

Detection of Principal Directions in Unknown Environments for Autonomous Navigation

Dmitri Dolgov
AI & Robotics Lab
Toyota Research Institute
Ann Arbor, MI 48105
Email: ddolgov@ai.stanford.edu

Sebastian Thrun
Stanford Artificial Intelligence Lab
Stanford University
Stanford CS 94305
Email: thrun@ai.stanford.edu

Abstract—Autonomous navigation in unknown but well-structured environments (e.g., parking lots) is a common task for human drivers and an important goal for autonomous vehicles. In such environments, the vehicles must obey the standard conventions of driving (e.g., passing oncoming vehicles on the correct side), but often lack a map that can be used to guide path planning in an appropriate way. The robots must therefore rely on features of the environment to drive in a safe and predictable way. In this work, we focus on detecting one of such features, the principal directions of the environment.

We propose a Markov-random-field (MRF) model for estimating the maximum-likelihood field of principal directions, given the local linear features extracted from the vehicle’s sensor data, and show that the method leads to robust estimates of principal directions in complex real-life driving environments. We also demonstrate how the computed principal directions can be used to guide a path-planning algorithm, leading to the generation of significantly improved trajectories.

I. INTRODUCTION

Autonomous navigation in outdoor environments is an active area of research in robotics, with extensive work being done in two distinct areas: i) the creation of robotic vehicles capable of driving on streets and highways [7, 24, 25, 23], and ii) the development of robots for off-road navigation [14, 22, 4, 21, 1, 16]. In the former case of on-street driving, it is sensible and commonplace to assume that there is a map of the global road network available, and the robot must obey the standard rules of driving (drive on the appropriate side of the street, not straddle the lane boundaries, etc). In the other mode of off-road driving, a detailed map of the environment is typically not available to the robot *a priori*, but the robot is not confined by any rules of the road and is free to choose any drivable and safe path to its goal.

There is also a large middle ground between these two areas, consisting of well-structured environments where it is not reasonable to assume knowledge of a detailed map, but the robotic vehicle must nonetheless obey the common rules of driving. Such conditions arise, for example, in parking lots, shopping malls, and construction zones. In such areas—even in the absence of typical road markings and signs—human drivers are usually able to partition the space into drivable lanes and drive appropriately. For instance, when driving in a parking lot such as the one shown in Figure 1a, most people




Fig. 1. Detecting structure of the driving environment. a) Aerial imagery of a parking lot; b) A typical structure imposed on the environment by human drivers.

will not have any difficulty detecting the main drivable lanes, similar to the ones shown in Figure 1b.

Similarly, a robot operating in such an environment needs to be able to—by using data from its on-board sensors—recognize features of the environment that will allow it to drive in accordance with common rules. Compliance to established rules and driving conventions is an especially important safety requirement if robots are to share the environment with human drivers, because it then becomes an issue of predictability.

In this work, we take a step towards this goal by considering one of the most basic features of environments: their principal directions. Knowing the principal directions of the environment is a prerequisite condition for implementing many high-level driving behaviors, such as driving on the correct side of the road (which is critical for collision avoidance) and avoiding diagonal paths across parking lots (which is often considered rude and unsafe).

Finding principal directions based purely on data from the robot’s on-board sensors is a challenging task due to the

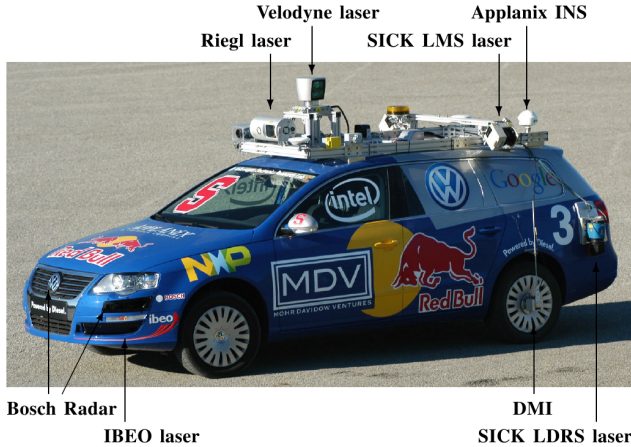


Fig. 2. Stanford Racing Team’s robotic vehicle, Junior. The results presented in this work were obtained using the Applanix GPS+IMU system and the Velodyne 3D LIDAR.

following issues: i) Many environments have several principal directions that vary from point to point (e.g., Figure 1), ii) Sensor data is noisy and can give rise to conflicting local hypotheses regarding the global principal directions, and iii) Computation of principal directions has to be done efficiently to be useful for navigation.

We approach the problem of computing the field of principal directions in a probabilistic framework. Assuming that at every point in the two-dimensional space of the driving environment there is a main preferred direction of motion, we formulate a Markov-random-field (MRF) [6] model that allows us to infer a high-likelihood field of principal directions, corresponding to the observed evidence provided by the vehicle’s on-board sensors. This approach is inspired by the extensive use of MRFs for image processing [18, 26].

The main steps of our approach are as follows. First, a map of the environment is computed from the robot’s on-board sensors. In this work, we present results based on LIDAR data, but our approach can be applied to other sensor modalities, such as cameras (potentially even leading to improvements in accuracy). We then scan the resulting map for local linear features that provide evidence regarding the principal directions of the environment. Having thus obtained local evidence, we formulate an MRF whose nodes form a 2D grid and represent the global principal direction of the environment at the corresponding locations. A solution to the MRF is obtained using fast numeric-optimization techniques.

We also present results demonstrating the robust performance of the MRF approach on several real environments driven with our robotic vehicle (Figure 2), and show how the resulting field of principal directions can be used to guide path planning, leading to a significant improvement in generated trajectories.

II. LOCAL PRINCIPAL DIRECTIONS

The core of our MRF-based approach is independent of the method used to obtain features of the environment that serve

as local evidence. However, for completeness and continuity, in this section, we briefly describe the specifics of our method for obtaining evidence for the MRF.

We used a vehicle equipped with a 3D Velodyne LIDAR as the sole environmental sensor. Below, we outline the techniques we used to find lines in the environment, given the range data produced by the LIDAR. The main steps of the process are illustrated in Figure 3.

The Velodyne LIDAR outputs a 3D point cloud, as shown in Figure 3a. In the first step, we filter out the ground plane, integrate data over time, and project the points onto a plane, leading to a 2D obstacle map shown in Figure 3b. Any line segments in this 2D map (e.g., curbs, other cars) are evidence of the principal driving directions, and we can detect those local lines using common computer-vision approaches.

There are several good known techniques for detecting lines in images. We found that the following sequence of standard transforms leads to satisfactory results for our application. First, we smooth the data using a symmetric 2D Gaussian kernel (Figure 3c) and then apply a binary threshold (25%) to the result (leading to the data shown in Figure 3d). The effect of these two steps is to remove small noise from the data, smooth out jagged lines, and “fill in” small gaps within objects (such as cars or trees), thereby eliminating extra edges. We then apply the Canny algorithm [5] for edge detection, yielding the image shown in Figure 3e. Finally, we use a randomized Hough transform (RHT) [27] to find line segments in the data (Figure 3f).

III. MARKOV RANDOM FIELD FOR ESTIMATING GLOBAL PRINCIPAL DIRECTIONS

The lines obtained in the previous section can be viewed as local evidence regarding the principal directions of the environment. Our goal is therefore to estimate the maximum likelihood field of principal directions $\theta(x, y)$, which specifies the principal orientation at every point x, y in the region of interest. Note that most complex driving environments—e.g., Figure 1—do not have a single global orientation, but rather have several principal directions that differ from point to point.

We formulate a discrete version of the inference problem on a regular 2D grid, associating with each point (x_i, y_i) an MRF node with a scalar continuous value θ_i . Figure 4 illustrates our MRF construction. Each MRF node $\theta_i \in [0, \pi/2)$ has associated with it a set of evidence nodes $\alpha_{ik} \in [0, \pi/2)$, one for the angle of each line segment $k \in [1, K_i]$ that crosses the grid cell i .

All angles are normalized to $[0, \pi/2)$, because orthogonal lines (e.g., edges of a car) support the same hypothesis for the principal direction at a point.

Our MRF uses two sets of potentials associated with the nodes and edges of the graph in Figure 4. The first potential (Ψ) is defined on the nodes of the MRF and ties each θ_i to its local evidence α_{ik} . It is defined as follows:

$$\Psi(\theta, \alpha) = \sum_i \sum_{k=1}^{K_i} \lambda_{ik} \psi(\theta_i, \alpha_{ik}), \quad (1)$$

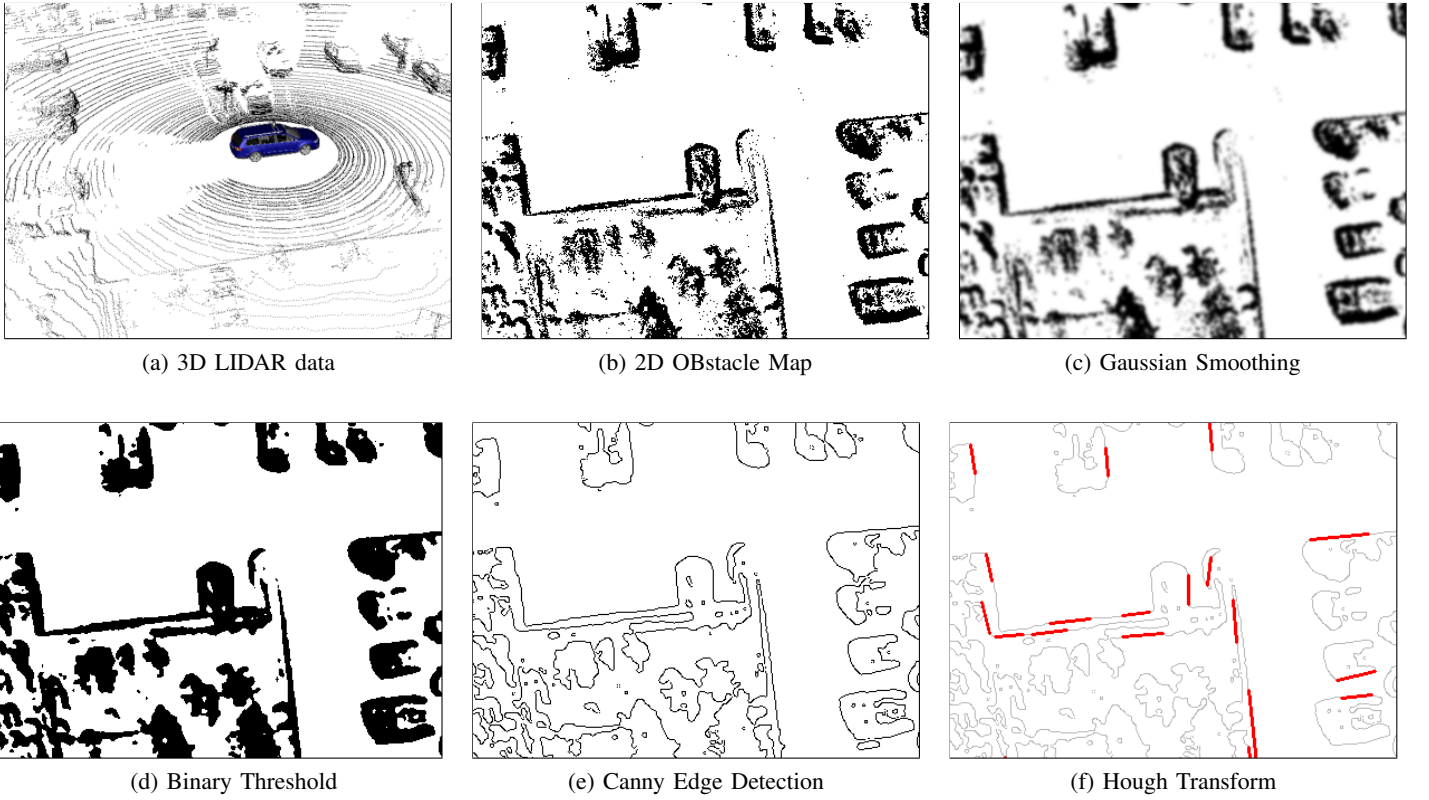


Fig. 3. Line detection in a parking lot: a) 3D LIDAR data; b) Obstacle map; c) After smoothing with a Gaussian kernel; d) After binary thresholding; e) After Canny edge detection; f) After Hough transform.




Fig. 4. MRF for inferring the field of principal directions of the environment. MRF variables θ_i are shown as red circles, α_{ik} are the input nodes, corresponding to the angles of the observed lines.

where λ_{ik} is the weight associated with line k , and ψ is a distance measure between two angles; both are defined below.

The second potential (Φ) is defined on the edges of the MRF and encodes a prior that enforces a smoothness relationship between the principal directions at neighboring nodes:

$$\Phi(\theta) = \frac{1}{2} \sum_i \sum_{j \in \mathcal{N}(i)} \phi(\theta_i, \theta_j), \quad (2)$$

where $\mathcal{N}(i)$ is the set of neighbors of node i , and ϕ is a distance measure between two angles (defined below).

There are many reasonable choices for the distance measures ψ and ϕ , as well as the weights λ_{ik} . We evaluated several options for each, and experimentally settled on the following. For the weights λ_{ik} , we used the length of the corresponding line segment (the longer the segment, the stronger the evidence). The choice of the distance measures ψ and ϕ is an interesting topic in itself: some distances favor smoother fields (e.g., L_2 norm), others (e.g., L_1 norm) have better discontinuity-preserving properties; see [9] for an applicable discussion of different norms in optimization. We empirically investigated several functions, and found the following to be a good choice of a norm for both evidence and smoothness potentials:

$$\psi(\beta, \gamma) = \phi(\beta, \gamma) = \sin^2(2(\beta - \gamma)). \quad (3)$$

This measure behaves quadratically for small $\beta - \gamma$, and has natural wrap-around properties for our angles in the interval $[0, \pi/2)$.

Finally, the distribution of the MRF variables θ for a specific set of observed α is given by a Gibbs distribution:

$$\mathbb{P}(\theta|\alpha) = \frac{1}{Z} \exp\left(- (w_\psi \Psi(\theta, \alpha) + w_\phi \Phi(\theta))\right), \quad (4)$$

where w_ψ and w_ϕ are weights, and Z is a normalizer or the partition function.

Our goal is to find the maximum-likelihood field of principal

directions θ , given the observed evidence α :

$$\theta_\alpha^* = \arg \max_{\theta} \mathbb{P}(\theta|\alpha), \quad (5)$$

or, in other words, find θ that minimizes the Gibbs energy $U = w_\psi \Psi + w_\phi \Phi$.

IV. OPTIMIZATION

For computational reasons, it is infeasible to compute an exact maximum-likelihood solution to the MRF defined in the previous section for anything but the simplest problems (a typical MRF for a realistic environment will have several thousand to tens of thousands nodes). Similarly to the approach of Diebel and Thrun [8, 9], we therefore settle for a high-probability mode of the posterior, which we compute using conjugate-gradient (CG) optimization [12, 20].

CG works best when an analytical gradient of the energy is specified, which is easily computed for our MRF potentials:

$$\frac{\partial \Psi}{\partial \theta_i} = 4 \sum_k \lambda_{ik} \sin \left(2(\theta_i - \alpha_{ik}) \right) \cos \left(2(\theta_i - \alpha_{ik}) \right),$$

$$\frac{\partial \Phi}{\partial \theta_i} = -\frac{\partial \Phi}{\partial \theta_j} = 2 \sin \left(2(\theta_i - \theta_j) \right) \cos \left(2(\theta_i - \theta_j) \right).$$

Given the above potentials and the gradient, the implementation of conjugate gradient is standard [20]. The output of the optimization is a high-likelihood field $\theta(x_i, y_i)$ that corresponds to the observed lines α .

V. RESULTS

In this section, we present results on the performance of our MRF-based approach to computing principal directions. The use of these directions in path-planning is discussed and evaluated in the next section.

We tested our algorithm using a vehicle equipped with an Applanix pose estimation GPS-IMU system and a Velodyne 3D LIDAR. Some representative examples of executing our method are shown in Figure 5. The left column shows the 2D obstacle map obtained from the vehicle’s sensors. The center column shows the lines detected in the obstacles map using the method described in Section II; these lines serve as the input to our MRF. The right column of Figure 5 shows the resulting field of principal directions computed by conjugate-gradient optimization.

The top row of Figure 5 shows a nearly ideal scenario: a parking lot with two main orthogonal principal directions, which are easily computed by our algorithm. The second row shows data for another, more complex, parking lot. Notice the presence of trees and a second, differently oriented, parking lot in the bottom-left part of the map. Despite these challenges, the MRF computes a very good estimate of the preferred driving directions for this environment. The third row of Figure 5 demonstrates the ability of our approach to handle environments with gradually-changing orientations. Notice that the field of principal directions correctly follows the curved road segment. The fourth row of Figure 5 depicts another challenging situation with a curved street, an intersection,

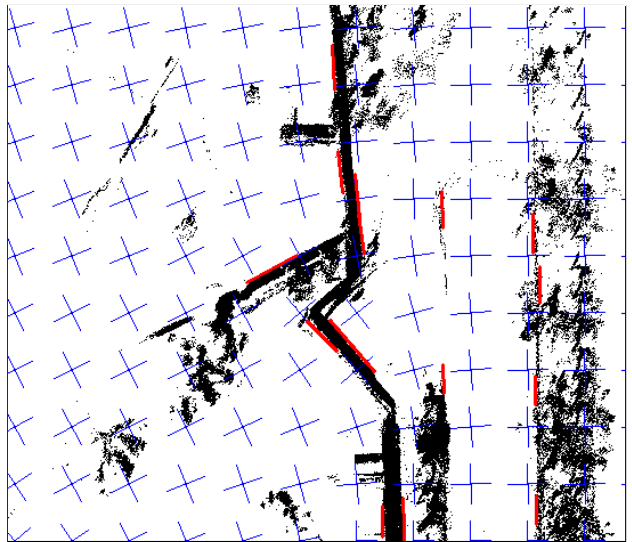


Fig. 6. The MRF computes the correct orientation for the drivable region of the parking lot, despite bad evidence provided by an adjacent building corner.

and an adjacent parking lot. It also highlights an interesting challenge for our method: parking lots with diagonally-parked cars (upper-right corner of the map). In this situation, the parked cars—whose sides are usually detected as lines—and features of the parking lots themselves (e.g., curbs) present conflicting evidence regarding the principal orientation of the environment.

Another complex situation is shown in Figure 6, where a building is located very close to a parking lot, but is oriented at a different angle, providing bad evidence for the MRF. However, as can be seen from the vector field in Figure 6, the MRF is able to compute the correct orientation for the drivable area, despite the fact that the curb separating the corner of the building from the parking lot is not detected.

Additional examples illustrating the computation of principal directions in real driving environments are shown in Figure 9.

Figure 7 shows the running time of our algorithm as performed on a 3Ghz Intel Core-2 PC. The data presented in Figure 7 is for the parking lot shown in the second row of Figure 5).

Figure 7a compares the running times of the main components of our approach. The timing results in Figure 7a are for experiments run with the following parameters: the obstacle grid was 260m×260m with 15cm resolution, the MRF evidence grid was 260m×260m with 5m resolution (resulting in an MRF with 2704 nodes). As can be seen from the data, the performance of our algorithm is usually dominated by the running time of the conjugate-gradient algorithm, and with the average computation time of under 300ms, the method is well-suited for use in online path planning (since it is usually not necessary to update the field of principal directions during every planning cycle).

Figure 7b shows how our MRF inference scales with the size of the MRF grid (with a 5-meter discretization of the MRF

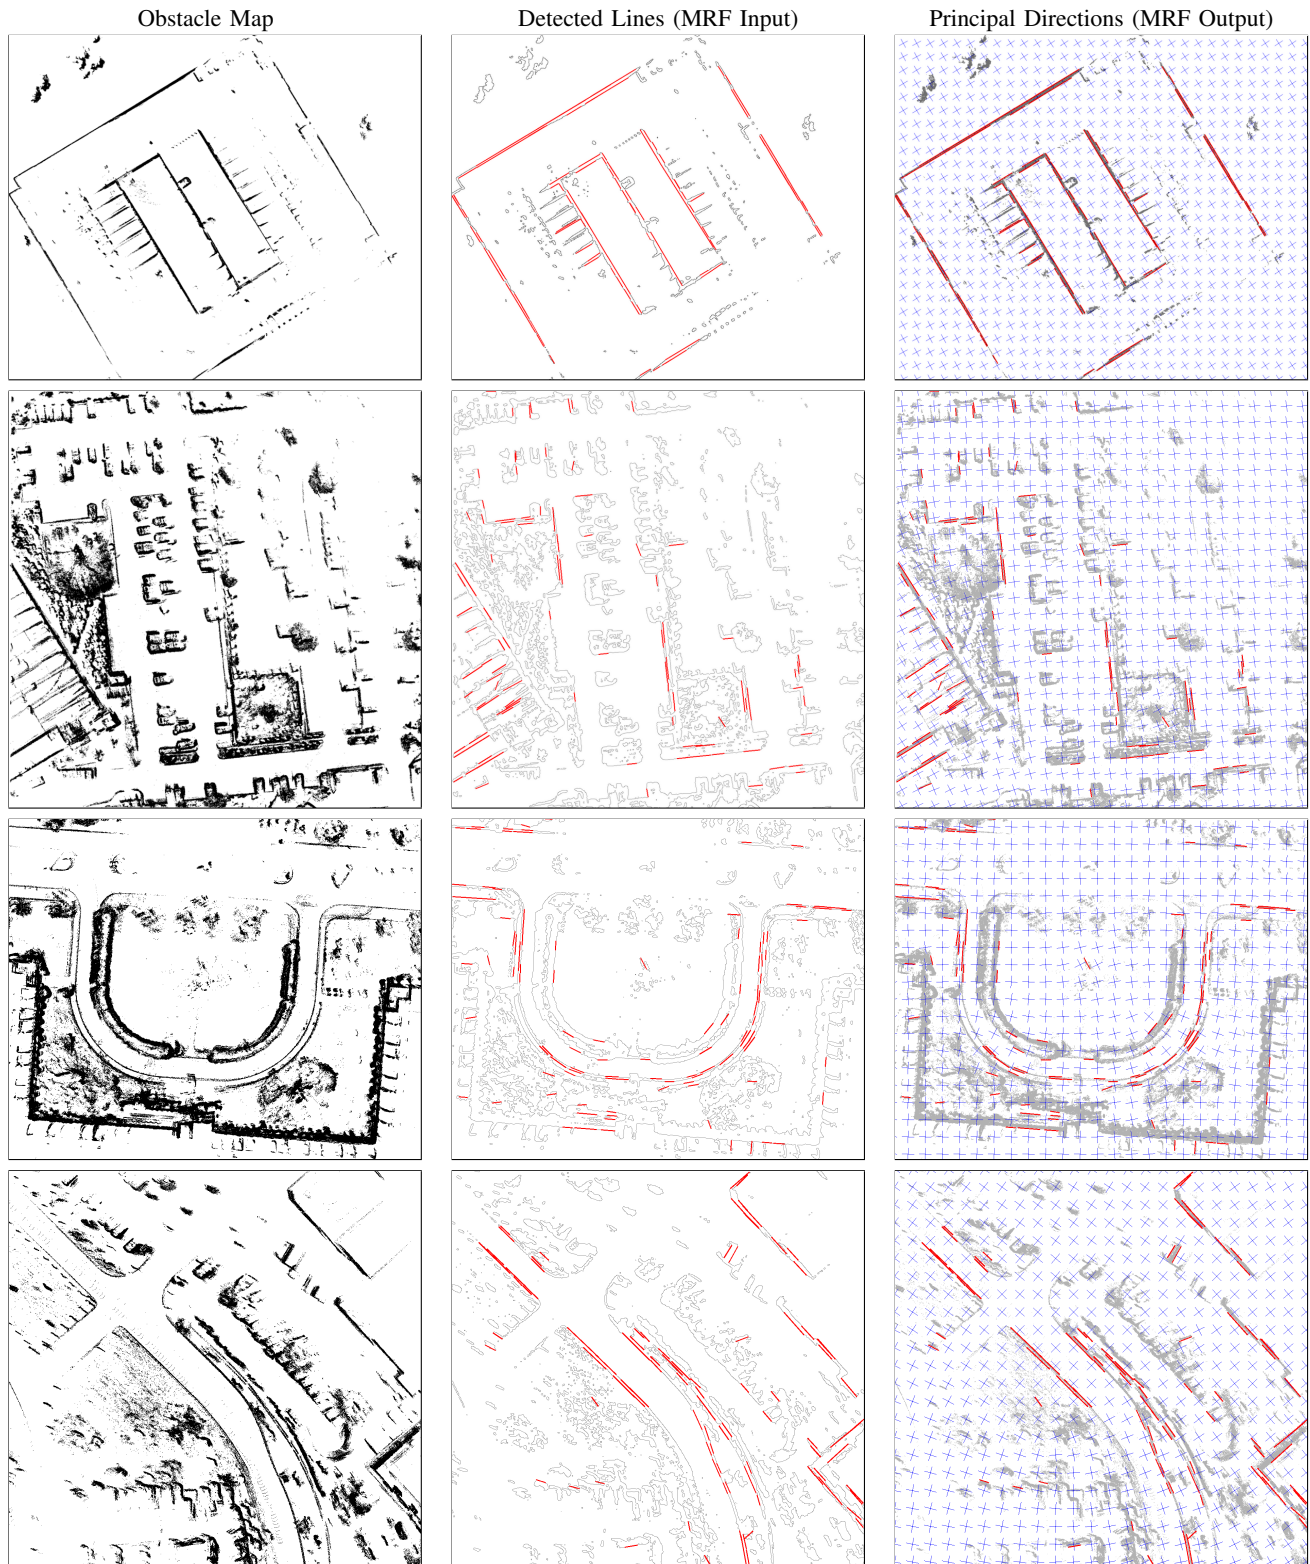


Fig. 5. Results of computing principal directions in several real driving environments.

grid). Figure 7c illustrates the scaling of the MRF inference as a function of the number of MRF variables θ_i (corresponding to MRF grids of size 40m to 300m, at 2.5m resolution).

We should note that a finer discretization of the MRF grid does not necessarily lead to better results. In fact, we found a resolution of around 5m to work best for typical driving areas.




Fig. 7. (a) Running time of the main components of our method on a sequence of real obstacle maps: i) Time for map pre-processing and Canny edge detection, ii) Time for the Hough Transform, iii) Time for the MRF conjugate-gradient optimization, and iv) Number of local linear features detected via the Hough transform. (b) MRF inference time as a function of the MRF grid size at a 5-meter resolution (number of MRF variables is quadratic in the size). (c) MRF inference time as a function of the number of MRF variables (grid cells) at a 2.5-meter discretization.

VI. PATH PLANNING

There are several ways in which knowledge of principal directions can be fruitfully utilized in path planning, some of which were outlined in the Introduction. Below, we describe a path-planning algorithm that favors trajectories aligned with the environment, and illustrate the improvements this leads to, compared to a standard planner with a bias towards smooth trajectories with arbitrary headings.

In what follows we leave aside the problem of global planning, which in practice can be accomplished via several existing algorithms, such as continuous forward search via path sampling (e.g., [3, 13, 17, 19]) or discrete search (e.g., [10, 15, 11]), just to name a couple classes of algorithms. Assuming that a rough global plan has been computed, we show how the field of principal directions can be used in a path-smoothing phase of planning.

Given a global plan represented as a sequence $\langle x_l, y_l \rangle$, $l \in [1, N]$, we can formulate the smoothing problem as a continuous optimization on the coordinates of the vertices of the trajectory. For illustrative purposes, below we focus on two terms: smoothness and compliance with principal directions; other important aspects of a realistic path smoother (such as kinematic constraints of the vehicle and collision-avoidance terms) are omitted for brevity.

For a given trajectory $\{\langle x_l, y_l \rangle\}$, let us define the displacement vector at vertex l :

$$\Delta \mathbf{x}_l = \begin{pmatrix} x_l \\ y_l \end{pmatrix} - \begin{pmatrix} x_{l-1} \\ y_{l-1} \end{pmatrix}, \quad (6)$$

and the heading of the trajectory at vertex l :

$$\sigma_l = \tan^{-1} \left(\frac{y_{l+1} - y_l}{x_{l+1} - x_l} \right). \quad (7)$$

The smoothness of the trajectory and the bias towards driving along principal directions can then be expressed as:

$$f = w_{sm} \sum_{l=1}^{N-1} (\Delta \mathbf{x}_{l+1} - \Delta \mathbf{x}_l)^2 + w_{pd} \sum_{l=1}^N \mu(\theta(x_l, y_l), \sigma_l), \quad (8)$$

where $\theta(x_l, y_l)$ is the principal direction at the MRF node closest to $\langle x_l, y_l \rangle$; w_{sm} and w_{pd} are the weights on smoothness and the principal-direction bias, respectively; μ is a potential on the difference of the two angles, which can be defined in a variety of ways, for example:

$$\mu(\theta(x_l, y_l), \sigma_l) = 1 - \cos(2(\theta(x_l, y_l) - \sigma_l)). \quad (9)$$

This smoothing problem can be solved using several efficient numerical techniques. One approach is to again use the conjugate-gradient algorithm as in Section IV.

Figure 8 shows the output of different variants of path smoothing on the same problem. Figure 8a and Figure 8c shows the output of the smoother that minimizes the quadratic-curvature term in (8), subject to constraints on collision avoidance and kinematics of the car, but $w_{pm} = 0$ and principal directions are ignored. The driving style exemplified by such trajectories might in some cases be considered undesirable. Turning on the bias for aligning the trajectory with the principal directions leads to solutions shown in Figure 8b and Figure 8d; such compliance with the orientation of the environment can be useful in many driving situations.

VII. DISCUSSION

We presented an MRF-based method for inferring principal directions of unknown environment using data from 3D LIDAR, and illustrated the usefulness of principal directions for path-planning in well-structured environments, such as parking lots. We evaluated the performance of our approach in real driving scenarios, demonstrating its ability to robustly estimate the field of principal directions in the presence of noise, conflicting local evidence, and regions with smoothly curved boundaries.

This work takes a step towards designing autonomous vehicles that can operate predictably in unknown environments and follow the standard conventions of human driving.

Still, the principal directions are a rather crude property of the environment, and it is necessary to detect more comprehensive features and use them in path planning to achieve predictable human-like driving. For instance, another high-level driving behavior that is challenging to implement is the

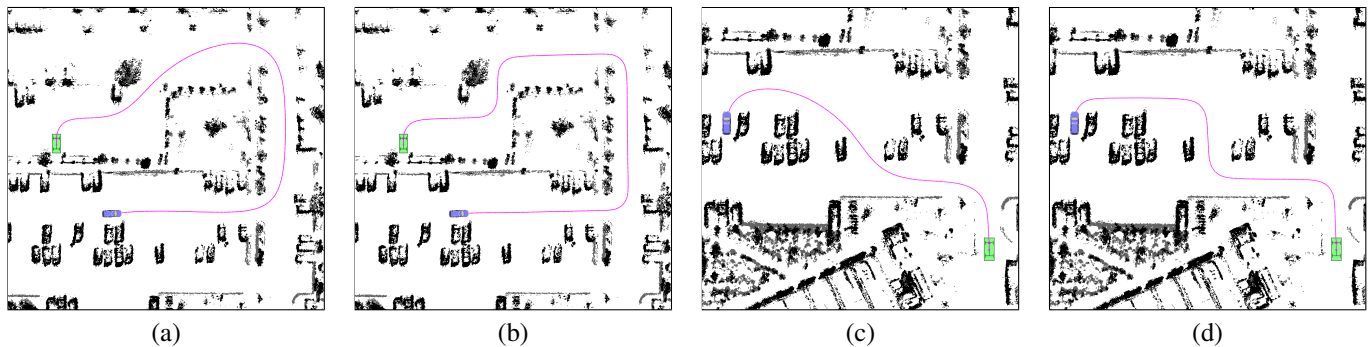


Fig. 8. Path planning using principal directions. a) and c) Principal directions are not used. b) and d) Planning favors trajectories that are aligned with principal directions of the environment.

adherence to the established convention of passing oncoming traffic on the correct side. A straightforward method of partitioning the space into “left” and “right” regions can be developed by using the Voronoi diagram [2] of the obstacle map and labeling points based on which side of the Voronoi edge they fall. However, a straightforward application of this method is brittle with respect to sensor noise, producing fragmented obstacle maps that can lead to highly irregular Voronoi diagrams. A more sophisticated method is needed for robust identification of drivable lanes in order for the robot to correctly handle situations involving other cars.

Another interesting thread of our current and future work lies along the direction of combining obstacle data with visual camera data as well as surface-reflectivity data from LIDARs for detecting and recognizing more advanced features in unknown environments (e.g., lane markings, curbs, other vehicles) and their use in path planning.

REFERENCES

- [1] The 2005 DARPA Grand Challenge: The Great Robot Race. Springer, 2005.
- [2] F. Aurenhammer. Voronoi diagrams – a survey of a fundamental geometric data structure. *ACM Computing Surveys*, 23(3):345–405, 1991.
- [3] J. Barraquand, B. Langlois, and J.-C. Latombe. Robot motion planning with many degrees of freedom and dynamic constraints. In *The fifth international symposium on Robotics research*, pages 435–444, Cambridge, MA, USA, 1990. MIT Press.
- [4] O. Brock and O. Khatib. High-speed navigation using the global dynamic window approach. In *International Conference on Robotics and Automation (ICRA)*, 1999.
- [5] J. Canny. A computational approach to edge detection. *IEEE Trans. Pattern Anal. Mach. Intell.*, 8(6):679–698, November 1986.
- [6] R. Chellappa and A. Jain. *Markov Random Fields: Theory and Applications*. Academic Press, Orlando, FL, 1993.
- [7] E. D. Dickmanns and A. Zapp. Autonomous high speed road vehicle guidance by computer vision. *Triennial World Congress of the International Federation of Automatic Control*, 4:221–226, July 1987 1987.
- [8] J. Diebel and S. Thrun. An application of markov random fields to range sensing. In Y. Weiss, B. Schölkopf, and J. Platt, editors, *Advances in Neural Information Processing Systems 18*, pages 291–298. MIT Press, Cambridge, MA, 2006.
- [9] J. R. Diebel, S. Thrun, and M. Bruenig. A bayesian method for probable surface reconstruction and decimation. *ACM Trans. Graph.*, 25(1):39–59, 2006.
- [10] T. Ersson and X. Hu. Path planning and navigation of mobile robots in unknown environments. In *IEEE International Conference on Intelligent Robots and Systems (IROS)*, 2001.
- [11] D. Ferguson and A. Stentz. Field d*: An interpolation-based path planner and replanner. In *Proceedings of the Int. Symp. on Robotics Research (ISRR)*, October 2005.
- [12] M. R. Hestenes and E. Stiefel. Methods of conjugate gradients for solving linear systems. *Journal of Research of the National Bureau of Standards*, 49:409436, 1952.
- [13] L. Kavraki, P. Svestka, J.-C. Latombe, and M. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. Technical Report CS-TR-94-1519, Department of Information and Computing Sciences, Utrecht University, 1994.
- [14] A. Kelly. An intelligent predictive control approach to the high speed cross country autonomous navigation problem, 1995.
- [15] S. Koenig and M. Likhachev. Improved fast replanning for robot navigation in unknown terrain. In *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2002.
- [16] S. Kolski, D. Ferguson, M. Bellino, and R. Siegwart. Autonomous driving in structured and unstructured environments. In *IEEE Intelligent Vehicles Symposium*, 2006.
- [17] S. LaValle. Rapidly-exploring random trees: A new tool for path planning, 1998.
- [18] S. Z. Li. Markov random field models in computer vision. In *ECCV '94: Proceedings of the Third European Conference-Volume II on Computer Vision*, pages 361–370, London, UK, 1994. Springer-Verlag.
- [19] E. Plaku, L. Kavraki, and M. Vardi. Discrete search leading continuous exploration for kinodynamic motion planning. In *Robotics: Science and Systems*, June 2007.
- [20] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C: The Art of Scientific Computing*. Cambridge University Press, New York, NY, USA, 1992.
- [21] S. Singh, R. Simmons, T. Smith, A. T. Stentz, V. Verma, A. Yahja, and K. Schwahr. Recent progress in local and global traversability for planetary rovers. In *Proceedings of the IEEE International Conference on Robotics and Automation, 2000*. IEEE, April 2000.
- [22] A. T. Stentz and M. Hebert. A complete navigation system for goal acquisition in unknown environments. In *Proceedings 1995 IEEE/RSJ International Conference On Intelligent Robotic Systems (IROS '95)*, volume 1, pages 425 – 432, August 1995.
- [23] C. Thorpe, T. Jochem, and D. Pomerleau. Automated highway and the free agent demonstration. In *Robotics Research – International Symposium*, volume 8, 1998.
- [24] S. Tsugawa, N. Watanabe, and H. Fujii. Super smart vehicle systemits concept and preliminary works. *Vehicle Navigation and Information Systems*, 2:269 –277, October 1991.
- [25] P. Varaiya. Smart cars on smart roads: Problems of control. *IEEE Transactions on Automatic Control*, 38(2), February 1993.
- [26] G. Winkler. *Image Analysis, Random Fields and Markov Chain Monte Carlo Methods: A Mathematical Introduction (Stochastic Modelling and Applied Probability)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2006.
- [27] L. Xu and E. Oja. Randomized hough transform (rht): basic mechanisms, algorithms, and computational complexities. *CVGIP: Image Understanding*, 57(2):131–154, 1993.

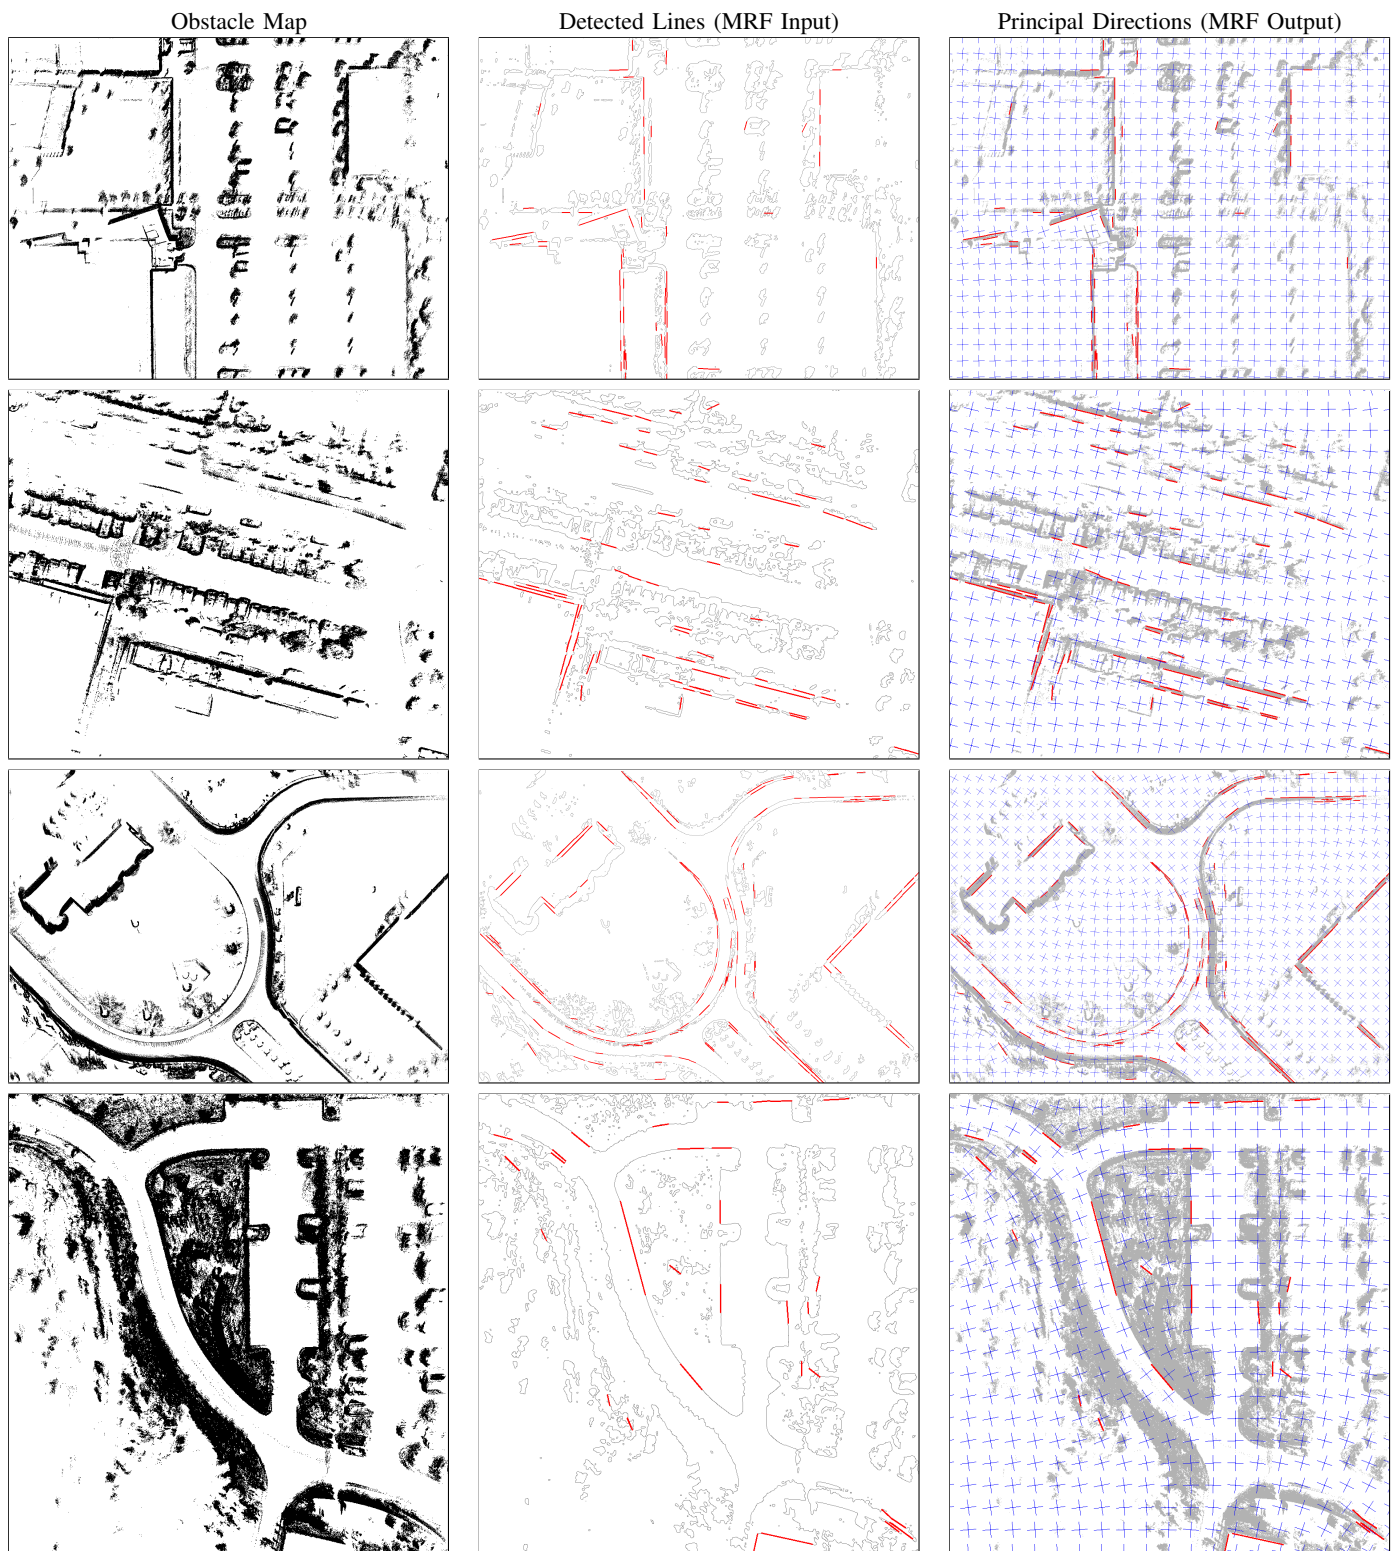


Fig. 9. Additional results of computing principal directions in several real driving environments.