Towards Safer Robot Motion: Using a Qualitative Motion Model to Classify Human-Robot Spatial Interaction*

Laurence Roberts-Elliott[1 [0000−0001−9877−0553]], Manuel Fernandez-Carmona[1 [0000−0002−0512−8594]], and Marc Hanheide[1 [0000−0001−7728−1849]]

University of Lincoln, Lincoln, LN6 7TS, UK
{laelliott, mfernandezcarmona, mhanheide}@lincoln.ac.uk

Abstract. For adoption of Autonomous Mobile Robots (AMR) across a breadth of industries, they must navigate around humans in a way which is safe and which humans perceive as safe, but without greatly compromising efficiency. This work aims to classify the Human-Robot Spatial Interaction (HRSI) situation of an interacting human and robot, to be applied in Human-Aware Navigation (HAN) to account for situational context. We develop qualitative probabilistic models of relative human and robot movements in various HRSI situations to classify situations, and explain our plan to develop per-situation probabilistic models of socially legible HRSI to predict human and robot movement. In future work we aim to use these predictions to generate qualitative constraints in the form of metric cost-maps for local robot motion planners, enforcing more efficient and socially legible trajectories which are both physically safe and perceived as safe.

Keywords: HRI · HRSI · spatial reasoning · Human-Aware Navigation · Hidden Markov Models · classification

1 Introduction

In industrial applications with environments shared between humans and robots, AMRs must move safely around humans and in a way which humans perceive to be safe. Physical human safety in robot navigation can be all but assured by simply stopping robot motion when anything is detected closer than a minimum safe distance to a robot's safety laser(s). This is highly inefficient and usually conflicts with personal space [9], thus perceived as unsafe.

The work discussed in this paper contributes to the safety stack of the human collaborative warehouse robots of the EU ILIAD Project [11], aiming to increase perceived safety by humans while preserving a critically physically safe system. The safety stack is implemented at four levels. At the highest level, global planning, robots take human flows into account to generate paths that

* Supported by the EU H2020 ILIAD Project (grant #732737).

don't interfere with human motion patterns [19]. The second level affects robot maximum speeds. If a robot detects a human in its vicinity, it will adapt its speed to the HRSI at hand. However, if the global path seems unfeasible given current HRSI, a new path will be triggered to account for the new interaction. Finally, a safety stop will be triggered if human gets within 0.5m of the robot's laser anyways. The ILIAD safety stack is described in greater detail in [7].

This work, involved in the second level of the ILIAD safety stack, addresses HRSI awareness with per-situation probabilistic models of qualitative abstractions of human and robot movement. Our probabilistic models use Qualitative Trajectory Calculus version C (QTC_C) [20] to encode the movements of two positions in space from one point in time to a subsequent point in time. This lets us represent pairs of trajectories in HRSI as sequences of qualitative states, where each state describes the relations of the movements of both human and robot.

The main contribution of this work is the multi-HMM classifier of the HRSI situation of sequences of qualitative descriptions of human and robot movement, extending [2] with multi-HMM classification and modelling of additional situations. The classifier is very fast to train, and fast and accurate enough at classification to be used in real-time in a safety critical situational HAN approach. The use of QTC, described in detail in Sec. 3.2, abstracts pairs of trajectories, decreasing the impact of sensor and tracking error. We train and test our classifier using data recorded from HRSIs with an automated Linde CiTiTruck pallet truck, making it more suitable for industrial applications than approaches which model robot movement after human-human interaction.

We plan to use the classifier, along with per-situation HMMs of QTC_C state transitions from socially legible HRSI, to predict the socially legible next state for a HRSI in realtime, given the HRSI's situation and QTC_C sequence. In future work, we aim to constrain the robot's local planner to enforce the transition to this state.

The rest of the paper is structured as follows. We first present other works within the field of HRSI in Sec. 2. Later on, we describe in Sec. 3 the theoretical foundations of our system. Implementation and analysis of the system is made in Sec. 4. Finally we summarise the results and draft future work in Sec. 5.

2 Related Work

Common approaches to robot navigation consider humans the same as any other obstacle [14], resulting in movements that are inefficient and perceived as unsafe. HAN is required for legible paths which are more likely to be perceived as safe [6]. Approaches for HAN often consider Hall's proxemic zones [9], but neglect to consider the intentions of human's movement [14].

Predictive models address this limitation by identifying and forecasting human trajectories. They rely on studying human motion in social environments, so that most likely paths are known when a similar situation is

matched. These works propose the use of qualitative domains and symbols to reduce the complexity of the task at hand [17] and include desired features (that increase social normativity) [13] with good performance in crowded scenarios.

However, these approaches do not capture in their motion models the presence of a robot. There is the implicit assumption that humans will move as if the robot was just another pedestrian. This is not appropriate for navigation of heavy industrial robots, which cause humans to feel unsafe when the robot moves close to them. Also, using these models as path planners for the robot may not provide the safest route, but instead the most human alike. This is particularly relevant for the domain at hand: shared warehouse environments. Within industrial applications, robots need to ensure safety over any other requirement, so mimicking human trajectories (e.g. cutting through a crowd) may be discouraged.

Instead of just learning motion models, other authors choose instead to embed human awareness in more traditional robot planners. This was first proposed by Lu et al. [16], creating socially aware cost-maps for occupancy grid-based planners. Similarly, Sisbot et al. [1] uses the so called *"Social Force Model"* that can be embedded in a classic sampling motion planner while accounting for human proxemics, or Araceli et al. [18], who developed an adaptation of the elastic band algorithm that accounted for potential human activity spaces.

Our work is in line with those authors, aiming for a description of the human interaction based on QTC that can be introduced in any planner. QTC used in our HRSI model describes relative movements of both human and robot [20] in the same way that two humans walking on intersecting trajectories negotiate their movements without knowledge of their quantitative positions [10]. A set of per-situation HMMs of transitions between QTC states can be used to classify the HRSI situation from QTC sequences generated from pairs of human and robot trajectories [2]. Here we extend the HRSI situation classification of [10], modelling additional situations.

3 System definition

3.1 Human Position Tracking

Warehouses are a challenging scenario for people tracking (multiple occlusions, dim lights, people sitting, kneeling, etc.), which may be a limiting factor for qualitative HRSI analysis. Our probabilistic HAN model uses human trajectories obtained with the real-time first-person people tracking system described in [5]. This tracking fuses data from the robot's on-board RGB-D camera and laser, which has a limited reach and requires extra computing power, but does not require any sensors to be installed in the warehouse.

3.2 Probabilistic QTC Model

Our probabilistic QTC model, uses sequences of QTC_C states to develop a Hidden Markov Model (HMM) for each of a set of HRSI situations, defined as

classes in Sec. 3.3 and extending our previous work in [4]. We encode pairs of human and robot trajectories using QTC version C (QTC_C) [20]. In QTC_C, movements of two agents in space are represented by a 4-tuple of state descriptors (h_1, r_1, h_2, r_2). Each descriptor expresses a qualitative spatial relation using a symbol $\in \{-, 0, +\}$. With this 4-tuple of descriptors comprised of 3 symbols, there a total of $3^4 = 81$ possible QTC_C states.

The relations of the descriptor symbols are defined as follows:

h_1) movement of h w.r.t. r at time t:
- : h is moving towards r
0 : h is neither moving towards nor away from r
+ : h is moving away from r

r_1) movement of r w.r.t h at time t:
The same as h_1), but with h and r swapped

h_2) movement of h w.r.t the line $\overrightarrow{hr}$ at time t:
- : h is moving to the left side of $\overrightarrow{hr}$
0 : h is moving along $\overrightarrow{hr}$ or not moving at all
+ : h is moving to the right side of $\overrightarrow{hr}$

r_2) movement of r w.r.t the line $\overrightarrow{rh}$ at time t:
The same as h_2), but with h and r swapped

where t is the earlier of the two points in time, r is the robot's position, and h is the human's position.

For example, Fig. 1 shows an interaction: human is moving towards the robot ($h_1 = -$) and robot is approaching too ($r_1 = -$). Human is directly headed to the robot ($h_2 = 0$) but robot is to its left side ($r_2 = -$).

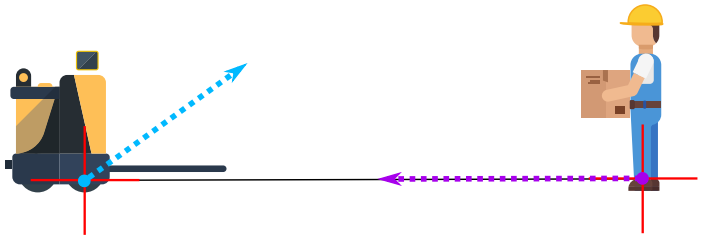

Fig. 1. QTC_C state $(-, -, 0, -)$: human is moving directly towards the robot, while the robot is moving toward and on its left side.

3.3 Classes of HRSI

In [4] there were two relevant classes in human robot spatial interaction. We extend this initial classification to account for the interactions in warehouse environments. Specifically we focus on interactions that require a robot to adjust its movements to accommodate the human's. In order to account for these, we model the following situations (see Fig. 2) with HMMs for our classifier:

- Passing By on the Left (PBL): Both actors pass each-other on the left side from their perspective, moving in opposite directions.
- Passing By on the Right (PBR): Both actors pass each-other on the right side from their perspective, moving in opposite directions.
- Robot Overtakes Left (ROL): The robot passes on the left of the human while both move in the same direction.
- Robot Overtakes Right (ROR): The robot passes on the right of the human while both move in the same direction.
- Path-Crossing on the Left (PCL): The robot has to slow or stop movement to allow the human to move across the robot's intended path from the robot's left side.
- Path-Crossing on the Right (PCR): The robot has to slow or stop movement to allow the human to move across the robot's intended path from the robot's right side.
- Rejection: Any situation in which a human is detected, but a qualitative constraint of the robot movement is not required. E.g., RMSH: The robot moves toward the stationary human, and stops when close. We record trajectory pairs from examples of RMSH situations to test our multi-HMM classifier's ability to reject such situations. For these situations we apply a default constraint, enforcing safe and comfortable distance between human and robot. Explanation of these constraints is outside of the scope of this work, but is detailed in the ILIAD 'Report on human-robot spatial interaction and mutual communication of navigation intent' [7].

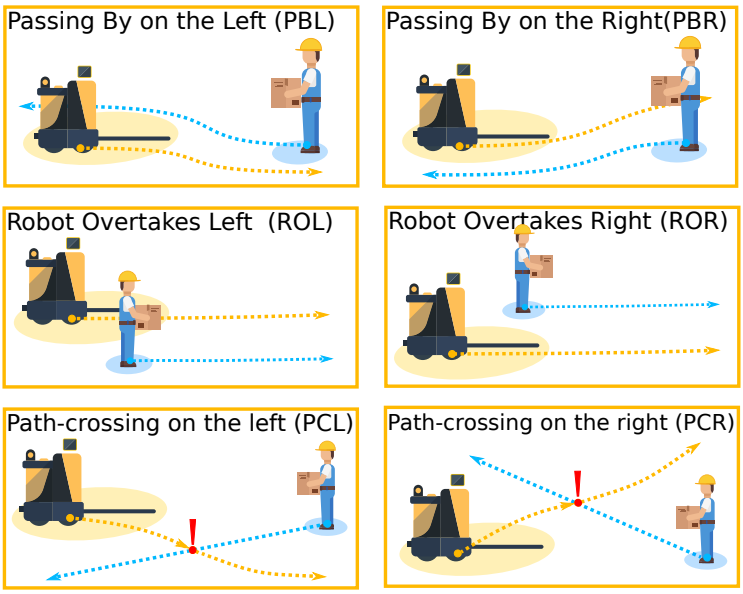

Fig. 2. HRSI Classes

3.4 Creating HMMs for our Multi-HMM Situation Classifier

We consider in this study 6 classes C = ('PBL', 'PBR', 'ROL', 'ROR', 'PCL', 'PCR'), exluding rejection, thus we will model 6 different HMMs, where the observation corresponds to a new state. Each HMM is comprised by $|Q| \times |Q|$ transition matrices listed in A_i, a $|Q| \times |Q|$ observation matrix B, and the $1 \times |Q|$ initial state vectors listed in I_i. These account for all possible transitions in QTC_C states $Q = ((- - \ - -) \dots (+ + \ + +))$, as in [20] on each class.

Our system is composed by the collection of transition matrices $A = (A_1 \dots A_{|C|})$, and initial state vectors $I = (I_1 \dots I_{|C|})$, indexed by class number. Each element of the list of per-class HMM $H = (H_1 \dots H_{|C|})$ is a tuple composed by (A_c, B, I_c) with c indexing the class's name in C, fully describes our system.

All of the classifier's HMM share B as their observation matrix. We use B to account for the possibility of generating incorrect QTC_C states due to sensor and tracking error, assuming a probability $t = 0.95$ that the true (hidden) QTC_C state matches the emitted QTC_C state generated from tracked human and robot positions. So, we initialise matrix B almost as an identity matrix with some noise, with diagonal $B[i, i] = t$ and the rest of elements $B[h, o] = \frac{1-t}{|Q|-1} \mid (b \neq o)$.

A list S of recorded QTC_C state sequences, generated from human-robot trajectory pairs, is used to obtain A, and I. First, we map each QTC_C state in each sequence S_i to its index in Q. Each state sequence in S will have a class label assigned $l_i \in [1 \dots |C|]$, so that the list of labels will be $L = (L_1 \dots L_{|S|})$ and L_s is the class label for sequence S_s. Initially, A and I are assigned uniform probabilities. Then we use the recorded state sequence list S to model the probabilities

$$I_{L_n}[S_n[1]] = I_{L_n}[S_n[1]] + 1 \text{ for } n = 1 \text{ to } |S|. \tag{1}$$

$$A_{L_n}[S_n[q], S_n[q{+}1]] = A_{L_n}[S_n[q], S_n[q{+}1]]{+}1 \text{ for } n = 1 \text{ to } |S|, \text{ and } q = 1 \text{ to } |S_n|{-}1. \tag{2}$$

Finally we normalise matrix B, the matrices of A, and the vectors of I, such that each row sums to 1. With these HMM, we can classify a QTC_C sequence as the class of the HMM that estimates the highest log-likelihood of the given sequence being observed [21]. If the KL divergence [12] of these log-likelihoods, normalised to sum to 1, from a uniform distribution of the same size is greater than a given threshold then the sequence is instead rejected.

4 Experiments

4.1 Laboratory Setup for Recording HRSI Situations

Our robot is a modified Linde CitiTruck equipped with a front-facing safety laser that automatically stops motors if anything is detected in front of the robot closer than 0.5m, prioritising human safety over any software command. We have additional sensors used for human detection, such as a LiDAR and

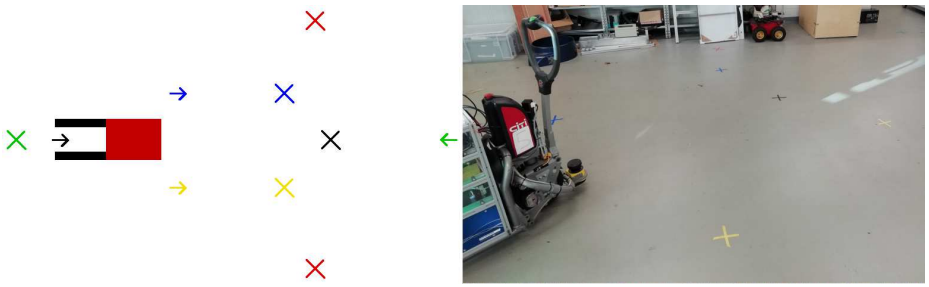

Fig. 3. Illustration (left) and photograph (right) of laboratory setup for recording HRSI situations.

Kinect 2. The robot is controlled using a standard PC running ROS, interfacing with the sensors, and the pallet truck's motor controllers. It is one of several automated pallet trucks belonging to the ILIAD Project. In our robotics lab we placed coloured tape on the floor, marking start and end positions for the human and the robot, as pictured in Fig. 3. In the diagram on the left of Fig. 3, arrows indicate start positions, and crosses indicate end positions. The robot follows the path between the black positions. In HRSI situations where the robot is stationary, the robot stays at the start position. When the robot begins moving it emits a click sound, which we use to signal the human to move. In conditions where the robot does not move, the experimenter speaks the signal 'go'. We record the robot and nearest human position on the robot's metric map, from when the robot begins moving, to when the human reaches their end position. Human positions are tracked within an 'active area' to reduce the risk of the experimenter being tracked instead of the interacting human.

The human moves as follows for the different situations:

PBL – The human moves from the green arrow to the green cross, moving to their left to pass the robot.
PBR – The same as PBL, but the human moves to their right to pass the robot.
ROL – The human moves from the yellow arrow to the yellow cross, moving as slowly as possible to allow the robot to overtake at a safe speed.
ROR – The same as ROL, but the human moves from the blue arrow to the blue cross.
PCL – The human moves from the topmost red cross to the other red cross.
PCR – The human moves from the bottom red cross to the topmost red cross.
RMSH – The human stands stationary at the yellow cross, the robot moves from the black arrow to the black cross.

Training dataset To train our multi-HMM classifier we recorded 35 interactions between a robotics expert and the robot for each of the 6 situation classes of Sec. 3.3, and 15 interactions for the rejection situation RMSH. A total of 225 HRSI. We create the multi-HMM classifier's HMM as described in

Sec. 3.4, using QTC_C sequences generated from human and robot trajectories using *QSRlib* [8]. We use *QSRlib*'s 'collapse' feature when generating all of our QTC_C sequences from HRSI, to remove repeating states, reducing the variance between sequences from HRSI of differing length. We used 3-fold cross validation for preliminary estimation of our classifier's performance, to measure the likely impact of changes to our model to its ability to classify HRSI situations beyond those recorded in this training set.

Test dataset To test the ability of our multi-HMM classifier to classify the situation of spatial interactions between the robot and non-experts of varied age, gender and cultural background, we conducted a study. In this repeated measures study, participants enacted HRSI situations using the methods described in Sec. 4.1. These situations were enacted in a randomised order, with each participant also performing 1 of a set of 5 rejection situations, chosen at random. We realised in hindsight that only 1 of these situations, RMSH, would need to be rejected by our classifier as the others did not involve the robot moving. Thus with 11 participants, we recorded a total of 66 interactions, including 2 RMSH interactions. We created the multi-HMM classifier's HMM using QTC_C sequences from the training set, and evaluated its performance in classifying the HRSI situation of QTC_C sequences from the study's HRSI. Training of the classifier took only 50 ms to execute. Each classification took 60 ms to execute on average. The number of interactions recorded per class is detailed in the study's confusion matrix in Sec. 4, which has its statistics explained in Sec. 4.2.

4.2 Results and Discussion

Fig. 4 is a confusion matrix containing metrics of the performance of the classifier at predicting the HRSI situation from QTC_C sequences in the test set described in Sec. 4.1, trained on sequences from the training set described in Sec. 4.1. The cells of the confusion matrix with a green or pale red background contain the count of classifications for the predicted class and actual class given by the cell's row and column respectively. Below each of these counts is the count as a percentage of the total number of classifications. The rightmost column of the matrix contains, in this order, the count of QTC_C sequences classified as the row's class, and the precision and False Positive Rate of classifications as the row's class. The bottom row of the matrix contains, in this order, the count of QTC_C sequences that are labelled with the column's class, and the recall and False Negative Rate of classifications as the column's class. The bottom-right cell contains, in this order, the total count of all classifications, the overall accuracy, and the overall misclassification rate.

The classifier's overall accuracy is high at 95.45%, as it must be as a component of the safety focused HAN approach in ILIAD. The ability of our qualitative probabilistic model to accurately classify HRSI with 11 non-experts, trained on HRSI with 1 robotics expert, demonstrates the benefit

Fig. 4. Confusion matrix for validation of our multi-HMM classifier using QTC_C sequences from study HRSIs as test data.

of abstracting HRSI to a qualitative description. It should be noted that while no rejection situations were misclassified, the precision of rejections is relatively low at 66.67%. This could be due to the small number of RMSH interactions recorded in our study, as explained in Sec. 4.1, as 3-fold Cross-Validation using the training dataset with its 15 RMSH interactions showed much higher rejection precision at 78.95%. The classifer is much less likely to confuse any of the other 6 classes, with precision and recall $> 90\%$ for all of these classes. We hope that we can improve the rejection precision and in turn the overall accuracy of our classifier by recording more examples of RMSH interactions and other rejection situations.

Our classifier is able to identify an interaction from a QTC sequence in less than 100ms. Each QTC state in our system is obtained almost instantaneously given current robot and human positions, which means that the limiting factor in our pipeline is the human tracker module itself. The current tracker [5] has an average rate of 25 Hz which sufficiently accounts for socially normative human behaviours and speeds, e.g. the mean human speed of 0.58 m/s and maximum robot speed of 0.91 m/s recorded in our study. The maximum human speed we recorded was 2.28 m/s, close to jogging speed, but these high speeds are not

represented frequently in our study. When an interaction is not classified or is given the rejection class, a default constraint described in Sec. 3.5.1 of [7] is applied to ensure safe and comfortable human-robot distance at all times, and a safety laser stops the robot motors when within 0.5m of the human.

We 'collapse' QTC sequences, removing consecutive repeats of states to account for variation in sequence length caused by varying temporal length of interactions. We plan to include a more robust tracker [15] capable of tracking humans at higher speeds, and we may consider QTC formulations that include speeds in their state description. We may also replace our HMMs with Continuous-Time Hidden Markov Models (CTHMMs) which better describe sequences where states are observed at irregular times [15]. All of these offer the possibility to extend our system to identify and react to non socially normative situations.

5 Conclusion and Future Work

We have presented here the framework used in the ILIAD Project to classify HRSI situation. If global planning is unable to create paths avoiding HRI, the classified situation will define how the intermediate safety layers will react in ILIAD. An efficient and accurate prediction will have a direct impact in the number of safety stops triggered by the lowest safety layer. Future work will evaluate the performance of this approach in the project's overall safety architecture.

The high accuracy of our multi-HMM HRSI situation classifier when tested on the HRSIs recorded in our experiment demonstrates its suitability for use in a situational HAN approach, with some room for improvement in the precision of rejection, which may be possible by taking the steps described in Sec. 4.2.

While this fully describes the applicability of our work in a qualitative HAN approach, we have not tested the systems in this paper with QTC_C sequences generated from partial interactions, needed to be applied to real-time robot navigation in the presence of humans. We will address this in future work, developing a continuous real-time implementation of our classifier that informs qualitative constraint of robot motion. The real-time HAN system should be tested on an industrial AMR in its ability to improve perceived safety and social legibility of robot movement in HRSI, while minimising safety-laser stops. For this real-time system, we aim to continuously classify the situation of a HRSI, making a new classification each time a QTC_C state is added to the sequence representing the interaction, and providing situational robot motion planner cost-maps enforcing a qualitative constraint of robot motion, i.e. [7], in order to improve perceived safety with socially normative movement, extending [3]. When multiple constraint cost-maps are to be applied simultaneously, the union of these is used to constrain the robot motion appropriately for simultaneous HRSIs with any number of humans.

For data to train the HMMs and validate this approach, we look to label and process trajectory pairs in larger datasets, then test in real-time on our robot in a real warehouse, interacting with human workers.

We hope to use the qualitative constraint cost-maps to communicate the robot's navigation intention, e.g. projecting on the floor in green on areas of high-cost, and red on areas of low-cost as conditioned symbols for 'go' and 'don't go'.

References

1. Cosgun, A., Sisbot, E.A., Christensen, H.I.: Anticipatory robot path planning in human environments. In: 2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN). pp. 562–569 (Aug 2016). https://doi.org/10.1109/ROMAN.2016.7745174
2. Dondrup, C., Bellotto, N., Hanheide, M.: A Probabilistic Model of Human-Robot Spatial Interaction Using a Qualitative Trajectory Calculus. In: AAAI Spring Symposium Series. AAAI, Palo Alto, California, USA (Mar 2014), https://www.aaai.org/ocs/index.php/SSS/SSS14/paper/view/7714
3. Dondrup, C., Hanheide, M.: Qualitative constraints for human-aware robot navigation using Velocity Costmaps. In: 2016 IEEE RO-MAN. pp. 586–592. IEEE, Piscataway, New Jersey, USA (Aug 2016). https://doi.org/10.1109/ROMAN.2016.7745177
4. Dondrup, C., et al.: A Computational Model of Human-Robot Spatial Interactions Based on a Qualitative Trajectory Calculus. Robotics **4**(1), 63 (2015). https://doi.org/10.3390/robotics4010063
5. Dondrup, C., et al.: Real-time multisensor people tracking for human-robot spatial interaction. ICRA / IEEE, Seattle, WA (May 2015), http://eprints.lincoln.ac.uk/id/eprint/17545/
6. Fernandez Carmona, M., Parekh, T., Hanheide, M.: Making the Case for Human-Aware Navigation in Warehouses. In: Althoefer, K., et al. (eds.) TAROS 2019. pp. 449–453. LNCS, Springer International Publishing, Cham (2019). https://doi.org/10.1007/978-3-030-25332-5_38
7. Fernandez-Carmona, M., et al.: Report on human-robot spatial interaction and mutual communication of navigation intent. Tech. rep., ILIAD Project (May 2020), http://iliad-project.eu/wp-content/uploads/Deliverables/D3_3.pdf
8. Gatsoulis, Y., et al.: QSRlib: a software library for online acquisition of qualitative spatial relations from video. In: Qualitative Reasoning29th International Workshop on Qualitative Reasoning. pp. 36–41. IJCAI, California, USA (Jul 2016), https://ivi.fnwi.uva.nl/tcs/QRgroup/qr16/pdf/05Gatsoulis.pdf
9. Hall, E.T., et al.: Proxemics. Current Anthropology **9**(2/3), 83–108 (1968), https://www.jstor.org/stable/2740724
10. Hanheide, M., Peters, A., Bellotto, N.: Analysis of human-robot spatial behaviour applying a qualitative trajectory calculus. In: 2012 IEEE RO-MAN. pp. 689–694. IEEE, Paris, France (Sep 2012). https://doi.org/10.1109/ROMAN.2012.6343831
11. ILIAD Project: ILIAD Project – An EU-funded research project on Intra-Logistics with Integrated Automatic Deployment for safe and scalable fleets in shared spaces. (Feb 2020), http://iliad-project.eu/
12. Joyce, J.M.: Kullback-Leibler Divergence, pp. 720–722. Springer Berlin Heidelberg, Berlin, Heidelberg (2011). https://doi.org/10.1007/978-3-642-04898-2_327

13. Kretzschmar, H., et al.: Socially compliant mobile robot navigation via inverse reinforcement learning. The International Journal of Robotics Research **35**(11), 1289–1307 (2016). https://doi.org/10.1177/0278364915619772
14. Kruse, T., et al.: Human-aware Robot Navigation: A Survey. Robotics and Autonomous Systems **61**(12), 1726–1743 (Dec 2013). https://doi.org/10.1016/j.robot.2013.05.007
15. Liu, Y.Y., et al.: Efficient Learning of Continuous-Time Hidden Markov Models for Disease Progression. Advances in neural information processing systems **28**, 3599–3607 (2015), https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4804157/
16. Lu, D.V., Allan, D.B., Smart, W.D.: Tuning cost functions for social navigation. In: Herrmann, G., Pearson, M.J., Lenz, A., Bremner, P., Spiers, A., Leonards, U. (eds.) Social Robotics. pp. 442–451. Springer International Publishing, Cham (2013)
17. Mavrogiannis, C.I., Knepper, R.A.: Multi-agent path topology in support of socially competent navigation planning. The International Journal of Robotics Research **38**(2-3), 338–356 (2019). https://doi.org/10.1177/0278364918781016
18. Vega, A., et al.: Socially aware robot navigation system in human-populated and interactive environments based on an adaptive spatial density function and space affordances. Pattern Recognition Letters **118**, 72 – 84 (2019). https://doi.org/https://doi.org/10.1016/j.patrec.2018.07.015, cooperative and Social Robots: Understanding Human Activities and Intentions
19. Vintr, T., et al.: Time-varying pedestrian flow models for service robots. In: 2019 European Conference on Mobile Robots (ECMR). pp. 1–7 (Sep 2019). https://doi.org/10.1109/ECMR.2019.8870909
20. Van de Weghe, N., et al.: A Qualitative Trajectory Calculus and the Composition of Its Relations. In: Hutchison, D., et al. (eds.) GeoSpatial Semantics, vol. 3799, pp. 60–76. Springer Berlin Heidelberg, Berlin, Heidelberg (2005). https://doi.org/10.1007/11586180_5
21. White, A.M.: Sequence Classification - with emphasis on Hidden Markov Models and Sequence Kernels (Sep 2009), http://www.cs.unc.edu/ lazebnik/fall09/sequence_classification.pdf