Walk With Me

Socially Acceptable Speed and Distance Control for Mobile Wayfinding Robots

Hansen, Karl Damkjær; Clement, Anders; Larsen, Hans Christian Østgård; Andersen, Jonas; Beyer Lauritsen, Oliver; Møller-Rahbek, Katrine

Published in:
IEEE International Conference on Imaging Systems and Techniques (IST 2023)

DOI (link to publication from Publisher):
10.1109/IST59124.2023.10355692

Publication date:
2023

Document Version
Early version, also known as pre-print

Walk With Me: Socially Acceptable Speed and Distance Control for Mobile Wayfinding Robots

Karl D. Hansen[1], Anders Clement[2], Hans Christian Østgård Larsen[2], Jonas Andersen[2], Oliver Beyer Lauritsen[2] and Katrine Møller-Rahbek[2]

***Abstract*—Robots can help people find their way around in complex buildings and environments. Robots are embodied, which makes it more intuitive for users to follow them around compared to, e.g., colored lines on the floor. However, they need to be socially intelligent for people to accept them. Like people, robots should keep a comfortable distance to its users when working.**

We investigate the optimization of a guiding robot's speed and distance when guiding to improve the user experience. Our findings indicate a correlation between the user's speed and the distance they keep from the robot, which can be utilized to control the robot's speed.

A person tracking system is implemented on a Spot quadruped mobile robot platform from Boston Dynamics, using the on-board depth and 2D cameras plus an Intel NUC and a Jetson Nano for processing. The system uses the Mobilenet-SSDv2 CNN for user detection and Kalman filtering for tracking.

Data from human-robot interaction tests on footpaths at Aalborg University is analyzed to create a linear model of the speed-distance relationship. Based on this, a control law is proposed and tested, demonstrating the ability to build a controller that allows the person following the robot to set the desired velocity of the robot.

Fig. 1: The Boston Dynamics robot Spot guiding a person on a sidewalk.

I. INTRODUCTION

People get lost in large and complex building such as airports and hospitals. Indeed, airports have increasing amounts of passenger boardings and hospitals are growing in size [1][2]. Getting lost in these places, in turn, leads to delayed flights, missed doctor's appointments, frustration for the visitors and additional costs for everyone involved [3]. To help people navigate, different wayfinding methods are used. These include: Signs, 'You Are Here'-maps, and painted guiding lines on the floor [4].

Robots can be used as an alternative or complement to the existing wayfinding technologies. Personal guiding in, e.g., hospitals is often infeasible because of lack of personnel; robots can alleviate this pain. Having guide robots located at the entrances of buildings ready to escort people to their destinations could potentially help solve some of these problems. Employing mobile assistant robots in these dynamic and busy environments calls for research and further development in many areas of robot navigation, one of these is the social capabilities of robots. When people are guided to an unknown place, it is important that they feel the robot acts in a way that is predictable and comfortable. To achieve this, the robot has to act socially acceptable [5, 6, 7]. The study of inter-personal relations, proxemics [8], can guide the design of new robot systems. Specific studies at airports, museums, hospitals, etc. show how people interact with wayfinding robots [9, 10, 11], and how this is a complex and challenging topic.

In this work we describe the design of a control system to maintain a comfortable speed of the wayfinding robot, in order to improve the experience of being guided by a robot. The main novelty of this contribution is the identification of a mapping between comfortable distance and walking speed and the associated control structure. Figure 1 shows a participant being guided by the robot.

The remainder of this paper is organized as follows: Related works will be presented in section II where the state of the art research relating to guiding robots will be looked into. Our design is described in section III, which includes: The hardware used, the people tracking system, the test setup for data gathering, and control design for the guiding robot. Section IV shows the experimental results of the implemented controller design when guiding participants. In section V the results are discussed and future work is presented.

[1]K. D. Hansen is with the Department of Materials and Production, Aalborg University, 9220 Aalborg, Denmark kdh@mp.aau.dk

[2]A. Clement, H. C. Ø. Larsen, J. Andersen, O. B. Lauritsen, K. Møller-Rahbek are students of the Robotics MSc program at Aalborg University, 9220 Aalborg, Denmark [acleme18, hlarse16, jan18, olauri18, kmalle18]@student.aau.dk

II. Related Work

An early instance of robots guiding is the museum-guide robot RHINO [12]. The RHINO robot is used in order to guide museum guests through the Deutsches Museum Bonn. Their main focus was to ensure safe movement through crowded areas. Later, the same team developed another museum guide robot; Minerva [13]. This robot had more sophisticated navigation and an emotion system was implemented, which made the human-robot interaction better.

A newer project addressing guiding robots is the SPENCER project [14]. The purpose of the SPENCER project is to assist people in large and busy airports. Through detection and tracking of individuals and groups of people using 2D laser scanners and RGB-D cameras, the SPENCER robot can perform socially acceptable navigation which abides to the social rules of the surrounding people, while tracking people who are being guided to their destination. The tracking is used to determine if the user is still following the robot. These three works focus on guiding a person, using socially acceptable navigation through crowded locations. Although they track the person following the robot they do not consider the social dynamics of the guiding itself.

Walters et al. [15] found that comfortable distances between human and robot are comparable to human-human social distances. In $60\ \%$ of cases, the most comfortable distance, between participants and a stationary robot, was found to be from $45\ \mathrm{cm}$ to $360\ \mathrm{cm}$, depending on the person's preference. The remaining participants preferred to be closer. The same results were observed when the robot approached a stationary participant, and they had to stop it. Gockley et al. [16] created a robot that follows a human at the distance $(120 \pm 10)\ \mathrm{cm}$, to keep the robot just outside of the human's personal space. The participants of that study found that this was a little too far away. Boladeras et al. [11] look into how people follow robots. They found that when a single person is guided, then all persons walked behind the robot rather than to the side of it.

Based on these works, we believe that it is important to understand the distance between the person and robot, the walking speed, and possibly where the person positions themselves relative to the robot. Then, it is believed that this knowledge can be used to control the robot in a more socially acceptable manner.

III. Design

This section details 1) the hardware; 2) the people tracking system which uses Mobilenet-SSDv2 and a Kalman filter for position estimation; 3) how a behavior model is created and used for estimating the preferred walking speed and distance; and finally, 4) the design of a controller that regulates the speed of a robot according to the behavior of the person following it.

A. Hardware

The robot used in this study, is the Spot robotic platform created by Boston Dynamics. This platform was chosen because of its versatility and speed. However, our solution should be usable on most mobile platforms used for guidance. Spot is a quadruped robot, designed to walk on difficult terrain. It has a size of $1.1\ \mathrm{m} \times 0.5\ \mathrm{m} \times 0.6\ \mathrm{m}$ and a battery capacity of 564 Wh giving approximately 90 minutes of continuous walking. During the testing it was found that although Spot is reported to have a maximum speed of $1.6\ \mathrm{m/s}$, it is actually $1.4\ \mathrm{m/s}$.

Spot has five camera modules, mounted in a configuration with one camera on the back and each side of Spot, and two in the front. This gives 360 degrees vision of the surroundings of the robot. The cameras module each has an infrared active stereo camera which provides depth images. Although unspecified by Boston Dynamics, this module resembles the Intel Realsense D430 depth module (87°x58°, 1280x720 px). The modules also feature a black and white camera with approximately the same FOV as the depth module.

The robot is controlled using Clearpath Robotics' Spot driver [1] for ROS. The driver runs on a small form factor computer, an Intel NUC with an Intel Core i7-5557U Processor and $16\ \mathrm{GB}$ RAM, which is mounted on top of Spot and connected by Ethernet. The computer also collects and stores data and logs for offline processing.

An NVIDIA Jetson Nano is also mounted on Spot for running the visual object detector. This is a small computer with a 128-core NVIDIA Maxwell GPU, a Quad-Core ARM Cortex-A57, $4\ \mathrm{GB}$ RAM and $16\ \mathrm{GB}$ storage, designed for running neural networks. It is connected to the Intel NUC using Ethernet over USB.

B. Tracking

Since Spot is a mobile platform, a detection algorithm is preferred to be efficient and low power, and must be good at detecting people in the proximity of Spot. According to the review of convolutional neural network object detectors from 2017 by Huang et al. [17] the MobileNet architecture was the fastest architecture available with respect to time and accuracy. Since then, Mobilenet-SSDv2 has been released, which is a combination and improvement of MobileNet and SSD [18]. In a comparison by Rios et al. [19] though, it was found that Mobilenet-SSDv2 has lower recall than YoloV3. However, Mobilenet-SSDv2 is better at detecting large objects in the image compared to other detectors [20]. Therefore, Mobilenet-SSDv2 is used in this project.

The Jetson Nano and receives images from the cameras of Spot and uses the Mobilenet-SSDv2 to detect people. The output is a bounding box surrounding the detected person and a confidence of the detected class. In this case we use the detection when the confidence is greater than 0.5.

Given the detection, the next step is to find the center of mass (CoM) of the detected persons. From the bounding box given by Mobilenet-SSDv2, the upper half and the center $50\ \%$ of the bounding box width is cut out. This is done to ensure that most of the background is removed from the detection as the upper middle of the bounding box locates the upper legs

[1] https://github.com/clearpathrobotics/spot_ros

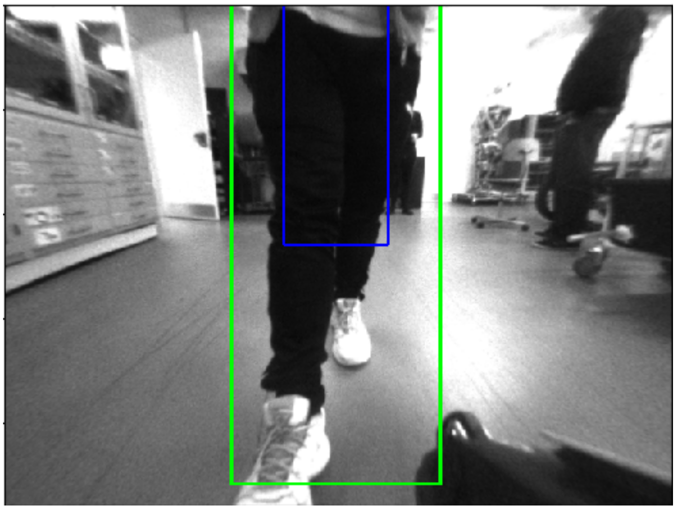

Fig. 2: Mobilenet-SSDv2 detecting a human in an image from the rear-facing camera on Spot. The large green bounding box is the detection box given by Mobilenet-SSDv2 and the small blue bounding box is the cutout used for the histogram and deprojection.

or torso of the user, as seen on Figure 2. This cutout from the visual image is applied to the corresponding depth image from the stereo camera pair. Then a histogram of depth values is created using bins of size $200\ \mathrm{mm}$. This is done to remove effects of fore- and background objects still present in the view. All depth pixels in the largest bin are then deprojected to a point cloud from which the CoM is calculated as the average position of the points. This CoM is then considered the position of the person. Finally, the CoM is transformed into Spot's base coordinate frame.

Our implementation of the tracker reaches 14 fps when running with 2 people in view of the robot and 10 fps when there are 4 people in view. However, running all five cameras, it was only possible to get 2 fps from each camera on Spot. Limiting the use to the rear-facing camera gives an frame rate of 10 fps.

When the system has detected a person, a Kalman filter with a constant velocity motion model is used in order to keep track of the person following the robot. Assuming constant linear velocity is a simple, but accurate, method compared to more advanced models [21].

Detections are tracked in $[x, y, z]$ coordinates with the state vector defined as:

$$x(k) = [\bar{x}\ \bar{y}\ \bar{z}\ \dot{\bar{x}}\ \dot{\bar{y}}\ \dot{\bar{z}}]^T \tag{1}$$

where $\bar{x}, \bar{y}, \bar{z}$ are the positions of the filter, and $\dot{\bar{x}}, \dot{\bar{y}}, \dot{\bar{z}}$ are the velocities respectively. A distance threshold of $1.5\ \mathrm{m}$ between the predicted position of the Kalman filter and new detections is used to remove spurious detections of people far away.

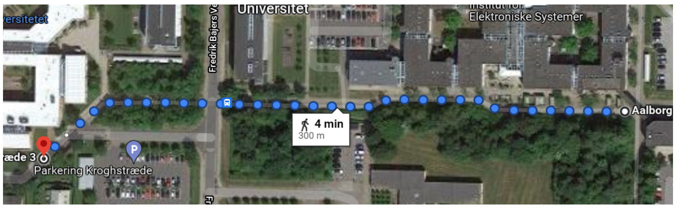

Fig. 3: Route used for testing. A footpath, approximately $300\ \mathrm{m}$ with an underpass. Located around 57.0136N, 9.9855E

The predicted state is given as:

$$x(k+1|k) = F \cdot x(k) + v(k) \tag{2}$$

$$x(k+1|k) = \begin{bmatrix} \mathbf{I} & \mathbf{I} \cdot T_s \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \cdot x(k) + v(k) \tag{3}$$

Where T_s is the sampling time and $v(k)$ the process noise with covariance $V(k) \in \mathbb{R}^{6\times 6}$ which is set to 1 along the diagonal, and 0.1 elsewhere.

The predicted measurement is given as:

$$y(k) = H(k+1)x(k) + w(k) \tag{4}$$

$$y(k) = \begin{bmatrix} \mathbf{I} & \mathbf{0} \end{bmatrix} \cdot x(k) + w(k) \tag{5}$$

Where the measurement noise $w(k)$ has covariance $W(k) \in \mathbb{R}^{3\times 3}$ defined as 10 along the diagonal, and 1 elsewhere.

$W(k)$ and $V(k)$ were chosen empirically. Note that the large scale of $W(k)$ is necessary in order to handle changes in speed of the tracked person.

C. Test setup

Three different tests were conducted to collect data on human behaviour when following Spot, so that a human behavior model could be created.

The tests were conducted on a 300-meter straight outdoor path on the campus of Aalborg University, see Figure 3. The entire test covered $600\ \mathrm{m}$ as the path was traveled both ways. The straight path was chosen to reduce variability in the tests.

Step response test: This test was done three times per participant using three different speeds: $1.4\ \mathrm{m/s}$, $1.0\ \mathrm{m/s}$ and $0.5\ \mathrm{m/s}$. There is a $10\ \mathrm{s}$ pause in between each step, which lasts $60\ \mathrm{s}$.

Step up test: A second test was performed where the robot starts with a step from $0\ \mathrm{m/s}$ to $0.8\ \mathrm{m/s}$, and increases after $60\ \mathrm{s}$ from $0.8\ \mathrm{m/s}$ to $1.6\ \mathrm{m/s}$.

Random staircase test: In this test the speed of Spot randomly increase or decrease with a step size of $\pm 0.25\ \mathrm{m/s}$. The time interval between the steps was randomized in the range $5\ \mathrm{s}$ to $15\ \mathrm{s}$. The participants were asked to evaluate the speed during each interval. The evaluation was based on a scale of 1 to 10 where 1 is too slow compared to the participant's preferred speed, 10 is too fast, thereby, making 5 the preferred following speed.

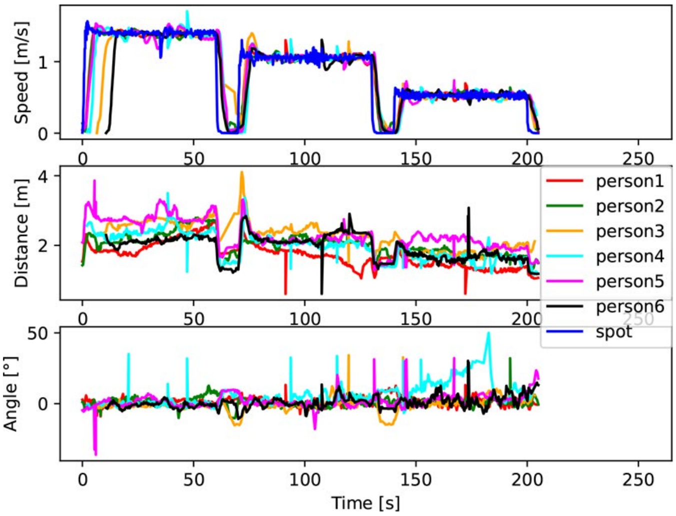

Fig. 4: Measured speed, distance and angle relative to Spot during the 'step response' test for the six participants. Speed is calculated from Spot's odometry. Distance is the Euclidean distance from the person to the center of Spot. Angle is measured relative to Spot's body, and is 0 when the person is directly behind Spot.

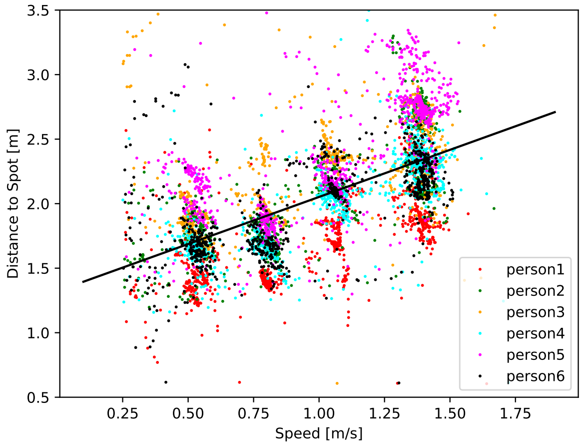

Fig. 5: Figure with all participants' speeds and distances to Spot for 'step response' and 'step up' tests. The data points are cropped to be above 0.25 m/s on the first axis and below 3.5 m on the second axis.

D. Results on preferred speed, distance, and angle

Six (n=6) people, unrelated to the project, participated in the initial tests. Their speed, distance and angle in relation to Spot, based on the output of the Kalman filter output, is shown in Figure 4.

The bottom plot of Figure 4 shows the angle of the participants in relation to Spot. With the exception of stops and turns, it was found that all six participants walked behind Spot, which is in line with what was found by Díaz-Boladeras et al. [11].

Figure 5 shows the six participants' distances to spot in relation to their speed in the three tests. Although noisy, there is an upward trend in preferred distances when asked to follow Spot at higher speeds. Their ideal distance to Spot (d_i) in relation to their absolute speed (V_p) is modelled with the affine function:

$$d_i(V_p) = 0.73 \cdot V_p + 1.32 \tag{6}$$

The results indicates that there is a shift of this line based on personal preference, however, more data is needed to generalize.

E. Preferred speed

The responses from each participant in the "random staircase" test can be seen in Table I. From the test participants' ratings of the different speeds, it was found that the preferred speed is around 1.4 m/s. This supports the results from [22], estimating the average preferred walking speed of healthy people to be 1.4 m/s. As Spot is not able to walk faster than 1.4 m/s, responses where the speed of Spot was uncomfortably fast for the participants were not achieved.

Table I: Scores from 1 to 10 from the six test participants during the 'random staircase' test, indicating the satisfaction of the speed. 1 being too slow, 5 being appropriate, and 10 being too fast.

Speed [m/s]	Test participant						
	1	2	3	4	5	6	Mean
0.25	1	1	1	1	1	1	1
0.50	1	1	2	1	2	2	1.5
0.75	3	1	3	3	3	3	2.7
1.00	4	2	3.5	3	4	3	3.25
1.25	5	4	5	4	5	4	4.5
1.40	5.5	4.5	5	5.5	5	4.5	5

F. Controller

Based on the above findings, it is believed that in order to be more socially acceptable, the robot should strike a balance between the preferred and actual speed while keeping an appropriate distance. We propose the following control law on the forward speed of Spot:

$$V_s(k+1|k) = V_s(k) + \alpha(k) + \beta(k) + \gamma(k) \tag{7}$$

Where $V_s(k)$ is the speed of Spot at time k, with:

$$\alpha(k) = (V_i - V_s(k)) \cdot K_{p1} \tag{8}$$

$$\beta(k) = (d_i(V_p(k)) - d(k)) \cdot K_{p2} \tag{9}$$

$$\gamma(k) = (V_p(k) - V_s(k)) \cdot K_{p3} \tag{10}$$

Where V_i is the ideal speed chosen to be 1.4 m/s based on the scores in Table I, $d_i(V_p(k))$ is equation 6, $V_p(k)$ is the speed of the person, and $d(k)$ is the distance between Spot and the person.

The α term pulls the speed of the robot, V_s, towards 1.4 m/s, V_i. It is implemented to have quicker acceleration during starts and not solely rely on the distance term (β). The β term regulates the distance, d, to maintain the ideal

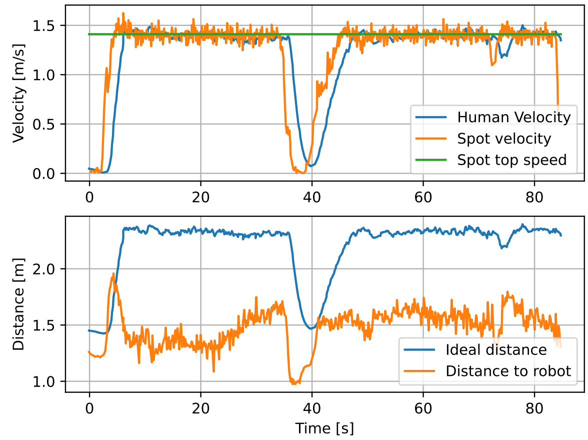

Fig. 6: Example of a participant following Spot, first at Spot's top speed ($1.4\ \mathrm{m/s}$), then after $30\ \mathrm{s}$, using the proportional controller. In this case, the participant pushes Spot to its maximum velocity, which the controller then maintains.

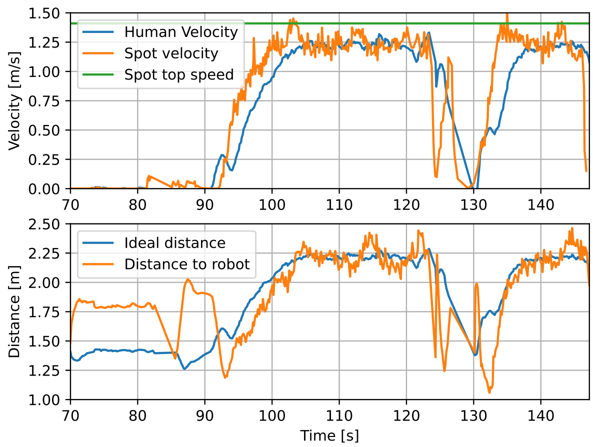

Fig. 7: Example of a participant following Spot using the proportional controller. It can be seen that the participant controls Spot to achieve a speed of approximately $1.2\ \mathrm{m/s}$ from $90\ \mathrm{s}$ to $120\ \mathrm{s}$. Unfortunately tracking failed for a period, leading Spot to stop, before starting moving again at $130\ \mathrm{s}$.

distance, d_i, which in turn is dependant on the actual speed of the person following. The γ term maintains the same speed as the person, V_p. Based on empirical testing, the following gains were chosen: $K_{p1} = 0.5$, $K_{p2} = 1.0$, $K_{p3} = 0.5$.

IV. Results

The system was tested by guiding high school students (n=8), visiting Aalborg University, between lecture rooms. The tests were performed outside. First, Spot was set to top speed for a period, then Spot was stopped, and set to use the control law from equation 7. The participant was asked afterwards to rate the two types of control, on a scale from 1 to 10, 10 being best. Figure 6 shows an entire run for a single participant.

The participant in Figure 6 rated the run with the proposed controller design higher than the one with constant speed. It is believed that this is due to the smoother velocity curve and the distance to Spot goes to $1.4\ \mathrm{m}$ and stays there more consistently.

Figure 7 illustrates a run with the proposed controller where a participant walks slower than $1.4\ \mathrm{m/s}$. Spot managed to keep the ideal distance to the participant, while following the speed of the participant. Here, the participant rated the proposed controller lower than the constant speed, however, this was mostly due to the abrupt stop around $125\ \mathrm{s}$.

Further tests (n=4), were conducted where the participants were told to walk at three different paces; slow, medium and fast. These tests found that when walking at a constant $0.5\ \mathrm{m/s}$ pace, the speed of Spot oscillates with up to $\pm 0.6\ \mathrm{m/s}$ at an approximate frequency of 0.2 Hz. This is mainly due to the α term, as it is far from the working point of $1.4\ \mathrm{m/s}$. Lowering the K_{p1} gain, however, results in a more sluggish and less responsive robot.

V. Discussion and Conclusion

Generally, the socially enabled robot design was better perceived as the open-loop controlled robot, which supports the use of proxemics in robot control.

As found by Díaz-Boladeras et al. [11], when guiding people individually, they stay mostly right behind the robot, and the average preferred speed was $\sim 1.4\ \mathrm{m/s}$, which also is the average preferred walking speed found by Samson et al. [22]. However, as was seen on Figure 4, one person drifted to the side of the robot when the walking speed was perceived as too slow.

During testing, Spot was forced to its maximum speed which removes the option of increasing speed to obtain the ideal distance. It would be interesting to replicate the initial tests on a mobile platform which is able to move at speeds that are uncomfortably fast for the participants, in order to model how people behave in these situations.

A key finding is, as one may expect, that the distance to the robot increases as the speed increases. I this work, the correlation of the distance and speed was approximated using a linear model, however, some participants had difficulties getting the robot to start moving when from a standstill, as they had to get close enough for the β term to become positive. A second order function could possibly prove to be a better fit than the linear model, using the data from participants standing still, since this would give an increased ideal distance at standstill, and a mostly linear region once the robot starts going. This, however, has been left for future work.

The sample size for the human behavior modelling was quite small. A larger sample size would aid in choosing a (possibly non-linear) model, as well as ensure that the observed behavior generalizes. We also observed that it seems that individual people have different model parameters, see Figure 5, which could possibly be learned on-the-fly.

References

[1] Federal Aviation Administration. *FAA National Forecast FY 2019-2039*. 2019.

[2] H. Scott Latimer, Hillary Gutknecht, and Kimmey Hardesty. "Analysis of Hospital Facility Growth: Are We Super-Sizing Healthcare?" In: *HERD* 1.4 (2008).

[3] Andrew Cave, Alethea Blackler, Vesna Popovic, and Ben Kraal. "Examining Intuitive Navigation in Airports". In: 2014.

[4] Saman Jamshidi, Mahnaz Ensafi, and Debajyoti Pati. "Wayfinding in Interior Environments: An Integrative Review". In: *Frontiers in Psychology* 11 (Nov. 2020). ISSN: 1664-1078. DOI: 10.3389/fpsyg.2020.549628.

[5] Xuan Tung Truong and Trung Dung Ngo. "Dynamic Social Zone based Mobile Robot Navigation for Human Comfortable Safety in Social Environments". In: *International Journal of Social Robotics* 8 (5 2016), pp. 663–684. ISSN: 18754805. DOI: 10.1007/s12369-016-0352-0.

[6] Thibault Kruse, Amit Kumar Pandey, Rachid Alami, and Alexandra Kirsch. "Human-aware robot navigation: A survey". In: *Robotics and Autonomous Systems* 61 (12 Dec. 2013), pp. 1726–1743. ISSN: 09218890. DOI: 10.1016/j.robot.2013.05.007.

[7] Anca D. Dragan, Kenton C.T. Lee, and Siddhartha S. Srinivasa. "Legibility and predictability of robot motion". In: IEEE, Mar. 2013, pp. 301–308. ISBN: 978-1-4673-3101-2. DOI: 10.1109/HRI.2013.6483603.

[8] Edward Twitchell Hall. *The Hidden Dimension*. Doubleday, 1966. ISBN: 0-385-08476-5.

[9] Sara Nielsen, Emil Bonnerup, Andreas Kornmaaler Hansen, Juliane Nilsson, Lucca Julie Nellemann, et al. "Subjective Experience of Interacting with a Social Robot at a Danish Airport". In: IEEE, Aug. 2018, pp. 1163–1170. ISBN: 978-1-5386-7980-7. DOI: 10.1109/ROMAN.2018.8525643.

[10] Sara Nielsen, Mikael B. Skov, Karl Damkjær Hansen, and Aleksandra Kaszowska. "Using User-Generated YouTube Videos to Understand Unguided Interactions with Robots in Public Places". In: *ACM Transactions on Human-Robot Interaction* (Aug. 2022). ISSN: 2573-9522. DOI: 10.1145/3550280.

[11] Marta Díaz-Boladeras, Dennys Paillacho, Cecilio Angulo, Oriol Torres, Jonathan González-Diéguez, et al. "Evaluating Group-Robot Interaction in Crowded Public Spaces: A Week-Long Exploratory Study in the Wild with a Humanoid Robot Guiding Visitors Through a Science Museum". In: *International Journal of Humanoid Robotics* 12 (04 Dec. 2015), p. 1550022. ISSN: 0219-8436. DOI: 10.1142/S021984361550022X.

[12] Wolfram Burgard, Armin B Cremers, Dieter Fox, Dirk Hähnel, Gerhard Lakemeyer, et al. "Experiences with an interactive museum tour-guide robot". In: *Artificial Intelligence* 114.1-2 (Oct. 1999), pp. 3–55. ISSN: 00043702. DOI: 10.1016/S0004-3702(99)00070-3.

[13] Sebastian Thrun, Maren Bennewitz, Wolfram Burgard, Armin Cremers, Frank Dellaert, et al. "MINERVA: A second-generation museum tour-guide robot". In: *Proceedings - IEEE International Conference on Robotics and Automation* 3 (Feb. 1999), pp. 1999 –2005.

[14] Rudolph Triebel, Kai Arras, Rachid Alami, Lucas Beyer, Stefan Breuers, et al. "SPENCER: A Socially Aware Service Robot for Passenger Guidance and Help in Busy Airports". In: *Field and Service Robotics: Results of the 10th International Conference*. Springer International Publishing, 2016, pp. 607–622.

[15] M.L. Walters, K. Dautenhahn, R. te Boekhorst, Kheng Lee Koay, C. Kaouri, et al. "The influence of subjects' personality traits on personal spatial zones in a human-robot interaction experiment". In: *ROMAN 2005. IEEE International Workshop on Robot and Human Interactive Communication, 2005.* 2005, pp. 347–352.

[16] Rachel Gockley, Jodi Forlizzi, and Reid Simmons. "Natural person-following behavior for social robots". In: *Proceedings of the 2007 ACM/IEEE Conference on Human-Robot Interaction* (Jan. 2007), pp. 17–24.

[17] Jonathan Huang, Vivek Rathod, Chen Sun, Menglong Zhu, Anoop Korattikara, et al. "Speed/Accuracy Trade-Offs for Modern Convolutional Object Detectors". In: IEEE, July 2017, pp. 3296–3297. ISBN: 978-1-5386-0457-1. DOI: 10.1109/CVPR.2017.351.

[18] Yu-Chen Chiu, Chi-Yi Tsai, Mind-Da Ruan, Guan-Yu Shen, and Tsu-Tian Lee. "Mobilenet-SSDv2: An Improved Object Detection Model for Embedded Systems". In: IEEE, Aug. 2020, pp. 1–5. ISBN: 978-1-7281-5960-7. DOI: 10.1109/ICSSE50014.2020.9219319.

[19] Adriana Carrillo Rios, Douglas Henke dos Reis, Rodrigo Mattos da Silva, Marco Antonio de Souza Leite Cuadros, and Daniel Fernando Tello Gamarra. "Comparison of the YOLOv3 and SSD MobileNet v2 algorithms for identifying objects in images from an indoor robotics dataset". In: *2021 14th IEEE International Conference on Industry Applications, INDUSCON 2021 - Proceedings* (2021), pp. 96–101.

[20] Jonathan Hui. *Object detection: speed and accuracy comparison*. URL: https://jonathan-hui.medium.com/object-detection-speed-and-accuracy-comparison-faster-r-cnn-r-fcn-ssd-and-yolo-5425656ae359.

[21] Christoph Schöller, Vincent Aravantinos, Florian Lay, and Alois Knoll. "What the Constant Velocity Model Can Teach Us About Pedestrian Motion Prediction". In: *IEEE Robotics and Automation Letters* 5 (2 Apr. 2020), pp. 1696–1703. ISSN: 2377-3766. DOI: 10.1109/LRA.2020.2969925.

[22] M. M. Samson, A. Crowe, P. L. de Vreede, J. A. G. Dessens, S. A. Duursma, et al. "Differences in gait parameters at a preferred walking speed in healthy subjects due to age, height and body weight". In: *Aging Clinical and Experimental Research* 13 (1 Feb. 2001), pp. 16–21. ISSN: 1594-0667. DOI: 10.1007/BF03351489.