

Sensor Space Planning

with applications to construction environments

Edward Latimer, DeWitt Latimer IV, Rajiv Saxena, Catherine Lyons, Lisa Michaux-Smith, Scott Thayer
Robotics Institute
Carnegie Mellon University
Pittsburgh, USA

edlatimer@cmu.edu, dlatimer@ieee.org, rajiv1@cmu.edu, col@andrew.cmu.edu, lmm278@psu.edu, sthayer@cmu.edu

Abstract—Outlined is a new approach to sensor space planning and its application to the construction industry. The software planning tool described here generates sensor placements automatically for use in assessing deviations in construction environments. The first step is to separate construction information goals into clusters, simplifying the planning space in order to reduce computational complexity. For each cluster, the planner generates the space of potential sensor placements for a set of information goals and selects a minimal set of subspaces to take advantage of views that can achieve multiple goals simultaneously. Sensing locations are chosen that maximize the probability of achieving each goal and a path is generated to minimize the transit cost between the various sensing locations within each cluster. Finally, paths are generated that minimize the transit cost between clusters. This method is demonstrated on a desktop computer and shown to support LIDAR information goal sensor planning within a construction site.

Keywords—sensor planning; vision planning; set covering; cluster planning

I. INTRODUCTION

Sensor planning serves to optimize the use of sensors to achieve a particular set of information goals. It differs from information goal planning, which focuses on whether the set of goals is sufficient to attain more profound conclusions about the environment (e.g. “is it correct?” or “will it suffice?” or “what is the optimal path?”).

In the context of construction inspections, the information goals typically are well established. Certain quantities must be inspected for compliance with regulatory codes or to ensure work conforms to best practices, e.g. the Lean Construction principles [14]. The window of opportunity during which data is acquired is often limited and the desired information must be collected efficiently.

Light Detection and Ranging (LIDAR) may be used for detection of various geometric features in a construction site [2], [15]. Scanning LIDAR sensors typically exhibit data represented easily as a 2-D image, and report range and intensity on line of sight phenomena. Such scans would be used to determine such properties as alignment and relative position, and compared with civil engineering specifications.

Whereas saturation of the construction environment may seem like a potential solution, it is expensive to provide

sensor coverage for areas that do not require frequent or even occasional inspections. Instead, the goal of the system described here is to optimize sensor configurations while minimizing cost and time to acquire the necessary data. This approach, coupled with a metric to determine the net-present value of sensor placements (including the initial and recurring cost of the sensor during the construction process and the processing of the data) will result in a more cost-efficient system than naïve saturation.

Scanning itself is a costly measure, as it takes time to scan. Field tests average 4-5 minutes to set up a LIDAR sensor, take a scan, and visually verify the data. In order to minimize the length of interruptions to the construction process when validating site conditions, the fewest number of scans should be used to achieve a desired set of goals. Thus a longer planning time is acceptable in order to limit the amount of physical scanning time. Algorithmically this involves finding a minimum number of sensor locations that achieve all goals; however, maximizing the frequency of goals achieved for each of the locations leads to more scan points per goal, and higher quality data. The algorithms defined here are built around these objectives.

II. BACKGROUND / RELATION TO CURRENT WORK

Traditional vision and range sensor (e.g. LIDAR) planning focuses on fixed sphere planning about a single object [9]. In such systems, only the placement of the sensors along the surface of a sphere is considered. Thus, for environments with multiple discrete objects with discrete information goals, the fixed sphere representation drastically limits the overlap to circles, curves, or points. In reality, depending on the sensor's properties, there is a volumetrically larger region of overlap that mutually achieves both goals. Our work focuses on expanding the notion of the configuration space of the sensor (also called the viewpoint space [9]) to achieve more overlap of the subspaces of the configuration spaces for examining multiple objects. Such overlap represents achieving multiple goals in a single scan or view.

Coverage planning in vision systems has received some attention [13]. The approach used for modeling an object, adapted here for multiple object / feature space, is to partition the configuration space into views that cover different occluded areas. These partitions are then searched using

greedy algorithms for a set of views to minimize the number of sensor placements.

The work in model acquisition provides a framework in which large scale sensor planning can take place. Reference [12] demonstrates how to achieve an information goal with a single point range information goal. Further work in [3], [8] and [11] extends the notion of information goals into the desired "Next Best Views" where these planned views represent the information desired. Explicit research into the placement of the sensors [3] and optimizing the placement for model creation extends naturally to selection of poses for use in global sensor placement and path optimization.

The AVENUE [1] system for automating site-modeling measurements focuses on acquiring photo-realistic models as opposed to general assessments of specific site conditions. However, a complete, photo-realistic model is not needed in many cases for construction assessment.

Our system is a component within a larger construction assessment research project at Carnegie Mellon [2]. Our software planner is designed to interface with modules that designate inspection items and with the assessment & recognition systems that interpret the sensory data. The civil engineering system provides appropriate site models, sensor resources, and information goal designations. The recognition system provides constraints on placement based on algorithmic capabilities and can re-designate goals for further scanning iterations in cases in which the target is not achieved by the current goals (e.g. when the target is not in the goal region or a tighter tolerance is needed to interpret the data).

As seen in Fig. 1, inputs to the sensor planning process are the information goals and a construction model along with sensor attributes from the civil engineering system. The output (sensing plan) is used to scan the construction site, and data from the scans is passed to the assessment & recognition systems to validate site conditions. The process feeds back on itself if the assessment system designates new goals for another scanning iteration.




Figure 1. Scan Planning Process Plan, as relates to construction assessment

III. SENSOR PLANNING SYSTEM

The process of sensor planning involves taking a series of 3-D information goals and generating a series of 2-D sensor configurations and paths between sensor configurations that will achieve the goals. Because of the quantity of data generated for each sensor scan and the scanning time for each, the system first minimizes the number of sensor configurations needed to cover the space and then seeks to

minimize the cost of transitions between sensor configurations. This process is accomplished in 6 steps:

- A. Separating information goals into clusters (cluster generation)
- B. Generating an initial set of sensor configuration spaces (sensor configuration spaces)
- C. Optimizing the minimum number of configuration spaces (coverage planning)
- D. Selecting sensing locations (sensing location selection)
- E. Path planning within each cluster (motion path planning within a cluster)
- F. Path planning between clusters (cluster path planning)

These six steps are described in detail below.

A. Cluster Generation

Cluster analysis is the study of algorithms and methods for grouping or classifying objects and data [4]. The term clustering is often used to describe methods for grouping of unlabeled data [6]. As applied to sensor planning, the clustering module takes the information goals and generates a set of clusters which lower the overall computation cost of generating the final sensing plan. Using the clustering approach, the space divides into a collection of distinct clusters (each of which is comprised of similar goals of interest) to be planned in subsequent steps of the algorithm.

Clustering makes the planner scalable by enabling it to plan for very large sets of information goals. The large sets can be broken into logical clusters which reduce planning complexity and computation time in later algorithmic steps. It also allows for fast re-plans when new information goals are added since unchanged clusters do not need to re-compute steps B-E of the algorithm.

Cluster generation is not done in isolation, but requires civil engineering knowledge. The clustering algorithm currently implemented is based on a nearest neighbor approach. For any information goal within a cluster, it is at most a distance d to another goal within the same cluster, with the value d as an input depending on the site dimensions, sensor parameters, and user preference. This approach is well-suited for our task because most information goals exist in long chains in the construction environment, such as along a wall or series of columns. Clustering these goals allows coverage planning, described in step C, to take advantage of their proximity.

Several other clustering approaches may be appropriate for scanning in an active construction site. These include a means of analyzing the site to group goals by floor or by room, or k-means clustering, which would group goals around a set number of center points. These methods would limit the disruption to any particular area of construction to the time to scan a single cluster.

B. Sensor Configuration Spaces

The sensor configuration space indicates sensing locations capable of attaining an information goal. For a typical LIDAR in 3-D this includes both position and orientation. For example, if a user wanted to scan a small volume with a LIDAR having no minimum range, the configuration space would be a solid sphere, excluding obstacles and regions occluded from line-of-sight with the goal by obstacles. The notion of occlusions is well addressed in vision planning research [8].

The LIDAR parameters of interest are minimum and maximum range, height, and vertical and horizontal angular range. These parameters bound the 3-D sensor configuration space. The scanner is further assumed to be at a fixed height, and thus limited to 2-D motion. Thus, the sensor configuration space for a given LIDAR system would be a line-of-sight, bounded, 2-D set of all points within the minimum and maximum scanning ranges of the information goal, at the height of the sensor. To avoid the complexities of continuous space, the configuration space is uniformly sampled to create a set of configuration space points.

Attributes of the information goal such as height and desired resolution determine the dimensions of the configuration space. As the sensor and the goal may not be located at the same height, the 2-space minimum and maximum ranges of the configuration space are not necessarily the same as that of the information goal in 3-space. The final dimensions of the 2-dimensional configuration space are dependent on the height and range of the sensor when coupled with the height of the goal, as seen in Fig. 2.




Figure 2. Configuration spaces for goals at heights 1 and 10 meters with sensor height 1, range [2...25] meters

Some goals may have specific minimum density and noise requirements to ensure quality data as specified by the assessment & recognition system [15]. A sensor placed closer to the goal will capture data at a higher resolution, and better satisfy these constraints, as LIDAR point density and noise is a function of distance. The data density constraint is met through the concept of resolution, which decreases the maximum range of the sensor configuration space for a specific information goal.

After the configuration space is initially created with a fixed number of sample points, the need for line of sight to all non-obstacle portions of the information goal reduces the number of c-space points, yielding the occlusions illustrated in Fig. 3.




Figure 3. Sampled configuration space with occlusions for a LIDAR sensor about a point goal

C. Coverage Planning

Coverage planning involves dividing the configuration space sets into non-overlapping regions called intersection sets and assuring all goals are satisfied while minimizing the number of sets to visit. During the minimization, the planner maximizes the frequency of each goal being satisfied to ensure the maximum amount of data for all information goals while still taking a minimal time to achieve all goals. This means the planner must take advantage of overlaps in the sensor configuration spaces for different goals, where the sensor can see more than one goal, as seen in Fig. 4.




Figure 4. Overlapping Configuration Spaces

This is achieved by taking a set of sensor configuration spaces and generating the set of intersection sets labeled for each unique intersection of goals. Formally, G is defined as the set of all goals g_j in a site with N total goals, in (1).

$$G = \{g_j \mid \forall j \in J = \{0 \dots N - 1\}\} \quad (1)$$

Equation (2) gives the configuration space set c_j for goal g_j , where each point c_{ji} is a valid sensing point in goal g_j 's configuration space set. Equation (3) gives us an intersection set S_K from a given set of goals. The complement of (3) is (4), giving a set of goals that can be seen by a given overlap of configuration space sets.

$$f(g_j) = c_j = \{c_{ji}, \forall i\} \quad (2)$$

$$F(\{g_k | k \in K\} \subset G) = s_K = \bigcap_K c_k, K \subset J \quad (3)$$

$$F'(s_K) = \{g_k | \forall k \in K, s_K \subset F(g_k)\} \quad (4)$$

The set of all unique intersection sets S is given by (5). The covering set C , defined in (6), is a possible grouping of intersection sets that can see all goals.

$$S = \{F(\{g_k | k \in K\}) | \forall K \subset J\} \quad (5)$$

$$C \subset S, \bigcup_C (F'(s_{c \in C})) = G \quad (6)$$

The planner attempts to minimize the size of the covering set C so that all goals are achieved in a minimum number of intersection sets. Maximally, this solution will generate a plan to scan each goal once if there is a one-to-one mapping between intersection sets and goals. However, if some intersection sets enable multiple goals to be achieved, a fewer number of configurations can be visited.

When creating intersection sets for a sensor with a 360° field of view, any intersection or subtraction of configuration space sets is a valid intersection set. For sensors with less than a 360° field of view, angles to goals corresponding to the configuration space overlaps must fall within the field of view of the sensor to be in the valid intersection set.

The task of finding the minimal covering set is aided by use of an intersection set tree, whereby a node's depth in the tree denotes the number of intersecting configuration space sets it contains. The intersection set tree creates minimally N nodes and maximally N -factorial labeled intersection sets. In Fig. 5, the node labeled 012 corresponds to the intersection set $s_{\{0,1,2\}}$, created by the intersection of the configuration spaces of 3 goals $\{g_0, g_1, g_2\}$.




Figure 5. Intersection set tree derived from Fig. 4

Frequently, the planner can obtain a minimized configuration by depth-first traversal of the tree, choosing only the leaves which include new goals (goals not covered by leaves that have already been chosen). In Fig. 5, sets $s_{\{0,1,2\}}$, $s_{\{3,4\}}$, and $s_{\{3,5\}}$ satisfy a minimal covering set C . However, this depth-first traversal is not always optimal. Solving for the minimum is an NP-complete set-covering problem, thus there is no known polynomial time algorithm for a minimal solution. Instead, the planner approximates a minimal set with a simulated annealing algorithm [7].

Simulated annealing is used to approach a minimal covering set C in a finite number of iterations, starting with an initial set, by mutating the set and accepting the mutation with some probability dependant on the results. Selecting a set from the leaves of the intersection set tree leads to the property that goals will be satisfied with a higher frequency. Thus only the leaves of the tree are considered for submission into the simulated annealing algorithm.

The initial covering set is taken from the depth-first traversal of the intersection set tree, including only the leaves which cover new goals. When a leaf is chosen for the set, it is assigned to cover those new goals, and so a coverage list is made of the intersection sets assigned to each goal. A set of all the leaves of the tree is taken for the search space. The simulated annealing algorithm alters the current set at each iteration step by taking an intersection set from the search space and assigning it to all goals that it covers, replacing any intersection sets previously assigned to those goals. If this eliminates some intersection sets from the coverage list, it has reduced the covering set and the change is kept. Otherwise, the change is still kept with some probability.

Table 1 shows an iteration of the simulated annealing algorithm from the i^{th} iteration to the $i+1^{\text{th}}$ iteration. In this example, the search space R contains the intersection sets denoted by the leaves of an intersection set tree and the goal set G encompasses all goals. Intersection sets are chosen to satisfy each of the goals. In the i^{th} iteration, the intersection set $s_{\{0,1\}}$ satisfies both goals g_0 and g_1 , and the intersection set $s_{\{1,2\}}$ satisfies goal g_2 .

TABLE I. EXAMPLE OF SIMULATED ANNEALING MINIMIZATION ITERATION

Iteration	Search space		
	Goal coverage satisfying goal in position $G = \{g_0, g_1, g_2, g_3, g_4\}$	Covering set size	Mutation
i	$C = \{s_{\{0,1\}}, s_{\{0,1\}}, s_{\{1,2\}}, s_{\{2,3\}}, s_{\{3,4\}}\}$	4	$F'(s_{\{2,3\}}) = \{g_2, g_3\}$
$i+1$	$C = \{s_{\{0,1\}}, s_{\{0,1\}}, s_{\{2,3\}}, s_{\{2,3\}}, s_{\{3,4\}}\}$	3	

As a sample mutation, $s_{\{2,3\}}$ is chosen to replace all instances of intersection sets satisfying goals specified by F' , thus it replaces $s_{\{1,2\}}$ to satisfy goal g_2 (while still

satisfying g_3), and we are left with a minimally-sized covering set.

D. Sensing Location Selection

Once the minimized intersection sets are generated, the actual sensing locations for each intersection are chosen, one from each of the final minimized sets. For an ideal scanner, sensing locations would be as close as possible to the minimum scanning range for each goal, but since intersections can satisfy multiple goals, the sensing location is chosen as the location within the intersection set that is closest to the centroid for the corresponding goals satisfied. For non-360° FOV sensors, a sensing location contains both a location and a direction that encompasses all goals satisfied for the set.

E. Motion Path Planning within a Cluster

Motion path planning takes a set of sensing locations and a site plan for a cluster and optimizes a path within that cluster based on traversal distance between sensing locations while avoiding obstacles. The algorithm returns a motion plan consisting of points to be traversed and an order for the sensing locations.

The first step is to create a visibility graph for the obstacles, to include the sensing locations. An A* heuristic search algorithm [5] searches the visibility graph for the shortest path between all pairs of sensing locations and an edge cost matrix of path lengths is created. Finally, permutations of all possible orderings of the sensing locations are generated and the shortest path that visits all sensing locations is returned.

A permutation algorithm works well when the number of sensing locations is small, but is grossly inefficient otherwise. Thus, for larger numbers of sensing locations, a simulated annealing algorithm is used to determine the order in which to visit the locations. Unlike the permutation algorithm, this method does not guarantee an optimal path between sensing locations, but performs well given suitable annealing parameters [7].

F. Cluster Path Planning

An overall motion plan is generated that links together the individual cluster plans. A* is again used to generate an edge cost matrix for all pairs of motion path endpoints between each cluster, and the simulated annealing is used again to order the clusters for a minimal path.

IV. IMPLEMENTATION

This system is being applied to planning sensing activity for construction inspections. Most construction inspection research seeks to characterize the applicability of specific instruments to specific inspection tasks, e.g. applied vision, LIDAR systems, position sensors for applications such as model construction, part tracking, or site modeling, however there has been little research into providing a framework and system into which all sensing activity can be coordinated, planned, and optimized.

The system described in the previous section was developed as a planning library and tested on a desktop interface (Figs. 6-7) from which the screen shots were taken. The desktop application supports creation and fine-tuning of sensing plans and integrates with a hand-held application. Since site conditions can change between the office plan and the field, such as temporary work areas or unexpected obstacles like equipment, the hand held interface provides the ability to view sensing plans in the field, update for site conditions, and create new sensor paths based on the field conditions.

The hand-held system does not need to create an entire sensing plan, but instead make small changes quickly to an existing plan. To do this, the data produced at any stage of planning can be saved by the desktop application and accessed by the hand-held. For example, if a new obstacle is added, and it does not affect the current scan locations, the hand-held would only have to repeat the motion planning steps in order to bring the plan up to date.




Figure 6. Desktop system showing project case study construction site




Figure 7. Desktop system with configuration space points and plan

V. RESULTS AND ANALYSIS

The following results are based on five different test sets, including a construction site case study. The first four are samples to test various aspects of the planner, and the last test set is on a construction site with goals specified by the civil engineering system. We start with a runtime analysis of the algorithm.

A. Runtime Analysis

Order of magnitude analysis gives us runtime bounds on the growth rate of a function with respect to the size of the

input. The most common is big-O analysis, which gives an upper-bound on the runtime of the algorithm.

There are six system components which merit big-O analysis. They are cluster generation, configuration space generation, configuration space occlusion reduction, intersection set creation, intersection set minimization, A* motion planning, and shortest-path selection. For the following, c is the number of configuration space points, G is the number of goals, V is the number of obstacle vertices, and S is the number of sensing points. We assume worst-case runtime of a single cluster for planning.

Cluster generation is $O(G^2)$ because we need to check the distance between all goals. Configuration space generation is $O(cG) = O(G)$ because c is the constant number of points generated for the randomly sampled configuration space, so c does not affect big-O analysis. Occlusion reduction is $O(cG * V) = O(GV)$ to verify that each configuration space point can see its own goal. The number of configuration space points will typically be less than c after occlusion reduction, but here we will assume the worst case, with no occluded points. Thus, intersection set creation is approximately $O(cG * G * V) = O(G^2V)$ to check, for each valid c -space point, which of the other goals it can see. Intersection set minimization is $O(G)$ because simulated annealing can be linear [7]. Computation of the A* motion plan is $O(S^2V)$ to create a fully-connected graph between all sensing locations and avoid obstacles between the locations. Finally, computation of the shortest path is $O(2^S)$ for small values of S and $O(S)$ for larger values when the system switches from using brute-force to simulated annealing to find the shortest path.

The largest runtimes for our algorithm are $O(G^2V)$ and $O(S^2V)$. Thus, worst-case we have square dependence on the number of goals and a linear dependence on the number of obstacle vertices, since we expect the number of goals to be larger than the final number of sensing locations.

B. Scattered Goals Without Motion Plan

A set of widely spaced goals with no obstacles serves as the base case. Since there is no overlap of the c -spaces in this case, there is no minimization necessary when selecting the minimal coverage set. The case was run with for varying number of goals without motion planning to get a runtime base. As seen in the first data set in Fig. 8, the runtime is roughly proportional to the square of the number of goals.




Figure 8. Scattered vs. overlapping goals - Tests B, C and E

C. Scattered Goals With Motion Plan

The previous test case was run again, with motion planning. An example scan plan for this case, with 20 goals, is shown in Fig. 9. Without any obstacles, A* does very little work to generate the visibility graph, and the motion planner finds the shortest path with the constant-time simulated annealing algorithm. The runtime, shown as the second data set in Fig. 8, is not much different from the previous case.




Figure 9. Scanning and Motion Plan for Scattered Goals

D. 20 Scattered Goals with Obstacles

To test runtime dependence on obstacles, the previous case was used, with 20 goals and varying numbers of obstacles added to the environment. With a constant number of sensing locations, the runtime depends only on the obstacle vertices, and shows a linear dependence as seen in Fig. 10. Creation of the motion plan does not increase the runtime appreciably for this small numbers of goals.




Figure 10. Obstacles with 20 goals

E. Line of Overlapping Goals

To test intersection set minimization, a line of overlapping goals was created where each goal's configuration space overlaps with exactly two others. An example with 10 goals is shown in Fig. 11. The runtime, the 3rd data set in Fig. 8, shows an approximately linear increase over the scattered goals, as expected by big-O analysis.




Figure 11. Scanning and Motion Plan for Overlapping Goals.

F. Construction Site Case Study

Figures 6-7 depict the planning system for a construction site. The site has 23 information goals and 246 obstacle vertices. Given the 4-5 minute estimate to scan, it would take approximately 1.5 hours to scan each goal once. Sensor planning without the motion plan took 43 seconds and with motion plan took 58 seconds, generating a sensing plan with 7 sensing locations. Scanning from the sensor plan would take approximately 30 minutes, saving the user an hour in scanning time. The runtime on this real-world data is very similar to that expected from the 20 goal test, which predicts approximately 60 seconds for the runtime with motion plan.

G. Hand-held Application

The hand-held system was tested on an iPAQ PocketPC, with a small number of goals and obstacles. Though timed test scripts were not available for the hand-held application, observed run times were considerably slower than with the desktop environment. A simple decrease in the number of points sampled in each configuration space considerably reduced runtime.

VI. CONCLUSIONS AND FUTURE WORK

As part of a larger site assessment system, this sensor planning system creates sensing and motion plans for sensor deployment to search for defects at a construction site. The system brings robotic optimization tools to users in the field, improving productivity in construction areas by reducing rework. The dual interfaces (desktop and palm) provide a front end for construction professionals to achieve lower cost measurement plans by working with the sensor planning structures presented in this paper.

All of our sample tests, including the case study, match our analysis of the expected runtimes. Further work will test the speed and effectiveness of various clustering algorithms.

Some improvements considered for the planning system relate to integration within the larger construction assessment research project. The civil engineering system does not always specify information goals as points, but may specify a wall (line) or a surface (area) as an information goal. A revised planning system would accept goals in a range of shapes. The sensor configuration space model allows these

different goal types to integrate easily into the planning process, since the configuration space representation, once generated, is independent of the type of goal

Given the time-critical nature of the construction environment, scan and motion time is costly, hence the need to minimize the time to achieve necessary information goals as outlined in this paper. Crucial to this work is the sensor configuration space model. We have made a case for continued development of the algorithms and interfaces to support the models in the construction arena. Further field evaluations will allow refinement of the algorithms and the interfaces.

ACKNOWLEDGMENT

The project is funded by a grant from the National Science Foundation, CMS #0121549. NSF's support is gratefully acknowledged. Any opinions, findings, recommendations, or conclusions presented in this paper are those of the authors and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

- [1] Allen, Stamos, Gueorguiev, Gold, Blaer, "AVENUE: automated site modeling in urban environments", IEEE Transactions on Digital Imaging & Modeling, 2001.
- [2] Gordon, Boukamp, Huber, Latimer, Park, Akinci, "Combining reality capture technologies for construction defect detection: a case study", EIA-9, 2003
- [3] Banta, Wong, Dumont, Abidi, "A next-best-view system for autonomous 3-d object reconstruction", IEEE Transactions on Systems, Man, and Cybernetics, Vol. 30, No. 5, 2000.
- [4] D. Chaudhuri and B. B. Chaudhuri, "A novel multi-seed nonhierarchical data clustering technique", IEEE Transactions on Systems, Man, and Cybernetics, Vol. 27, No. 5, 1997
- [5] Hart, Nilsson, Raphael, "A formal basis for the heuristic determination of minimum cost paths", IEEE Transactions on Systems Science and Cybernetics, SSC-4(2), 1968
- [6] Jain, "Data clustering: a review", ACM Computing Surveys, Vol. 31, No. 3, 1999.
- [7] Kirkpatrick, Gelatt, Vecchi, "Optimization by simulated annealing", Science, Number 4598, 1983.
- [8] Maver, Bajcsy, "Occlusions and the next view planning", Proceedings of the IEEE International Conference on Robotics and Automation, 1992.
- [9] Plantinga, Dyer, "Visibility, occlusion, and the aspect graph", 1987.
- [10] MacQueen, "Some methods for classification and analysis of multivariate observations", Proc. 5th Berkeley Symp. Math. Statist. Prob., 1967, Vol. 1
- [11] Reed, Allen, "3-d modeling from range imagery: an incremental method with a planning component", Image and Vision Computing, issue 17, pp. 99-111, 1999.
- [12] Reed, Allen, "A Robotic System for 3-D Model Acquisition from Multiple Range Images", IEEE Journal on Robotics & Automation, 1997.
- [13] Strobel, Illmann, Prassler, "Intuitive programming of a mobile manipulator system designed for cleaning tasks in home environments", Proc. of Int. Conf. on Field and Service Robotics FSR, 2001.
- [14] Koleska, Lauri, "Application of the new production philosophy to construction", Stanford University, CIFE Technical Report, 1992.
- [15] Huber, Hebert, "Fully automatic registration of multiple 3D data sets," Image and Vision Computing, Vol 21, No. 7, July, 2003, pp. 637-650.