Research Article

Open Access

Akansel Cosgun* and Henrik I. Christensen

Context-aware robot navigation using interactively built semantic maps

<https://doi.org/10.1515/pjbr-2018-0020>

Received October 24, 2017; accepted July 12, 2018

Abstract: We discuss the process of building semantic maps, how to interactively label entities in them, and how to use them to enable context-aware navigation behaviors in human environments. We utilize planar surfaces, such as walls and tables, and static objects, such as door signs, as features for our semantic mapping approach. Users can interactively annotate these features by having the robot follow him/her, entering the label through a mobile app, and performing a pointing gesture toward the landmark of interest. Our gesture-based approach can reliably estimate which object is being pointed at, and detect ambiguous gestures with probabilistic modeling. Our person following method attempts to maximize future utility by search for future actions assuming constant velocity model for the human. We describe a method to extract metric goals from a semantic map landmark and to plan a human aware path that takes into account the personal spaces of people. Finally, we demonstrate context-awareness for person following in two scenarios: interactive labeling and door passing. We believe that future navigation approaches and service robotics applications can be made more effective by further exploiting the structure of human environments.

Keywords: mobile robot navigation, semantic mapping, human-robot interaction

1 Introduction

Millions of robots around the world are in operation today, however, most of them operate in factories and are physically separated from humans. In the future, robots could be deployed in human environments, such as hotels, hospitals, offices and homes, and they could be used for el-

Figure 1: An example of the type of map produced by our system. Planar features are visible by the red convex hulls and red normal vectors. The small red arrows on the ground plane show the robot's trajectory. The point clouds used to extract these measurements are shown in white and have been rendered in the map coordinate frame by making use of the optimized poses from which they were taken.

derly care, cleaning, welcoming guests, and object delivery. There are two observations we make for this problem domain. First, human environments are designed for human convenience. For example, rooms offer privacy and compartmentalization of activities, doors are easy to open for humans, people put their things on planar surfaces, and door signs help people distinguish different rooms. Contemporary approaches to robot navigation typically do not take advantage of such human-made structures. Second, robots will be in close proximity to humans and interacting with us on a daily basis. Standard robot path planning algorithms do not distinguish people from obstacles, thus ignoring the social aspect of navigation. In this paper, we focus on how to utilize human-made structures to improve reasoning capabilities of service robots, as well as complying with social conventions to navigate efficiently among people.

We aim to develop intelligent mobile robots that understand the semantics of human environments and the spatial relationships with and between humans. Our mapping approach leverages our prior knowledge of semi-structured human environments to provide a rich representation for service robotics tasks. Specifically, our maps contain high-level features such as objects, planar sur-

*Corresponding Author: Akansel Cosgun: Monash University, Australia; E-mail: akansel.cosgun@monash.edu

Henrik I. Christensen: University of California San Diego, USA; E-mail: hichristensen@ucsd.edu

faces and signs that contain text in addition to metric coordinates. For example, planar landmarks enable the robot to know the locations of tables, counters, rooms, and doors. Door sign landmarks enable room locations and numbers to be automatically added to the map, while object landmarks can be used for fetch and carry tasks. Our approach additionally supports manual annotation of these high-level landmarks, so that they can be referenced by name in interactions with users.

We use the user-annotated landmarks and people tracking in four ways to enable context-aware robot navigation: First, high-level features, such as planar surfaces, is used for robust localization. Second, users and robots refer to the same landmarks by name, which enable users to provide human-friendly navigation goals instead of goals in metric coordinates. Third, knowledge of nearby landmarks is used to infer the user's intention, which helps the robot to move appropriately in certain tasks such as door passing and person following. Finally, our path planner treats humans differently than obstacles, predicts future trajectories of people and takes into account human safety and comfort.

This paper integrates our previous work on interactive semantic map building and introduces navigation behaviors that use the information contained in the semantic maps. The contributions of this paper are as follows:

- A rich map representation that contains high-level features, such as planar surfaces, static objects, and door signs that are grounded in metric coordinates;
- A multi-modal interaction model for annotating semantic landmarks based on natural gestures and an app; and
- Use of the annotated landmarks and person tracking for demonstrating context-aware navigation behaviors.

We make use of external algorithms including GoogleGoggles image recognition engine, GTSAM mapping [1], and OpenNI NITE skeleton tracker.

The rest of this paper is organized as follows: A literature survey of related work is given in Section 2, followed by our semantic mapping approach in Section 3. Section 4 describes how users interactively label semantic elements, as well as sub-components, such as person following and pointing gestures, that make interactive labeling possible. In Section 5, we evaluate these sub-components. Section 6 discusses context-aware navigation behaviors and we conclude in Section 7.

2 Related work

Research on semantic mapping and context-aware robot navigation has been ongoing for several years, and a large body of work exists that is related to this paper.

We first survey mapping techniques for mobile robotics in Section 2.1, and do a deep dive on semantic mapping in Section 2.2 and human-augmented mapping in Section 2.3. We then review literature on person detection and tracking techniques in Section 2.4. Section 2.5 is concerned with context-aware navigation where we provide related work on human-aware navigation in Section 2.5.1 and navigation using semantic information in Section 2.5.2.

2.1 Mapping in robotics

For most tasks, mobile robots need to keep a representation of the environment based on sensor readings and possibly prior knowledge. The probabilistic formulation of creating this representation and building a map, called Simultaneous Localization and Mapping (SLAM), was first addressed by Smith and Cheeseman *et al.* [2], and Leonard and Durrant-Whyte *et al.* [3]. There are usually two types of SLAM approaches: filter-based and graph-based. Early works used filter-based SLAM approaches that focused on the temporal aspect of the sensor measurements. Graph-based SLAM approaches, instead of solving for only the current robot pose, typically maintain a graph of the entire robot trajectory in addition to the landmark positions [1]. Another area of interest is feature-based SLAM, which uses landmarks to solve the SLAM problem, such as the M-Space model [4]. A central challenge to SLAM is the data association problem, especially when the robot revisits a location [5].

Three types of map representations are commonly used in robotics: metric, topological, and semantic. Metric maps typically use low-level representations, such as raw sensor measurements (i.e. point clouds [6]), positions of salient features, or occupancy grids [7, 8]. In topological maps [9, 10], the environment is represented as a graph where nodes represent the distinct places in the environment and edges represent the connections between the places. Semantic maps aims to build richer, more useful maps that include objects, their categories, and common-sense knowledge.

In our work, we use a hybrid representation: a semantic map for task-level goal assignment and human-robot

interaction, and a metric map for motion planning. Below, we review the literature on semantic mapping.

2.2 Semantic mapping

Semantic mapping uses high-level modalities such as object recognition, optical character recognition and interaction with humans. Kuipers *et al.* [11] proposed the Spatial Semantic Hierarchy (SSH), which is a qualitative and quantitative model of knowledge of large-scale space consisting of multiple interacting representations. This map also informs the robot of the control strategy that should be used to traverse between locations in the map. Martinez-Mozos *et al.* [12] introduce a semantic understanding of the environment creating a conceptual representation referring to functional properties of typical indoor environments. Ekvall *et al.* [13] integrated an augmented SLAM map with information based on object recognition, providing a richer representation of the environment in a service robot scenario. Nüchter *et al.* [14] investigated semantic labeling of points in 3D point cloud based maps. Semantic interpretation was given to the resulting maps by labeling points or extracted planes with labels such as floor, wall, ceiling, or door. Pronobis *et al.* [15] proposed a joint spatial-semantic environment model by fusing multi-modal data including natural language and object classifiers. Recent work in semantic mapping include object-oriented semantic mapping [16, 17]. These methods use objects as landmarks and can create maps that are meaningful to humans. Semantic maps can be useful for describing spatial relations with natural language [18, 19], such as understanding commands like “get the mug on the table”. For more related work on this topic, the reader is referred in-depth surveys on perception approaches to semantic mapping by Kostavelis *et al.* [20] and on spatial reasoning by Landsiedel *et al.* [21].

In our approach, we utilize multiple modalities of semantic features, including household objects, door signs and labeled planar surfaces.

2.3 Human-augmented mapping

Human-augmented mapping was first introduced by Topp *et al.* [22], where a human assists the robot in the map building process. This is motivated by the scenario of a human guiding a service robot on a tour of an indoor environment and adding relevant semantic information to the map throughout the tour for later reference. Users could ask the robot to follow them throughout the environment

and provide labels for locations, which could later be referenced in commands, such as “go to label”. This means of providing labels seems quite intuitive, as users are co-located in the environment with the robot platform.

One of the key concepts in semantic mapping is that of “grounding”, or establishing “common ground” [23]. Of particular interest for mapping is grounding references, in order to ensure that the human and robot have common ground when referring to regions of a map, structures, or objects. Many spatial tasks may require various terms to be grounded in the map. Dialog in human augmented mapping has been investigated in [24]. Clarification dialogs were studied in order to resolve ambiguities in the mapping process, for example, resolving whether or not a door is present in a particular location. This was applied to the Cosy Explorer system, described in [25], which includes a semantic mapping system that build multi-layered maps, including a metric feature based map, a topological map, as well as detected objects.

Gemignani *et al.* [26] presents evaluation of an interactive semantic mapping system. In contrast to our work, their approach does not utilize semantic features as landmarks during SLAM. Their approach, however, extracts a topological map from the semantic map in order to facilitate task planning. This work and many others use natural language as the modality to provide labels for the semantic map whereas we use a smartphone app.

2.4 Person detection and tracking

The applicability of person detection and tracking is wide ranging, including congestion analysis in crowded places, security, diagnostics of orthopedic patients, autonomous vehicles and human-computer interfaces. A large body of work exists in the computer vision area; an extensive survey is given in [27]. Popular methods in image-based person detection include using temporal templates [28], histogram-based methods [29], deformable part-based methods [30, 31] and multi-modal methods [32]. Depth cameras are commonly used for body pose estimation [33]. More recently, convolutional neural networks [34] and deep learning [35] became the dominant method for image-based object detection. These techniques have been applied to person detection [36] and tracking [37].

For mobile robotics, laser scanners remain the most commonly used sensor for person detection and tracking, because, as opposed to monocular cameras, they can more easily determine the distance to the detections and their higher field of view makes it possible for a single sensor to cover the surroundings of the robot. Legs in laser scans are

typically distinguished using a multitude of geometric features [38]. Schulz *et al.* [39] uses particle filters and statistical data association. Topp *et al.* [40] demonstrates that leg tracking in cluttered environments is possible, but prone to false positives. Bellotto *et al.* [41] combine leg detection and face tracking in a multi-modal tracking framework. Zanlungo *et al.* [42] utilize the Social Forces Model to describe pedestrian motions, where parameters are trained with real pedestrian data. Leigh *et al.* [43] track multiple people with laser scanners. Dondrup *et al.* [44] present a framework that utilizes multiple sensor modalities for real-time tracking.

Person tracking provides the robot with the position, and potentially orientation of the humans. However, richer information is typically needed for Human-Robot Interaction (HRI) applications. Pointing gestures are commonly used in HRI, such as for object references [45] and providing navigation goals [46]. After deciding if a pointing gesture occurred or not, typically the direction of pointing is also estimated. A commonly used method is to extend a ray from a body part to another and assume this ray is aimed toward the object of interest. The two of most commonly used methods are elbow-hand [47] and head-hand rays [45].

We use a laser-based torso detection approach and track each person individually using a Kalman Filter. Our approach to data association is nearest neighbors. Our pointing gesture approach can take as input both the elbow-hand and head-hand rays, and uses pointing statistic priors to determine the target object.

2.5 Context-aware navigation

Path planning for mobile robotics is traditionally seen as a shortest-path problem and doesn't utilize semantic information. While such approaches generate collision-free paths, the resulting robot behavior may not be preferable to humans. For example, the robot would make people feel unsafe by getting too close to them, or it won't be able to predict the intentions of people if it doesn't recognize gestures. Context-aware navigation has found interest in two fronts: human-aware navigation and navigation using semantic information.

2.5.1 Human-aware navigation

Human-aware navigation algorithms are concerned with planning a motion for a mobile robot given obstacles and people around.

A common situation in human environments is when the robot encounters bystanders on the way to its goal position. An approach to encode mobility constraints for navigating around humans is through costmaps [48, 49]. These approaches typically model personal spaces by assigning costs according to distance and orientation of the robot with respect to humans. Walters and Dautenhahn *et al.* [50] show that people's personal spaces can differ according to their personality, gender and preferences. Predicting the future movements of people is important for planning robot motion. Luber *et al.* [51] trains a model to estimate the future relative motion of people and plan a path. Kidokoro *et al.* [52] simulates hypothetical situations using real data to anticipate how pedestrians' walking comfort would be affected. Bordallo *et al.* [53] and Köeckemann *et al.* [54] first explicitly estimate the goal of the people, and then plan for the robot motion accordingly. Understanding the predictability and legibility of robot motion by human observers is a relevant factor in designing robot behaviors [55, 56]. There has been efforts to extend the aforementioned ideas to navigation among crowds [57, 58].

Another type of an application is when a person is part of the goal definition, such as when the robot is following, guiding [59] or moving alongside a specific person [60]. Our work involves person following and here we review related works on that topic. Ohya *et al.* [61] present a following method that escorts a target on the side while avoiding obstacles. It was assumed that the target would move with the same acceleration and velocity. Murakami *et al.* [62] present a method to first estimate the sub-goal of the leading person and then following as if the robot knows the goal. Park *et al.* [63] model the problem as a control problem and offer an algorithm based on Model Predictive Control. Granata *et al.* [64] present behaviors such as going towards, following and searching a user. Gockley *et al.* [65] compared two elementary following methods: direction following, where the robot always attempts to drive towards the tracked person, and path following, in which the robot follows the exact path the person took. It was shown that direction following behavior was perceived as more human-like and natural than path following. More detailed surveys on human-aware navigation can be found in [66, 67].

We use a costmap-based approach similar to Sisbot *et al.* [48] and Kirby *et al.* [49]. Similar to Bordallo *et al.* [53], the path is planned by taking into consideration the future movements of humans. Our person following approach involves a limited-horizon search and allows different robot positioning around the human.

2.5.2 Navigation using semantic information

The robot can exploit the information contained in semantic maps and possibly prior domain knowledge from the environment to increase the effectiveness of its navigation capabilities.

Regier *et al.* [68] present a planner that predicts traversal costs of potential paths by considering the amount of clutter in the environment. Pacchierotti *et al.* [69] adjust the robot's speed when the robot is in a hallway setting. Wilde *et al.* [70] learns the cost function weights for path planning from users who choose the path for the robot in a map that contains semantic information. Galindo *et al.* [71] generate goals for the robot when there are violations of semantic knowledge.

Natural gestures and spoken language are often used to boost HRI: Lu *et al.* [72] show that using gaze cues makes robot-human hallway passing more efficient. Loper *et al.* [73] presents a system that is capable of responding to verbal and non-verbal gestures and following a person. Anderson *et al.* [74] present a method that interprets visually-guided navigation instructions using deep learning. Tellex *et al.* [19] address the same problem but use a graphical model.

Zender *et al.* [75] considers context-awareness for person following, specifically for handling of door and corridor passages. To handle door passages, the robot increases its following distance and that leads the robot to wait for a while. Our approach to navigation using semantics is similar to this work, as we also use objects such as doors for case-based behavior generation.

3 Semantic mapping

As service robots become increasingly capable and are able to perform a wider variety of tasks, we believe that new mapping systems could be developed to better support these tasks. Towards this end, we developed a SLAM approach that uses planar surfaces and objects as landmarks, and maps their locations and extent. We chose planar surfaces because they are prevalent in indoor environments in the forms of walls, tables, and other surfaces. We also utilize door signs, and use this information to enhance robot navigation behavior.

Non-technical users prefer human terms for objects and locations when assigning tasks to robots instead of whatever indices or coordinates the robot uses to represent them in its memory. Semantic mapping offers an advantage for robots to understand task assignments given

to them by human users. We allow humans to label planar landmarks that are automatically acquired during the SLAM process, as described in Section 4. Users provide navigation goals in terms of these labeled landmarks. Our approach of finding goal points for a given planar landmark will be discussed in Section 6.1.1.

Planar landmarks provide semantic information about the space, as vertical planes correspond to walls, showing how space is partitioned, while horizontal planes correspond to tables and shelves, where objects of interest may occur. We describe in Section 3.2 how higher level objects, specifically door signs, can be used as landmarks in SLAM. We further explore in Section 6.2.2 how detection of door signs and therefore the existence of doors, can be used for robot navigation.

3.1 Plane landmarks

We believe that feature-based maps are suitable for containing task-relevant information for service robots. For example, a home service robot might need to know the locations of kitchen tables, countertops, cupboards and shelves. Structures such as walls could be used to better understand how space is structured and partitioned. We describe a SLAM approach capable of creating maps of the locations and extents of planar surfaces in the environment using both 3D and 2D landmarks.

Our SLAM implementation makes use of the GTSAM library [1]. This library represents the graph SLAM problem with a factor graph which relates landmarks to robot poses through factors. GTSAM builds a factor graph of nonlinear measurements. Our approach involves using multiple types of landmark measurements as factors of nonlinear measurements. Planar surfaces are detected in point cloud data generated by a Asus Xtion RGB-D camera. An example of a map produced by our system is shown in Figure 1.

A plane in $\mathbb{R}^3$ has the equation

$$ax + by + cz + d = 0, \quad (1)$$

where a, b, c, d are parameters that define the plane and x, y, z are cartesian coordinates of a point that lies on the plane. We use this representation for planes, while additionally representing the plane's extent by calculating the convex hull of the observed points. While only the plane normal and perpendicular distance are used to correct the robot trajectory in SLAM, it is essential to keep track of the extent of planar patches, as many coplanar surfaces can exist in indoor environments, and we would like to represent these as distinct entities. We therefore represent

planes as

$$p = [n, \text{hull}], \quad (2)$$

where

$$n = [a, b, c, d] \quad (3)$$

and *hull* is a point cloud consisting of the vertices of the plane's convex hull. As planes are re-observed, their hulls are extended with the hull observed in the new measurements. That is, the measured hull is projected onto the newly optimized landmark's plane using its normal, and a new convex hull is calculated for the sum of the vertices in the landmark hull and the measurement's projected hull. In this way, the convex hull of a landmark can grow as additional portions of the plane are observed.

We use a Joint Compatibility Branch and Bound (JCBB) technique for data association [76]. JCBB works by evaluating the joint probability over the set of interpretation trees of the measurements seen by the robot at one pose. The output of the algorithm is the most likely interpretation tree for the set of measurements. We are able to evaluate the probability of an interpretation tree quickly by marginalizing out the irrelevant portions of the graph of poses and features. The branch and bound recursion structure from the EKF formulation is used in our implementation.

Given a robot pose X_r , a transform from the map frame to the robot frame in the form of $(R, \vec{t})$, a previously observed feature in the map frame $(\vec{n}, d)$ and a measured plane $(\vec{n}_m, d_m)$, the measurement function h is

$$h = \begin{pmatrix} R^T * \vec{n} \\ \langle \vec{n}, \vec{t} \rangle + d \end{pmatrix} - \begin{pmatrix} \vec{n}_m \\ d_m \end{pmatrix}. \quad (4)$$

The Jacobian with respect to the robot pose is

$$\frac{\partial h}{\partial X_r} = \begin{bmatrix} 0 & -n_a & n_b & 0 \\ n_a & 0 & -n_c & 0 \\ -n_b & n_c & 0 & 0 \\ 0 & 0 & 0 & \vec{n}^T \end{bmatrix}. \quad (5)$$

The Jacobian with respect to the landmark is

$$\frac{\partial h}{\partial n_{map}} = \begin{bmatrix} [R_r] & \vec{0} \\ X_r^T & 1 \end{bmatrix}. \quad (6)$$

Using this measurement function and its associated Jacobians, we can utilize planar normals and perpendicular distances as landmarks in our SLAM approach. During optimization, the landmark poses and robot trajectory are optimized.

3.2 Object landmarks: door signs

The previous section introduced how we use planar landmarks for SLAM. In this section, we present a method for using a learned object classifier in a SLAM context to provide measurements suitable for mapping.

First, walls are extracted from straight lines in the laser scan. We use a RANSAC technique to extract lines from the laser data. Only lines which are longer than a certain threshold are passed to the mapper as measurements.

Figure 2: This sign is recognized and a measurement is made in the mapper. GoogleGoggles has read both the room number and the text, so this sign can be used for data association.

The door-sign-detector module makes use of a Support Vector Machine (SVM) classifier, trained on Histogram of Oriented Gradient (HOG) features. If an image region is classified as a sign by the SVM then a query is made from this image region to the GoogleGoggles server. If GoogleGoggles is able to read any text on the sign then it will be returned to us in a response packet. Detected signs with decoded text are then published as measurements that can be used by the mapper. The measurements consist of the pixel location in the image of the detected region's centroid, the image patch corresponding to the detected region, and the text string returned from GoogleGoggles. An example detection of a door sign is shown in Figure 2.

Detected lines in the laser scan and by the door sign detector are added as non-linear measurements to the factor graph. At the time of this study, we did not have a RGB-D sensor on the robot. Therefore, measurements were made on the 3D coordinates of the back-projected image location directly. Range is recovered by finding the laser beam from the head laser which projects most closely to the image coordinates of the sign. This technique approximates the true range. This factor also incorporates an additional variable which corresponds to the transformation between the robot base and the camera.

To implement this factor in GTSAM, we must specify an error function and the error function's derivatives in terms of all of the variables which contribute to it. The error function is the difference in the 3D position of the predicted location of the sign and the measured value given by the recognition module.

4 Interactive map labeling

There are several methods to support the annotation of entities in a robot map. For example, while the robot is building its representation of the environment, it can recognize objects or landmarks, such as doors, tables, rooms, and automatically add these features to its map. Even though such a system would be useful, it may wrongly label some objects. In that case, the correct label can be provided by a human with an interactive system. Custom labels would also allow custom annotations such as “Joe’s Room”.

We presented our method for building semantic maps in Section 3. We will assume that we have a metric and semantic map for the rest of this paper for simplicity. To enable a common ground between the humans and the robot, we developed an interactive procedure to annotate landmarks in the semantic map. In this procedure, the person guides the robot to the landmark of interest first and refers to the landmark of interest by pointing at it. The steps for labeling a landmark is shown in Figure 3. The robot follows the user around between labeling of landmarks. This allows the user to guide the robot to virtually any location in the environment. Our interactive map labeling approach was previously described in [77].

Various modalities could be used for the interaction model, such as speech or GUI-only. The reason why we combine the GUI with pointing gestures is that we think natural gestures would play an important role for HRI in the future.

For the rest of this section, we describe the sub-components necessary to realize interactive labeling, namely person following (Section 4.1), pointing gestures (Section 4.2), and object labeling (Section 4.3).

4.1 Person following

The robot has to continuously estimate the position of the user in real-time for robust person following. We focus on tracking people who are either walking or standing, as these are the two most common human poses around a mobile robot. Below are the brief descriptions of the per-

son detection methods used on the robot as shown in Figure 6(b):

1) Leg Detector: A front-facing laser scanner at ankle height (Hokuyo UTM 30-LX) is used. We trained a leg classifier using three geometric features: width, circularity and inscribed angle variance [78]. We find a distance score for each candidate segment using the weighted sum of the distance to each feature and then threshold the score for detection.

2) Torso Detector: A back-facing Hokuyo laser scanner placed at torso-level is used for this detector. We model the human torso as an ellipse and fit each segment in the laser as an ellipse. The ellipse fitting approach always returns a result, even for bad data. In addition to the geometric features we use for legs, we use two additional features for torso detection: the horizontal and vertical axes of the fitted ellipse. Similar to leg detection, we use a threshold test for the detection result.

The output of the detectors are input to a state estimation module. Using a state predictor for human movement have two advantages. First, the predicted trajectories are smoother than raw detections. Smooth tracking helps the robot maintain consistent trajectories for person following. Second, it provides a posterior estimate that can be used for data association when there is a lack of matching detections. This allows the tracker to handle temporary occlusions. We use a linear Kalman Filter (KF) with a constant velocity model to estimate the position and velocity of a person. We used a KF for tracking because it has acceptable tracking performance and is computationally cheap, which is important in real-time applications.

For person following, the robot uses the most probable location of the KF, which is the mean of the Gaussian distribution. We use Dynamic Window Approach (DWA) [79] at the core of our planner to sample velocity and acceleration-bounded trajectories with the modification of using time as an additional dimension. DWA forward-simulates allowable velocities and chooses an action that optimizes a function that will create a goal-directed behavior while avoiding obstacles. Our approach projects the future locations of the target, creates a tree of trajectories, and scores each tree node according to a goal function which is a function of the relative pose of the robot with respect to the human. Our planner takes the laser scan measurement, predicted positions of the person, and the number of time steps to plan as input and outputs a sequence of actions. A robot configuration at time t is expressed as $q^t = (x^t, y^t, \theta^t, v^t, \omega^t)^T$, where x^t and y^t denote positions, θ^t is the orientation, v^t and w^t are the linear and angular velocities at time t . The person configuration p^t is defined the same way. An action of the robot is de-

Figure 3: Steps for interactively labeling a landmark in the semantic map. First, the user activates person following using the app. The user stops nearby the target landmark, and enters the requested label using the app. Then the user performs a pointing gesture towards the target and waits for acknowledgement. The robot assesses the likelihood of nearby objects being the intended target, and asks for confirmation if ambiguity is detected. Finally, a string label is attached to the corresponding landmark in the semantic map.

defined as a velocity command for some duration: $a(t, \Delta t) = (v_a^t, \omega_a^t, \Delta t)$. We assume a unicycle kinematics model for trajectory sampling. Using this model, we generate a tree up to a fixed depth, starting from the current configuration of the robot. A tree node consists of a robot configuration as well as the information about the previous action and parent node. Every depth of the tree corresponds to a discretized time slice. Therefore, every action taken in the planning phase advances the time by a fixed amount. This enables the planner to consider future steps of the person and simulate what is likely to happen in the future. The planner uses depth-limited Breadth First Search (BFS) to search all the trajectories in the generated tree and determines the trajectory that will give the robot the maximum utility over a fixed time in the future. Given the robot and person configuration at some particular time, goal function $0 \leq g(q^t, p^t) \leq 1$ determines how desirable the situation is for the task. Goal function can be defined in any way and provides flexibility to the navigation behavior designer. We assume that it is desirable for the robot to follow from behind as it would give the robot a better chance to predict the human's motions and implicitly mimic human's path, which is known to be obstacle free. We report our results on this person following method in Section 5.1. In previous work, we applied this following method on an autonomous telepresence robot [80].

Another important capability to enable interactive labeling is to be able to refer to the same landmarks in the environment. Our design involves the person extending his/her arm and point at the landmark of interest. The method for detecting the pointing target is described next.

4.2 Pointing gestures

In this section, we present an uncertainty model for estimating pointing gesture targets based on previous work [81]. Estimating this uncertainty allows us to interpret whether a pointing gesture is ambiguous or not and when objects are too close to one another. We model the uncertainty of pointing gestures using a spherical coordinate system. We use this model to determine the correct pointing target and detect when there is ambiguity. As reviewed in Section 2.4, a common method in inferring pointing gesture directions is to extend a virtual ray from a body part to another. We evaluate two of the most commonly used rays, elbow-hand and head-hand, using a 3rd party skeleton tracking algorithm, OpenNI NITE in Section 5.2. We use a simple gesture detection algorithm as our focus is to estimate the gesture target given that a pointing gesture was performed. Using the skeleton data, pointing gestures are recognized if a human's forearm makes more than a fixed angle with the vertical axis and elbow and hand joints stay almost stationary for a fixed duration. Gesture detection is activated only after the user requests a labeling action. This design is intended to reduce false positive detections.

We represent a pointing ray in two angles: a "horizontal" sense we denote as θ and a "vertical" sense we denote as ψ . We first attach a coordinate frame to the hand point, with its z-axis oriented in either Elbow-hand $\vec{v}_{eh}$ or Head-Hand $\vec{v}_{hh}$ directions. The hand was chosen as the origin for this coordinate system because both of head-hand and elbow-hand pointing methods include the user's hand. The transformation between the sensor frame and the hand frame ${}^{sensor}T_{hand}$ is calculated by using an angle-axis rotation method. An illustration of the hand coordi-

nate frame for Elbow-Hand method and corresponding angles are shown graphically in Figure 4.

Figure 4: Vertical (ψ) and horizontal (θ) angles in spherical coordinates are illustrated. A potential intended target is shown as a star. The z-axis of the hand coordinate frame is defined by either the Elbow-Hand (this example) or Head-Hand ray.

Given this coordinate frame and a potential target point P , we first transform it to the hand frame

$${}^{hand}p_{target} = T_{hand} p_{target}. \quad (7)$$

We calculate the horizontal and vertical angles for the target point ${}^{hand}p_{target} = (x_{targ}, y_{targ}, z_{targ})$ as

$$[\theta_{target} \ \psi_{target}] = [\text{atan2}(x_{targ}, z_{targ}) \ \text{atan2}(y_{targ}, z_{targ})] \quad (8)$$

where $\text{atan2}(y, x)$ is a function that returns the value of the angle $\arctan(\frac{y}{x})$ with the correct sign.

We estimate the likelihood of objects being the target using statistical data from previous pointing gesture observations. We observed that head-hand and elbow-hand methods returned different angle errors depending on the target location. Our approach relies on finding error statistics of these approaches, and compensating the error when the target object is searched for. First, given a set of prior pointing observations, we calculate the mean and variance of the vertical and horizontal angle errors for each pointing method. This analysis will be presented in Section 5.2. Given an input gesture, we apply correction to the pointing direction and find the Mahalanobis distance to each object in the scene.

When a pointing gesture is recognized and the angle pair $[\theta_{target} \ \psi_{target}]$ is found then a correction is applied by subtracting the mean terms from measured angles

$$[\theta_{cor} \ \psi_{cor}] = [\theta_{target} - \mu_{\theta} \ \psi_{target} - \mu_{\psi}]. \quad (9)$$

We also compute a covariance matrix for angle errors in this spherical coordinate system:

$$S_{type} = \begin{bmatrix} \sigma_{\theta} & 0 \\ 0 & \sigma_{\psi} \end{bmatrix}. \quad (10)$$

We get the values for μ_{θ} , μ_{ψ} , σ_{θ} , σ_{ψ} from Table 2 for the corresponding gesture type and closest target location. We then compute the Mahalanobis distance to the target

$$D_{mah}(target, method) = \sqrt{[\theta_{cor} \ \psi_{cor}]^T S_{method}^{-1} [\theta_{cor} \ \psi_{cor}]} \quad (11)$$

We use D_{mah} to estimate which target or object is intended. We consider two use cases: the objects are represented as a point or a point cloud. For point targets, we first filter out targets that have a Mahalanobis distance larger than a threshold $D_{mah} > D_{thresh}$. If none of the targets has a D_{mah} lower than the threshold then we decide the user did not point to any targets. If there are multiple targets that has $D_{mah} \leq D_{thresh}$ then we determine ambiguity by employing a ratio test. The ratio of the least D_{mah} and the second-least D_{mah} among all targets is compared with a threshold to determine if there is ambiguity. If the ratio is higher than a threshold then the robot can resort to additional action, such as initiating a dialogue to ask or confirm the intended object.

4.3 Labeling object models

Our method supports labeling two types of landmarks: planar surfaces and objects. The UI shows two labeling buttons “Label Object” and “Label Planar Surface”, so that the robot knows what the user is intending to label. Section 4.2 demonstrated how point targets or point cloud targets can be referenced via pointing gestures. For labeling planar surfaces, once the pointing gesture is performed, we check whether any planar feature in the semantic map intersects with the corrected gesture direction. For labeling object models, first the large planar surface corresponding to the table is detected. This is removed from the point cloud, and point clusters above this are detected. The cluster with a centroid nearest to the reference point is selected as the object to be modeled. The cluster’s points are projected into the camera image and are used to generate a region of interest. SURF features are detected for the region of interest, and are stored as an object model along with the provided label.

We assume that the objects are unique and will remain static throughout, which is obviously a strong assumption for real operation. The object consistency problem is tackled by object-based mapping research [16, 17], but is not the focus of this paper.

Once the intended landmark is determined it is annotated with the label entered by the user and can then be recognized later as described in our previous work [82]. Figure 5 shows the steps the robot executes for this task. Service robotic tasks that use such a map can then reference the object by label, rather than generating a more complex referring expression (e.g. “the large object” or “the object on the left”).

Figure 5: a) A user pointing at an object; b) A detected pointing gesture (blue and green spheres) and the referenced object (red sphere); c) Features are extracted from an image patch corresponding to the cluster and annotated with the provided label.

5 Evaluation of sub-components

In this section, we evaluate two of the core sub-components that enable building interactive semantic maps. In Section 5.1 we analyze the person following behavior and in Section 5.2 we evaluate our pointing gesture target estimation method.

The algorithms presented in this paper were implemented on three robot platforms, as shown in Figure 6.

5.1 Evaluation of person following

In this section, we report on an analysis of the person following sub-component. In our experiments, the robot fol-

lowed seven people for three laps and we logged the total distance robot followed the person and the average distance to the person. Users did not have prior experience with the robot and were asked to walk in a corridor while the robot is following them. Subjects were encouraged to adjust to the robot’s speed which was slightly lower than regular walking speed. A lap consisted of leaving the starting point, going to an intermediate point at the end of the corridor, and coming back to the starting point from the same path. The corridor was L-shaped and did not have any clutter or obstacles other than a couple of tall columns. Occasionally, the robot lost track of people due to fast motions of people, and the experiment was continued after re-initializing the tracking. The goal function for the robot was chosen such that there were two global minimums: a region about 1m behind the human, and another region that 1m behind and 0.8m to the right of the human. These goal regions are designed to encourage the robot to either stay behind the human or shift a bit to one side accordingly to the obstacles around. Table 1 shows the data pertaining to this study. In total, the robot followed people for about 1.2km. The robot did not come into contact with any obstacles or people during the experiments and the average distance between the person and the robot across all runs was 1.16m.

Table 1: Performance of the person follower on seven subjects. Each row shows a run where the robot followed the subject on a course. The distance traversed per run is given in the second column. The average distance between the robot and the human was provided in the third column. The average distance to the person across runs are relatively consistent, and slightly higher than 1 meter because the robot got higher rewards by keeping that much of distance to the human.

Subject #	Dist. traversed (m)	Avg. dist. to human (m)
1	171.4	1.1
2	161.8	1.13
3	160.4	1.14
4	169.9	1.25
5	174.2	1.04
6	166.2	1.3
7	171.5	1.2

Figure 7 plots the time versus following distance for a sample run that consists of 1 lap. At $t = 0$, the following is initiated and robot leaves the starting point. Around $t = 8s$, the robot and the person start making a right turn. The sudden drop in following distance at $t = 57.2$ signifies that the robot lost track of the person and the person

Figure 6: Robot platforms used in this paper. The robot shown in a) is a telepresence robot fitted with a Microsoft Kinect sensor. It was used for the person following experiments presented in Figure 5.1. The robot shown in b) has a Segway base with caster wheels, two laser scanners (one at torso height), and is used in human-aware path planning in Section 6.1.2. The robot shown in c) has a Segway base with caster wheels, has a UR5 robot arm, Ocular 3D rotating sensor and an Asus Xtion RGB-D sensor. This robot is used for the rest of the experiments including pointing gestures in Section 5.2 and context-aware person following in Section 6.2. All robot platforms have laser scanners for navigation.

Figure 7: Following distance as a function of time for one of the person following experiments.

tracker is reinitialized. At $t = 65.6s$, the intermediate point at the end of the corridor is reached so the person makes a 180° turn. The robot is close to the person (about $0.75m$) around this time because the robot is rotating around itself while the person is turning back. Between $t = 70$ and $t = 80s$, the person walks faster than the robot, so the following distance reaches to a maximum of $1.91m$. At $t = 79.8s$, the person is lost again. Around $t = 115s$, the robot and the person make a left turn. The lap ends at $t = 124s$. Note that the high frequency fluctuation of the following distance is a result of tracking only a single leg. There could be other contributors to this error, including inconsistent human walking speeds and noise in sensor data.

Our person following approach is applied to a telepresence robot where there is a remote user connected to the robot, as seen in Figure 6(a). In a study with then subjects, the motions of the robot was found natural, evidenced by getting an average score of 5.4 on 7-point Likert scale. The reader is referred to [80] for more details on the user study.

5.2 Evaluation of pointing gestures

To evaluate the accuracy of pointing gesture target detection, we first find the error statistics for pointing gestures and then apply it to a scenario for distinguishing two objects with varying separation.

We collected data from six users with seven targets where people pointed at each target with their right arms (Figure 8). Our use case is on a mobile robot platform capable of positioning itself relative to the user (Figure 6). For this reason, we can assume that the user is always centered in the image as the robot can easily rotate to face the user and can position itself at a desired distance from the user. The ground truth points that are represented in the camera frame are found by first finding the pixel values of targets using a corner detector, extending a virtual ray from the camera's origin to a target, and finding the intersection point to the supporting plane which is extracted from the point cloud data. We computed the mean and standard de-

Figure 8: Our study involved six users that pointed to seven targets being recorded using 30 frames per target. Four targets were placed horizontally on the table (indicated in blue) and three targets were placed vertically on the wall (indicated in yellow).

viations of the angular errors in the spherical coordinate system for each pointing gesture method and target.

The error statistics for Target 2 and across all targets are given in Table 2. The reason for reporting Target 2 only is that we use the error statistics of that target for the object separation study. From the data, we can tell that for the elbow-hand pointing method users typically point about 11° to the left of the intended target direction, and about 9° above the target direction. Similarly, the data from the head-hand pointing method reports that users typically point about 2° to the left of the intended pointing direction, but with a higher standard deviation than the elbow-hand method. On average, the vertical angle ψ was about 5° below the intended direction with a higher standard deviation than the elbow-hand method. The horizontal angle θ has a higher variation than the vertical angle ψ . Examining the errors for individual target locations shows that this error changes significantly with the target location. Therefore, for a given target location, we first choose the closest target category in our data set, and use the corresponding mean and standard deviation values.

Next, we conducted an experiment to determine how our approach distinguished two potentially ambiguous pointing target objects. The setup consisted of a table between the robot and the person and two coke cans on the table (Figure 9) where the separation between objects was varied. The center positions of objects were calculated in real-time by a point cloud segmentation with supporting plane assumption. The separation between objects were varied with 1 cm increments from 2 cm to 15 cm and with 5 cm increments between 15 cm – 30 cm. We could not conduct the experiment below 2 cm separation because of the limitations of our perception system. The experiment was conducted with one user who was not in the training dataset. For each separation the user performed five

Figure 9: Example scenarios from the object separation test are shown. Our experiments covered separations between 2 cm (left images) and 30 cm (right images). The object is comfortably distinguished for the 30 cm case, whereas the intended target is ambiguous when the targets are 2 cm apart. Second row shows the point cloud from the RGB-D camera's view. Green lines show the Elbow-Hand and Head-Hand directions whereas green circles show the objects that are within the threshold $D_{mah} < 2$.

pointing gestures to each object. The person pointed to one of the objects and the Mahalanobis distance D_{mah} to the intended object and the other object is calculated. Error statistics of Target 2 (Figure 8) was used for this experiment.

The results of the object separation experiment is given for Elbow-Hand (Figure 10(a)) and Head-Hand (Figure 10(b)) methods. The graphs plot object separation versus the Mahalanobis distance for the intended unintended objects for corrected and uncorrected pointing gestures. First, the Mahalanobis distance D_{mah} for the intended object was always lower than the other object. The corrected D_{mah} for both Elbow-Hand and Head-Hand methods for the intended object was always below 2. Because of that, we chose the threshold $D_{thres} = 2$. We notice that some distances for the unintended object at 2 cm separation is also below $D_{mah} < 2$. Therefore, when the objects are 2 cm apart the pointing target becomes ambiguous for this setup. For separations of 3 cm or more, D_{mah} of the unintended object is always over the threshold so there is no ambiguity. Second, correction significantly improved Head-Hand accuracy at all separations, slightly improved Elbow-Hand between 2-12 cm but slightly worsened Elbow-Hand after 12 cm. Third, the Mahalanobis distance stayed generally

Table 2: μ and σ of angular errors (in degrees) are given for Target 2 and across all targets. Error statistics of Table 2 was used for the object separation evaluation. The reader is referred to [81] for the complete table.

	Target 2				All Targets			
	θ		ψ		θ		ψ	
	μ	σ	μ	σ	μ	σ	μ	σ
Elbow-Hand	-3.8	6.6	11.3	10.9	-11.2	7.6	9.6	6.3
Head-Hand	10.2	6.7	-5.7	8.0	-2.4	9.6	-5.3	6.4

Figure 10: Resulting Mahalanabis distances of pointing targets from the Object Separation Test is shown for a) Elbow-Hand and b) Head-Hand pointing methods. Plots for intended objects are shown in green and for the unintended objects are shown in red. Solid lines show distances after correction is applied. Less Mahalanabis distance for intended object is better for reducing ambiguity.

constant for the intended object, which was expected. It linearly increased with separation distance for the other object. Fourth, patterns for both methods are fairly simi-

lar to each other, other than Head-Hand uncorrected distances being higher than Elbow-Hand.

6 Robot navigation using semantic maps

Autonomous navigation is one of the most fundamental capabilities for a mobile robot. There are many approaches that achieve point-to-point autonomous navigation thanks to the advances in mapping, localization and motion planning research. Many of these algorithms are optimized to find the least-cost path or the shortest path. However, often there are additional social factors to consider for navigation among humans.

First, it is not natural for humans to provide the goals in exact coordinates. Instead, the robot should be able to understand goals that are expressed in natural language and grounded with shared references. In our approach users provide annotated landmarks as goals to the robot using a mobile app. Our method of goal calculation from a user query is discussed in Section 6.1.1.

Second, robots should pay special attention to their motions when there are humans in the environment, or when there is a probability of encountering a human. For example, while it is acceptable for a robot to get very close to a wall, doing so to a human is socially unacceptable and unsafe. Similarly the sudden appearance of a robot can surprise humans and cause discomfort. There are many other social scenarios where the shortest path may not be optimal. Therefore, context-aware path planning algorithms should treat humans and obstacles differently to enable intelligent robot behaviors. Our human-aware path planning method is described in Section 6.1.2.

Third, semantic maps could be exploited to enhance the navigation behaviors. Robot navigation behaviors could be tailored to the task at hand. For example, when the robot is following a person during the map labeling task the robot chooses its sub-goals to facilitate interaction. Similarly, semantic maps could be useful to negotiate

passing in bottlenecks, such as door passages. We present our context-aware person following approach in Section 6.2 built on top of the person following method presented in Section 4.1.

6.1 Navigating to labeled landmarks

6.1.1 Finding the goal point

In our current application, the goals are either labeled landmarks or objects using a phone app. When the user enters a landmark as a navigation goal, the robot first finds a goal point in the metric map, then plans and executes a socially acceptable path toward this goal.

When a planar landmark is entered as the goal, the metric goal is chosen towards the closest edge of the plane. We first select the closest vertex on the landmark's convex hull to the robot's current position, and projects it to the floor plane. We find the line between the closest vertex on the convex hull and the robot's current pose. The goal position is selected to be on this line, a meter away from the vertex. With this design, the robot would be in close proximity of the desired planar surface and be oriented towards it. This method is suitable for horizontal planes, such as tables, and vertical planes alike, such as doors. An example for finding a goal pose for a uniquely labeled planar landmark is shown in Figure 11(a).

When there are multiple planes associated with the same label we interpret this landmark as a region or space, such as a room or corridor. In this case, we project the points of all planes with this label to the ground plane and compute the convex hull. The goal position is chosen as the centroid of the convex hull and the goal orientation is unspecified, meaning the robot would not change its orientation upon reaching the goal position. An example goal position where the goal landmark label "hallway" represents two walls enclosing a hallway is shown in Figure 11(b).

6.1.2 Human-aware path planning

After a goal position is calculated, a socially acceptable path is planned starting from the current position of the robot. Most approaches divide the robot path planning problem into two: global and local planning. Our approach adopts this template and further divides global planning into two parts: static and dynamic planners. The static planner finds a path on the map of the environment by considering the safety and disturbance of humans, as

(a)

(b)

Figure 11: a) Top down point cloud view of a room. A planar landmark with the label *Table* has previously been annotated by a user. The convex hull for the planar landmark is shown in red lines. When asked to navigate to *Table*, the robot calculates a goal position, which is shown as the yellow point; b) Top down point cloud view of a hallway. The user has previously annotated two planar landmarks with the same label: *Hallway*. When asked to navigate to *Hallway*, the robot chooses a goal position in the middle of the planar landmarks which is shown as the yellow point.

well as the path length but does not consider the future movements of humans. The dynamic planner simulates the future motions of humans by using a social motion model and refines the static path. The dynamic planner takes the static path and refines it by considering the predicted temporary goals of humans and their reaction to the robot's future movements. The predicted goals are used to forward-simulate the human trajectories and generate 'social forces' for the social motion model. A part of the robot's path is then recomputed in compliance with the model.

Figure 12: System overview. When a map and goal is provided to the static planner obstacles and humans are detected and then a path is planned using A* search. The dynamic planner refines the static plan by simulating human reactions to the robot motion. The local planner receives the result, and computes the linear and angular velocities necessary to follow the path. The controller applies these velocities to the robot, which in turn acts in the world. The sensors generate new data, and the loop restarts.

Our local planner is a trajectory planner that computes the linear and angular velocities that would allow the robot to follow the dynamic path. The navigation system overview is shown in Figure 12. The obstacles are differentiated from humans in two modules:

1. Static planner: Approaching humans add a safety cost and traversing between humans add a disturbance cost.
2. Dynamic planner: Future motion of the humans are simulated, which is then used to refine the static path.

More information about this work can be found in [83].

The static planner takes the start and goal positions and a 2D grid map as input and aims to find a set of waypoints that connects the start and goal cells. The output path has the minimum cost using a linearly weighted cost function with three components: path length, safety and disturbance. We use A* search with Euclidean heuristics on a 8-connected grid map to find the minimum cost path. The path length cost is the total length of a path. The safety cost aims to model personal spaces of people. A gaussian cost function is attached to each person in the environment. The safety cost of a cell is the maximum safety cost value among all humans in the environment. The disturbance cost aims to represent the cases where the robot potentially disturbs the interaction of a group of humans. For example, if two people are facing each other and talking, then the robot should not cross between them. The disturbance cost is a non-zero cost if the robot's path crosses between two people who are in reasonable proximity to each other. We do not detect if there actually is conversation between the people but estimate the disturbance cost using body poses of agents. This cost increases if the body orien-

tations of two people are facing each other and is inversely proportional to the distance between a pair of humans. Detecting human formations is a challenging task that has been addressed in the literature [84], however we compute this cost for each pair of humans in the scene without explicitly detecting the formation. This works because the disturbance cost becomes zero after a cut-off distance threshold.

The dynamic path refinement processes the static plan by simulating parts of the path where group of humans are closeby. We use the Social Forces Model (SFM) [85] to simulate the motions of humans and the robot. Interactions between people are modeled as attractive and repulsive forces in SFM, similar to potential fields. The forces are recomputed iteratively and the resulting simulated path sections replaces the corresponding path sections in the static plan. We use DWA as the local trajectory planner. The refinement step allows considering future motions of humans due to the robot's future motions.

We demonstrate our approach with an example in simulation (Figure 13). The goal of the robot is to navigate to a goal position in an office environment where there are four people present in the environment. In this scenario, we show how the path changes significantly when only poses of humans are varied. There are three main ways the robot can navigate to its goal: left, center or right corridor.

In the first configuration in Figure 13(a), two people are grouped together as they are looking at each other and likely conversing. The robot decides to take the center corridor. First, it slightly disturbs the speaking duo, then switches sides in the corridor, and reaches its goal. In the figure, the dynamic path (pink line) is overlaid on the static path (green line).

In the second configuration in Figure 13(b), the third person at the center corridor joins the conversation. Now we have two group regions (rectangles) in the scene. Since passing through a group of three people would introduce a high disturbance cost in addition to the safety cost, the robot decides to take a longer route (left corridor). Since this path does not intersect any group regions the dynamic simulation was not conducted.

In the third configuration in Figure 13(c), the group of three hasn't moved, but the fourth person has changed its position. In this case, if the left corridor is taken again, an additional safety cost would be incurred. Therefore the robot decides to take the longest route (right corridor). Again, since the robot travels far from humans the dynamic simulation was not conducted.

(a)

(b)

(c)

Figure 13: Path planner's output differ given the poses and grouping of humans. a) The robot takes shortest route, traveling in the vicinity of a group of two and another individual; b) third individual joins the group. Robot takes a longer path that doesn't have humans on path; c) fourth person changes his position, leading the robot to take the longest route.

6.2 Context aware person following

As briefly reviewed in Section 2.5, most person following methods in the literature have the same underlying principle: a target position is calculated given the human's position each iteration and a control method finds actions iteratively to navigate towards that position. This results in reactive robot behaviors where the robot follows the human blindly irrespective of the task and context. Our person following method presented in Section 4.1 also falls under this category.

Although reactive methods are sufficient for some scenarios, it can easily lead to deadlock scenarios. For example, consider the case that the followed person goes through a door and stops just outside the doorway. In this case, the robot would occupy the doorway, blocking other people's passage, however does not know it caused an undesirable social situation. If the robot knows what the user intends to do, it can anticipate those actions and suitably adjust its behavior. Person following can be used in different contexts, such as for carrying luggage in airports or groceries in a supermarket. We showed in previous sections that semantic information could be used to communicate goals between the robot and the user. The stored semantic information can also be used to facilitate robot navigation.

We model the task scenario during person following as a state machine, where transitions are triggered via events. A general scenario during person following is implemented as a sequence of four phases:

1. **Signal:** The robot detects an event using perceptual cues.
2. **Approach:** The robot moves to a position better suited to the task.
3. **Execution:** The robot and/or the human executes the task.
4. **Release:** The robot detects the end of event and continues with the basic following behavior.

We focus on two specific scenarios of context-aware person following using the 4-phase model: following for labeling in Section 6.2.1 and passing doors in Section 6.2.2. For both of the scenarios, we demonstrate the capability with a user who is knowledgeable of the robot's capabilities.

6.2.1 Following for interactive labeling

We first examine the person following scenario for interactive labeling of semantic landmarks as described in Section 4. For this scenario the robot follows the user as

he/she moves between the landmarks or objects of interests. Sometimes when a user wants to label an object, undesirable social situations can occur because the robot does not have the task context. In this example, the context is defined as the understanding of being a part of a collaborative task: interactive labeling. When a user stops in front of a landmark or object to label it, if the robot stays behind it can not perceive the pointing gesture and the landmark at the same time. This situation is illustrated in Figure 14.

Figure 14: a) A common problem encountered during person following for interactive labeling. The user wants to label an object on the table, however, the robot does not know the user's intention and stays behind at a fixed distance to the user; b) Our solution is for the robot to navigate to a location that gives the robot a better chance to observe the user and the object simultaneously.

The robot can behave more intelligently if the robot can predict ahead of time when the user is going to label a landmark. When the robot detects that the user intends to label a landmark or object our approach is to position the robot base so it has a better chance to perceive both the pointing gesture and the object/landmark of interest.

We follow the 4-phase behavior design for person following for interactive labeling. The Signaling phase is triggered whenever the user is close to a unlabeled landmark in the semantic map. The user must have close to zero speed to enable signaling for this behavior, because the user may walk past the landmark. After the robot detects the signal we sample positions around the group to locate a “suitable” goal pose for the robot. A pose that is collision free but that gives the robot highest chance of interaction

is favored. A suitable goal position should be at an equal distance to the landmark and the user, and the goal orientation should be selected so the robot faces in between the landmark and the user. Moreover, the goal point should not be very close to an obstacle. We linearly sample points around the “group” formed by the user and the landmark's centroid. The points are sampled from the p -space of this group, which is a circle that includes the landmark and user center positions. This is influenced by Kendon *et al.* [86] on how people form groups in interactive settings. Every sampled position p has a score of

$$\text{Score}(p) = 1.0 - \text{Cost}_{\text{vis}}(p) - \text{Cost}_{\text{obs}}(p),$$

where we define the costs as:

$$\begin{aligned} \text{Cost}_{\text{vis}}(p) &= (\text{dist}(p, LM) - \text{dist}(p, usr)) / \text{dist}(usr, LM), \\ \text{Cost}_{\text{obs}}(p) &= \max(\text{local_cost}(p), \text{global_cost}(p)), \end{aligned} \quad (12)$$

where $\text{dist}()$ is a function that returns the Euclidean distance between two 2D points, $\text{local_cost}(p)$ and $\text{global_cost}(p)$ are the cost values calculated at point p from the local and global costmap, respectively.

The local and global costs are fetched from the normalized local costmap which is formed by the laser scanner readings. The sample with the highest non-negative score is chosen as the goal position. The orientation of the robot is chosen as looking toward the center of all the people in the group.

When the robot completes its move to the goal position, the user labels the landmark or objects via pointing gestures. After the task is completed, the robot waits until the user leaves the vicinity of the landmark. When that happens, the robot continues following the user. If, during any of the phases, the person tracking fails, it informs the user so following can be restarted. The phases and conditions for this behavior are summarized in Table 3. Images from a demonstration for this behavior is shown in Figure 15.

6.2.2 Door passing

The second behavior we inspect during person following is door passing. In our experience, the reactive person following behavior can cause problems while passing doors. For example, if the user intends to close an open door or open a closed door, the robot might end up blocking the movement of the door. Moreover, a deadlock situation occurs when the user wants to go through a door with spring-loaded hinges. In that case, the user would need to hold to

Figure 15: Demonstration of context-awareness for interactive labeling. The robot is following the user throughout the environment and keeping a fixed distance of $1.2m$ to the user. a) Signal phase: The user has stopped and is in the close proximity to the convex hull of the table; b) Approach phase: The robot calculates and navigates to a goal position, so it can perceive the pointing gesture and target. Execution phase: The user points out to the object on the table; c) Release phase: the user moves away from the table d) Basic following behavior continues.

Table 3: Conditions to trigger phases when the user is involved with the Landmark Labeling Event during following.

Signal	$\text{dist}(\text{user}, \text{convex hull}(\text{landmark})) < \text{threshold}$ $\text{speed}(\text{user}) \sim 0$ person roughly facing landmark
Approach	Optimal goal: Close to both the landmark and person, facing in between
Execution	User points and labels landmark
Release	$\text{dist}(\text{user}, \text{convex hull}(\text{landmark})) > \text{threshold}$

door to keep it open, and because the distance between the robot and the user is less than the following threshold, the robot would stay still and won't pass the door.

The robot can assume that the user might be intending to open, close or pass through a door when the user is approaching the door. In our approach, the robot continu-

ously monitors the user's proximity to the doors using the semantic map if the door signs were detected and added to the semantic map beforehand as explained in Section 3.2. The distance check between the user and each door sign is executed each iteration by projecting the centroid of the door sign feature to the ground floor.

Figure 16: Demonstration of context-awareness for door passing during person following. This is a swing door with spring loaded hinges, so it would close if not kept open actively. a) The robot is following the user by keeping a fixed distance to the user; b) Signal Phase: The user has stopped, is in close proximity to the door and performed a pointing gesture toward the other room; c) Approach Phase: The robot passes the door while the user is holding the door; d) Release Phase: The user has more than a threshold distance to the door, and the robot continues with the basic following.

Table 4: Conditions to trigger phases when the user is passing through a door during following.

Signal	$\text{dist}(\text{user}, \text{door}) < \text{threshold}$ $\text{speed}(\text{user}) \sim 0$ User performs pointing gesture towards the passage
Approach	Optimal Goal: A position on the other side of the door that doesn't block the doorway
Execution	Robot and user meet at the same side of the door
Release	$\text{dist}(\text{user}, \text{door}) > \text{threshold}$

The phases and conditions for door passing situation are summarized in Table 4. The robot takes action when the user is nearby a door and performs a pointing gesture towards it to signal that the robot should pass from the door (Signal Phase). If the action is not signaled, the robot continues with basic following during the door passage. After the detection of a pointing gesture, a goal position is calculated (Approach Phase). The goal positions are sampled on the other side of the door indicated by the pointing gesture ray. A collision-free position with the least obstacle cost sample is chosen as the goal point. Note that while the robot is moving, it does not aim to keep fixed distance to

the user anymore. After the robot reaches the goal, it waits for the person to pass the door (Execution Phase). After the user moved away from the door, the standard following behavior takes over. Images from the demonstration of context-aware person following for door passing is shown in Figure 16.

There are several limitations in this work that are worth mentioning. First, we don't detect whether the door signs are to the left or right of the door and rely on the metric map to when we are sampling points for the robot. Second, when sampling goal points around the door, we make assumptions about the maximum size of the door

and use a simple distance heuristic for sampling. Third, we rely on the correctness of the pointing gesture direction to figure out when sampling points from the other side of the door. Finally, we don't currently handle cases when multiple nearby doors are involved, but a data association step would be needed for those cases.

Currently, our approach to door passing only applies to doors with door signs, however, it would be possible to detect doors as planar features and label them as a door category instead of an object or planar surface.

7 Conclusion

In this paper we discussed the process of building semantic maps, how to interactively label entities in them, and use them to enable new navigation behaviors for specific scenarios. We utilize planar surfaces, such as walls and tables, and static objects, such as door signs as features to our semantic SLAM approach. Users can interactively annotate these features by having the robot follow him/her, entering the label through a mobile app and performing a pointing gesture toward the landmark of interest. These landmarks can later be used to generate context-aware motions.

Our pointing gesture approach can reliably estimate the target object using human joint positions and detect ambiguous gestures with probabilistic modeling. Our person following algorithm maximizes future utility by searching future actions, assuming constant velocity model for the human. We showed that our person following method can keep a near-constant distance to the human. We described a simple method to extract metric goals from a semantic map landmark and presented a human-aware path planner that considers the personal spaces of people to generate socially-aware paths. Finally, we demonstrated context-awareness for person following in two scenarios: interactive labeling and door passing. For interactive labeling, the robot utilizes the task knowledge and moves to a favorable position to facilitate interaction if an unlabeled landmark is detected near the person. For door passing, the robot utilizes the existence of a door passage by querying detected door signs in the semantic map and to execute a door passage behavior.

Semantic maps would facilitate communication of goals from an HRI perspective and enable navigation behaviors that are not feasible with metric maps. We showed proof of concept for enabling context-aware navigation behaviors using semantics and believe that there is much to explore in this research area. We think as the sensing tech-

nology improves and maps with richer semantic information become common, it would make intelligent navigation algorithms possible.

One limitation of our work is that there is only implicit interaction between the robot and human in our interaction design. The robot signaled its intention only through motion and did not explicitly communicate with people. As future work, dialogue and gaze could be utilized to complement the motions of the robot.

We think implementation and qualitative validation of robot behavior is a critical first step for path planning algorithms among humans. In this paper, we showed that our approach produced sound solutions in a number of example scenarios. As future work, we think effectiveness of context-aware navigation could be evaluated with usability studies with users who are not familiar with the robot. Moreover, scenarios could be performed under different conditions to test the generality and validate the robustness of the system.

References

- [1] F. Dellaert, M. Kaess, Square root SAM: Simultaneous localization and mapping via square root information smoothing, *The International Journal of Robotics Research*, 2006, 25(12), 1181-1203
- [2] R. C. Smith, P. Cheeseman, On the representation and estimation of spatial uncertainty, *The International Journal of Robotics Research* (IJRR), 1986, 5(4), 56-68
- [3] J. J. Leonard, H. F. Durrant-Whyte, Simultaneous map building and localization for an autonomous mobile robot, In: *Workshop on Intelligence for Mechanical Systems at IEEE/RSJ Intelligent Robots and Systems (IROS)*, 1991, 1442-1447
- [4] J. Folkesson, P. Jensfelt, H. I. Christensen, The M-Space feature representation for SLAM, *IEEE Transactions on Robotics*, 2007, 23(5), 1024-1035
- [5] B. Williams, M. Cummins, J. Neira, P. Newman, I. Reid, J. Tardós, A comparison of loop closing techniques in monocular SLAM, *Robotics and Autonomous Systems*, 2009, 57(12), 1188-1197
- [6] P. Henry, M. Krainin, E. Herbst, X. Ren, D. Fox, RGB-D mapping: Using depth cameras for dense 3d modeling of indoor environments, In: *12th International Symposium on Experimental Robotics (ISER)*, 2010
- [7] A. Elfes, Using occupancy grids for mobile robot perception and navigation, *Computer*, 1989, 22(6), 46-57
- [8] G. Grisetti, C. Stachniss, W. Burgard, Improved techniques for grid mapping with Rao-Blackwellized particle filters, *IEEE Transactions on Robotics*, 2007, 23(1), 34-46
- [9] E. Remolina, B. Kuipers, Towards a general theory of topological maps, *Artificial Intelligence*, 2004, 152(1), 47-104
- [10] J. Boal, A. Sánchez-Mirallas, A. Arranz, Topological simultaneous localization and mapping: a survey, *Robotica*, 2014, 32(5), 803-821

- [11] B. Kuipers, The spatial semantic hierarchy, *Artificial intelligence*, 2000, 119(1), 191-233
- [12] O. M. Mozos, C. Stachniss, W. Burgard, Supervised learning of places from range data using AdaBoost, In: *IEEE International Conference on Robotics and Automation (ICRA)*, 2005, 1730-1735
- [13] S. Ekvall, D. Kragic, P. Jensfelt, Object detection and mapping for service robot tasks, *Robotica*, 2007, 25(02), 175-187
- [14] A. Nüchter, J. Hertzberg, Towards semantic maps for mobile robots, *Robotics and Autonomous Systems*, 2008, 56(11), 915-926
- [15] A. Pronobis, P. Jensfelt, Large-scale semantic mapping and reasoning with heterogeneous modalities, In: *2012 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2012, 3515-3522
- [16] S. Choudhary, L. Carlone, C. Nieto, J. Rogers, Z. Liu, H. I. Christensen, F. Dellaert, Multi robot object-based SLAM, In: *International Symposium on Experimental Robotics*, 2016, 729-741
- [17] N. Sünderhauf, T. T. Pham, Y. Latif, M. Milford, I. Reid, Meaningful maps with object-oriented semantic mapping, In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017, 5079-5085
- [18] J. Fasola, M. J. Mataric, Using semantic fields to model dynamic spatial relations in a robot architecture for natural language instruction of service robots, In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2013, 143-150
- [19] S. Tellex, T. Kollar, S. Dickerson, M. R. Walter, A. G. Banerjee, S. J. Teller, N. Roy, Understanding natural language commands for robotic navigation and mobile manipulation, In: *25th AAAI Conference on Artificial Intelligence*, 2011
- [20] I. Kostavelis, A. Gasteratos, Semantic mapping for mobile robotics tasks: A survey, *Robotics and Autonomous Systems*, 2015, 66, 86-103
- [21] C. Landsiedel, V. Rieser, M. Walter, D. Wollherr, A review of spatial reasoning and interaction for real-world robotics, *Advanced Robotics*, 2017, 31(5), 222-242
- [22] E. A. Topp, H. I. Christensen, Topological modelling for human augmented mapping, In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2006, 2257-2263
- [23] H. H. Clark, S. E. Brennan, Grounding in communication, *Perspectives on socially shared cognition*, 1991, 13(1991), 127-149
- [24] G.-J. M. Kruijff, H. Zender, P. Jensfelt, H. I. Christensen, Clarification dialogues in human-augmented mapping, In: *Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction*, 2006, 282-289
- [25] H. Zender, P. Jensfelt, Ó. M. Mozos, G.-J. M. Kruijff, W. Burgard, An integrated robotic system for spatial understanding and situated interaction in indoor environments, In: *AAAI*, 2007, 7, 1584-1589
- [26] G. Gemignani, D. Nardi, D. D. Bloisi, R. Capobianco, L. Iocchi, Interactive semantic mapping: experimental evaluation, In: *Experimental Robotics*, 2016, 339-355
- [27] T. B. Moeslund, A. Hilton, V. Krüger, A survey of advances in vision-based human motion capture and analysis, *Computer Vision and Image Understanding*, 2006, 104(2-3), 90-126
- [28] A. F. Bobick J. W. Davis, The recognition of human movement using temporal templates, *IEEE Transactions on pattern analysis and machine intelligence*, 2001, 23(3), 257-267
- [29] N. Dalal, B. Triggs, Histograms of oriented gradients for human detection, In: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2005, 1, 886-893
- [30] K. Mikolajczyk, C. Schmid, A. Zisserman, Human detection based on a probabilistic assembly of robust part detectors, In: *European Conference on Computer Vision (ECCV)*, Springer, 2004, 69-82
- [31] G. Shu, A. Dehghan, O. Oreifej, E. Hand, M. Shah, Part-based multiple-person tracking with partial occlusion handling, In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012, 1815-1821
- [32] T. Darrell, G. Gordon, M. Harville, J. Woodfill, Integrated person tracking using stereo, color, and pattern detection, *International Journal of Computer Vision*, 2000, 37(2), 175-185
- [33] J. Shotton, A. Fitzgibbon, M. Cook, T. Sharp, M. Finocchio, R. Moore, et al., Real-time human pose recognition in parts from single depth images, In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011, 1297-1304
- [34] Y. LeCun, Y. Bengio, et al., Convolutional networks for images, speech, and time series, *The handbook of Brain Theory and Neural Networks*, 1995, 3361(10)
- [35] A. Krizhevsky, I. Sutskever, G. E. Hinton, ImageNet classification with deep convolutional neural networks, In: *Advances In Neural Information Processing Systems*, 2012, 1097-1105
- [36] Y. Tian, P. Luo, X. Wang, X. Tang, Deep learning strong parts for pedestrian detection, In: *Proceedings of the IEEE International Conference on Computer Vision*, 2015, 1904-1912
- [37] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, S. Savarese, Social LSTM: Human trajectory prediction in crowded spaces, In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, 961-971
- [38] K. O. Arras, Ó. M. Mozos, W. Burgard, Using boosted features for the detection of people in 2d range data, In: *IEEE International Conference on Robotics and Automation (ICRA)*, 2007, 3402-3407
- [39] D. Schulz, W. Burgard, D. Fox, A. B. Cremers, Tracking multiple moving targets with a mobile robot using particle filters and statistical data association, In: *IEEE International Conference on Robotics and Automation (ICRA)*, 2001, 2, 1665-1670
- [40] E. A. Topp, H. I. Christensen, Tracking for following and passing persons, In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2005, 2321-2327
- [41] N. Bellotto, H. Hu, Multisensor-based human detection and tracking for mobile service robots, *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 2009, 39(1), 167-181
- [42] F. Zanlungo, T. Ikeda, T. Kanda, Social force model with explicit collision prediction, *EPL (Europhysics Letters)*, 2011, 93(6), 68005
- [43] A. Leigh, J. Pineau, N. Olmedo, H. Zhang, Person tracking and following with 2D laser scanners, In: *IEEE International Conference on Robotics and Automation (ICRA)*, 2015, 726-733
- [44] C. Dondrup, N. Bellotto, F. Jovan, M. Hanheide, et al., Real-time multisensor people tracking for human-robot spatial interaction, *Workshop on Machine Learning for Social Robotics at IEEE International Conference on Robotics and Automation (ICRA)*, 2015
- [45] J. Schmidt, N. Hofemann, A. Haasch, J. Fritsch, G. Sagerer, Interacting with a mobile robot: Evaluating gestural object references, In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2008, 3804-3809

- [46] M. Van den Bergh, D. Carton, R. De Nijs, N. Mitsou, C. Landsiedel, K. Kuehnlenz, et al., Real-time 3D hand gesture interaction with a robot for understanding directions from humans, In: IEEE International Workshop on Robot and Human Interactive Communication (ROMAN), 2011, 357-362
- [47] A. G. Brooks, C. Breazeal, Working with robots and objects: Revisiting deictic reference for achieving spatial common ground, In: Proceedings of the 1st ACM SIGCHI/SIGART Conference on Human-Robot Interaction, 2006, 297-304
- [48] E. A. Sisbot, L. F. Marin-Urias, R. Alami, T. Simeon, A human aware mobile robot motion planner, IEEE Transactions on Robotics, 2007, 23(5), 874-883
- [49] R. Kirby, R. Simmons, J. Forlizzi, Companion: A constraint-optimizing method for person-acceptable navigation, In: The 18th IEEE International Symposium on Robot and Human Interactive Communication, (RO-MAN), 2009, 607-612
- [50] M. L. Walters, K. Dautenhahn, R. Te Boekhorst, K. L. Koay, C. Kaouri, S. Woods, C. Nehaniv, D. Lee, I. Werry, The influence of subjects' personality traits on personal spatial zones in a human-robot interaction experiment, In: IEEE International Workshop on Robot and Human Interactive Communication (RO-MAN), 2005, 347-352
- [51] M. Luber, L. Spinello, J. Silva, K. O. Arras, Socially-aware robot navigation: A learning approach, In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2012, 902-907
- [52] H. Kidokoro, T. Kanda, D. Bršćic, M. Shiomi, Simulation-based behavior planning to prevent congestion of pedestrians around a robot, IEEE Transactions on Robotics, 2015, 31(6), 1419-1431
- [53] A. Bordallo, F. Previtali, N. Nardelli, S. Ramamoorthy, Counterfactual reasoning about intent for interactive navigation in dynamic environments, In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2015, 2943-2950
- [54] U. Köeckemann, F. Pecora, L. Karlsson, Inferring context and goals for online human-aware planning, In: IEEE 27th International Conference on Tools with Artificial Intelligence (ICTAI), IEEE, 2015, 550-557
- [55] T. Kruse, P. Basili, S. Glasauer, A. Kirsch, Legible robot navigation in the proximity of moving humans, In: IEEE Workshop on Advanced Robotics and its Social Impacts (ARSO), 2012, 83-88
- [56] A. D. Dragan, K. C. Lee, S. S. Srinivasa, Legibility and predictability of robot motion, In: 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI), 2013, 301-308
- [57] P. Trautman, J. Ma, R. M. Murray, A. Krause, Robot navigation in dense human crowds: Statistical models and experimental studies of human-robot cooperation, The International Journal of Robotics Research, 2015, 34(3), 335-356
- [58] P. Henry, C. Vollmer, B. Ferris, D. Fox, Learning to navigate through crowded environments, In: IEEE International Conference on Robotics and Automation (ICRA), 2010, 981-986
- [59] R. Philippsen, R. Siegwart, Smooth and efficient obstacle avoidance for a tour guide robot, In: IEEE International Conference on Robotics and Automation (ICRA), 2003
- [60] L. Y. Morales Saiki, S. Satake, R. Huq, D. Glas, T. Kanda, N. Hagita, How do people walk side-by-side?: using a computational model of human behavior for a social robot, In: ACM/IEEE International Conference on Human-Robot Interaction, 2012, 301-308
- [61] A. Ohya, T. Munekata, Intelligent escort robot moving together with human-interaction in accompanying behavior, In: Proceedings FIRA Robot World Congress, 2002, 31-35
- [62] R. Murakami, L. Y. Morales Saiki, S. Satake, T. Kanda, H. Ishiguro, Destination unknown: walking side-by-side without knowing the goal, In: Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction, 2014, 471-478
- [63] J. J. Park, B. Kuipers, Autonomous person pacing and following with model predictive equilibrium point control, In: IEEE International Conference on Robotics and Automation (ICRA), 2013, 1060-1067
- [64] C. Granata, P. Bidaud, A framework for the design of person following behaviors for social mobile robots, In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2012, 4652-4659
- [65] R. Gockley, J. Forlizzi, R. Simmons, Natural person following behavior for social robots, In: 2nd ACM/IEEE International Conference on Human-Robot Interaction (HRI), 2007, 17-24
- [66] K. Charalampous, I. Kostavelis, A. Gasteratos, Recent trends in social aware robot navigation: A survey, Robotics and Autonomous Systems, 2017, 93, 85-104
- [67] T. Kruse, A. K. Pandey, R. Alami, A. Kirsch, Human-aware robot navigation: A survey, Robotics and Autonomous Systems, 2013, 61(12), 1726-1743
- [68] P. Regier, S. Oßwald, P. Karkowski, M. Bennewitz, Foresighted navigation through cluttered environments, In: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2016, 1437-1442
- [69] E. Pacchierotti, H. I. Christensen, P. Jensfelt, Human-robot embodied interaction in hallway settings: a pilot user study, In: IEEE International Workshop on Robot and Human Interactive Communication (ROMAN), 2005, 164-171
- [70] N. Wilde, D. Kulic, S. L. Smith, Learning user preferences in robot motion planning through interaction, In: IEEE International Conference on Robotics and Automation (ICRA), 2018
- [71] C. Galindo, A. Saffiotti, Inferring robot goals from violations of semantic knowledge, Robotics and Autonomous Systems, 2013, 61(10), 1131-1143
- [72] D. V. Lu, W. D. Smart, Towards more efficient navigation for robots and humans, In: IEEE/RSJ International Conference On Intelligent Robots and Systems (IROS), 2013, 1707-1713
- [73] M. M. Loper, N. P. Koenig, S. H. Chernova, C. V. Jones, O. C. Jenkins, Mobile human-robot teaming with environmental tolerance, In: Proceedings of the 4th ACM/IEEE International Conference on Human-Robot Interaction, ACM, 2009, 157-164
- [74] P. Anderson, Q. Wu, D. Teney, J. Bruce, M. Johnson, N. Sünderhauf, et al., Vision-and-language navigation: Interpreting visually-grounded navigation instructions in real environments, In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2018
- [75] H. Zender, P. Jensfelt, G.-J. M. Kruijff, Human-and situation-aware people following, In: The 16th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN), 2007, 1131-1136
- [76] J. Neira, J. D. Tardós, Data association in stochastic mapping using the joint compatibility test, IEEE Transactions on Robotics and Automation, 2001, 17(6), 890-897
- [77] A. J. Trevor, A. Cosgun, J. Kumar, H. I. Christensen, Interactive map labeling for service robots, In: IROS Workshop on Active Semantic Perception, 2012
- [78] J. Xavier, M. Pacheco, D. Castro, A. Ruano, U. Nunes., Fast line, arc/circle and leg detection from laser scan data in a player

- driver, In: IEEE International Conference on Robotics and Automation (ICRA), 2005, 3930-3935
- [79] D. Fox, W. Burgard, S. Thrun, The dynamic window approach to collision avoidance, *IEEE Robotics & Automation Magazine*, 1997, 4(1), 23-33
- [80] A. Cosgun, D. A. Florencio, H. I. Christensen, Autonomous person following for telepresence robots, In: IEEE International Conference on Robotics and Automation (ICRA), 2013, 4335-4342
- [81] A. Cosgun, A. J. Trevor, H. I. Christensen, Did you mean this object?: Detecting ambiguity in pointing gesture targets, In: Towards a Framework for Joint Action Workshop at 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI 2015), IEEE Press, 2015
- [82] A. J. Trevor, J. G. Rogers III, A. Cosgun, H. I. Christensen, Interactive object modeling & labeling for service robots, In: Proceedings of the 8th ACM/IEEE International the Conference on Human-Robot Interaction (HRI), 2013, 421-422
- [83] A. Cosgun, E. A. Sisbot, H. I. Christensen, Anticipatory robot path planning in human environments, In: 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN), 2016, 562-569
- [84] M. Cristani, L. Bazzani, G. Paggetti, A. Fossati, D. Tosato, A. Del Bue, et al., Social interaction discovery by statistical analysis of F-formations, In: The 22nd British Machine Vision Conference (BMVC 2011), 2011, 2
- [85] D. Helbing, P. Molnar, Social force model for pedestrian dynamics, *Physical review E*, 1995, 51(5), 4282
- [86] A. Kendon, Conducting interaction: Patterns of behavior in focused encounters, In: *Studies in interactional sociolinguistics*, vol. 7, Cambridge University Press, 1990