

Qualitative Design and Implementation of Human-Robot Spatial Interactions

Nicola Bellotto¹, Marc Hanheide¹, and Nico Van de Weghe²

¹ School of Computer Science, University of Lincoln, LN6 7TS, United Kingdom
{nbellotto,mhanheide}@lincoln.ac.uk

² Department of Geography, Ghent University, 9000, Belgium
nico.vandeweghe@ugent.be

Abstract. Despite the large number of navigation algorithms available for mobile robots, in many social contexts they often exhibit inopportune motion behaviours in proximity of people, often with very “unnatural” movements due to the execution of segmented trajectories or the sudden activation of safety mechanisms (e.g., for obstacle avoidance). We argue that the reason of the problem is not only the difficulty of modelling human behaviours and generating opportune robot control policies, but also the way human-robot spatial interactions are represented and implemented. In this paper we propose a new methodology based on a *qualitative* representation of spatial interactions, which is both flexible and compact, adopting the well-defined and coherent formalization of Qualitative Trajectory Calculus (QTC). We show the potential of a QTC-based approach to abstract and design complex robot behaviours, where the desired robot’s motion is represented together with its actual performance in one coherent approach, focusing on spatial interactions rather than pure navigation problems.

1 Introduction

In the context of this paper, human-robot spatial interaction (HRSI) is defined as a set of relative motion events between two or more (possibly coordinated, cooperative and/or communicative) agents, which are executed according to particular social rules, agents objectives and safety constraints. In this paper we focus particularly on the 2D free-motion case, i.e. the trajectories followed by humans and robots on a planar space without obstacles, generally associated with the actions of walking towards something or someone, standing still, moving away, etc. The interpretation of such motion behaviours, as well as the capability of performing them in a social context, are essential skills for a mobile robot aiming at interacting and providing services to humans.

Typically, the trajectories of HRSIs are treated as geometrical relations in a metric frame, which is often complicated or dependent on specific training datasets. Such an approach has proven only partially effective so far, preventing the implementation of more complex and meaningful spatial behaviours. In contrast to the majority of existing approaches, we propose a *qualitative* approach

to represent and implement HRSIs, which offers the following advantages with respect to the typical quantitative solutions: it is possible to design and implement social rules of spatial behaviours without necessarily having to learn them from huge datasets; human-human and human-robot spatial interactions can be easily mapped into semantic descriptions close to natural language. The key novelty presented in this paper, founded on our previous work on qualitative robot control and human behaviour analysis [1,2], is the combination of coarse and fine QTC representations for a comprehensible ad-hoc definition of HRSIs, which is both flexible and compact. This approach allows for the rapid design of complex spatial behaviours with varying resolution of qualitative description.

Numerous solutions for HRSI have been developed in the past. For example, in [3] the authors developed a geometrical-based algorithm for a mobile robot to enter a queue of people according to their (static) position and orientation; in [4], a solution to learn typical motion patterns of people in an office environment was proposed, which was used to estimate the metrical location of a person tracked across different rooms. Recent works have considered the motion activity of people in relation to their spatial location, so that a social robot can predict the position of potential users and approach them more effectively [5]. None of these approaches, however, take into account the robot's influence in estimating and modifying current human trajectories. Among the solutions that have considered explicitly the effect of a robot's action on human motion behaviours, the "social force model" proposed by [6] is often used to provide a quantitative description of pedestrian behaviours [7]. This model describes human motion according to forces driven by internal objectives, such as the desire of reaching a target or avoiding an obstacle, although some recent work suggests the model is not suitable for dealing with individual pedestrians during evasive manoeuvres [8]. The authors proposed instead to learn a model of human motion, based on the principle of maximum entropy, from pedestrian observations. A probabilistic framework was also proposed in [9] to generate collision-free trajectories with a robot in dynamic human environments. Differently from others, the last two solutions take into account mutual human-robot interactions to estimate and plan joint trajectories. However, all these works are associated with a numerical representation of the agents position, which might be not the most accessible approach for programming social behaviours with robots.

A qualitative interpretation of motion activities seems to be a more tractable way of dealing with HRSI, as shown by relatively simple but effective QTC representations implemented in our previous work [1]. There, QTC was adopted as a formalism for representing and implementing HRSIs. Initial simulation results using QTC Basic (QTC_B) suggested it was possible to abstract simple human and robot trajectories to generate motion commands based on qualitative terms. The work in [2] extended the spatial behaviours representation to analyse more complex trajectories with QTC Double Cross (QTC_C). To this end, however, no solutions have been implemented that exploit both QTC_B and QTC_C representations and generate, from real observations, robot control policies within the same qualitative framework. This is explored by the current research.

2 The Qualitative Trajectory Calculus

When implementing HRSI behaviour in autonomous robots, the designer is usually not interested in the exact trajectory of the robot, but rather on how it *qualitatively* moves in relation to the human, obeying implemented rules and conventions. To accommodate that need, we propose a qualitative framework based on the analysis of *relative* position and movement direction between two interacting agents on a 2D environment. In order to reduce the space domain and focus only on those terms relevant to HRSI, we adopt the well-defined set of symbols and relations provided by QTC, which is an elegant formalism to deal with the relative motion of two points in space [10]. QTC belongs to the broad research area of qualitative spatial representation and reasoning, inheriting some of its properties and tools [11]. There are several versions of QTC, depending on the number of factors considered (e.g. distance, speed, direction, etc.) and on the dimensions, or constraints, of the space where the points move.

QTC_B represents the relative motion of two points *k* and *l* (Fig. 1), with respect to the reference line connecting them, with a 3-tuple of qualitative relations (a b c), where each element can assume any of the values {−, 0, +} as follows:

- a) movement of *k* with respect to *l*
 - − : *k* is moving towards *l*
 - 0 : *k* is stable with respect to *l*
 - + : *k* is moving away from *l*
- b) movement of *l* with respect to *k* as above, swapping *k* and *l*
- c) relative speed of *k* with respect to *l*
 - − : *k* is slower than *l*
 - 0 : *k* has the same speed of *l*
 - + : *k* is faster than *l*

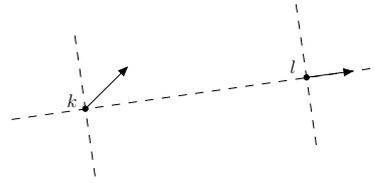


Fig. 1. Example of moving points *k* and *l*. The respective QTC_B and QTC_C relations are (−+) and (−+−0).

Depending on the application, a simplified version of QTC_B without the speed relation can be adopted, considering only the 2-tuple (a b). All the different combinations and relative motion description for two points are illustrated in Fig. 2. In this case there are 9 (3²) possible states, the transitions of which can be represented by a Conceptual Neighbourhood Diagram (CND) [10]. In the QTC framework, a CND restricts the number of legally possible transitions. This helps to reduce the complexity in building temporal sequence of QTC states.

A variant of QTC_C extends the previous calculus to specify which side the two points are moving on, with respect to the reference line connecting them (see Fig. 1). In addition to the previous relations, the following ones are included:

- d) movement of *k* with respect to $\overline{k\ l}$
 - − : *k* is moving to the left side of $\overline{k\ l}$
 - 0 : *k* is moving along $\overline{k\ l}$
 - + : *k* is moving to the right side of $\overline{k\ l}$
- e) movement of *l* with respect to $\overline{l\ k}$ as above, swapping *k* and *l*

QTC_C has 81 (3⁴) states in total and 1088 possible transitions in the relative CND. It can be combined with QTC_B to represent and reason about HRSIs.

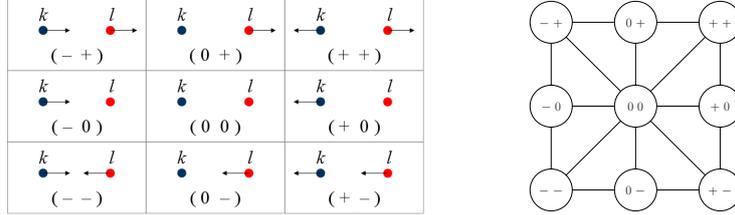


Fig. 2. Graphical representation of QTC_B (without speed) and relative CND. Due to the original formulation [10], there are no direct transitions in the CND between some of the states that, at a first glance, appear to be adjacent (e.g. (-0) and $(0-)$).

3 Application Case Studies

The proposed approach facilitates the design of a variety of HRSIs. Here we will study two cases that stem from real-world scenarios of robots moving and interacting with people in public spaces (i.e. to approach potential users or share narrow passageways) extending and improving our previous models in [1,2]. The considered HRSIs are not intended to be comprehensive, but just a means to more complex robot behaviours. In particular, we consider the scenario in which a person k and a robot l are in front of each other, just a few meters away and without obstacles between them. The robot is programmed to proactively engage with the person in response to two possible actions: Case I) the person approaches the robot, but then he/she stops and moves away from it before being reached; Case II) the person move towards the robot, but then he/she deviates from the initial trajectory to pass on its left-hand or right-hand side. Both the situations, illustrated in Fig. 3, terminate with the robot standing still.

3.1 Approach and Withdraw (Case I)

The first scenario refers to the following temporal sequence, which extends a previous QTC_B -based only example discussed in [1]:

$$(-000) \rightsquigarrow (- - 00) \rightsquigarrow (0-) \rightsquigarrow (+-) \rightsquigarrow (+0) \tag{1}$$

The person initially approaches the robot, triggering the same response on it. A QTC_C representation is necessary to specify that the person is moving straight towards the robot, and not aside of it. The initial state of the sequence is therefore (-000) , which reads “(person) k moves towards (robot) l , while l stands still;



Fig. 3. HRSI in Case I (left) and II (right). The arrows indicate the trajectories of the agents. The crosses correspond to the locations where the robot stops.

neither k nor l moves laterally”. The robot should then approach the person and stop only when this moves away (see Fig. 3). A command must be issued to the robot to facilitate the first transition $(-000) \rightsquigarrow (-00)$, as discussed in Sec. 4. In the remaining part of the sequence, the robot moves and stops regardless of the human (and robot) motion side. A simpler QTC_B representation is therefore sufficient to generate the opportune robot behaviour. In particular, in the following $(0-)$ state, the robot keeps approaching the person as long as he/she does not move away. Should that happen, the robot stops, as indicated by the transition $(+0) \rightsquigarrow (+0)$.

The shift in resolution from a detailed to a coarser QTC representation is particularly interesting. In reasoning language, this can be formally presented as follows [10]: there is a change from a fine relation in QCT_C to a coarse relation in QTC_B ; coarse relations are specific unions of fine relations; in this case, the union of the QTC_C relations $(0-??)$ gives the QTC_B relation $(0-)$. The possibility to switch between resolutions is important in order to deal with the computational complexity arising from the many possible interactions between two or more moving agents. In the QTC_C space and relative CND , this is somehow equivalent to switch between neighbouring subsets. The transitions between QTC_C subsets correspond to the transitions between the associated QTC_B states, and the choice between one or the other representation depends on the level of accuracy required to model the particular spatial interaction.

3.2 Approach and Avoid (Case II)

In the second case, we consider a variant of the scenario discussed in [2], where the robot followed predefined trajectories to let a person pass in a narrow corridor, but without reacting to the human movements. Our qualitative approach can encode the same situation, more flexibly, accommodating the situative behaviour of the human. The interaction is described by the following temporal sequence, in which the agents keep the left-hand or right-hand side, following the respective top or bottom branch after the second state (see also Fig. 3):

$$(-000) \rightsquigarrow (-00) \begin{cases} \nearrow (-00) \rightsquigarrow (-000) \searrow \\ \searrow (-00) \rightsquigarrow (-000) \nearrow \end{cases} (00) \rightsquigarrow (++) \rightsquigarrow (+0) \quad (2)$$

The complexity of the manoeuvre, in this case, is reflected by the increased length of the sequence. It differs from the previous scenarios also by the fact that the robot, besides executing “towards” and “stop” actions, has to perform additional “left” and “right” movements, depending on the convention chosen by the person. These actions are represented in Eq. (2) by the transitions $(-0) \rightsquigarrow (-00)$ and $(-000) \rightsquigarrow (-000)$ on the top and bottom branch of the sequence respectively. In practice, the robot gives way to the person, who has the priority in deciding which side of the corridor to carry on. Two more robot commands are generated next, independently of the motion side: one to move away the robot from the person when they are side-by-side, i.e. $(00) \rightsquigarrow (++)$;

another one to stop the robot when the person walks away, i.e. $(++)\rightsquigarrow(+0)$. For this last part of the sequence, a QTC_B representation is sufficient.

It is important to note that, although in the above cases the human always does “the first move”, it does not imply that the HRSI designer has to specifically program robot commands *in response* to particular human actions. With our qualitative approach, indeed, robot commands “emerge” naturally from the QTC description of the interaction scenario. One could design HRSIs in which the robot is the first agent taking the initiative, for example, to prompt some desired human motion behaviour. The command generation process is transparent from the designer’s point of view.

4 Implementation

As in a previous work [1], the implemented system consists of three intercommunicating modules: a laser-based people tracker; a high-level reasoner for QTC-based inference and robot commands; and a control module that converts high-level commands to low-level instructions for the robot. The last module feeds a simple motion planning algorithm with obstacle avoidance provided by the robot’s middleware. In a typical run of the system, the tracking module provides positions and velocities of the agents (i.e. person and robot) in form of string messages for the reasoner. The information is converted in QTC relations by the reasoner and included as new evidence, upon which an inference process is run. The output is a status message (e.g. a QTC string ‘(-000)’) and, if available, a command for the robot control module (e.g. ‘moveRightOf(robot, Agent)’). The tracking and control modules are based on standard algorithms available in the literature [12,13]. The high-level reasoner, instead, is based on an inference engine, which is discussed next.

4.1 QTC Relations

The QTC relations described in Sec. 2 are implemented with the logic constructs of F-Limette¹, an inference engine based on rules that are written in a Prolog-like language for Fuzzy Metric-Temporal Horn Logic (FMTHL). In particular, we extended and improved the QTC_B implementation in [1] to include also QTC_C , following the specifications suggested in [14]. The “fuzzyness” is used to accommodate some of the ambiguities in the qualitative characterization of the motion properties. It is implemented assigning *degree of validities* (i.e. functions of the difference between current speed/orientation and some nominal value) to QTC relations, which are utilized by F-Limette during the inference process.

4.2 Situation Graph Trees

Knowledge about particular motion behaviours, expressed in terms of QTC sequences, is encoded in F-Limette using a schematic representation called Situation Graph Tree (SGT) and created with a dedicated editor [15]. During the

¹ http://cogvisys.iaks.uni-karlsruhe.de/Vid-Text/f_limette/

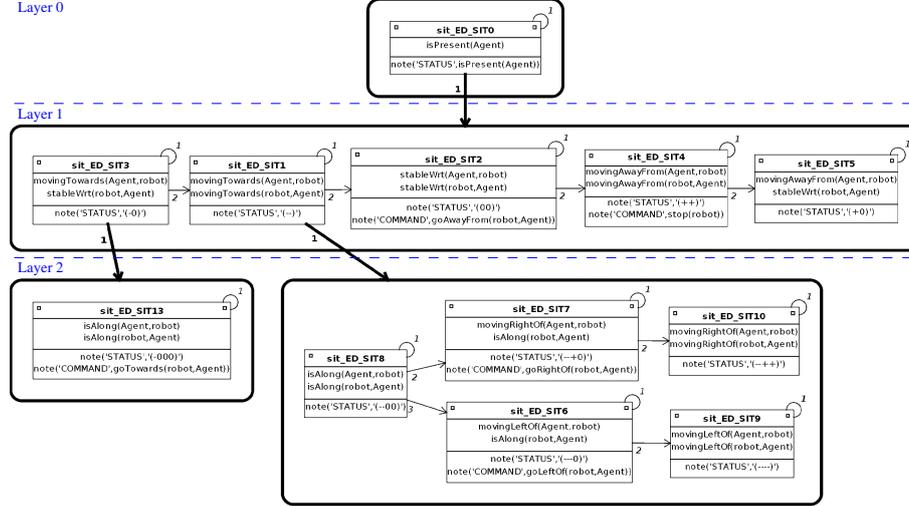


Fig. 4. SGT of Case II. Small boxes are situations with QTC relations and associated actions. Situations are temporally connected to others by thin arrows to form a graph. Situation graphs at the bottom are specializations of situations at the top.

inference process, the SGT is traversed in a depth-first fashion to find instanciable situations, each one corresponding to a particular QTC state (except the root, which is used simply to instantiate the presence of a human `Agent`). When available in a situation, the action `COMMAND` of a successful traversal is sent to the robot controller, enabling the potential instantiation of the next situation. Details about robot commands using F-Limette and SGTs are discussed in [1].

Since the SGTs are conceptually similar, we describe only the one relative to Case II, which is shown in Fig. 4. The nine QTC states of the sequence in Eq. (2) are encoded by the situations in the middle and bottom layer (Layer 1 and 2 respectively) of the SGT. Thin arrows between situations indicate prediction edges, while thick arrows from Layer 1 to Layer 2 point to specializations, respectively, from QTC_B to QTC_C . When a new inference process starts, the SGT traversal tries to instantiate the logic predicates in the first situation of Layer 1 (i.e. `movingTowards(Agent, robot)` and `stableWrt(robot, Agent)`), which correspond to the QTC_B state (-0) . An attempt is also made to satisfy the respective specialization in Layer 2 (i.e. `isAlong(Agent, robot)` and `isAlong(robot, Agent)`), which, combined with the previous, gives the QTC_C state (-000) . If successful, the command `goTowards(robot, Agent)` is issued, enabling the following transition $(-000) \rightsquigarrow (-00)$ at the next traversal.

The traversal proceeds with the next situation in Layer 1. Since to the new QTC_B state $(--)$ may have five QTC_C extensions, according to the sequence in Eq. (2), the respective specialization in Layer 2 is another temporal graph with five situations. Note that two of them, $(--0)$ and $(--0)$, include also the new robot commands `goRightOf(robot, Agent)` and `goLeftOf(robot, Agent)` respectively. The process continues in a similar way to the next situations.

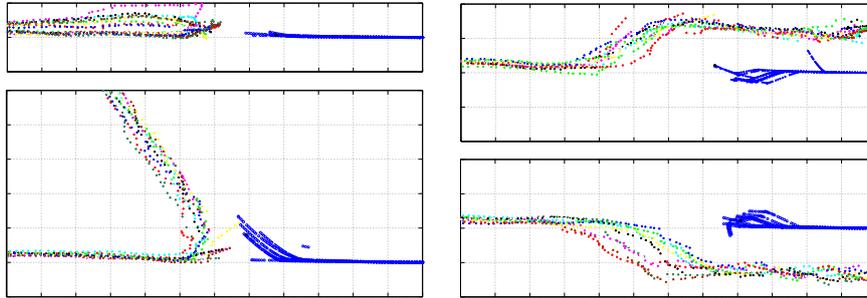


Fig. 5. Trajectories of Case I (left) and II (right). The human started on the left side, the robot on the right side. The grid is 0.5m x 0.5m. [Top left] The person walks towards the robot, turns back and walks away. [Bottom left] The person walks towards the robot, turns left and walks away. [Top right] The person walks towards the robot and passes on its left-hand side. [Bottom right] As before, but on the right-hand side.

5 Experiments

To validate the proposed approach, a set of human trajectories have been collected in an office environment using a laser-based tracking algorithm for mobile robots [12]. The trajectories, consisting of 2D coordinates and velocities, were performed according to the descriptions of Case I and II, and used to feed a real-time robot simulation. The robot was controlled by the system described in Sec. 4. Four trajectory types have been recorded, two for Case I and two for Case II (see Fig. 5), each one consisting of 10 runs (i.e. 20 trajectories per case).

The robot trajectories in Fig. 5 show that the system performed mostly correctly, stopping or driving the robot towards the desired direction according to the planned HRSI. Only in a couple of occasions, during the execution of the avoid behaviour (Case II), the interaction was unsuccessful, with the robot failing to stop after passing the human. The errors were caused by the rotation in place (no translation) of the robot, which unfortunately is not captured by the current versions of QTC. The problem is sensitive to the particular robot motion planner and it should be addressed in future implementations of the system.

To analyse the actual performance of HRSIs, the observed QTC sequences have been modelled using Markov chains (Fig. 6), the transition probabilities of which were derived from the total number of trials per case (i.e. 20). It is interesting to note that, although the behaviours generally followed the expected ones, in both cases a QTC_B state that was not in the original sequence (i.e. $--$) for Case I and (-0) for Case II), enabled the successful completion of the interaction replacing missing transitions between other states, mostly QTC_C . This evidence further supports the advantage of using hybrid QTC representations.

Real-world case studies have also been carried out with our MetraLabs SCITOS G5 platform using a Sick S300 laser scanner restricted to 180° field of view to detect the subject's pair of legs and obstacles. The robot employed DWA local path planning [13] to implement safe QTC movement patterns. A subject

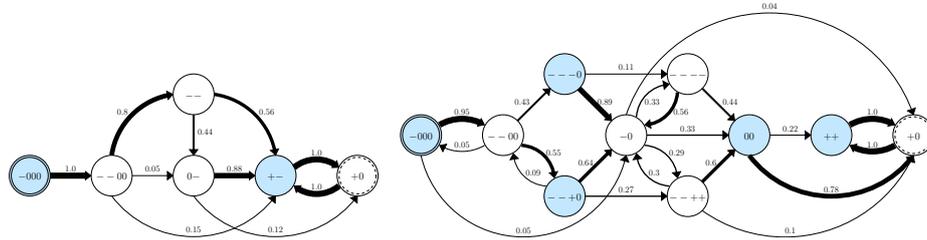


Fig. 6. Markov chains of Case I (left) and II (right). Thicker edges correspond to higher transition probabilities. Initial and final states are marked, respectively, by continuous and dashed inner circles. Blue states indicate the robot execution of motion commands.

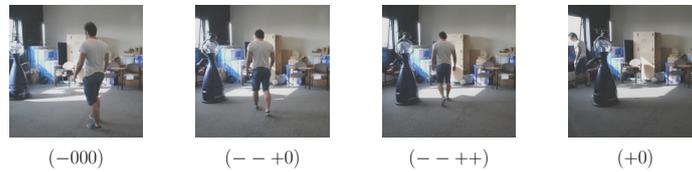


Fig. 7. Case II with real robot. The robot gives space to the subject shifting on its right-hand side. This is illustrated also by the QTC states under each snapshot.

was instructed to either approach and turn around for the withdraw behaviour (Case I) or to pass the robot on its left side for the avoid behaviour (Case II). Five runs of each case were analysed. In all of them, the robot behaviour was according to our model, evidencing that both tracking and reasoning are applicable in a real-world setting. However, the actual sequence of QTC states varied in a similar fashion as it did in the simulation runs. One prototypical sequence of the Case II interaction taken from our study is illustrated in Fig. 7 with photos of the subject augmented by the corresponding QTC states as they have been recognised by the robot. Supplementary material is available on our website: <http://robots.lincoln.ac.uk/research/icr13>

6 Conclusions

In this paper, we proposed a novel approach for the qualitative design of HRSIs, which is based on a hybrid QTC representation of human and robot trajectories. Our solution allows for compact and flexible descriptions of spatial interactions, which are enabled, within the same qualitative framework, by implicit robot control policies. Several experiments with real data validated the proposed solution, but highlighted also some limitation of the current implementation. Although fuzzy sets are useful to represent ambiguities of (top-down) HRSI models, they seem unsuitable, at least in our case, to deal with real-world uncertainties. Indeed, F-Limette strongly relies on correct tracking estimates and reliable robot motion planners. Also, the hierarchical structure of SGTs does not allow for a hybrid representation of QTC transitions always consistent with the CNDs.

Therefore, they cannot fully exploit the properties of the calculus, in particular for the recovery of missing transitions caused by sensor noise or occlusions. Other solutions based on robust probabilistic frameworks, e.g. POMDP [16], should be able to overcome some limitations in future implementations of our work.

Finally, at the present, it is not possible to create very complex models of HRSI: each QTC sequence has to be hand-crafted or learned by datasets for the specific application. This poses severe constraints in real-world scenarios with large numbers of potential HRSIs. Further research should investigate incremental approaches, where initial QTC models are provided as skeletons to be refined and extended autonomously by the robot, based on actual experience.

References

1. Bellotto, N.: Robot control based on qualitative representation of human trajectories. In: AAAI Spring Symposium – Designing Intelligent Robots (2012)
2. Hanheide, M., Peters, A., Bellotto, N.: Analysis of human-robot spatial behaviour applying a qualitative trajectory calculus. In: Proc. of IEEE RoMan, pp. 689–694 (2012)
3. Nakauchi, Y., Simmons, R.: A social robot that stands in line. *Autonomous Robots* 12(3), 313–324 (2002)
4. Bennewitz, M., Burgard, W., Cielniak, G., Thrun, S.: Learning motion patterns of people for compliant robot motion. *IJRR* 24(1), 31–48 (2005)
5. Kanda, T., Glas, D.F., Shiomi, M.F., Hagita, N.F.: Abstracting people’s trajectories for social robots to proactively approach customers. *IEEE Trans. on Robotics* 25(6) (2009)
6. Helbing, D., Molnár, P.: Social force model for pedestrian dynamics. *Physical Review E* 51(5), 4282–4286 (1995)
7. Ikeda, T., Chigodo, Y., Rea, D., Zanlungo, F., Shiomi, M., Kanda, T.: Modeling and prediction of pedestrian behavior based on the sub-goal concept. In: Proc. of RSS (2012)
8. Kuderer, M., Kretschmar, H., Sprunk, C., Burgard, W.: Feature-based prediction of trajectories for socially compliant navigation. In: Proc. of RSS (2012)
9. Althoff, D., Kuffner, J., Wollherr, D., Buss, M.: Safety assessment of robot trajectories for navigation in uncertain and dynamic environments. *Autonomous Robots* 32(3), 285–302 (2012)
10. Van de Weghe, N.: Representing and Reasoning about Moving Objects: A Qualitative Approach. PhD thesis, Ghent University (2004)
11. Cohn, A.G., Renz, J.: Qualitative Spatial Representation and Reasoning. In: van Harmelen, F., Lifschitz, V., Porter, B. (eds.) *Handbook of Knowledge Representation*, vol. 3, ch. 13, pp. 551–596. Elsevier (2008)
12. Bellotto, N., Hu, H.: Multisensor-based human detection and tracking for mobile service robots. *IEEE Trans. on SMC – Part B* 39(1), 167–181 (2009)
13. Fox, D., Burgard, W., Thrun, S.: The dynamic window approach to collision avoidance. *IEEE Robotics and Automation Magazine* 4(1), 23–33 (1997)
14. Delafontaine, M., Cohn, A., de Weghe, N.V.: Implementing a qualitative calculus to analyse moving point objects. *Expert Systems with Applications* 38(5) (2011)
15. Arens, M., Nagel, H.-H.: Behavioral knowledge representation for the understanding and creation of video sequences. In: Günter, A., Kruse, R., Neumann, B. (eds.) *KI 2003. LNCS (LNAI)*, vol. 2821, pp. 149–163. Springer, Heidelberg (2003)
16. Ong, S., Png, S., Hsu, D., Lee, W.S.: POMDPs for robotic tasks with mixed observability. In: Proc. of RSS (2009)