Advances in Probabilistic Modeling: Applications of Stochastic Geometry

By Martin Adams, Ba-Ngu Vo, and Ronald Mahler

Stochastic geometry is an established branch of mathematics that studies uncertainty in geometric structures [1], [2] and is, therefore, a befitting framework for autonomous robotic mapping and the well-known simultaneous localization and mapping (SLAM) problem, where the fundamental concern is succinctly captured in the title of the 1988 seminal paper by Hugh Durrant-Whyte, "Uncertain Geometry in Robotics" [3]. The theory of random sets has long been used by statisticians in many diverse applications, including agriculture, geology, and epidemiology [1], [2], [4]–[6]. In addition, there has been substantial work by probabilists and statisticians in point process filtering, such as those by Singh et al. [7] and Caron et al. [8]. Applications of this are random point pattern methods for multiple object recognition in image analysis [9], [10], and recent work based on random set analysis by Vo et al. [11], which laid the foundations for set-based multiobject visual tracking by Hoseinnezhad et al. [12], [13]. The application of random sets in multitarget tracking has led to the development of finite set statistics (FISST), which provides the basis for novel filters, such as the probability hypothesis density (PHD) filter [14]–[16] and the cardinalized (C)-PHD filter [17], which recently attracted considerable research

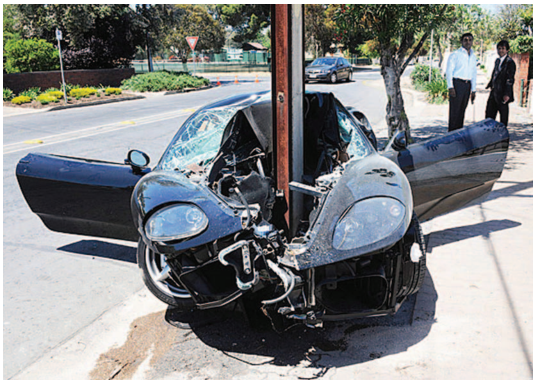

Figure 1. The importance of object number estimation in navigation tasks.

interest as well as deployment in commercial applications.

As noted in the field of multitarget filtering by Mahler [17, p. 571]:

> Having a good estimate of target number is half the battle in multitarget tracking. If one has 1,000 measurements but we know that roughly 900 of them are false alarms, then the problem of detecting the actual targets has been greatly simplified.

The articles in this special issue advocate that the same principle applies to feature detection and autonomous mapping in robotics, where, instead of referring to the problem of target estimation, the problem of map feature or environmental object estimation are of concern. From here on, map features, targets, and environmental objects of interest will simply be referred to as "features." In the case of robotic mapping and SLAM, realistic feature detection algorithms produce false alarms and missed detections, and estimating the true number of map features is, therefore, central to these problems.

A philosophy often encountered within the SLAM community is that the number of estimated map features is not important in SLAM, provided that enough are estimated to provide successful robot location estimates. In response to this, the reader is referred to Figure 1, in which a human driver has clearly not estimated the correct number of objects within his/her environment. Unfortunate accidents aside, failing to correctly estimate the true number of objects or features that have passed through the field(s) of view of a vehicle's sensor(s) can only be detrimental to the location estimation performance of any SLAM algorithm. This special issue, therefore, addresses the concept of Bayes optimality for estimation with an unknown feature number by formulating autonomous mapping, SLAM, and general tracking algorithms as random finite set (RFS) estimation problems.

An RFS is simply a finite-set-valued random variable. Similar to random vectors, the probability density (if it exists) is a very useful descriptor of an RFS, especially in filtering and estimation. However, the space of finite sets does not inherit the usual Euclidean notion of integration and density. Hence, the standard tools for random vectors are not appropriate for random finite sets. Mahler's FISSTs provide practical mathematical tools and principled approximations for dealing with RFSs

Digital Object Identifier 10.1109/MRA.2014.2314018
Date of publication: 10 June 2014

[14], [18] based on a notion of integration and density that is consistent with point process theory [19]. Therefore, what are the advances in the applications of stochastic geometry advocated in this special issue? In contrast to the state-of-the-art vector- based implementations of Bayesian filters, which require separate filters/routines to manage and associate measurements to features, the use of RFSs unifies the independent filters adopted by previous solutions through the recursive propagation of a distribution of an RFS of features. This allows the joint propagation of the estimated feature density to take place and, in the case of SLAM, leads to optimal map estimates in the presence of unknown map size, spurious measurements, feature detection, and data association uncertainty. The proposed framework further allows for the joint treatment of error in feature number and location estimates as it jointly propagates both the estimate of the number of features and their corresponding states and, consequently, eliminates the need for feature management and association algorithms.

Summary of Random Set-Based Implementations in Robotics

Mullane et al. [14] first applied the random set concept to the SLAM problem, in which a first-order random set statistic—the PHD—was used, from which the joint vehicle and feature-based map estimates could be extracted. Referred to as a *brute force* approach, it demonstrated a viable RFS-based SLAM solution. For environments with a significant number of features, it is, however, computationally intractable. Therefore, a more elegant and computationally tractable RFS solution, based on Rao–Blackwellization, was published in 2010 and further elaborated on in 2011, in which the CPHD and multitarget multi-Bernoulli SLAM filtering concepts were presented [21]–[23]. A simplified version of these publications comprises the first article of this special issue.

Lee et al. [24] addressed the SLAM problem with a single cluster (SC) PHD filter, which also utilized Rao–Blackwellization but generated a measurement likelihood function for trajectory weighting in a different manner. An extension of this article, applied to underwater SLAM, is the focus of the second article in this special issue. Also in 2012, Moratuwage et al. [25] extended the RFS concept to multivehicle SLAM, providing demonstrations with two vehicles that collaborate to estimate a global map based on PHD filtering along with their own trajectories. An extension of this work comprises the fourth article in this special issue. Finally, also in 2012, Adams et al. [26] demonstrated constant false alarm rate and scan integration feature detection techniques to provide principled detection statistics for a short-range radar in a Rao–Blackwellized PHD filter SLAM framework. To complement the multiple publications on the applications of FISST to robotic-based problems, the first international workshop on stochastic geometry in SLAM was held at the 2012 IEEE International Conference on Robotics and Automation (ICRA 2012), in Minneapolis St. Paul, Minnesota. This full-day workshop was opened by the founder of FISST, Ronald Mahler, who presented the foundations behind many of the FISST-based filtering concepts. It also provided a forum for various FISST-based robotic mapping, navigation, and control presentations, some of which are extended in this special issue. During this workshop, many of the presenters and members of the audience indicated the importance of publicizing the recent advances in the application of stochastic geometry to robotic problems, which has instantiated this special issue.

Overview of the Special Issue

The first of the six articles in this special issue, "SLAM Gets a PHD," focuses on a SLAM implementation that uses the most basic Bayesian set-based estimator—the PHD filter. The article first demonstrates the random nature of detections in sensing modalities as diverse as radar, laser range finders, and vision. The information, referred to as *features*, provided by any feature detection algorithm and based on any sensing type, is prone to randomness both in the quantity of the detected features and their attributes, such as range and bearing, or image-based quantities, such as contrast levels. It shows that realistic measurement uncertainty comprises detection uncertainty in the form of false alarms and missed features as well as the usually considered spatial (e.g., range and bearing) uncertainty. The ability to account for all of these in a joint manner provides the motivation for remodeling the SLAM problem as a set- rather than a vector-based framework. The concepts of RFSs are introduced, and the implementation of the PHD filter, in the form of manipulating sums of Gaussians, is demonstrated through the use of simple block diagrams. A marine environment, in which an autonomous kayak estimates the number and location of objects on the sea surface as well as its own trajectory, provides the complex setting for SLAM trials. These are based on the presented sum of the Gaussians PHD SLAM algorithm, with performance comparisons being made with state-of-the-art multiple-hypothesis (MH) FastSLAM.

The SLAM problem is treated as an SC process in the second article titled "SLAM with SC-PHD Filters." Here, SLAM is defined as a particular type of cluster process in which the configuration of the map features is a daughter process conditioned on the state of the vehicle, which is represented as a single parent process. The SC-PHD filter approach can be separated into a parent and a conditional daughter term, allowing a hybrid particle filter and Gaussian mixture approach to be used for SLAM in a manner similar to that proposed in the first article. However, in contrast to the first article, an SC process, rather than a Poisson process, is assumed on the prior map feature cardinality distribution. An implementation of the concept is demonstrated on an underwater robotic vehicle, the Girona 500, which utilizes stereo imagery and a speeded-up robust feature (SURF) detector to detect key points in images for underwater SLAM.

The third article, "Playing Fetch with Your Robot," is based on a Segway robotic platform and again uses vision for the detection of an unknown number of objects, this time based on shape and color matching. These are scattered about an environment for the robot to locate, collect, and return to the user. This work uses a grid with cells containing occupancy probability values. An RFS is used to represent a set of labels of occupied cells together with a further set, which contains every combination of the RFSs from zero to the maximum number of objects that can be tracked. In contrast to the first two articles, a Bayesian filter iterates over the distribution of the RFSs to estimate the varying number of object locations. This information is then used to automatically instruct the robot to move to maximize its immediate information gain—a technique referred to as *information surfing*. The robot is controlled to move in the maximum gradient of mutual information between the sensor readings and the cell-based object position estimates based on a quantity known as the Rényi divergence. This moves the robot and the camera's field of view into the direction of objects to be fetched.

The focus of the fourth article, "RFS Collaborative Multivehicle SLAM," is collaborative multivehicle (CM) SLAM under an RFS framework. Formulated for two vehicles, which collaborate to build a single global map and estimates of both trajectories, this article introduces the concept of the general multisensor PHD update. This update requires the union of the set-based measurements from each vehicle to be partitioned into binary subsets. The article highlights the computational problem that results in the case of many robots, since all possible combinations of these subsets would be necessary to form the Bayes optimal CM-SLAM measurement update. To test the concepts with two robots, simulations demonstrate the robustness of the RFS-based solution under varying degrees of clutter (false alarms). In a real experiment, a parking lot with moving people provides the scene for the dynamic environment where the proposed RFS based multivehicle SLAM solution is compared against a state-of-the-art CM-SLAM solution that depends on external feature association and management routines.

The fifth article titled "Group Mapping," again addresses multirobot applications and focuses on grid-based map fusion. Occupancy grid maps from multiple vehicles are merged without prior knowledge of their relative transformations. This is achieved through graph matching, in which the graph is a topological representation of the map, and is based on a generalized Voronoi diagram (GVD). Referred to as *map fusion*, the approach demonstrated in this article exploits the uniqueness of GVDs to combine large maps. The confidence values associated with certain areas of each robot's map is further encoded into the topological structure by building a probabilistic (P)GVD. Map matching takes place via a two-dimensional correlation to match laser range finder-based edges. The probabilistic nature of the PGVDs allows areas of the maps with higher certainty to be preferentially matched. This technique is verified through four experiments: the first is based on a publicly available data set; the second is based on two real indoor vehicles communicating map information between them; the third is based on three vehicles operating in a larger indoor environment; and the fourth is based on a simulated, highly cluttered environment.

The estimation of the location and number of objects/features that can each generate multiple sensor detections due to their large size and/or occlusions is the subject of the final article "Random Set Methods." This article demonstrates how the assumptions of the earlier articles, in which single objects are assumed to yield single detections, can be relaxed, allowing for multiple detections per object. This scenario naturally lends itself to the RFS concept since a subset of all of the detections can now result from each object. This, in turn, requires a partitioning of the full set of detections into subsets, the union of which comprises the full detection set. During the PHD filter measurement update, each partition requires a likelihood corresponding to how probable that group of measurements stems from a single object. There are many ways in which such partitions can be formed, all of which should be theoretically considered for Bayes optimality. This article provides algorithms for limiting the number of partitions to computationally manageable levels without sacrificing estimation performance. By considering different alternatives for the measurement model, based, for example, on assumed geometric extended object shapes, a multiobject tracking PHD filter implementation is demonstrated. The experimental results demonstrate the ability of the filter to estimate the quantity and location of an unknown number of pedestrians based on laser range finder data, even when pedestrians occlude each other. This article models the probability of detection as nonhomogeneous in the sensor surveillance region based on object location estimates, yielding good estimates of pedestrian numbers even when they are completely occluded.

Looking Forward

Stochastic geometry has been applied in diverse engineering fields for many decades, but only in the last decade have the tools of FISST become available for set-based estimation applications. This special issue is largely a collection of robotic applications based on these recently formulated tools. Within the field of robotics, many avenues exist for further research based on FISST. These include improved sensor models that take into account object occlusions; generalized mapping concepts, such as semantic maps; and active navigation, in which vehicles are autonomously commanded to maximize their information gain. We hope that this special issue provides motivation for further advances in the use of stochastic geometry in SLAM and general robotics applications.

References

[1] D. Stoyan, W. S. Kendall, and J. Mecke, *Stochastic Geometry and its Applications*, 2nd ed. New York: Wiley, 1995.

[2] I. Molchanov, "Probability and its applications," in *Theory of Random Sets*, London, U.K.: Springer, 2005.

[3] H. F. Durrant-Whyte, "Uncertain geometry in robotics," *IEEE J. Robot. Automat.*, vol. 4, no. 1, pp. 23–31, 1988.

[4] G. Matheron, *Random Sets and Integral Geometry* (Probability and Statistics series). New York: Wiley, 1975.

[5] D. J. Daley and D. Vere-Jones, "Probability and its applications," in *An Introduction to the Theory of Point Processes: Volume I: Elementary Theory and Methods*, 2nd ed. New York: Springer, 1988.

[6] J. F. C. Kingman, *Poisson Processes* (Oxford Studies in Probability 3). London, U.K.: Oxford Univ. Press, 1993.

[7] S. S. Singh, B.-N. Vo, A. Baddeley, and S. Zuyev, "Filters for spatial point processes," *SIAM J. Control Optim.*, vol. 48, no. 4, pp. 2275–2295, 2009.

[8] F. Caron, P. del Moral, A. Doucet, and M. Pace, "On the conditional distributions of spatial point processes," *Adv. Appl. Probab.*, vol. 43, no. 2, pp. 301–307, 2011.

[9] A. J. Baddeley and M. N. M. Van Lieshout, "Stochastic geometry models in high level vision," *J. Appl. Stat.*, vol. 20, nos. 5–6, pp. 231–256, 1993.

[10] A. Dasgupta and A. E. Raftery, "Detecting features in spatial point processes with clutter via model based clustering," *J. Amer. Statist. Assoc.*, vol. 93, no. 441, pp. 294–302, 1998.

[11] B.-N. Vo, B.-T. Vo, N.-T. Pham, and D. Suter, "Joint detection and estimation of multiple objects from image observations," *IEEE Trans. Signal Processing*, vol. 58, no. 10, pp. 5129–5141, Oct. 2010.

[12] R. Hoseinnezhad, B.-N. Vo, and B.-T. Vo, "Visual tracking in background subtracted image sequences via multi-Bernoulli filtering," *IEEE Trans. Signal Processing*, vol. 61, no. 2, pp. 392–397, 2013.

[13] R. Hoseinnezhad, B.-N. Vo, B.-T. Vo, and D. Suter, "Visual tracking of numerous targets via multi-Bernoulli filtering of image data," *Pattern Recognit.*, vol. 45, no. 10, pp. 3625–3635, 2012.

[14] R. Mahler, "Multi-target Bayes filtering via first-order multi-target moments," *IEEE Trans. Aerospace Electron. Syst.*, vol. 4, no. 39, pp. 1152–1178, Oct. 2003.

[15] B.-N. Vo and W.-K. Ma, "A closed form solution for the probability hypothesis density filter," in *Proc. Int. Conf. Information Fusion*, Philadelphia, PA, 2005, pp. 1–8.

[16] B.-N. Vo and W.-K. Ma, "The Gaussian mixture probability hypothesis density filter," *IEEE Trans. Signal Processing*, vol. 54, no. 11, pp. 4091–4104, 2006.

[17] B. T. Vo, B.-N. Vo, and A. Cantoni, "Analytic implementations of the cardinalized probability hypothesis density filter," *IEEE Trans. Signal Processing*, vol. 55, no. 7, pp. 3553–3567, July 2007.

[18] R. Mahler, *Statistical Multisource Multitarget Information Fusion*. Norwood, MA: Artech House, 2007.

[19] B. Vo, S. Singh, and A. Doucet, "Sequential Monte Carlo methods for multi-target filtering with random finite sets," *IEEE Trans. Aerospace Electron. Syst.*, vol. 41, no. 4, pp. 1224–1245, Oct. 2005.

[20] J. Mullane, B.-N. Vo, M. Adams, and W. S. Wijesoma, "A random set formulation for Bayesian SLAM," in *Proc. IEEE/RSJ Int. Conf. Intelligent Robots Systems*, France, Sept. 2008, pp. 1043–1049.

[21] J. Mullane, B.-N. Vo, and M. Adams, "Rao-Blackwellised PHD SLAM," in *Proc. IEEE Int. Conf. Robotics Automation*, Alaska, May 2010, pp. 5410–5416.

[22] J. Mullane, B.-N. Vo, M. D. Adams, and B. T. Vo, "A random-finite-set approach to Bayesian SLAM," *IEEE Trans. Robot.*, vol. 27, no. 2, pp. 268–282, Apr. 2011.

[23] J. Mullane, B.-N. Vo, M. Adams, and B.-T. Vo, *Random Finite Sets for Robot Mapping and SLAM* (Springer Tracts in Advanced Robotics 72). Berlin, Germany: Springer, 2011.

[24] C. S. Lee, D. E. Clark, and J. Salvi, "SLAM with single cluster PHD filters," in *Proc. IEEE Int. Conf. Robotics Automation*, St. Paul, Minnesota, May 2012, pp. 2096–2101.

[25] D. Moratuwage, B.-N. Vo, and D. Wang, "A hierarchical approach to the multi-vehicle slam problem," in *Proc. 15th Int. Conf. Information Fusion*, Singapore, July 2012, pp. 1119–1125.

[26] M. Adams, J. Mullane, E. Jose, and B.-N. Vo, *Robotic Navigation and Mapping with Radar*. Boston, MA: Artech House, 2012.