

Rapid Multi-Robot Exploration with Topometric Maps

Anthony Cowley and Camillo J. Taylor
GRASP Laboratory
University of Pennsylvania, USA
{acowley, cjtaylor}@seas.upenn.edu

Ben Southall
Sarnoff Corporation, USA
bsouthall@sarnoff.com

Abstract—Multi-robot map building has advanced to the point where high quality occupancy grid data may be collected by multiple robots collaborating with only intermittent connectivity. However, the tasking of these agents to most efficiently build the map is a problem that has seen less attention. Unfamiliar, highly cluttered environments can confound exploration strategies that rely solely on occupancy grid frontier identification or even semantic classification methods keyed on geometric features. To reason about partial maps of novel, highly cluttered locations, hypotheses about significant structure in the disposition of free space may be used to guide exploration task assignment. A parsing of map data into *places* with semantic significance to the exploration task provides a foundation from which one may infer an efficient exploration strategy.

I. INTRODUCTION

An encouraging phenomenon of modern probabilistic mapping techniques [1] is just how good the resulting maps look to a person. A combination of filtering techniques to correct for local pose drift and larger scale loop closure operations that maintain global consistency have led to a surfeit of entirely legible occupancy grids collected via SLAM approaches. Mapping techniques have crossed the threshold at which a human looking at an output map would say that it is “good.”

However, this value judgement is predicated on the availability of the kinds of high-level reasoning capabilities humans can bring to bear on the problem of map interpretation. For example, a map of an open field with a large tree in the middle suggests a sparse representation of geometry: you are free to navigate the field howsoever you please, so long as you avoid the tree. On the other hand, a forest packed with trees and unnavigable brush may be more usefully represented by focusing on the available trails: one wishes to avoid losing the trail for the trees.

This latter example, the navigation of a cluttered environment cut through by continuous stretches of traversable terrain, may be seen as an instance of the problem of reasoning about maps whose structure is reflected in the distribution of free space, rather than geometry. The specific instantiation of this problem considered here is the exploration of cluttered office-like environments consisting of junk-lined hallways connecting furnished rooms. Such an environment could, perhaps ideally for exploration purposes, be represented by a floor plan diagram. Such a diagram would not include anything beyond the architectural design of the building: the clutter would be abstract.

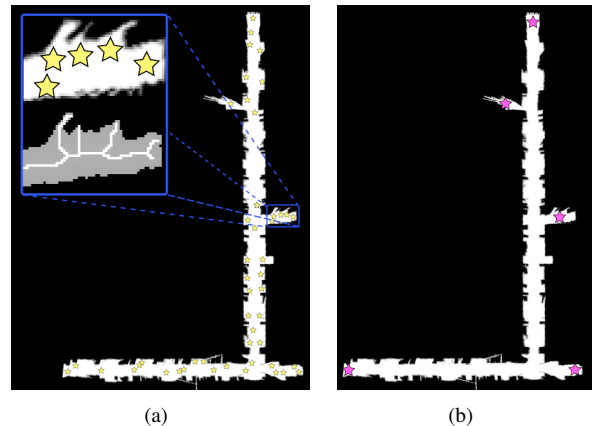


Fig. 1. Occupancy grid with goals, marked by yellow stars, identified by Karto 1.1’s exploration module (a). The inset highlights the complex contour of the occupancy grid and the resulting topological map. Occupancy grid with desired goals, marked by violet stars, that reflect the overall structure of the map (b).

A. A Focus on Exploration

We define exploration to be the incremental creation of a map that approaches a state of being complete and consistent with respect to an idealized floor plan of an environment. Given this goal, we can expect intermediate values to take the form of partial floor plans of the environment. These intermediate values must be rationalized into a form suitable for the tasking of robots to efficiently build out the new map by pushing outward into unexplored territory without getting distracted by the nooks and crannies created by clutter.

Thus we desire a system that, when presented with a partial floor plan whose distinct features may be as of yet only partially observed, is able to produce a set of locations corresponding to the architectural frontiers of the environment that we would like a robot to visit. This discrete set of locations suggests the need for a decomposition of the map into distinct locations, and it is from this division that the frontier representatives may be chosen. In practice, the frontier set will consist of locations farther down halls or into rooms than any robot has yet sensed.

The difficulty in computing such a semantically significant decomposition of cluttered free space is illustrated in Figure 1. The goals chosen for a real map of a cluttered environment by version 1.1 of the Karto mapping and exploration software produced by SRI [2] are shown, as they tend to represent the output of an exploration strategy focused on the occupancy grid frontier. The multitude of



Fig. 2. An environment whose architectural geometry is barely visible.

goals suggested by the occupancy grid frontier bears little relation to the abstract structure of the building, and are not ideally chosen if the goal is to expand the map as quickly as possible.

In contrast, a decomposition of the map into a hierarchical, graph representation directly follows, in particular, from Kuipers’ approach of applying the Spatial Semantic Hierarchy [3]–[5] to map understanding. The exploration strategy espoused in that work is one of opportunistically identifying the place containing an unexplained terminal node in the topological representation of the map. This representation style focuses on the connectivity between discrete places in the map.

Focusing on rapid exploration, we wish for the robot allocation strategy to be initially biased toward “Exploration” over “Wandering” as defined in Brooks’ subsumption architecture [6]: find the overall disposition of free space, then fill in the gaps. While bare occupancy grids may present too noisy an estimate of the map frontier to reliably make this distinction due to unstructured geometry, purely topological methods can struggle with accidental complexity induced by clutter and furniture fracturing the free space. Both of these phenomena are highlighted by the inset image in Figure 1(a).

We propose a fusion of the two approaches that leverages a stack of processing stages to incrementally winnow the set of potential goals using *both* topological and metric-based evaluations of the available map data. This parsing process utilizes what may be referred to, with a slight abuse of terminology, as a topometric map interpretation.

II. RAPID MAPPING

Rapid exploration should take advantage of the maximal available sensing range of all sensors, and must be able to operate in unknown environments in an on-line fashion. To this end, we have developed an analysis procedure for processing range data gathered in highly cluttered environments consisting of rooms, clutter, and corridors hundreds of meters in length. The analysis procedure can run in under a second (processing time is around 200ms on a modern laptop for the included examples), resulting in a system that remains




Fig. 3. 3D point cloud data rendered over an occupancy grid generated by Karto. Exploration goals selected by Karto are shown in yellow, while exploration goals selected by the proposed method are marked by green octahedra and attached integer labels.




Fig. 4. One of the platforms used in the experiments. The two laser range finders generate 3D point cloud data as the robot moves through the environment.

responsive to newly discovered exploration frontiers. A representative hallway scene from the experimental environment is shown in Figure 2.

The input data for all processing was gathered by two Pioneer platforms outfitted with dual Hokuyo UTM-30LX laser range finders mounted at right angles such that one forward-facing sensor scans a plane parallel to the ground, while a second, coronal sensor scans a plane whose normal is the forward motion direction of the robot, as shown in Figure 4. These two scanners in combination offer both long range sensing (30m) and dense 3D point geometry. Occupancy grids, and baseline exploration goal identification, are computed by version 1.1 of the Karto mapping and exploration library produced by SRI [2]. An example of the data generated by the system is shown in Figure 3.

A. Entropy Compass

The map parsing process is bootstrapped by the observation that office-like environments tend to be aligned to a pair of orthogonal axes. This alignment defines a notion of cardinal directions (up/down, left/right), with a quarter-revolution ambiguity: any 90° rotation of the map is just as good as another. While this alignment is useful when presenting maps to human users, it also serves to provide a strong prior to geometry recognition tasks such as wall




Fig. 5. Entropy vs. orientation of occupancy grid projections showing a dominant orientation at 47° .

extraction.

The proposed method of determining the dominant orientation of an environment presumed to be significantly rectilinear is to consider histograms of projections of occupancy grid data. For a given orientation, θ , 2D point data from the occupancy grid is projected onto a line and binned into a histogram, H_θ , with bin extents, bin_i .

$$H_{\theta,i} = \sum_{p \in points} count_{\theta,i}(p)$$

$$count_{\theta,i}(p) = \begin{cases} 1 & \text{if } \sin(\theta)p_x + \cos(\theta)p_y \in bin_i \\ 0 & \text{otherwise} \end{cases}$$

To determine whether this projection of the map data is aligned with a dominant direction of the building geometry, the entropy of the histogram, $-\sum_i H_{\theta,i} \log(H_{\theta,i})$, associated with each projection angle is summed with its orthogonal partner, $\theta + 90^\circ$, yielding a measure whose minimum coincides when the projection and building orientations (with respect to an arbitrary coordinate frame established by the robots) coincide.

The intuition behind this approach is that most walls in a building lie on lines that are either parallel or orthogonal to each other. Hallways then provide strong reinforcement for a specific orientation, while adding unstructured noise to the orthogonal projection. A representative plot of the described measure, shown in Figure 5, displays the characteristic local minima at a pair of orientations separated by a quarter revolution at 47° and 137° .

B. Wall Extraction

The canonical orientation provided by the entropy compass informs a straightforward mechanism for extracting line segments representing wall geometry. The histogram for a given orientation is considered, for instance the histogram associated with the 47° projection in Figure 5, and the local maxima of the histogram bins are identified, as these bins correspond to likely walls in the mapped environment. The coordinates of occupancy grid cells that project into the locally maximal histogram bins may then be sorted along the axis of projection and broken into contiguous runs. These

contiguous runs represent collinear wall segments, and can be filtered by requiring that a meaningful wall segment be above some minimum length (e.g. one meter).

The robustness of this method may be improved by scanning in a plane perpendicular to the ground plane. Such scan data may be used to identify occupancy grid cells for which a vertical column of points have been detected. An example is to require that a particular 2D occupancy grid cell be sensed as occupied at three different heights separated by at least half meter intervals. The resulting wall occupancy information can be efficiently represented using a data structure tuned for sparse tenancy. The implementation chosen for this system is a Patricia tree [7] storing an integer for each partially occupied grid cell. Individual bits of the stored integer are flipped for each detected height at that location. This persistent data structure allows for easy snapshotting of the map state for analysis and output of wall segments while concurrently inserting new observation data.

The wall extraction component of the topometric map parsing technique is not required to directly detect every wall in the environment. Instead, the aim is to provide sparse, yet confident, support for divisions between free space regions. Since the walls in the environments considered here are seldom visible, the entropy compass is relied upon to provide a strong prior for detecting the suspected planar geometry. The raw range data is then used merely to register the few places where easily recognizable geometry may be directly sensed. This recognition is insufficient by itself to reconstruct a floor plan suitable for effective exploration, but its output is still valuable for subsequent analyses.

C. Place Segmentation

Decomposition of the map into regions that are semantically significant to the exploration task begins with a skeletonization of the occupancy grid. This procedure thins free space regions to one-pixel-width lines whose intersections and terminal endpoints are classified as nodes in a topological map. The thinning method implemented here is an iterative procedure that produces a medial axis transform of an original binary occupancy grid while preserving connectivity. The output is a distance-to-boundary measure for every free space cell along with the topological map whose edges represents the central skeleton of the map. While clutter and obscured lines of site tend to fracture the free space, as shown by the inset in Figure 1(a), the number of nodes computed by the medial axis transform is a small fraction of the number of unoccupied occupancy grid cells and represents the first significant complexity reduction of the original map data into meaningful *places*.

This first set of places is used to perform a labeling of the free space identified by the occupancy grid. The initial labeling represents a significant over-segmentation of the map, and is subsequently fed into a graph reduction process that collapses the labels of adjacent nodes in the skeleton map. This reduction stage is governed by two primary concerns: (1) an axis-aligned bounding box containing all the free space assigned a particular label should contain

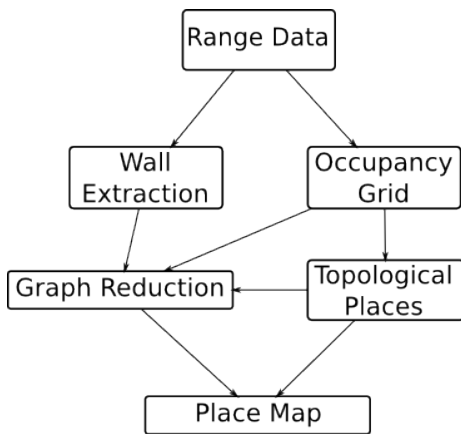


Fig. 6. Map analysis architecture.

mostly free space (as opposed to occupied cells, represented by black pixels in the occupancy grid bitmaps); (2) merging the free space attached to two nodes under consideration for merging should not involve crossing a registered wall segment. The former consideration limits the complexity of the free space assigned a particular label, while the latter prevents a place label from leaking into a room hanging off of a corridor.

The reduction process is iteratively applied to the place graph until a fixed point is reached. In practice, the resulting sparse graph includes nodes for rooms that are distinctly labelled from their connecting corridors, and corridor segments that are distinctly labelled from each other when their connectivity is mediated by a sharp turn or significant narrowing or widening. Note that such features are not topologically meaningful, but may be respected by pairing occupancy grid-derived metric information to the graph structure. The overall architecture of the analysis procedure is shown in Figure 6.

D. Place Map Example

A synthetic map is considered to highlight some of the hard-to-classify features encountered when mapping cluttered environments. Where an architectural floor plan shows the smooth walls of rooms and corridors, Figure 7(a), human considerations lead to real environments whose smooth surfaces are often hidden by what we refer to as clutter (e.g. shelving units, stacks of miscellaneous items of varying size, obscured views through partially open doors, etc.), as in Figure 7(b). The skeleton associated with this map, shown in Figure 7(c), rationalizes some of the irregular perimeter features of the occupancy grid into spurs of the topological map, but suggests an over-segmentation of the free space. This is most visible in the large, approximately central room that contains two items of furniture not placed against a wall (represented by the black holes in the occupancy grid).

During graph reduction, the scale of free space topological complexity is determined by the scale of the free space attached to the relevant regions of the topological graph. This adaptive scale, resulting from the stipulation that adjacent graph nodes may only be merged if the union of their free




Fig. 7. An unnaturally clean synthetic map (a); the same map displaying some of the clutter found in real, lived-in environments (e.g. door alcoves, furniture, shelves against walls) (b); a skeletonization of the map showing the essential topology (c); a segmentation of the map into “places” (d).

space labeling is dense (i.e. a bounding polygon contains mostly unoccupied occupancy grid cells), results in the unification of the large central room in Figure 7(d) despite the fact that it contains several nodes in the topological map. The extents of large places are ultimately bounded by the sparsely detected walls, shown as red lines. In this way, the two restrictions used by the graph reduction algorithm provide an aggressive adaptability to the scale of what is treated as clutter while remaining faithful to the observed delimiting geometry.

E. Exploration Goal Identification

The selection of locations that should be visited to expand the map is based on the previously described place segmentation. Since places are bounded by axis-aligned bounding rectangles that contain a limited amount of topological complexity, an analysis of a graph of boxes is conducted to produce an initial set of exploration goals.

While it is relatively straight forward to generate many possible exploration goals (e.g. occupancy grid frontier cells, or leaves of the topological map), efficient, rapid exploration requires the identification of a comparatively small set of goals that capture the structure of the free space of the environment. The approach taken here is to begin with a set of goals believed to contain all the desired goals, then compose a stack of filters that can winnow that initial guess down into a set of *essential* exploration targets.

We begin by considering the place map, which induces a Voronoi labeling of the free space in the occupancy grid. The Voronoi decomposition of the map yields a connected graph

of places from which we generate a first set of exploration goals that represent all the ways to enter or leave a place. Since places in the reduced graph contain a limited amount of topological complexity and abut walls in the environment, we approximate each place with a bounding rectangle (this can be generalized to more complex polygonal hulls if building geometry is expected to be highly non-rectangular). The initial set of exploration goals are placed along the sides of these bounding polygons (e.g. four points for each place: one generated at the midpoint of each side of the bounding rectangle). This set of exploration goals is believed to contain all the essential goals, but also contains many internal or insignificant goals.


The first filtering step is to remove goals associated with internal place-place boundaries, leaving only goals corresponding to place boundaries that are true frontiers of the map. The remaining points are then steered using a controller that first checks if the initial dense topological map includes a nearby node that hasn't already had a goal assigned to it. If a nearby free node is found, the goal point is relocated there; if not, the goal point climbs the cost function created by the medial axis transform described in section II-C until it is safely clear of any occupied or unknown cells in the occupancy grid or it hits the spine of the topological map.

The points associated with a given place are then compared to determine which may safely be discarded. Any points that have been driven to nearby their place's centroid relative to other goal points associated with the same place are dropped. The intuition behind this step is the example of long corridors: goals tentatively placed along the long sides of the corridor tend to end up much closer to the bounding rectangle's centroid than those that were initially placed along the short sides of the corridor's associated bounding rectangle. Rooms and other free space leading off of a corridor will generate their own exploration goals; a given place should only contribute exploration goals when there is no other place that could lead a robot to a given location.


The goals produced by each place are concatenated and compared with the pose histories of all robots. Any goals nearby a visited location are rejected, thus dealing with areas of the environment that, for whatever reason, are not cleanly bounded by geometry apparent in either the occupancy grid or wall extraction output. This compensates for difficulty in dealing with partial views through narrow openings that a robot can not navigate. Once a robot has visited what seemed to be an attractive location, we do not wish to return there.

F. Cooperative Exploration

In order to assign multiple robots located in a shared map to distinct exploration goals, a mapping from robots to goals is needed. The approach taken here is to perform a k -means clustering on the goals with k set to the number of robots to be tasked (if $k \leq N_{robots}$ then the clustering step is unnecessary). A greedy assignment is then performed that matches a robot to a cluster by considering the navigation cost from a robot to the nearest representative of the cluster.



(a)



(b)

Fig. 8. Robots, represented by numbered orange cubes, are assigned to exploration goals marked in green and given a numeric label corresponding to the robot assigned to that goal. The robots begin on the left at the end of a hallway, and only one robot is initially tasked with exploring the frontier of the known map (a). Places are highlighted by colored, translucent polygons; a robot's pose history is shown in blue, while its planned path is shown in gold. When a junction is encountered, two exploration goals are identified, leading to both robots being tasked (b).

The least expensive robot-cluster assignment is made, and that cluster is removed from consideration for the remaining assignments.

The greedy nature of this assignment is not guaranteed to find an assignment that is optimal in travel distance, but instead is designed to expand the map frontier as quickly as possible. The reason for this prioritization is that the exploration task will perform better when more is known about the map. For example, we wish to discover a long corridor leading to another wing of a building as soon as possible so that at least one robot may immediately be tasked with extending the frontier in that direction.

An experiment involving two robots is shown in Figure 8. The two robots begin at one end of a corridor, and know nothing of the world other than the corridor in front of them, a situation that induces a single exploration goal as shown in Figure 8(a). The first robot, identified by the numeral 2, proceeds down the hallway; as more of the map is uncovered, the single exploration goal recedes down the hallway, pulling the robot with it. As soon as a junction is encountered, two exploration goals are generated, one for each branch of the newly discovered corridor. The first robot is greedily assigned to the nearest exploration goal, while the second robot, labeled 3, is assigned to take the opposite branch of the intersection, Figure 8(b).




Fig. 9. An occupancy grid for a cluttered environment dominated by intersecting corridors. The partial map shown here includes a 45m length of one primary corridor, and a 70m length of the other. The 16 colored regions indicate the derived place segmentation; the blue trajectory represents the pose history of the robot building the map which starts in the lower-left portion of the map and ends near the top-center; the red lines indicate identified wall segments; the yellow crosses indicate exploration goals identified by Karto 1.1; the circled green markers indicate the exploration goals identified by the proposed exploration mechanism.

III. ANALYSIS

The motivating environment considered in this project is a basement consisting primarily of long corridors laced with steam and water pipes, draped with electrical and network cables, and lined with unused industrial equipment, Figure 2. The interesting characteristics of this environment are that the ceiling is completely obscured by irregular geometry; the walls are seldom visible; and there are apparent alcoves between the palettes of detritus that can easily be mistaken for doorways, or even small rooms.

A comparison of Karto 1.1’s exploration goal identification with topometric planning is shown in Figure 9. The key point is that reducing the set of frontier-based goals can be very ambiguous. While the frontier-based exploration strategy has few false negatives (i.e. useful exploration goals that have not been identified), the several dozen false positives (goals that are redundant) can only be reasoned about by placing them into the context of the overall flow of the map.

For a human, parts of the map without a directly observed frontier that correspond to the nooks between the stacks of clutter along the walls are easily disregarded in favor of the clear topological structure of the simple floor plan. Critically, this intuitive filtering of what is important and what is not uses a combination of the topological structure of the map with the metric information: a small nook is probably nothing, but a big enough space jutting off from a corridor is worth exploring. This reasoning process is implemented in the graph reduction process described in section II-E, and demonstrated by the four goals queued up by the system for the robot to explore, circled in Figure 9.

Fully autonomous map collection experiments demonstrated that task allocation would reliably reflect newly




Fig. 10. Occupancy grid constructed by one robot (a). Occupancy grid of the same area collaboratively constructed by two robots (b).

discovered hallway intersections, directing distinct robots down each corridor, and that such an allocation applied to two robots increased mapping speed when compared to a single robot experiment. Figure 10 shows two maps: one collected by a single robot beginning in the upper-left corner of the map, and the other by two robots beginning in the same location. For the two robot experiment, the second robot was started approximately 1 minute after the first. Both types of experiment were deemed to begin when the first robot began moving, and declared over when the last robot returned to the start location. The two-robot map was generated in 4 minutes, 43 seconds, while the single robot was able to map the same environment in 6 minutes and 10 seconds.

Goal locations evolved over the course of each experiment by receding down each hallway ahead of an exploring robot before splitting into a pair of goals: one at the end of the hallway and another in the side room accessible from that hallway. Once both hallways and both side rooms had been visited, the system, exhausted of exploration goals, commanded each robot to return to the start location. Further aspects of the experimental setup are described in Southall et al. [8].

IV. RELATED WORK

The wall extraction procedure’s usage of the prior produced by the entropy compass is reminiscent of the virtual scans technique developed by Lakaemper [9]. The method presented here differs in that the prior is used only to suggest orientations of points rather than specific geometry.

Frontier-based exploration has received significant attention, usually in the context of distributed robots constructing a shared map [10]–[12]. However, the primary inspiration and foundation for our work is Vincent’s description of the *Centibots* project [13]. Our goal is to produce a more aggressive simplification of the task graph described by Vincent, in order to rapidly identify and allocate exploration goals to robots fed into a cluttered environment from a deployment point. While Vincent dealt with robot-robot localization and map merging, the task allocation problem was focused on object of interest discovery and coverage maintenance. We found that the speed of exploration of a new environment is massively hindered by robots inspecting crevices and

channels that have little to no bearing on the desired floor plan-style output. The aim here is to provide a process tuned for an initial, rapid exploration of an environment that can serve as a prelude to subsequent search and coverage activities.

Place classification [14], [15], and exploration driven by place classification [16] represents an exciting direction for semantically-driven autonomous map building. Ideally, a system integrating geometry- and imagery-driven semantic classification with the free space segmentation described here can be developed. Such a system would be able to draw from each approach to rationalizing map data in order to cope with both recognizable locations and unfamiliar, unstructured environments.

V. CONCLUSION

Reasoning about maps at a higher level of abstraction more appropriately tuned for exploration tasks than occupancy grids opens the door to several powerful, adaptive strategies. By understanding the disposition of available resources with respect to the significant structure of a map, one is able to make the early decisions necessary to get robots where they are needed before having a complete picture of the world. The structure of an environment as represented by the place map provides a rapidly acquirable base for future map-based activities.

REFERENCES

- [1] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*. The MIT Press, September 2005.
- [2] Karto SDK. [Online]. Available: <http://www.kartorobotics.com/>
- [3] B. Kuipers and T. Levitt, "Navigation and mapping in large scale space," *AI Magazine*, vol. 9, no. 2, 1988.
- [4] B. Kuipers and Y.-T. Byun, "A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations," *Journal of Robotics and Autonomous Systems*, vol. 8, pp. 47–63, 1991.
- [5] B. Kuipers, "The spatial semantic hierarchy," *Artificial Intelligence*, vol. 119, no. 1-2, pp. 191 – 233, 2000.
- [6] R. Brooks, "A robust layered control system for a mobile robot," *Robotics and Automation, IEEE Journal of*, vol. 2, no. 1, pp. 14 – 23, mar. 1986.
- [7] C. Okasaki and A. Gill, "Fast mergeable integer maps," in *In Workshop on ML*, 1998, pp. 77–86.
- [8] B. Southall, M. Bansal, B. Matei, A. Das, J. Eledath, H. Sawhney, A. Cowley, and C. J. Taylor, "A system for rapid exploration and mapping with 3d structure characterization," in *Submitted for Publication*, 2011.
- [9] R. Lakaemper and N. Adluru, "Using virtual scans for improved mapping and evaluation," *Autonomous Robots*, vol. 27, pp. 431–448, 2009.
- [10] B. Yamauchi, "A frontier-based approach for autonomous exploration," in *In Proceedings of the IEEE International Symposium on Computational Intelligence, Robotics and Automation*, 1997, pp. 146–151.
- [11] R. Simmons, D. Apfelbaum, W. Burgard, M. Fox, D. an Moors, S. Thrun, and H. Younes, "Coordination for multi-robot exploration and mapping," in *Proceedings of the AAAI National Conference on Artificial Intelligence*. Austin, TX: AAAI, 2000.
- [12] D. Fox, J. Ko, K. Konolige, B. Limketkai, D. Schulz, and B. Stewart, "Distributed multi-robot exploration and mapping," in *In Proceedings of the IEEE*, 2006.
- [13] R. Vincent, D. Fox, J. Ko, K. Konolige, B. Limketkai, B. Morisset, C. Ortiz, D. Schulz, and B. Stewart, "Distributed multirobot exploration, mapping, and task allocation," in *Special Issue on Multi-Robot Coverage, Search, and Exploration*, D. A. Shapiro and D. G. A. Kaminka, Eds. Annals of Math and Artificial Intelligence Journal (AMAI), 2008, vol. 52, no. 2-4, pp. 229–255.
- [14] A. Pronobis, O. M. Mozos, B. Caputo, and P. Jensfelt, "Multi-modal semantic place classification," *International Journal of Robotics Research*, vol. 29, no. 2–3, pp. 298–320, February-March 2010.
- [15] O. M. Mozos, "Semantic place labeling with mobile robots," Ph.D. dissertation, Dept. of Computer Science, University of Freiburg, July 2008.
- [16] C. Stachniss, Ó. Martínez Mozos, and W. Burgard, "Efficient exploration of unknown indoor environments using a team of mobile robots," *Annals of Mathematics and Artificial Intelligence*, vol. 52, pp. 205–227, 2008.