Study, Design, and Evaluation of Exploration Strategies for Autonomous Mobile Robots

ALBERTO QUATTRINI LI, Politecnico di Milano

My research activity deals mainly with the problem of defining exploration strategies that drive autonomous mobile robots around environments to be discovered. One of my contributions is towards bridging the gap between theory and practice for exploration strategies, modeling problems related to exploration to theoretically analyze them, and applying obtained results to more realistic scenarios. Another contribution is the use of semantic information in multirobot systems for more informed decisions during exploration. Further contributions provide insights for a more mature experimental assessment of exploring robotic systems.

Categories and Subject Descriptors: I.2 [**Artificial Intelligence**]: Robotics

General Terms: Design, Algorithms, Performance

Additional Key Words and Phrases: Autonomous mobile robots, exploration strategies, coordination methods

ACM Reference Format:
Alberto Quattrini Li, 2015. Study, Design, and Evaluation of Exploration Strategies for Autonomous Mobile Robots. *ACM* V, N, Article A (January YYYY), 4 pages.
DOI: http://dx.doi.org/10.1145/0000000.0000000

Recently, autonomous mobile robotics has seen a wide spread progress, due to its possible applications (e.g., search and rescue). Designers face several challenges during the development of systems of autonomous mobile robots, from low level issues (e.g., actuators) to high level issues (e.g., navigation) [Siegwart et al. 2011]. One of the most important aspects that affect system's performance is the set of techniques that let robots decide next locations to reach, possibly in a coordinated way, according to their current knowledge of the world for autonomously executing some tasks.

My research activity mainly focus on *exploration*, namely the process of exploring an initially unknown environment by one or more mobile robots, in order to discover and map its features, for relevant applications like map building and search and rescue. Multirobot exploration basically involves the following steps: (a) perceive the surrounding environment, (b) integrate perceived data in a map representing the environment known so far, (c) decide where to go next (*exploration strategies*) and who goes where (*coordination methods*), (d) go to the destination locations chosen. Specifically, my focus is on step (c). Note that, to obtain a large amount of results that are easily reproducible, I usually perform experiments with realistic simulators.

Given the importance of the exploration problem, in the following I outline three main research goals and, for each, I provide a brief overview of the gaps in the literature, my contributions so far, and interesting future research directions.

To contribute to bridge the gap between theory and practice. Exploration strategies are usually defined following two rather different approaches. On the one

Author's address: A. Quattrini Li, Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Via Ponzio 34/5, 20133 Milano, Italy.

DOI: http://dx.doi.org/10.1145/0000000.0000000

hand, they are defined in practical contexts of real (or realistically simulated) robots and are empirically assessed by testing them in some environments, as for example in [Julia et al. 2012]. On the other hand, exploration strategies are defined in theoretical settings. The environments could be represented geometrically, as usually done by the computational geometry community. In this kind of works, some assumptions that are not fully realistic are made (e.g., line-of-sight visibility [Fekete and Schmidt 2010]). Another type of representation is based on graphs, as typically assumed by theoretical computer science community to disregard environments' geometry for focusing on their topological and combinatorial aspects (e.g., [Tovey and Koenig 2003]). In both approaches, proposed methods are assessed using theoretical tools as worst-case bounds and competitive ratio in some classes of environments.

Against this background, I contribute to define the problem of calculating the optimal offline exploration paths under some realistic assumptions, namely, robot with time-discrete and limited perception and environment represented as a grid (some of the results are presented in [Quattrini Li et al. 2012]). I analyze the relation between the solutions found in such discretization and its continuous counterpart and formulate the discrete problem as a search problem. Thus, I develop the first algorithm to find the (approximated) optimal exploration path. Simulation results show the viability of this approach for realistic environments. Moreover, I theoretically study the worst and average cases on the traveled distance to explore graph-based environments, of some exploration strategies that consider distance and information gain in selecting the next destination location. The obtained theoretical results show that, in the worst case, considering also information gain does not provide any advantage over considering only distance, while it does in the average case on graphs modeling realistic indoor environments. This kind of results contributes to strengthen the experimental results obtained with real (and realistically simulated) exploring robots.

Further research directions that could be worth investigation are to address other realistic cases, like constrained exploration, in which a robot should stay within a certain distance from the base station [Duncan et al. 2006], or that of multiple robots [Brass et al. 2011]. Both cases can be analyzed using artificial intelligence techniques to find optimal strategies and bounds on the performance of some state-of-the-art exploration strategies on graph- or geometrically-based environments.

To improve exploration strategies and coordination methods. Most of the exploration strategies and coordination methods proposed in literature (e.g., [Gonzáles-Baños and Latombe 2002; Basilico and Amigoni 2011]) base their decisions only on the current *metric map*, which represents the spatial features of the environment, like the position of obstacles. In the last years, several methods (e.g., [Mozos et al. 2005]) have been proposed to build *semantic maps* that associate semantic labels (e.g., 'corridor' or 'room') to portion of the underlying metric map. Despite the great effort in *constructing* semantic maps, the study of their *use* for exploration is still rather limited.

To fill this gap, I define exploration strategies and coordination methods that embed information coming from semantic maps. This allows to privilege some specific areas of the environment. For example, if robots know that an area of an environment is labeled as 'corridor', then that area should be privileged and more than one robot should be allocated to that area, so that the exploration of the environment is speeded up, as rooms are typically attached to corridors. I experimentally show that there is a significant improvement about the exploration of *relevant* and *total* areas of indoor environments within a given time interval, when *a priori* information about the relevant areas of the environment is available [Cipolleschi et al. 2013].

One possible extension that could be interesting pursuing is taking into account the distribution of probability about the location of the victims, starting from results of [Aydemir et al. 2013]. Also, an interesting study is to derive a theoretically-grounded way to determine the number of robots that will work best for a given environment, starting from the insights given by the work in [Nieto-Granda et al. 2014]. Further, it is relevant to study scenarios in which connectivity should be maintained over time. An interesting approach to start with is that of [Rekleitis et al. 2004], which describes the multirobot coverage problem. Another direction of interest is the investigation of how to include a lookahead more than one of the next candidate locations (i.e., each robot plans how to reach a sequence of candidate locations) or path optimization, starting from results of [Tovar et al. 2006], so that chosen paths are the most informative ones.

To improve the experimental assessment of multirobot exploration systems. A lively debate on good experimental methodologies is currently ongoing in the autonomous robotics community, as they have not reached yet a maturity level comparable to that of other disciplines [Amigoni and Schiaffonati 2014].

With the current experimental evaluation almost exclusively based on relative comparisons in some test environments, it is difficult to assess how much room for improvement an exploration system has. The method I propose in [Quattrini Li et al. 2012] could serve as a tool for complementing the evaluation of online exploration strategies for autonomous mobile robots by having an optimal reference against which is possible to compute the competitive ratio of an exploration strategy in an environment.

Another issue is the difficulty in reproducing experiments as parameters are usually not reported in the descriptions of experiments, and thus it is not clear what factors impact the performance of exploration. I contribute to fill this gap by quantitatively analyzing some controllable factors on exploration, providing some insights that could be useful for a roboticist who has to set these parameters. Specifically, I quantitatively assess the relative influence of exploration strategies and coordination methods on the exploration performance [Amigoni et al. 2012]: experimentally the role of exploration strategies dominates that of coordination methods in determining the performance of an exploring multirobot system in a highly structured indoor environment, while the situation is reversed in a less structured indoor environment. I also analyze the impact of different perception/decision timings [Amigoni et al. 2013]. The simulations' results confirm the intuitive idea that the best performance is obtained with fast-paced perceptions and decisions, but also suggest some trade-offs for the values of perception and decision frequencies in some settings.

Future interesting research on the improvement of the evaluation includes the development of systems that automatically generate realistic environments so that robotic systems can be easily tested in simulations. An initial attempt in this direction has been presented in [Amigoni et al. 2014]. This feature can be used to extend, for example, the system for robotic system verification shown in [O'Brien et al. 2014].

Some of my contributions and some of the future research directions could foster the achievement of the long-term goal towards the theoretical and practical definition and the evaluation of exploration strategies and coordination methods for increasing autonomy of mobile robots.

The exploration problem could be intended in a broader perspective in the sense that robots can "explore" other features of the environment. There are several interesting domains that are worth investigation. One interesting problem is that of continuously monitoring temporal-spatial phenomena in partially known environments, in applications like environmental or ocean monitoring [Parker et al. 2013; Smith et al. 2011], surveillance [Carpin et al. 2013; Thakur et al. 2013], olfactory exploration [Marjovi and Marques 2013], and tactile exploration of object properties [Pezzementi et al. 2011]. In this direction, the objective is to design a general framework in order to define general navigation strategies and coordination methods considering general domains, so that the development of a robotic system is eased.

REFERENCES

Francesco Amigoni, Nicola Basilico, and Alberto Quattrini Li. 2012. How much worth is coordination of mobile robots for exploration in search and rescue?. In *Proc. RoboCup*. 106–117.

Francesco Amigoni, Matteo Luperto, and Alberto Quattrini Li. 2014. Towards more realistic indoor environments for the Virtual Robot Competition. In *RoboCup2014 Team Description Papers*.

Francesco Amigoni, Alberto Quattrini Li, and Dirk Holz. 2013. Evaluating the impact of perception and decision timing on autonomous robotic exploration. In *Proc. ECMR*. 68–73.

Francesco Amigoni and Viola Schiaffonati. 2014. Autonomous mobile robots as technical artifacts: A discussion of experimental issues. In *Model-Based Reasoning in Science and Technology*, Lorenzo Magnani (Ed.). Springer, 527–542.

Alper Aydemir, Andrzej Pronobis, Moritz Göbelbecker, and Patric Jensfelt. 2013. Active visual object search in unknown environments using uncertain semantics. *IEEE T Robot* 29, 4 (2013), 986–1002.

Nicola Basilico and Francesco Amigoni. 2011. Exploration strategies based on multi-criteria decision making for searching environments in rescue operations. *Auton Robot* 31, 4 (2011), 401–417.

Peter Brass, Flavio Cabrera-Mora, Andrea Gasparri, and Jizhong Xiao. 2011. Multirobot tree and graph exploration. *IEEE T Robot* 27, 4 (2011), 707–717.

Stefano Carpin, Derek Burch, Nicola Basilico, Timothy Chung, and Mathias Kölsch. 2013. Variable Resolution Search with Quadrotors: Theory and Practice. *Journal of Field Robotics* 30, 5 (2013), 685–701.

Riccardo Cipolleschi, Michele Giusto, Alberto Quattrini Li, and Francesco Amigoni. 2013. Semantically-informed coordinated multirobot exploration of relevant areas in search and rescue settings. In *Proc. ECMR*. 216–221.

Christian Duncan, Stephen Kobourov, and Anil Kumar. 2006. Optimal constrained graph exploration. *ACM T Algorithms* 2, 3 (2006), 380–402.

Sándor Fekete and Christiane Schmidt. 2010. Polygon exploration with time-discrete vision. *Comp Geom-Theor Appl* 43, 2 (2010), 148–168.

Héctor Gonzáles-Baños and Jean-Claude Latombe. 2002. Navigation strategies for exploring indoor environments. *Int J Robot Res* 21, 10-11 (2002), 829–848.

Miguel Julia, Arturo Gil, and Óscar Reinoso. 2012. A comparison of path planning strategies for autonomous exploration and mapping of unknown environments. *Auton Robot* 33, 4 (2012), 427–444.

Ali Marjovi and Lino Marques. 2013. Multi-robot topological exploration using olfactory cues. In *Proc. DARS*. 47–60.

Oscar Mozos, Cyrill Stachniss, and Wolfram Burgard. 2005. Supervised learning of places from range data using AdaBoost. In *Proc. ICRA*. 1742–1747.

Carlos Nieto-Granda, John Rogers, and Henrik Christensen. 2014. Coordination strategies for multi-robot exploration and mapping. *Int J Robot Res* 33, 4 (2014), 519–533.

Matthew O'Brien, Ronald Arkin, Dagan Harrington, Damian Lyons, and Shu Jiang. 2014. Automatic verification of autonomous robot missions. In *Proc. SIMPAR*. 462–473.

Lonnie Parker, Richard Coogle, and Ayanna Howard. 2013. Estimation-informed, resource-aware robot navigation for environmental monitoring applications. In *Proc. ICRA*. 1041–1046.

Zachary Pezzementi, Erion Plaku, Caitlin Reyda, and Gregory Hager. 2011. Tactile-Object Recognition From Appearance Information. *IEEE T Robot* 27, 3 (2011), 473–487.

Alberto Quattrini Li, Nicola Basilico, and Francesco Amigoni. 2012. Searching for optimal off-line exploration paths in grid environments. In *Proc. AAAI*. 2060–2066.

Ioannis Rekleitis, Vincent Lee-Shue, Ai Peng New, and Howie Choset. 2004. Limited communication, multi-robot team based coverage. In *Proc. ICRA*, Vol. 4. 3462–3468.

Roland Siegwart, Illah Nourbakhsh, and Davide Scaramuzza. 2011. *Introduction to autonomous mobile robots*. MIT Press.

Ryan Smith, Mac Schwager, Stephen Smith, Daniela Rus, and Gaurav Sukhatme. 2011. Persistent ocean monitoring with underwater gliders: Towards accurate reconstruction of dynamic ocean processes. In *Proc. ICRA*. 1517–1524.

Dinesh Thakur, Maxim Likhachev, James Keller, Vijay Kumar, Vladimir Dobrokhodov, Kevin Jones, Jeff Wurz, and Isaac Kaminer. 2013. Planning for opportunistic surveillance with multiple robots. In *Proc. IROS*. 5750–5757.

Benjamin Tovar, Lourdes Munoz, Rafael Murrieta-Cid, Moisés Alencastre, Raùl Monroy, and Seth Hutchinson. 2006. Planning exploration strategies for simultaneous localization and mapping. *Robot Auton Syst* 54, 4 (2006), 314–331.

Craig Tovey and Sven Koenig. 2003. Improved analysis of greedy mapping. In *Proc. IROS*. 3251–3257.