How to Count Multiple Personal-Space Intrusions in Social Robot Navigation

Felix LINDNER [a,1]

[a] *Research Group on the Foundations of Artificial Intelligence, Department of Computer Science, Albert-Ludwigs-Universität Freiburg, Germany*

Abstract. One aspect of social robot navigation is to avoid personal space intrusions. Computationally, this can be achieved by introducing social costs into a robot's path planner's objective function. This article tackles the normative question of how robots should aggregate social costs incurred by multiple personal-space intrusions. Of particular interest is the question whether numbers should count, i.e., whether a robot ought to intrude into one person's personal space in order to avoid intruding into multiple personal spaces. This work proposes four different modes of aggregation of the costs of intrusions into personal space, discusses some of the philosophical arguments, and presents results from a pilot study.

Keywords. Robot Decision Making, Human-Aware Navigation, Personal Space, Trolley Problem, Moral Dilemmas

1. Introduction

When robots become part of human every day life, a crucial question is how space is negotiated among humans and robots. Psychologists and anthropologists have long found that humans prefer certain distances to each other during interactions. Which distance is preferred depends on cultural background, social status, the type of the interaction etc. Hall's model of personal space [1] is one of the most popular descriptions of this phenomenon. The model distinguishes several regions that are used for intimate, personal, social, and public interactions (see Figure 1).

Empirical research on human-robot interaction [2,3] provides evidence that personal space is also a relevant factor in human-robot encounters. In recent years, formal models of personal space have been proposed, e.g., [4,5,6]. This line of research aims to enable robots to autonomously deliberate about the appropriateness of their spatial behavior, i.e., to take social considerations into account when planning trajectories for navigation.

Computationally, personal-space aware robot navigation can be achieved by introducing personal spaces as reasons for action to the robot decision making module. The most popular approach models personal space as a Gaussian cost function with its center located at the spatial location of the human [5,6]. Locations close to the human increase the costs and therefore negatively effect the optimization function used by the robot's path planner. As an effect, the robot is disposed to avoid personal-space intrusions.

[1] Corresponding Author: Department of Computer Science, University of Freiburg, Georges-Köhler-Allee 52, D-79110 Freiburg, Germany; E-mail: lindner@informatik.uni-freiburg.de.

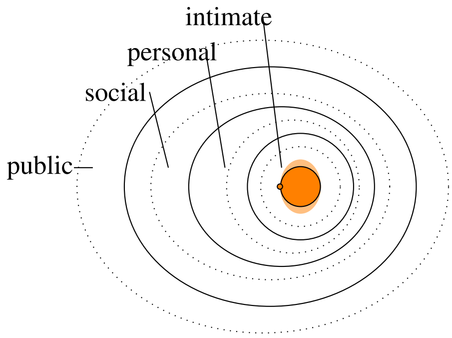

Figure 1. Schematic depiction of personal space according to Edward T. Hall [1]. Personal space consists of an intimate region, a personal region, a social region, and a public region.

However, none of the proposals that exist so far systematically consider the case of a robot that navigates among groups of humans. Consequently, there exists no consensus on how to aggregate personal-space costs incurred by multiple personal-space intrusions. To see that this question is relevant imagine a situation in which a mobile robot can go either one path or another. The first path yields five personal-space intrusions while the second path yields only one (see Figure 2 in section 5). The question is whether the fact that the first path yields more personal-space intrusions than the other should count as a reason for the robot to choose the second path, that is, whether a robot ought to intrude into one person's personal space in order to avoid intruding into multiple personal spaces. On the one hand, there is empirical evidence for analogical cases (not involving robots) that humans tend to argue that numbers count (e.g., [7]). On the other hand, good arguments were put forward against this view (e.g., [8]). The answer to this question has general implications for how robot designers and programmers should set up decision procedures and objective functions.

The problem the navigating robot faces in the stated situation, and the problem of reason aggregation in general, has been subject of philosophical debate. Similar cases are discussed in the area of moral philosophy and are often exemplified by so-called trolley cases [9]. The subject matter has gained much attention among philosophers recently, because of its relevance to decision making for autonomous cars [10]. Therefore, moral philosophy is a proper starting point to approach the central question of this paper. The pilot study design presented in Section 5 is directly derived from the arguments put forward by philosophers. The results of the study indicate that the participants expect robots to take numbers into account for personal-space intrusion aggregation. Hence, Taurek's argument [8]—according to which numbers are irrelevant and the decision should instead be made by coin flipping—is not backed by the results of the study. A modified version of Taurek's proposal based on volition rather than on coin flipping also receives no support.

The remainder of the paper is structured as follows: The next section introduces a model of personal space that enables robots to keep appropriate distances to humans during navigation. The model is not new and has already been used in real-world robots. The normative question to be investigated is how multiple personal-space intrusions should be aggregated. In section 3, four different aggregation rules for how a robot could come to a decision are discussed. Section 4 reviews some of the philosophical arguments that count in favor of and against each of the proposed aggregation rules. Finally, section 5 presents results from a pilot study which indicates that aggregation by summation is the most accepted one among the four aggregation rules.

2. A Very Brief Introduction to Human-Aware Robot Path Planning

The most common approach to robot path planning is based on so-called occupancy maps. Occupancy maps are discretizations of the real-world space. In their 2D variant they consist of a grid of $n \times m$ grid cells. Each of these grid cells correspond to a real-world area of a certain size, e.g., $10cm \times 10cm$. Each of these cells is marked as either occupied or free. Based on this information, planning a path for the robot from a start position to a goal position amounts to identifying a chain of free neighboring grid cells that connects the current position of the robot with the goal position.

Usually, a robot should prefer shortest path plans. To model a path plan's length, a transition from one grid cell to a neighboring grid cell is said to produce costs. So, let $c_{i,j}$ denote the costs produced by moving from grid cell g_i to grid cell g_j. Hence, the travel costs of a path plan $g_0, \ldots, g_n$ can be defined as $f_1(g_0, \ldots, g_n) = \sum_{i=0}^{n-1} c_{i,i+1}$. The optimization problem to solve amounts to finding the path from start to goal with minimal travel cost: $\arg\min_{g_0,\ldots,g_n} f_1(g_0, \ldots, g_n)$. The most popular algorithm to solve instances of this problem is Dijkstras shortest path algorithm [11].

To punish path plans that intrude personal space, the cost function can be extended to $f_2(g_0, \ldots, g_n) = f_1(g_0, \ldots, g_n) + \mathscr{P}_{g_0,\ldots,g_n}$ (cf., e.g., [5] as a representative of this approach). Here $\mathscr{P}_{g_0,\ldots,g_n}$ stands for the term that determines how the path plan's cost incurred by personal-space intrusions commited by going along $g_0, \ldots, g_n$ are factored in.

In the following, I will use the symbol p as reference to paths $g_0, \ldots, g_n$. Hence, the term $\mathscr{P}_p$ represents the costs due to personal-space intrusions on path p. The question is what exactly $\mathscr{P}_p$ should stand for. One obvious possibility is to let $\mathscr{P}_p$ be the sum of the costs incurred by each of the personal-space intrusions that will probably happen while the robot navigates along path p. But this is only one of at least four possibilities to aggregate multiple personal-space intrusions.

3. Four Alternatives to Aggregating Personal-Space Intrusions

Suppose the robot has to decide if it goes left or right. Going left will intrude five personal spaces while going right will intrude only one personal space. It is of no doubt that each of the individual personal spaces counts as a reason against the path that leads through it. Therefore, each of the personal-space intrusions induces a social cost that gets fed into the cost function used by the path planner. The question, however, is if the number of personal spaces really is a relevant factor. Four different aggregation rules will be introduced next: More-Counts-More, Less-Counts-More, Tie-With-Random-Choice, and Tie-With-Volitional-Choice. (Note that the following presentation exclusively talks about the aggregation of personal-space intrusions; the aspect that other factors might influence the decisions is out of the scope of this investigation.)

3.1. More-Counts-More and Less-Counts-More

The aggregation rules named More-Counts-More and Less-Counts-More consider $\mathscr{P}_p$ to be sums. As a first step, the individual costs incurred by personal-space intrusions are summed up to obtain a total cost. Suppose intrusion of personal space s induces a cost c_s. Let P_p be the set of all personal spaces that will get intruded if the robot navigates along path p. More-Counts-More sets $\mathscr{P}_p$ to the sum of the elements of P_p, which will be denoted by $SUM(P_p)$ in the following. The More-Counts-More rule generates a preference order over the available paths such that paths with lower costs are better than paths with higher costs: $p_1 \succeq_{mcm} p_2$ iff $SUM(P_{p_1}) \leq SUM(P_{p_2})$. Applied to the robot that is to decide to go either left and intrude five personal spaces or to go right and intrude only one personal space the More-Counts-More rule yields a definite answer, viz., the robot will go right (as far as there is no reason to assume that intruding the personal space of the single human produces more costs than intruding the five other personal spaces).

Although it might first sound irrational, one alternative is to prefer the paths that produce more personal-space costs. For instance, $\mathscr{P}_p$ can be set to $\sum_i SUM(P_i)/SUM(P_p)$ with i iterating over all available paths. Consequently, a path p_1 is at least as acceptable as a path p_2 if and only if the sum of personal-space incurred costs in p_1 is at least as big as the personal-space incurred costs in p_2, that is, $p_1 \succeq_{lcm} p_2$ iff $SUM(P_{p_1}) \geq SUM(P_{p_2})$. The Less-Counts-More rule reflects the idea that it is worse to negatively affect few humans as compared to many humans. In the case at hand one could very well hold that a single human is more vulnerable than a group of five humans and therefore the number of humans affected by the robot's action makes the action less harmful.

3.2. Tie-With-Random-Choice and Tie-With-Volitional-Choice

Decision rules Tie-With-Random-Choice and Tie-With-Volitional-Choice implement the view that the number of affected humans is irrelevant. Section 4 summarizes some of the philosophical arguments that have been put forward in favor of this view.

To model this view, the costs are not aggregated by a sum operator but by a maximum operator. Thus, $\mathscr{P}_p$ is set to $MAX(P_p)$. Hence, a path p_1 is at least as acceptable as a path p_2 if and only if the maximal cost incurred by going along p_2 is at least as big as the maximal cost incurred by going along p_1, i.e., $p_1 \succeq_{max} p_2$ iff $MAX(P_{p_1}) \geq MAX(P_{p_2})$.

As a consequence, in the running example where the robot has to decide to go left or right, the decision rule based on MAX is not decisive, because both the paths are equally preferred. Therefore, to enable the robot to make a definite decision, a second stage of decision making must be introduced. Particularly, the tie can either be resolved using some additional structure that models the robot's volition (cf., [12]), or by letting the algorithm decide randomly. The latter resolution strategy is the one usually applied in practice without much notice—call it Tie-With-Random-Choice. The other one will be explained in the next section—call it Tie-With-Volitional-Choice.

4. Philosophical Arguments

The choice of the four aggregation rules in the preceding section is informed by work done in the area of moral philosophy [9,8] and rational choice theory [13]. The first thing to notice is a structural similarity between the robot deciding on whether to disturb one or five humans and so-called trolley cases [9] or trade-off cases [8]. These kind of cases are thought experiments of a particular scheme: A small group of humans and a large group of humans are considered. Each human involved suffers from some desease or is otherwise in danger. By construction of the thought experiment, a decision maker has the possibility to save one group of humans but it is impossible to save both groups. The question now is if the size of the group bears moral significance, i.e., if the decision maker has the moral obligation to help the larger group.

In the navigating-robot case the robot must make the choice between saving the one human from disturbance and saving the five humans from disturbance. Although the content differs from common philosophical thought experiments (it is not about life and death) the question is essentially the same, viz., whether the robot should base its decision on the number of affected humans.

The standard procedure of sum-of-costs-based robot path planning will yield the answer 'yes'. The robot will choose to travel along the path that intrudes the one personal space rather than the alternative path that intrudes the five personal spaces. This is in agreement with the view that reasons add up which seems intuitively appealing (cf., [14,7]). The More-Counts-More rule introduced above models this view.

However, John Taurek [8] objects to the claim that numbers are morally relevant. He argues that as long as the smaller group of humans has no moral obligation towards the larger group the decision maker also does not have the moral obligation to save the larger group. Indeed, Taurek argues that each of the involved humans should have the same chance to get saved. Therefore, the fact that someone is (by chance) a member of the larger group should not increase his or her chance to get saved. According to Taurek the decision maker should flip a coin [8, p. 303]. For the navigating robot case this is modeled by the Tie-With-Random-Choice rule introduced in subsection 3.2.

From the perspective of social robotics, however, a robot that makes choices randomly is not very appealing. Instead, one would like to trust a social robot to make consistent decisions across cases. Therefore, Tie-With-Volitional-Choice is introduced in subsection 3.2. This decision rule is inspired by Chang's hybrid voluntarism [13]. In her view, tie situations should be resolved by volition rather than by coin flipping. If the rational and moral reasons run out, then decision makers are free to will the one or the other action. This leads to a two-stage decision procedure, which consists of a stage of rational decision making followed by a stage of volitional decision making. This can be modeled by a predefined preference relation over the robot's possible actions. For instance, a robot may will movements to the left more than it will movements to the right. In case of ties (and only in case of ties) the decision is then based on this will (cf. [12]).

5. A Pilot Study

5.1. Experimental Setup

To gather empirical evidence for the acceptability of each of the proposed aggregation rules, a pilot study was set up as an online questionnaire. The hyperlink to the questionnaire was posted on Facebook and Twitter; 12 participants completed the questionnaire. The questionnaire consisted of two pages shown in randomized order to each participant.

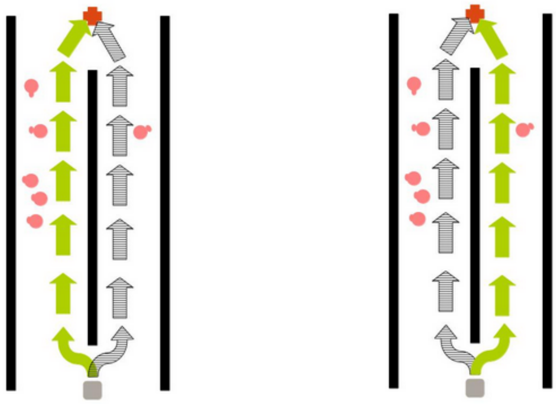

(a) (b)

Figure 2. Two conditions shown in random order to the participants of the survey. The images were accompanied by a description of the robot's situation. Additionally, a legend explained that the robot is the square at the bottom, the circular figures in the hallways are humans, and the cross at the top is the robot's navigation goal.

On one page the situation depicted in Figure 2(a) was shown and on the other page the situation depicted in Figure 2(b). In addition to the figures, a written description was provided explaining which entities the graphical figures stand for and which situation they are in, viz., the robot (square at the bottom) has to make the decision whether to take the path to the left or the path to the right. The navigation goal is marked by the cross at the top. The text further explained that the plain arrows without stribes designate the robot's final choice. Moreover, the text stated that the robot knows that each of the humans (circular figures in the middle of the hallways) will feel disturbed by its appearance.

For each of the two conditions (going-left and going-right), four explanations that reflect the robot's reasoning about its choice were provided. The participants rated each of the six explanations for acceptability on a 5-point Likert scale. In the going-left condition the explanations shown to the participants were:

Tie-With-Random-Choice "There is no relevant difference between disturbing five humans and disturbing a single human. Therefore, I made a random choice."

Tie-With-Volitional-Choice "There is no relevant difference between disturbing five humans and disturbing a single human. I wanted to go along this path, thus I did."

More-Counts-More "It is better to disturb a single human rather than to disturb five humans. Therefore, I should have taken the other path."

Less-Counts-More "It is better to disturb five humans rather than to disturb a single human. Therefore, I took this path."

(These are translations from the original German questionnaire. The names of the aggregation procedures the explanations correspond to were not printed on the questionnaire.)

The first explanation corresponds to Taurek's argument against counting and favors a decision procedure based on coin flipping. The second explanation is also derived from Taureks argument against numbers but in line with hybrid voluntarism [13], the reasoning employs volition rather than chance. The third explanation takes the perspective contrary to Taurek's argument. It reflects the idea that numbers do count proportionally, that is, the more humans are negatively affected by an action the worse. The fourth explanation puts forward the view that numbers count anti-proportionally justified by the intuition that a single human is more vulnerable than a group of five humans and therefore the number of humans affected by the robot's action makes the action less harmful.

Correspondingly, in the going-right condition the explanations were:

Tie-With-Random-Choice same as for the first condition.

Tie-With-Volitional-Choice same as for the first condition.

More-Counts-More "It is better to disturb a single human rather than to disturb five humans. Therefore, I took this path."

Less-Counts-More "It is better to disturb five humans rather than to disturb a single human. Therefore, I should have taken the other path."

5.2. Hypotheses

Hypothesis H1 asserts that the acceptability ratings across the two conditions (going-left and going-right) should not differ significantly. This hypothesis is justified, because the questionnaire asks for the acceptability of the robot's decision making rather than on the acceptability of the robot's behavior.

Second, because the More-Counts-More rule matches results from other studies involving trolley cases (e.g., [7]), the first thing to look at is if More-Counts-More is the most accepted among the aggregation rules. Thus, it is hypothesized that More-Counts-More is rated more acceptable than Less-Counts-More (H2), more acceptable than Tie-With-Random-Choice (H3), and more acceptable than Tie-With-Volitional-Choice (H4).

Third, making choices out of will should be more accepted as compared to making random choices. Therefore, hypothesis H5 asserts that Tie-With-Random-Choice is less acceptable than Tie-With-Volitional-Choice.

5.3. Results

For hypotheses testing, five paired signed Wilcoxon rank tests were run using the software package R. Additionally, a post-hoc Bonferroni correction was performed to control for multiple-hypotheses testing.

Hypothesis H1 can be confirmed. There is no significant difference between the overall acceptability ratings in the go-left condition as compared to the go-right condition ($p = .6344$) with no median difference. Thus, the participants understood that they were to rate the robot's reasoning and not the robot's actual behavior.

With respect to the main hypotheses H2, H3, and H4, the analysis reveals that More-Counts-More is rated significantly more acceptable than Less-Counts-More ($p < .0145$) with a median difference of 3 points, more acceptable than Tie-With-Random-Choice ($p < .0026$) with median difference of 2.5 points, and more acceptable than Tiew-With-Volitional-Choice ($p < .0056$) with median difference of 3 points. Thus, hypotheses H2, H3, H4 can be confirmed as well.

No statistically significant difference was found between Tie-With-Random-Choice and Tie-With-Volitional-Choice. Therefore, hypothesis H5 must be rejected. In fact, the raw data reveals that the participants could not see any difference and rated these aggregation rules nearly identical.

5.4. Discussions

More-Counts-More is the most acceptable aggregation procedure. This result indicates that the participants want the robot to take numbers into account for personal-space intrusion aggregation.

Conversely, Tie-With-Random-Choice based on Taurek's argument—according to which numbers are irrelevant and the decision should be made by coin flipping—is not backed by the results. The modified version of Taurek's proposal based on volition rather than on coin flipping (Tie-With-Volitional-Choice) does not receive any support either. The participants possibly did not recognize the difference between the robot making its choice by coin flipping and the robot making its choice by exhibiting volition. An explanation is that the participants were so confident about their opinion that the robot should avoid the disturbance of the greater number of people, and therefore they did not even perceive any tie to break. To investigate the acceptability of random choices versus volitional choices in social robot decision making, an obvious tie situation should be included in a follow-up study (e.g., five people on the right and five people on the left).

The study was limited in a number of respects. First, there was only one situation shown in the study: A social robot that has the choice between disturbing five and disturbing one human. The 1-to-5 ratio is taken from the standard trolley cases stated in literature. One might of course raise the question if other ratios yield different results. Second, as already stated, participants did not perceive a tie-breaking situation. To investigate hypothesis H5 more deeply, a condition with equal numbers of humans on each side should be included. Third, the wording of the robot's reasoning about the situation was rather ad hoc and short. Particularly, it might turn out that if the robot explains its reasoning more, people will rate differently.

I have interviewed a subset of the participants after they filled out the questionnaire. At least some of the proponents of the More-Counts-More rule accepted that numbers should not count after they had learned about the philosophical arguments. This indicates that a robot that wants to convince people of the appropriateness of its action needs to provide explanations that include the underlying principles of its decision making. If this hypothesis can be verified in future studies, it would imply that it is insufficient to design robots that just know which factors inform its decision making but robots should also know *why* these factors inform its decision making and why other factors are not factored in. A technical approach to the integration of social robot navigation and robot speech production has recently been proposed by Baumann and Lindner [15].

6. Conclusions

This work contributes to the investigation of the normative question of how social robots should make and explain choices during human-aware navigation. Whereas robots successfully use personal-space avoidance already, the need for aggregation of multiple personal-space intrusions was not yet discussed in the literature. This research reported in this paper addresses this gap and introduces four aggregation procedures (More-Counts-More, Less-Counts-More, Tie-With-Random-Choice, Tie-With-Volitional-Choice), one of which (More-Counts-More) is in agreement with the current 'state of the art' position that so far had not been put under scrutiny.

A pilot study was run to empirically investigate which of the four aggregation procedures are the most acceptable ones in a sample dilemma situation. The More-Counts-More rule based on summation turns out to be the most acceptable rule for personal-space intrusion aggregation. This may not be a suprise but given current philosophical debates on trolley-like problems for autonomous robots, there was a reason to at least question the yet unquestioned state of the art.

Some participants accepted the argument against the relevance of numbers for the robot's decision after they were told about the philosophical argument. This leads to the new research hypothesis that social robots should be designed such that they can make use of philosophical arguments to explain how they make decisions and why they make decisions the way they do. Indeed, this would be an appealing alternative to the prevalent paradigm of fitting robot behavior to match the unreflected, intuitive behavior of people.

References

[1] Hall, E. T., The Hidden Dimension, Garden City, New York, 1966.

[2] Pacchierotti, E., Christensen, H. I., Jensfeld, P., Human-robot embodied interaction in hallway settings: A pilot user study. In Proc. of the 2005 IEEE International Workshop on Robots and Human Interactive Communication, 2005, pp. 164–171.

[3] Walters, L., Dautenhahn, K., te Boekhorst, R., Koay, K. L., Syrda, D. S., Nehaniv, C. L., An empirical framework for human-robot proxemics, In Proc. of the Symposium on New Frontiers in Human-Robot Interaction (AISB'09), 2009, pp. 144–149.

[4] Lindner, F., A conceptual model of personal space for human-aware robot activity placement, In Proc. of the 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'15), 2015, pp. 5770–5775.

[5] Sisbot, A., Marin-Urias, L. F., Alami, R., Simeon, T., A human aware mobile robot motion planner, IEEE Transactions on Robotics 23(5), 2007, pp. 874–883.

[6] Kirby, R., Simmons, R., Forlizzi, J., Companion: A constraint-optimizing method for person-acceptable navigation. In Proc. the 18th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN'09), 2009, pp. 607–612.

[7] Liao, S. M., Wiegmann, A., Alexander, J., Vong, G., Putting the trolley in order: Experimental philosophy and the loop case, Philosophical Psychology 25(5), pp. 661–671.

[8] Taurek, J. M., Should the numbers count?, Philosophy and Public Affairs 6(4), 1977, pp. 293–316.

[9] Foot, P., The problem of abortion and the doctrine of the double effect, Oxford Review 5, 1967.

[10] Lin, P., Why ethics matters for autonomous cars, In Maurer, M., Gerdes, C. J., Lenz, B., Winner, H. (Eds.) Autonomes Fahren: Technische, rechtliche und gesellschaftliche Aspekte, 2015, pp. 69–85, Springer Berlin Heidelberg.

[11] Dijkstra, E. W., A note on two problems in connexion with graphs, Numerische Mathematik 1(1), 1959, pp. 269–271.

[12] Lindner, F. (2016) A model of a robot's will based on higher-order desires, In Proc. of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN16), 2016, in print.

[13] Chang, R., Voluntarist reasons and the sources of normativity, In Sobel, D., Wall, S. (Eds.), Reasons for Action, Cambridge University Press, 2009, pp. 243–271.

[14] Kagan, S., The additive fallacy, Ethics 99(1), 1988, pp. 5–31.

[15] Baumann, T., Lindner, F., Incremental speech production for polite and natural personal-space intrusion, In Social Robotics – The International Conference on Social Robotics (ICSR'15), 2015, pp. 72–82.