Detecting Region Transitions for Human Augmented Mapping

Elin A. Topp* and Henrik I. Christensen

***Abstract*—In this paper we describe a concise method for the feature based representation of *regions* in an indoor environment and show how it also can be applied for door passage independent detection of transitions between *regions* to improve communication with a human user.**

I. Introduction

With this paper we aim to make a case for a concise method for the segmentation of an indoor environment into a topological graph representation that is independent from particular ”transition indicators” like door passages and that allows to generate a human comprehensible environment representation for a mobile service robot. We assume such an environment representation as crucial to support meaningful interaction between a user and a supposed ”general purpose service robot” in and about the environment. We consider our framework of *Human Augmented Mapping* (HAM) [1, chapter 3] a possible way to approach the issue of integrating robotic and human environment representation in general. The framework subsumes different aspects of robotic mapping, spatial representation and human robot interaction. Within the context of HAM we assume an interactive scenario – a “home tour” – as the most natural way of providing the robot with the needed semantic information about the environment as it is seen by the user. The human user guides the robot and gives names to things and places according to her personal preferences, while the robot builds a suitable (hybrid) map that is augmented with this information. In such a tour it is not necessarily the case that the user will present all items actively [1, chapter 6], hence, the system driven detection of transitions, e.g., from one room into another, is essential to make sure that the representation generated by the robot corresponds to the user's understanding of the environment. An obvious way of detecting such changes is to find door passages. However, there are cases where the border between two structurally (and often also functionally) different rooms (or *regions*) is not described by an obvious separator like a door passage. Figure 1 visualises such a ”structural ambiguity”. There are of course also cases where one large room serves different functions, e.g., in very small studios with combined ”living room” and ”kitchen”, but in this work we want to focus on a segmentation based on structural features that can be observed in the environment. Such a feature based representation should be suitable for the generation of *region* nodes in a topological graph structure and support the detection of hypothesised transitions between *regions*. Ideally, the respective representation allows a robot to more or less immediately recognize a particular *region* as previously visited even when reaching it from a new ”entry point”. In this paper we describe our concise method for the feature based representation of *regions* as nodes of a topological graph representation and show how it also can be applied for door passage independent detection of transitions between *regions*. We show the applicability of the method in different (interactive) contexts and give one ”proof-of-concept” example for a successful ”loop closing” experiment.

Elin A. Topp was at the time this work was performed with the Centre for Autonomous Systems, Royal Institute of Technology (KTH), Stockholm, Sweden and is now with the Department of Computer Science, Lund University - LTH, 22100 Lund, Sweden, e-mail: Elin_Anna.Topp@cs.lth.se

Henrik I. Christensen is with the College of Computing, Georgia Institute of Technology, Atlanta, GA, 30332-0760, USA, e-mail: hic@cc.gatech.edu

* Corresponding author

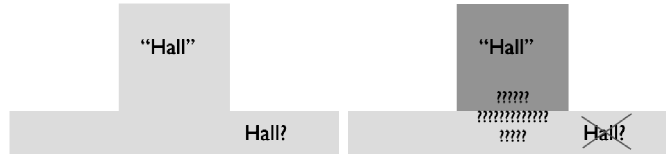

Fig. 1. (left) The ”hall” has been presented to the robot that now assumes the complete depicted area as ”hall” since no ”door passage” was passed while traveling. (right) The user wants the robot to understand that there is some part of the area that is NOT the hall, but, e.g., the ”corridor”. It seems natural to assume an ”unspecified area” in the transition.

II. Related Work

Since our presented work deals mainly with the issue of obtaining a topological partitioning of a given environment we give an overview of related work in this area. We are aware of several works using image analysis techniques and object recognition based representations or categorisations for rooms or regions, but due to the limited space we focus mainly on structure based methods in this brief overview.

One strategy is to predefine the topological structure of an environment and use this map for localisation and navigation purposes [2]. The limitations of such an approach in the context of and interactive framework and the arbitrary environment we assume are obvious: the complete possible working environment for the robot needs to be known in advance. Other, more adaptive methods that assume the robot to acquire a topological representation of its environment are based on (sensory) data obtained while travelling.

An unsupervised/autonomous method for the detection of *places* is suggested by Beeson *et al.* [3] based on earlier investigations in a related context [4], [5]. The definition of a “place” in these works suits the requirements and abilities of an autonomous system, but does not necessarily correspond to a personalised representation of a human user. This limitation can be observed also for other completely unsupervised methods of topology learning, as for instance proposed by Tapus *et al.* [6].

For the representation of convex areas Kröse showed that it is possible to represent such *regions* reliably by obtaining only one sample range data set and transform it to its centre point and bearing with the help of a principal component analysis to anticipate future scans [7]. Our representation for *regions* is closely related to this proposed approach, as to the one presented by Buschka and Saffiotti [8], who detect ”room-like” structures based on (sonar) range data, using a very similar method. However, due to the nature of the used range finder data, their approach is somewhat different to our method regarding necessary preparation steps.

Mozos *et al.* show, how the *category* of a certain area (room, doorway, or corridor) can be determined with the help of supervised learning [9], also used in another similar approach [10]. We adopt their idea of using a set of features to represent a (laser) range data set, that we obtain in *regions*, but use an even more concise set of features [1, chapter 4]. Further Mozos *et al.* label places in the complete environment into a fixed number of categories, while we do not rely on any previously defined categories for the *regions* that can be specified by the user. This allows us to concentrate on the transition from one *region* into the other, not regarding what category (in the sense of the mentioned work) the *regions* or the transition itself belong to.

III. Representing an indoor environment with regions

In the HAM framework we define a *region* as follows: *A region is a functionally and / or structurally delimited area of an indoor environment, that can be a container for one or several particular* *locations and objects. A region offers enough space to navigate (typically regions correspond to, e.g., rooms, corridors, delimited areas in hallways)*[1].

For this work we focus exclusively on *regions* and their structural properties. The *regions* that have been labeled in the assumed tour form the nodes of a topological graph structure that (among other entries according to the HAM model) contain subgraphs representing known, viable paths (navigation graphs). We introduce also a generic node, the "generic *region*" as a starting point and to cope with situations in which entities of other conceptual levels (e.g., locations) are specified without the surrounding *region* being named before.

The general assumption is that a "*region* node" in the topological graph is generated when the user shows a particular *region* to the robot. This can also happen when the robot detects a significant change in the environment – a hypothesised transition from one *region* into another – and asks for clarification of the situation, while the user did not (yet) introduce any new region actively.

A new *region* node is linked into a topological graph structure on two levels: A high level edge describes the topological link between two nodes, i.e., the fact that it is possible to somehow travel from one *region* to a neighboring one. The so far existing navigation graphs of those regions are rebuild so that the high level edge receives at least one concrete instantiation, describing *how* to travel. This concrete link is represented as a (metric) path vector relative to the *region* node's geometrical centre point. In addition to the topological links between the graph nodes, each *region* node is described with its centre point $\bar{X}$ and angle θ relative to the starting position of the tour. To derive these metric links we make use of a (corrected) position estimate. Since the metric links between *region* nodes are described relative to the corresponding node "origin", we believe it would be possible to decouple small local (metric) maps from the global metric one if necessary. Hence, we assume an arbitrary, "classic" simultaneous localization and mapping (SLAM) method as suitable for the purpose of retrieving a sufficiently correct pose estimation.

Representing regions and detecting transitions between them

To actually compute the representation for the *region* nodes in our topological graph structure, we rely on statistical features derived from laser range data sets. This is a very concise, computationally rather inexpensive and flexible method and we propose to use it not only for the description of the *region* nodes specified by the user, but also for continuous comparisons of hypothesised *region* representations for transition detection. The detection of transitions can then be handled independently from or as a complement to explicit cues like, for instance, door detectors as used by other approaches [11].

1) The region representation: We represent specified *regions* with the help of a number of statistical features computed from a 360° laser range data set [1, chapter 4]:

- Mass m: The free space surrounding the robot ("clutter index")
- Length $l1$ and $l2$: The length along the two principal components of the data set (overall "size")
- Excentricity e: The excentricity of the ellipse described by the two principal components (overall "shape")
- Centre point $\bar{X}$: The centroid of the data set
- Angle θ: The angle of the first principal component relative to the origin of the map / the starting point of the tour

The first three features describe the properties of the *region* while the latter two are used to link the corresponding *region* node into the graph structure as described previously. Although those features are related to each other we found in some empirical tests that they all contribute to the distinctive power of the description [1, chapter5]. The features over a range data set $\{X_i : 0 \leq i < N\}$, where N is the number of data points $X_i = (x_i, y_i)$ are computed as follows. To compensate for the distortion of the laser range data set[2] the centroid of the data set is computed as a range weighted average

$$\bar{X} = (\bar{x}, \bar{y}),$$

with

$$\bar{x} = \frac{1}{\sum_{i=0}^{N-1} r_i} \sum_{i=0}^{N-1} r_i x_i$$

and

$$\bar{y} = \frac{1}{\sum_{i=0}^{N-1} r_i} \sum_{i=0}^{N-1} r_i y_i$$

where $r_i = \sqrt{x_i^2 + y_i^2}$ is the distance of the data point from the origin of the data set, i.e., the position of the laser range finder. The data set is then transformed to the set $\{X'_i = (x_i - \bar{x}, y_i - \bar{y}) : 0 \leq i < N\}$ relative to the centroid. This allows estimation of the area m bordered by the data set to

$$m = \left(\sum_{i=0}^{N-2} m_i \right) + m_{N-1},$$

with

$$m_i = \frac{1}{2} tan(\alpha'_{i+1} - \alpha'_i) {r'_i}^2$$

and

$$m_{N-1} = \frac{1}{2} tan(\alpha'_{N-1} - \alpha'_0)(r'_{N-1})^2$$

where r'_i is the distance of the transformed point from the centroid. Since this estimated covered area is depending on objects in the vicinity it represents an index of clutter, which is helpful to differentiate between regions of the same basic layout, but with different furnishing.

We perform a principal component analysis (PCA) to obtain the two eigenvectors E_1 and E_2 of the data set. We then estimate the two features $l1$ and $l2$ as the maximum distances represented in the data set along the bearing angles of E_1 and E_2. To make sure that such a point is found, a tolerance threshold around the bearing angle is employed. The data set is now represented by the quadruple $regDesc = (name, m, l1, l2, e)$ and stored as properties of the corresponding *region*.

2) Detecting transitions while travelling: While travelling through the environment the available range data sets are continuously used to generate a "hypothesised *region*" representation of the surroundings, which is compared to a previously specified one to decide, whether the environment has changed significantly so that it appears likely to have entered a new *region* [1, chapter 4].

To compare two *region* representations we compute a distance measure d from the relative differences in each of the descriptive features:

$$d \quad = \quad \sqrt{\hat{m}^2 * \hat{l1}^2 * \hat{l2}^2 * \hat{e}^2}$$

with

$$\hat{f} \quad = \quad \left(1 - \frac{f_{hyp}}{f_{cur}} \right) \quad for \quad f \in \{m, l1, l2, e\},$$

[1] *Location* and *object* correspond to other spatial concepts used in the HAM framework, forming a conceptual hierarchy. We use the term *location* for a large, not as a whole manipulated object or a particular "workspace" (more or less static, e.g., a table, the fridge, the coffee-maker), and define an *object* as small and dynamic (manipulable) item (e.g., a cup or a remote control).

[2] as a result of the equidistant angular resolution with which a laser range finder scans the environment objects in the direct vicinity of the sensor are represented with considerably more data points than objects that are further away

with f_{cur} and f_{hyp} standing for the respective feature of the current and the hypothesised representation. We evaluated several distance measures in initial empirical tests and found the presented one most suitable to capture the changes being implied by the structure of the environment. If the distance measure d exceeds a threshold a significant change in the environment representation is hypothesised.

To improve stability, we assume that the change has to be stable over a number of data cycles. Additionally it is obvious that the robot cannot have entered a new *region* when it has not moved, hence we apply a minimum distance threshold between transition detections. Those two conditions allow to lower the computational effort and make the system more stable.

The hypothesised *region* representation is compared to the previously accepted current one. In the case that a significant change is detected, the hypothesised representation is checked against all other available representations (nodes in the graph) whether any of them matches sufficiently well and is not completely unlikely to have been entered, given its (metric) position. If none of the previously specified *region* representations match, the system assumes to be back in the "generic *region*". Given appropriate interaction capabilities, a transition detection with the hypothesis for the actually entered *region* can lead to a confirmation dialogue with the user, which then can result in the specification of a new *region*. The corresponding representation is then added to the graph and is used as the accepted current one for further comparisons.

IV. Implementation and Evaluation of the System

We investigate our method for the representation of *regions* and detection of transitions in the context of three different implementations. Although the used implementations were slightly different regarding their integration into an interactive system (or lack thereof) the general system setup for both online runs and data collection consisted of a mobile non-holonomic platform with one laser range data finder mounted at about 30cm (in one case 50cm, but still below the top level of most furniture) above the ground. In all cases time stamped odometer readings and laser data were made available and, depending on the interaction capabilities, also the user's labeling and the raw data sets used for the specification of *regions* together with the resulting feature based representation. As mentioned previously we consider two types of events relevant to generate a new region representation in the topological graph. One is the – user initiated – specification of a new region, the other is the – data driven and robot initiated – detection of a structural change in the environment. When the user through personal initiative or as a result of a clarification discourse specifies a *region*, the robot acquires a 360° range data set by turning around once, for the continuous comparisons we use "virtual scans" generated from a local map[3] to compensate for the fact that the used robots only have one laser range finder available.

Evaluation

We discuss our approach in the context of a number of data sets that were obtained in different indoor environments, representing the range from "laboratory conditions" in an office building to "real world conditions" in a small, actually inhabited, apartment. Also the settings range from explicit test runs, where data collections from tours with a remotely controlled robot [12] were evaluated, to a fully interactively controlled run conducted by a test user. These different settings had of course influence on how the system could handle a detected transition *after* it stated its hypothesis, but the main aspect was in all cases to use the *region* representation for the detection of transitions in the environment and evaluate the suitability of the method for the generation of a human comprehensible representation of an indoor environment. We evaluated the runs (guided tours) in those different environments with respect to the following criteria:

- Consistency of the generated separation of *regions* in the environment with the "common understanding" of this separation
- Detection of "obvious" transitions (doorways) and ambiguities where, e.g., a hallway opens up into a larger area.
- Loop closing ability on the conceptual / semantic level when coming back to a previously specified *region* through a new entry point
- Overall number n of detected ambiguities / transitions (and requests for confirmation from the system for the fully implemented systems), with $nCorr$ being the number of expected transition detections between structurally different areas given the path of the robot.
- Number $nSens$ of ambiguities detected in a sensible range (approximately 1 to 2 meters in a standard indoor / domestic environment) from an obvious transition in the environment (e.g., a doorway)
- Number $nSpurious$ of obviously spurious (erroneous) detections of ambiguities (e.g., in the middle of an open area)
- Number $nMiss$ of obviously missed transitions into a structurally different area

The generation of a new, explicitly specified, *region* was not considered as a detected change, but when this specified *region* was obviously left a detection should have occurred, otherwise a miss is counted.

In the following we describe the test scenarios, the results according to our evaluation criteria being summarised in table I as evaluation results for 1) two domestic settings (ap. 1 & 2), 2) test runs in the laboratory (lab 1 & 2), and 3) a test with a fully interactive system ("BIRON"[4]).

1) Domestic environment, mapping subsystem only: Two different domestic environments were considered, one being a rather small apartment with narrow passages and doorways, the second being a medium sized flat with partially rather wide passages and open spaces. In both apartments the living room, a bedroom and the kitchen were presented to the robot. In the larger apartment also two runs were conducted, both being actual guided tours in a user study setup of the "home tour" scenario.

2) Test runs in the laboratory / office environment: In the office environment we evaluated two runs one of which covered a large part of the corridor of one of our floors and two of the rooms. With the

TABLE I

counts & values	ap. 1	ap. 2	lab 1	lab 2	BIRON
min. movement	1m	1m	1m	1m	1m
data cycles	1	1	3	3	3
n	18	22	13	4	12
nCorr	18	24	13	4	5
nSens	18	20	13	4	9
% of nCorr	100	83	100	100	180
nSpurious	0	2	0	0	3
% of n	0	9	0	0	25
nMiss	0	4	0	0	0
% of nCorr	0	17	0	0	0

[3]part of the software tool package "CURE", courtesy of Patric Jensfelt and John Folkesson

[4]the "**Bi**elefeld **Ro**bot Companio**n**", that served as platform for the demonstrating Key Experiment 1 (the "Home Tour") of the integrated EU project COGNIRON, concluded in 2008 [13]

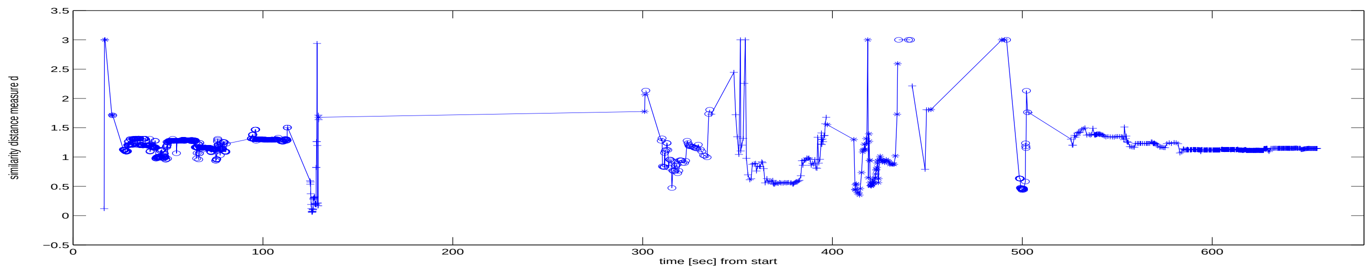

Fig. 2. *An illustration of the similarity measure for the "BIRON"-run as it is changing over time. A switch to a new marker type indicates a potential "change". Extreme values due to the user blocking the range finder are cut for readability of the plot. Jumps without markers indicate periods during which the robot was not moving (translating), hence the similarity distances were not regarded relevant.*

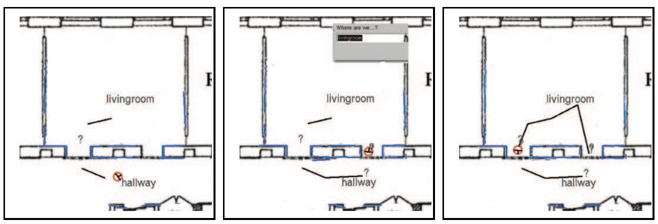

Fig. 3. *The run starting in the "living room" (left), leaving it and coming back in (centre, with the hypothesis "living room" in the dialogue box), and after merging the navigation graphs inside the room (right). Question marks indicate positions where the system asked for clarification and solid black lines represent the navigation graphs.*

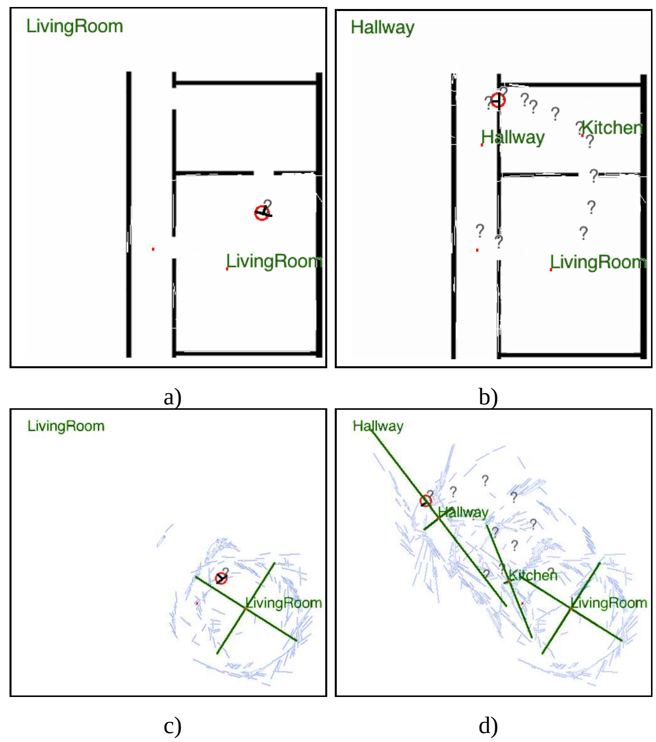

Fig. 4. The experiment with BIRON, visualised in post-hoc runs. Question marks indicate positions where the robot asked the user for confirmation. In the upper left corner of each image the system's hypothesis of the current *region* is shown. a) and b) Reconstruction with the help of a pose estimation module with the room labels marked at the positions where they were given to the robot by the user. c) and d) visualisation of the originally generated representation, based on raw odometer readings, the crosses indicating the ellipses for the three specified *regions*. a) / c) Starting in the "living room", b) / d) concluding the tour in the hallway after going through living room and kitchen twice.

other run the applicability for loop closing was tested by specifying the "living room" (one large laboratory room) and the connected hallway, where the robot was guided back into the "living room" through a different door than used when leaving it (fig. 3).

3) Test with a fully integrated interactive system: To test the applicability of the transition detection approach the third implementation, integrated into the communication framework of the robot BIRON, an experiment in a laboratory environment corresponding to a part of an apartment including living room, kitchen and hallway was conducted. An interesting aspect to the integrated system was that for this experiment no SLAM method was available to provide the corrected pose estimations usually assumed. We decided to use the experiment to investigate, in how far the now purely topological mapping subsystem, including the transition detection, would be capable of representing the environment in a way that allowed meaningful interaction with a user, relying only on the feature based representations.

Fig. 4 illustrates the guided tour with BIRON through the laboratory environment, conducted by a researcher acting as "user". Since the pose estimation error was obviously mostly depending on rotations of the robot platform (see the uncorrected illustration in fig. 4 c) and d)), the accumulated error was kept on a level that allowed to hypothesise the "hallway" correctly when it was re-entered, since no significant turning movements "on the spot" had been made after its specification. Figure 2 illustrates the similarity measures over time for the run, applying the same conditions for the detection of a transition as in the original run, i.e., a "new current *region* representation" is assumed (in this post-hoc run no confirmation question was actually posed) when a significant change ($d > 1.5$) is observed for more than three data cycles, and the robot is at least 1 meter away from the point were the previous representation was accepted as current one.

4) Summary: The results from the seven evaluated runs show, that most of the obvious transitions in our test environments are detected rather reliably. As "obvious transitions" we consider door passages, junctions of hallways (available in the office settings), and hallways opening into a room (available in the two domestic data sets for the "medium size apartment"). Most failures of the approach have to be counted regarding "false alarms". However, since we assume the user to assist the system, we consider this type of failure less critical than "false negatives". Those occurred significantly less often and only in one "apartment" setting. Adaptive setting of the threshold values to the type of environment ("narrow apartment" vs. "spacious laboratory") or the application of a more sophisticated change detection filter can be an option to cover such cases more appropriately. A number of spurious detections in one of the domestic settings can be explained with the user being very close to the robot (due to the interactive scenario) and thus covering larger parts of the laser range finder's "field of view". Such spurious detections can obviously be avoided by increasing the number of data cycles that a change needs to last, before a transisiton is hypothesised. This was done for the laboratory runs, where it seemed to have immediate impact in the sense that spurious detections did not occur that frequently, still being able to detect significant changes satisfyingly.

The second laboratory run (fig. 3) showed the advantage of using a feature based representation both for the detection of transitions and the representation of *regions*, so that it is not necessary to travel back to a previously observed path to hypothesise a loop closure, which would presumably be the case with a door detector in combination with pure metrical SLAM.

The aim of the integration of the mapping subsystem with the fully interactive framework on BIRON was to see if a meaningful interaction in and about the surroundings can be achieved with the proposed models and used representations. For this integrated system, it was decided to limit the functionality of the mapping subsystem to the rather basic situations described above, i.e., the specification of *regions* and the detection of transitions together with the resulting requests for confirmation. Within this limited context the question mentioned above can be positively answered at least for the discussed environment. The robot detected all expected transitions and produced only a very limited amount of surprising questions.

As an overall result we consider the approach to separating *regions* and detecting transitions between them a useful tool to support the acquisition of a usable and understandable representation of an arbitrary indoor environment, suitable for a meaningful communication with the user even without an underlying correction of the robot's pose estimation, as it could be demonstrated with the last presented experiment.

V. Conclusions and future work

In this article we presented our approach to the separation of *regions* (one central spatial concept in our framework for Human Augmented Mapping - HAM) in the environment and the detection of transitions between them. We assume an interactive guided tour in which a human user presents and explains a known environment to our robot.

We tested our implementation of the HAM framework, particularly its subsystem for topological graph building (region segmentation and transition detection), with off-line experiments as well as with a full interactive setup (graphical interface and tracking system included) and in one case integrated in a fully interactive framework, including dialogue abilities, in different on-line runs. With these experiments the applicability of our method for *region* segmentation and transition detection could be confirmed. No prior knowledge of spatial categories is needed to generate a topological graph representation of an arbitrary indoor environment that reflects the human user's conceptual understanding of the surroundings. This makes the approach very flexible. Our tests showed sufficiently good results in both office (or laboratory) and domestic environments. Other mentioned approaches aim to label an environment with spatial categories [9], [10], while our method can rather be considered to deal with transitions between any type of spatial categories. This makes it more flexible in situations where the spatial category is difficult to determine even for a human user. Thus, we consider our approach as a fast and easy-to-apply complement to such categorising methods.

So far the corrections made by the user are not persistent in the system – the robot simply "forgets" that it has already asked the user about a particular detected transition. Thus it has to be investigated, how the topological graph structure and respective representations of involved *regions* have to be changed persistently. Another interesting aspect for future investigations is whether personal differences can be handled with one instantiation of an environment representation or if several are needed.

On a more detailed level it would seem natural to investigate a more adaptive method to decide if in fact a transition has occurred. This should make the method better suitable to different types of environments (generally narrow or more open) without needing to adjust parameters manually. Another issue is to find out with the help of another study, where users actually would want the robot to react to a detected change and if the system acts in a comprehensible way when hypotheses about the current position are generated.

VI. Acknowledgements

The work described in this paper was mainly conducted at the Centre for Autonomous Systems, Royal Insitute of Technology (KTH), Stockholm, Sweden, partially funded by the European Commission Division FP6-IST Future and Emerging Technologies under Contract FP6-002020. We also thank the Applied Informatics group at the University of Bielefeld for their cooperation.

References

[1] E. Topp, *Human-Robot Interaction and Mapping with a Service Robot: Human Augmented Mapping*. Doctoral Dissertation, KTH School of Computer Science and Communication (CSC), Stockholm, Sweden, 2008.

[2] I. Nourbakhsh, A. Powers, and S. Birchfield, "Dervish an office-navigating robot," *AI Magazine*, vol. 16, no. 2, pp. 53–60, 1995.

[3] P. Beeson, N. Jong, and B. Kuipers, "Towards autonomous topological place detection using the extended voronoi graph," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Barcelona, Spain*, 2005.

[4] B. Kuipers and Y.-T. Byun, "A robust qualitative method for spatial learning in unknown environments," in *Proceedings of the National Conference on Artificial Intelligence (AAAI), Los Altos, CA, USA*, AAAI Press, 1988.

[5] H. Choset and K. Nagatani, "Topological simultaneous localization and mapping (slam): Toward exact localization without explicit localization," *IEEE Transactions on Robotics and Automation*, vol. 17, no. 2, pp. 125–137, 2001.

[6] A. Tapus, G. Ramel, L. Dobler, and R. Siegwart, "Topology Learning and Place Recognition using Bayesian Programming for Mobile Robot Navigation," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Sendai, Japan*, 2004.

[7] B. Kröse, "An efficient representation of the robot's environment," in *Proceedings of the International Conference on Intelligent Autonomous Systems (IAS), Venice, Italy*, IOS Press, 2000.

[8] P. Buschka and A. Saffiotti, "A virtual sensor for room detection," in *Proceedings of the IEEE International Conference on Intelligent Robots and Systems*, 2002.

[9] O. Martínez Mozos, C. Stachniss, and W. Burgard, "Supervised learning of places from range data using adaboost," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA), Barcelona, Spain*, 2005.

[10] S. Friedman, H. Pasula, and D. Fox, "Voronoi Random Fields: Extracting the Topological Structure of Indoor Environments via Place Labeling," in *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI), Hyderabad, India*, Springer, 2007.

[11] H. Zender, P. Jensfelt, Ó. Mozos, G.-J. Kruijff, and W. Burgard, "An Integrated Robotic System for Spatial Understanding and Situated Interaction in Indoor Environments," in *Proceedings of the National Conference on Artificial Intelligence (AAAI), Vancouver, BC, Canada*, 2007.

[12] Z. Zivkovic, O. Booij, B. Kröse, E. Topp, and H. Christensen, "From sensors to human spatial concepts: An annotated data set," *IEEE Transactions on Robotics*, vol. 24, no. 2, pp. 501–505, 2008.

[13] A. Haasch, S. Hohenner, S. Hüwel, M. Kleinehagenbrock, S. Lang, I. Toptsis, G. Fink, J. Fritsch, B. Wrede, and G. Sagerer, "Biron – the bielefeld robot companion," in *Proceedings of the International Workshop on Advances in Service Robotics (IWASR), Stuttgart, Germany*, Fraunhofer, IRB, 2004.