–198, May 1998.
- [11] R.A. Knepper and M.T. Mason. Empirical sampling of path sets for local area motion planning. In *International Symposium of Experimental Robotics*, Athens, Greece, July 2008.
- [12] A. Lacaze, Y. Moscovitz, N. DeClaris, and K. Murphy. Path planning for autonomous vehicles driving over rough terrain. In *Proc. of IEEE Intl. Symp. Intelligent Control*, 1998.
- [13] J. Leonard, J. How, S. Teller, M. Berger, S. Campbell, G. Fiore, L. Fletcher, E. Frazzoli, A. Huang, S. Karaman, O. Koch, Y. Kuwata, D. Moore, E. Olson, S. Peters, J. Teo, R. Truax, M. Walter, D. Barrett, A. Epstein, K. Maheloni, K. Moyer, T. Jones, R. Buckley, M. Antone, R. Galejs, S. Krishnamurthy, and J. Williams. A perception driven autonomous urban vehicle. *Journal of Field Robotics*, 25(10):727–774, October 2008.
- [14] M. Pivtoraiko, R.A. Knepper, and A. Kelly. Optimal, smooth, nonholonomic mobile robot motion planning in state lattices. Technical Report CMU-RI-TR-07-15, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, May 2007.
- [15] F. W. Rauskolb, K. Berger, C. Lipski, M. Magnor, K. Cornelsen, J. Effertz, T. Form, F. Graefe, S. Ohl, W. Schumacher, J. Wille, P. Hecker, T. Nothdurft, M. Doering, K. Homeier, J. Morgenroth, L. Wolf, C. Basarke, C. Berger, T. Gilke, F. Klöse, and B. Rumpe. Caroline: An autonomously driving vehicle for urban environments. *Journal of Field Robotics*, 25(9):674–724, 2008.
- [16] R. Simmons. The curvature-velocity method for local obstacle avoidance. In *International Conference on Robotics and Automation*, April 1996.
- [17] S. Thrun, M. Montemerlo, H. Dahlkamp, D. Stavens, A. Aron, J. Diebel, P. Fong, J. Gale, M. Halpenny, G. Hoffmann, et al. Stanley: The robot that won the DARPA grand challenge. *Journal of Field Robotics*, 23(9):661–692, 2006.
- [18] F. von Hundelshausen, M. Himmelsbach, F. Hecker, A. Mueller, and H. Wuensche. Driving with tentacles: Integral structures for sensing and motion. *Journal of Field Robotics*, 25(9):640–673, 2008.