Where can I park my robot? Modeling out-of-the-way parking spots in the home using room geometry

Dylan F. Glas[1] and William D. Smart[1,2]

***Abstract*— For social robots operating in home environments, identifying appropriate parking locations which are "out of the way" is a challenging and multi-faceted problem. This paper proposes a solution to one core aspect of that problem, specifically a model for estimating locations where the robot may block walking paths through narrow spaces. For generality, this model assumes no *a priori* knowledge about user behaviors or semantic features in the space, and is derived purely from spatial geometry based on a standard 2D occupancy map. An experimental validation based on self-reported parking spot preferences from long-term robot users demonstrates that the proposed model captures 74% of user preferences, outperforming a naive baseline condition in selecting user-preferred parking spots. The proposed method provides a basis for estimating socially-appropriate parking locations for robots operating in the home or other unstructured social spaces and serves as a foundation for developing more sophisticated parking spot preference models in the future.**

I. Introduction

As interactive mobile robots are being developed and tested for use in the home in academic studies [1]–[4], commercial home robot products are beginning to appear in industry as well. Whereas cleaning robots such as Roomba [5] have been widespread for years, social robots such as Kuri [6] and Astro [7] have more recently been developed to navigate and socially interact in home environments.

In this work, we focus on the challenge of determining acceptable parking spots in the home. Although it is typical for a robot to have an assigned dock for charging and storage, social robots like Astro (Fig. I) face the additional challenge of determining socially-acceptable places to park or wait in other rooms where they can be more accessible to users.

In particular, this research was conducted to enable two behaviors for the Astro robot: a "hangout" behavior, in which the robot will try to park near users in order to be accessible and encourage interaction, and a "self-park" behavior, where the robot searches for a nearby out-of-the-way location without considering user co-presence. Both problems require a model of good parking spots. Although work such as [8] has investigated ways to collect user preferences about parking spots, it is still necessary to develop computational models to select parking spots which conform to those preferences.

This is not a simple problem, as there is a high density of social and functional meaning associated with objects and locations in the home. Use of open floor space follows complex rules, with some areas off-limits due to high pedestrian traffic, others blocking access to important features in the home, and others impeding social activities in the space.

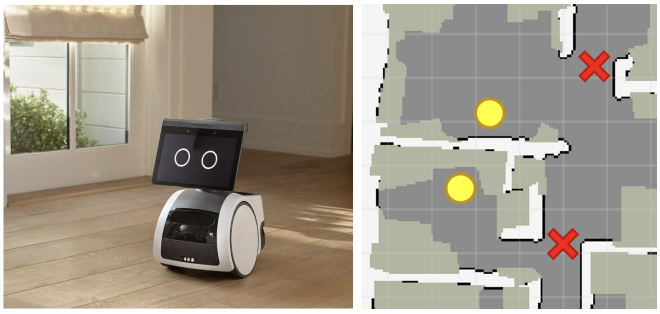

Fig. 1. Left: Amazon's Astro robot. Right: Examples of good and bad parking spots on an occupancy map, according to user annotations. Yellow circles indicate good spots, and red crosses show bad spots.

As a starting point for addressing this problem, we argue that the most fundamental factor in determining acceptability of a parking spot is whether it blocks any of the main walking paths in a space. If these walking paths can be estimated, then it will be possible to create a first-order estimate of parking spot availability. Additional information about the environment, if available, may later be combined with the proposed method to form a more sophisticated model.

In this work, we propose a method for using the geometry of the environment to infer primary pedestrian walking paths. Along these paths, width constraints including robot size and a human proxemics model are applied to to construct a parking spot availability map indicating where a robot would fully or partially obstruct those walking paths.

This approach is not specific to Astro or dependent on any specific sensing capabilities or additional knowledge of the environment such as object labeling or observations of trajectory patterns, and it relies only on a standard 2D occupancy map, which is available for most mobile robots.

Whether used alone or in combination with other cost functions, the proposed method provides a basis for the selection of out-of-the-way parking spots in different rooms in the home. This method has been patented in the US [9].

II. Related Work

Research into human-aware navigation often focuses on dynamic collision avoidance or human-relative positioning such as approaching or establishing social formations, rather than the determination of parking spots. Some surveys of human-aware robot navigation include [10]–[12].

[1]Both authors are with Amazon Lab126 Consumer Robotics, 905 11th Ave, Sunnyvale, CA, USA {dgglas, robobill}@amazon.com

[2]William D. Smart is also affiliated with the Collaborative Robotics and Intelligent Systems (CoRIS) Institute at Oregon State University.

A. Pedestrian trajectory modeling

Different approaches are taken towards modeling pedestrian trajectories, including empirical data collection and social force models of pedestrian flow. A survey of human motion trajectory prediction can be found in [13].

Many studies have modeled the typical paths taken by users in a space by collecting data sets of trajectory observations in an environment, building empirical pedestrian motion models to use as a basis for robot planning.

Kanda et al. built a model of observed pedestrian trajectories to select a robot's patrol path to avoid obstructing busy users [14]. Similar approaches have been used by Noguchi et al. in robot path planning [15], by Kidokoro et al. to improve "walking comfort" for pedestrians around a robot [16], and by Liu et al. for learning interactive robot behaviors from data [17]. Most similar to our work, Kitade et al. aggregated observed trajectories to develop a model of acceptable waiting spots for a robot in a shopping mall [18].

These techniques used statistical models of pedestrian trajectories observed by ambient sensor systems, which cannot be easily obtained in a home environment. Our approach assumes that a rich trajectory history is not initially available.

B. Proxemics

Much of the proxemics work in human-robot interaction has examined questions such as where a robot should position itself when interacting with a person [19] or how to comfortably navigate in the presence of humans, such as when passing people in a hallway [20]. Other studies have investigated path planning around individual pedestrians [21], approaching moving people with the intention of interaction [22], or positioning in group interactions [23].

Although the current work broadly concerns robot positioning related to humans, it differs from these studies in that its focus is not on considering the immediate locations of users, but rather on minimizing the risk of obstructing their paths at some time in the future.

C. Geometry-based approaches

Other work has modeled pedestrian walking paths through architectural spaces using the Medial Axis Transform [24], [25]. Our proposed method builds on this technique.

A related topic is the modeling of "choke points" in an environment in the field of gaming. Perkins et al. proposed an algorithm for identifying choke points in an unstructured environment map [26]. However, that work focuses on identifying individual choke points, such as doorways, which segment the spaces in the environment, whereas our work focuses on building a continuous costmap, *e.g.* modeling the narrow space along an entire hallway or passage.

III. Modeling parking spots

The goal of this work is to generate a parking spot availability map, where higher scores indicate spots where the robot is less likely to physically obstruct pedestrians.

From early user feedback we learned that a main driver of user frustration for poorly selected parking spots was Astro being "in the way". Although this phrase can admit interpretations such as blocking cabinets or appliances or lines of sight, we consider the main definition of "in the way" to specifically mean obstruction of walking paths.

Our strategy for finding out-of-the-way parking spots thus consists of two main steps. The first step is to establish a model of likely pedestrian walking paths in the environment. The next step is to geometrically model whether a given parking spot potentially obstructs any of these paths, in which case it can be considered to be "in the way".

A. Walking path estimation

The problem of determining pedestrian paths in the home is somewhat difficult, due to the large number of potential goals within the environment. Whereas large spaces like shopping malls or museums are designed around clear start and end goals such as shop entrances, exhibits, and external doors, the home contains many start and end goals such as sofas, desks, bookshelves, kitchen cabinets, closets, and so on. This makes path determination challenging, as not only does detection of these goal points require rich semantic understanding of home features, but it is also difficult to exhaustively recognize the large set of object types for different specific home situations (consider a piano, a cat box, a tool chest, a trampoline, a billiard table, and so on).

A simpler approach to modeling pedestrian paths is to use the skeleton, or medial axis, of a geometric representation of the navigable space in the home. This approach has been used in modeling pedestrian paths through architectural and urban spaces in works such as [24], [25]. As Fu et al. explain, "the medial axis-based graph network model is highly suitable for indoor navigation and route planning" due to its succinctness, consistency with common human cognition, and usefulness for spatial queries [24].

The skeleton can be defined as the "loci of centers of bitangent circles that fit entirely within the foreground region" [27], and it can be extracted via the Medial Axis Transform (MAT), a method proposed by Blum [28] as a way to capture the shape of a polygon for object representation. The medial axis $M(G)$ consists of points q within a simple polygon G where at least two boundary points are equidistant and closest to q [29]. The MAT is a common tool in image processing, as it is able to represent simplified shapes from raster data in a way robust to pixel-by-pixel boundary noise [30], [31].

Thus, the first step of our proposed algorithm uses the MAT to approximate a model of likely pedestrian paths through the home based only on geometric features.

B. Human and robot diameters

Once the main paths through the home have been determined, the next step of the analysis is to compute whether a robot would obstruct a human if parked near that path. To characterize this problem mathematically, we propose a simple model in which the robot and human are approximated as circles on the 2D plane. The model incorporates two parameters: d_R represents the physical diameter of the robot, and d_H represents the diameter of a human, that is, the

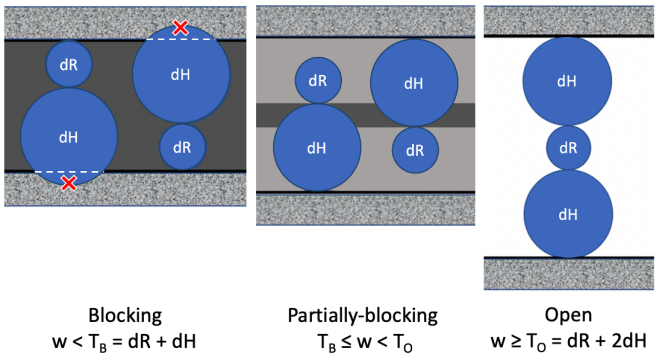

Fig. 2. Illustrations of blocking, partially-blocking, and open spaces based on gap width w. White: robot can park freely. Dark gray: robot cannot park. Light gray: allowable but non-preferred.

minimum amount of space a human would be comfortable occupying while passing by the robot. This is a somewhat subjective quantity which is a function of both a person's physical size and their proxemic preferences, which will be discussed in further depth in Section VI-B.

Because the MAT captures the minimum distance to equidistant obstacles along the skeleton line, it can be interpreted as representing the width of a gap between obstacles. Given d_R and d_H, it is then possible to characterize obstruction potential for gaps of differing widths according to the following thresholds, as illustrated in Fig. 2.

$$\text{Blocking threshold:} \quad T_B = d_R + d_H \tag{1}$$

$$\text{Open threshold:} \quad T_O = d_R + 2 \cdot d_H \tag{2}$$

Blocking For a gap of width $w < T_B$ (Eq. 1), there is no robot position which permits a human to pass, so these gaps are considered *blocking* and should be fully avoided.

Open By contrast, a gap of width $w \geq T_O$ (Eq. 2) is considered *open*, as no matter where the robot parks in the gap, the human can pass on the other side.

Partially Blocking For gap widths that lie between these two extremes, that is, where $T_B \leq w < T_O$, the ability of a human to pass by the robot is conditional. If the robot is close to one side or the other, a human can pass by, but a robot in the middle would be an obstruction. In this case, the center area of the gap should be marked as *blocking*. We additionally mark the edge areas outside of this center strip as *partially-blocking*, to indicate that they are permissible but not preferred for parking since the space is narrow.

The width of this *blocking* area in the center of the gap is a function of the overall gap width. Strictly speaking, this area has a width of $2 \cdot dH - w$, but for navigation we apply an inflation of $dR/2$ to the walls in order to model the position of the robot center, so the same inflation must be applied to the center strip, resulting in a strip width of $dR + 2 \cdot dH - w$, which is equal to $T_O - w$.

IV. Algorithm

The algorithm itself is implemented as a series of image-processing steps applied to a raster image of the robot's occupancy map. The OpenCV and scikit-image Python libraries were used for image-processing functions [32], [33].

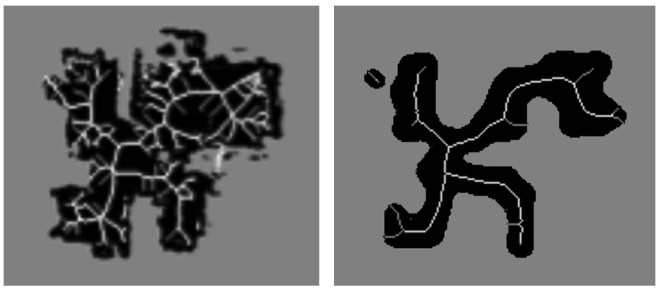

Fig. 3. Skeleton extraction from a raw, noisy map with jagged edges (left), compared with skeleton extraction from a smoothed version of the same map (right). The smoothed map enables a simpler skeleton to be extracted, which better represents the overall shape of the space.

The processing steps can be summarized as follows, with the details described in the following subsections:

1) Smooth map (IV-A)
2) Compute medial axis transform (IV-B)
3) Refine skeleton (IV-C)
4) Construct availability map (IV-D)
5) Apply smoothing (IV-E)

The output is a raster map containing parking availability scores for each navigable cell in the occupancy grid.

A. Smooth map

The first steps in the algorithm aim towards extracting a clean skeleton from the occupancy grid. When occupancy maps have noisy and jagged edges, it is difficult to extract a clean skeleton of the region. Since this noise does not often represent important features, the choice was made to trade off some map accuracy for clean structural representation.

The smoothing is performed in two steps. First, a morphological "closing" operation is performed. This has the effect of clearing out small specks of noise from the open regions in the occupancy map. Next, a Gaussian blur is applied to the entire image, blurring the edge around the open region, and a binary thresholding is applied to the resultant blurred image to restore a sharp edge. This results in smoother curves around the edges and a cleaner skeleton, as shown in Fig. 3. This tradeoff can be managed by tuning the size of the morphological kernel. However, important details can be lost. Notably, thin walls and similar features can be erased.

In order to restore some features erased by the smoothing step, we take the thresholded occupancy grid map and mask off the areas that are covered by the smoothed map (leaving the occupied cells in the thresholded map that are in the open areas of the smoothed map). Then we perform a connected components analysis of these remaining features, and restore any features over a given size (currently 10 pixels or larger) to the smoothed map. This effectively restores thin map features such as free-standing walls which may have been erased by the smoothing step.

Fig. 4. Computed skeleton, with problematic spurs highlighted in yellow. These spurs enter narrow spaces or corners but do not represent typical walking paths, so they need to be removed from the skeleton to avoid erroneously marking those areas as *blocking*.

B. Compute medial axis transform

As explained in Sec. III-A, the MAT is used to approximate the main walking paths through the space. This was calculated using the `morphology.medial_axis` function in the *scikit-image* Python library [33]. The output of this transform is equivalent to the distance transform of the space, masked by its skeleton. That is, the MAT will yield a set of pixels comprising the skeleton, where the value of each pixel equals the half-width of the gap traversed by the skeleton at that point.

C. Refine skeleton

One problem with this technique is that the generated skeletons often feature spurs leading to corners of rooms, as indicated in Fig. 4. Treating these spurs as walking paths can result in false labeling of areas as *blocking*, typically in the corners of rooms. Thus, it is necessary to clean up the skeletons by removing these spurs.

Starting from each endpoint and walking inward up the skeleton, we identify the first point where the MAT value (minimum distance to obstacles) is either no longer monotonically increasing or is greater than T_O, and trim off the segment between that point and the endpoint. This removes segments from the ends which often correspond to corners.

D. Construct Availability Map

The next step, described in Algorithm 1, is to generate a map representing availability of parking spots. In our method, a raster map is produced by coloring pixels on a grayscale range from 0 to 1. We define colors as follows. $C_B = 0$ represents areas unavailable for parking. $C_O = 1$ indicates open areas available for parking. $C_P = 0.25$ represents the available but less preferred areas in *partially-blocking* gaps, a value chosen to align with a score of 2 on a 1-5 scale.

Starting with a blank availability map initialized to C_O, points along the skeleton are classified as *blocking*, *partially-blocking*, or *open*, based on the thresholds in Eqs. 1 and 2.

Centered on these points, filled circles are drawn with a "darken" function on the availability map, meaning that the minimum of the current shade or painted shade is applied. For *blocking* points, a circle of color C_B is painted spanning the full gap width. For *partially-blocking* points, a circle of color C_B is painted spanning the region in the center of the gap to avoid, while a circle of color C_P is painted spanning the full gap width. The resultant map looks like Fig. 5, Left.

Algorithm 1 Construct availability map

Given: $M : (x, y) \mapsto d$ ▷ Refined MAT
$C_O \leftarrow 1.0$ (white) ▷ Available for parking
$C_P \leftarrow 0.25$ (gray) ▷ Less desirable
$C_B \leftarrow 0.0$ (black) ▷ Not available
Initialize $output\_map$ values to C_O
$T_B \leftarrow d_R + d_H$ ▷ Blocking threshold
$T_O \leftarrow d_R + 2 \cdot d_H$ ▷ Open threshold
$S_B \leftarrow \{(x, y) \mid 2 \cdot M(x, y) < T_B\}$ ▷ Blocking
$S_P \leftarrow \{(x, y) \mid T_B \leq 2 \cdot M(x, y) < T_O\}$ ▷ Partial
$S_O \leftarrow \{(x, y) \mid 2 \cdot M(x, y) \geq T_O\}$ ▷ Open
procedure CONSTRUCTAVAILABILITYMAP
  for each point $(x, y) \in M$ **do**
    if $(x, y) \in S_B$ **then**:
      DARKEN($output\_map, x, y, M(x, y), C_B$)
    else if $(x, y) \in S_P$ **then**
      $d_{avoid} \leftarrow T_O/2 - M(x, y)$
      DARKEN($output\_map, x, y, d_{avoid}, C_B$)
      DARKEN($output\_map, x, y, M(x, y), C_P$)
    else if $(x, y) \in S_O$ **then**
      Do nothing
    end if
  end for
  Return output_map
end procedure
procedure DARKEN($map, x, y, radius, value$)
  for each point $p \in map$ **do**
    if $distance(p, (x, y)) < radius$ **then**
      $map(x, y) \leftarrow min(map(x, y), value)$
    end if
  end for
end procedure

E. Apply smoothing

The map generated by the above procedure represents a utility function discretized into three levels. Smoothing the edges between these levels will provide a more continuous gradient for comparing the quality of potential parking spots within these discrete levels.

To achieve this, a Gaussian blur is applied to the map. This results in smoother gradients, as shown in Fig. 5.

V. EVALUATION

To evaluate the effectiveness of this algorithm, we performed an offline analysis using user-annotated maps.

As noted earlier, the setting of the human diameter dH is subjective and must be tuned, so the focus of this study was to investigate two questions:

RQ1: What value of dH is most effective at predicting the parking spot preferences of users?

Fig. 5. Left: Map before smoothing. Dark gray represents *blocking* and light gray represents *partially blocking* areas. Gray values adjusted and occupancy grid overlaid to show walls for clarity. Right: Final map after smoothing.

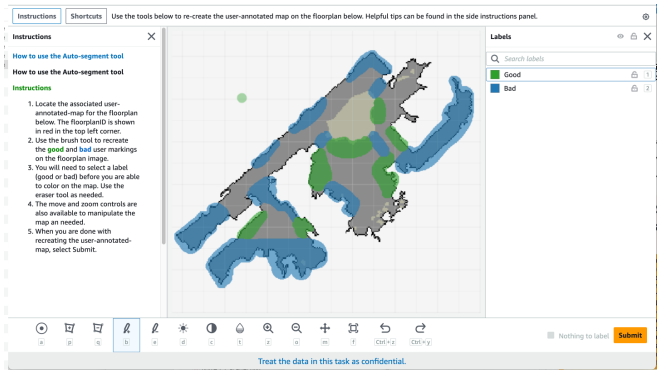

Fig. 6. Annotation interface used for coloring good and bad hangout spots.

RQ2: For this chosen value of dH, how effectively does the proposed algorithm predict parking spot preferences?

A. Participants and data collection

A pool of 34 users was recruited to provide annotations of good and bad hangout spots in their homes. All were Amazon employees who had used Astro in their homes for at least 2 weeks. Of the participants, 14 were unpaid volunteers internal to the Astro project, and 20 were paid participants recruited from outside the Astro project through our company's beta testing program.

Participants provided occupancy grids of their floorplans from their robot, and they annotated their preferences of good and bad parking spots by manually coloring over their map.

Using the semantic segmentation feature in Amazon Sagemaker GroundTruth, blue was used to represent a score of 0 (bad for parking), green a score of 1 (good for parking), and uncolored space a neutral score of 0.5. The navigable space in each map was then subsampled to a 0.25m grid. Fig. 6 shows an example of the annotation interface.

As a part of the initial robot setup, participants had divided their home into "regions" corresponding to rooms or semantically meaningful spaces. Each pose in the user preference map was thus associated with a region. In total, the collected data consisted of 34 floorplans containing 16643 unique poses grouped into 269 regions.

B. Conditions

To study the effect of the human diameter variable dH, multiple versions of the availability map were generated for each floorplan. Maps were generated for dH values at 5cm intervals ranging from 25cm to 70cm. As a baseline, we generated an additional blank map in which all navigable spots are considered *open*. This baseline represents random parking spot selection without using our model at all.

In all, 11 versions each of 34 floorplans were generated for this analysis, a total of 374 availability maps. Fig. 7 shows examples of user preference maps and generated availability maps for various dH values.

C. Measures

We define three metrics for comparison between the user preference map and our computed availability map. Each metric has advantages and drawbacks, as discussed below.

In this analysis, we define the user preference map U to contain a set of user preference scores $U(1)..U(n)$, defined over the set of n navigable poses sampled at 0.25m intervals. The preference scores are scaled to the interval $[0, 1]$, with higher scores representing more positive preferences.

We define the availability map A to contain a set of scores $A(1)..A(n)$ corresponding to the same set of poses as the user preference map. The availability scores computed by our algorithm are scaled to the interval $[0, 1]$, with higher scores representing better candidates for parking.

Finally, each map includes a set of user-defined regions R. Each pose p belongs to one region $r \in R$.

The three metrics are defined as follows:

1) Estimation Accuracy: This metric is a pose-by-pose comparison over all grid elements in the user preference map, and it answers the question, "How well does our model fit the user's preferences overall?"

Estimation accuracy is computed as the mean absolute difference between corresponding cell values in the user preference map and the availability map, subtracted from 1.

$$EA = 1 - \frac{1}{n}\sum_{i=1}^{n} |U(i) - A(i)|$$

where n is the total number of navigable poses in the map, $U(i)$ is the user preference score at pose i in the user preference map, and $A(i)$ is the availability score at the corresponding pose i in the availability map.

One limitation of this metric is that it considers all navigable space in the map equally, causing an implicit bias towards representing goodness-of-fit in larger spaces compared with smaller ones.

2) Spot Quality: When using the availability map to select a parking spot within a room, what matters most is the quality of the top selected pose. Spot Quality answers the question, "How good are the parking spots which are selected according to this model?"

For this analysis, the pose with the highest availability score is selected for each region. If the top score is shared by multiple poses, one is selected at random. The spot quality

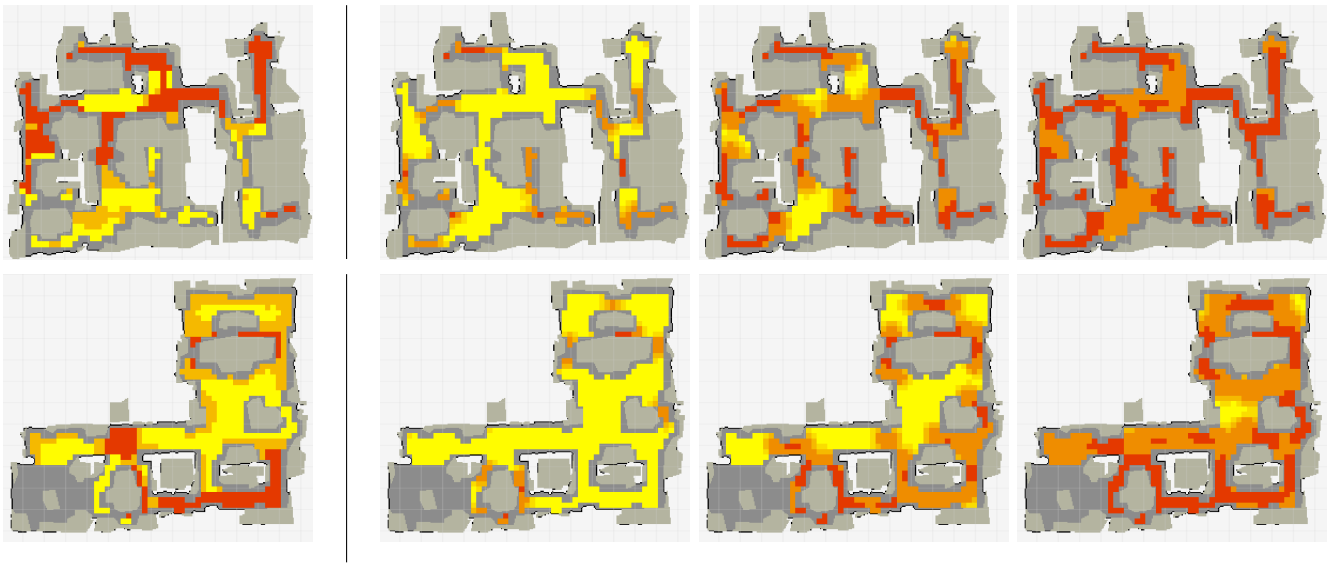

Fig. 7. Examples of two floorplans. Left: User-annotated preference map. Center to right: Availability maps generated for dH values of 25cm, 45cm, and 65cm. Red represents bad parking spots and yellow represents good. By inspection, the 45cm map gives the closest fit to the user-annotated map.

score for a region is defined as the user preference score corresponding to that region's selected pose. These quality scores are then averaged across all regions in the map. The Spot Quality metric SQ is thus defined as follows:

$$SQ = \frac{1}{|R|} \sum_{r \in R} U\left(\arg\max_{p \in P_r} A(p)\right)$$

where $|R|$ is the total number of regions in the map, P_r is the set of poses in region r, and $\arg\max_{p \in P_r} A(p)$ denotes the pose in region r with the highest availability score.

This metric provides a more practical view of the effectiveness of the model in selecting parking spots. However, it is a noisy measure due to the smaller number of data points and the use of random selection.

3) Region Agreement: Not every room in the home will necessarily contain a good parking spot. The availability map may be used to determine whether to park in a room at all. We thus propose the Region Agreement metric to answer the question, "How often does our model agree with the user on whether a region contains any good parking spots?"

For each region r, we evaluate the agreement between the highest availability score A_r and the highest preference score P_r found within that region. Since a score of 0.5 indicates no preference, 0.5 is used as an inclusive cutoff of positive or negative evaluation. The agreement condition RA is assigned a value of 0 or 1, expressed as follows:

$$\text{RA} = [(P_r \geq 0.5 \wedge A_r \geq 0.5) \vee (P_r \leq 0.5 \wedge A_r \leq 0.5)]$$

The RA scores are then averaged across all regions in the map to yield the Region Agreement score.

This metric attempts to capture a coarse-grained view of a user's parking spot preferences. However, it is a bit sensitive to the exact placement of region boundaries and it does not capture any of the nuances of parking spot placement.

TABLE I

METRICS RESULTS FOR HUMAN DIAMETER (DH) VALUES OF 25-70CM. BOLD VALUES INDICATE MAXIMUM VALUE FOR THAT METRIC.

dH (cm)	Estimation Accuracy	Spot Quality	Region Agreement
25	0.5102	0.3646	0.7871
30	0.5110	0.3686	0.7962
35	0.5507	0.3987	**0.8225**
40	0.6228	0.4952	0.8191
45	0.6374	**0.5568**	0.8014
50	**0.6416**	0.5074	0.7369
55	0.6396	0.4703	0.6909
60	0.6399	0.4760	0.6917
65	0.6236	0.4414	0.5777
70	0.6239	0.4426	0.5903
Baseline	0.4427	0.3116	0.7724

VI. RESULTS AND DISCUSSION

A. Results

The evaluation results are shown in Fig. 8 and Table I.

For **Estimation Accuracy**, we see a large jump between dH values of 35cm (55.1% accuracy) and 40cm (62.3%). The peak value was at 50cm (64.2%), with values fairly flat in the range of 40-60cm and possibly falling off slightly after that. All conditions outperformed the baseline case.

For **Spot Quality**, we see a clear peak value of 55.7% at a diameter of 45cm, with values dropping off sharply in both directions around that value. All conditions outperformed the baseline case.

For **Region Agreement**, we see another plateau for values up to 45cm, with the peak value of 82.3% for a diameter of 35cm. Region agreement dropped off substantially for diameters of 50cm and up, performing worse than the baseline case.

RQ1: What value of dH is most effective at predicting the parking spot preferences of users?

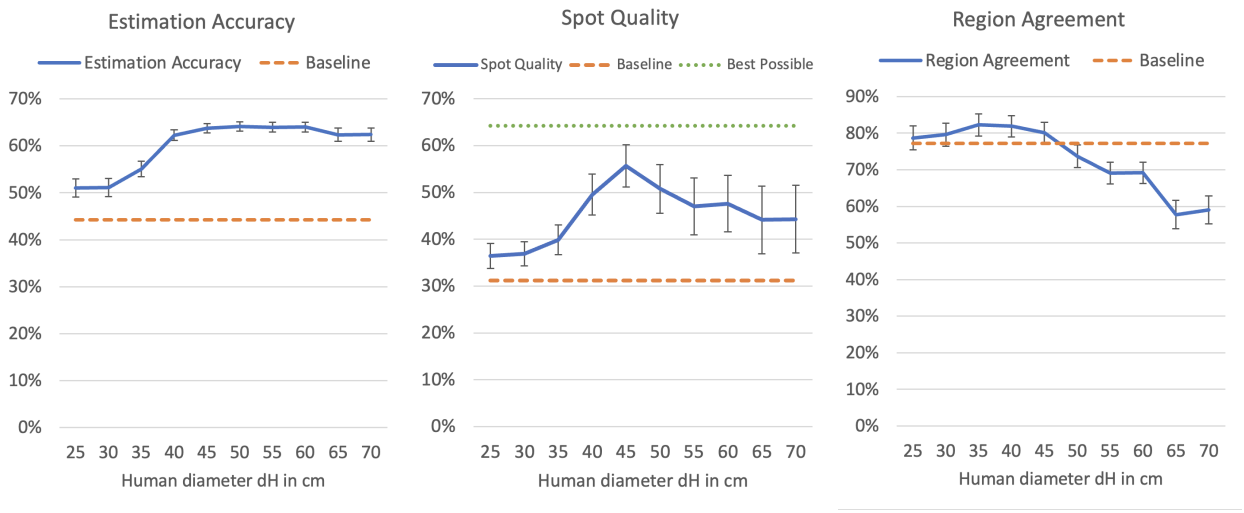

Fig. 8. Analysis results comparing model prediction with user-annotated preferences for varying values of dH. Left: Estimation accuracy, representing cell-by-cell agreement. Center: Spot quality, representing the score of the top selected spot in each region. Both the baseline and maximum achievable quality score are shown. Right: Region agreement, representing aggregate region-by-region agreement.

Considering all of these metrics together, it appears that a dH value of around 45cm would provide the best fit, as there is a clear peak in the Spot Quality results at this value and it is also within the plateau regions for the other two metrics, that is, $dH >= 40cm$ for Estimation Accuracy, and $dH <= 45cm$ for Region Agreement.

RQ2: For this chosen value of dH, how effectively does the proposed algorithm predict parking spot preferences?

For $dH = 45cm$, all measures showed improvement with respect to the baseline. Estimation Accuracy (63.7%, baseline 44.3%) and Spot Quality (55.7%, baseline 31.2%) showed large improvements, while Region Agreement showed a minor improvement (80.1%, baseline 77.2%).

To better understand the Spot Quality results, recall that Spot Quality represents an average of user preference scores for selected parking spots in each region in the home. However, not every region contains a desirable parking spot. Regions like hallways may be too narrow for parking, and regions like bathrooms or entryways may be unacceptable for other reasons. Of the 269 spaces, users only marked good parking spots in 151 (56%) of regions, and only 195 regions (72%) contained spots with neutral or better scores.

Averaging the top user preference scores across all 269 regions gives 64.2%. This represents the highest Spot Quality score achievable. Thus, considering the Spot Quality result of 55.7% between the baseline score of 31.1% and the maximum possible score of 64.2%, our proposed model has accounted for 74% of that gap, suggesting that the remaining 26% is likely due to factors outside of the proposed model.

B. Discussion

1) Interpretation of dH: The concept of "human diameter" merits some discussion - how should the 45cm result be interpreted? Whereas a single radius is often used for expressing distances such as Hall's zones of interpersonal distance [34], we may need to consider the width differences between the lower and upper body. It seems likely that human preferences for comfortable passing space would use a smaller diameter for low obstacles, and a larger diameter for high obstacles which could impede the arms or shoulders.

Our study involved a robot shorter than knee-height, but many of the obstacles in the map were walls extending from floor to ceiling. Future studies might investigate whether and how people's preferences change for low obstacles like a table or sofa, or taller robots. We expect that dH may be lower for shorter obstacles or higher for taller robots.

2) Possible applications: This algorithm was designed to enable a robot to select out-of-the-way parking spots in arbitrary rooms. It might also be useful for suggesting locations for permanent placement of objects such as a robot dock, or even in dynamic navigation, where the utility function defined in this map would naturally guide the robot to move along the edges of hallways and avoid bottlenecks.

C. Limitations and future work

1) Annotation Reliability: This work treats parking spot quality as a static property, but actual user preferences may depend on dynamic context such as the user's location, activity, or other factors. It would be worthwhile to follow up with real *in-situ* studies with a robot to better understand the degree to which dynamic context might cause user preferences to diverge from what they drew on their maps.

2) Entry and exit points: One limitation of the proposed approach is that the estimated walking paths do not consider entrances to the space such as doors and stairways. Mitigation strategies could include direct detection of these features by the robot, or augmenting the estimated walking paths with observed trajectories of people who enter or exit the space.

3) Unmodeled goal locations: The MAT provides a reasonable estimate of walking paths that go through rooms and hallways, but actual walking paths may have local goal locations such as chairs, sofas, or appliances, that are

not captured by this simple geometric approach. Again, mitigation strategies could include detection of these objects, or they could be inferred from pedestrian trajectories.

VII. Conclusion

In this work, we have proposed a method for estimating out-of-the way parking spots for robots in the home or other social spaces. The approach is highly generalizable, as it is based only on the geometry of the space as represented in a standard occupancy grid, and does not depend on specific sensing capabilities or *a priori* knowledge about the space.

In an evaluation based on preference maps collected from 34 users who had real experience with an Astro robot in their home, we demonstrated our technique to be successful in improving the quality of parking spot selection compared with a naive baseline. Our analysis showed a "human diameter" parameter of 45cm to be most effective, capturing 74% of user parking spot preferences.

As the trend of developing mobile robots for the home continues, we expect to see an increased interest in the modeling of user spatial preferences including robot parking spots. This work provides a useful technique that contributes toward modeling such preferences in real people's homes.

Acknowledgments

The authors would like to thank De'Aira Bryant and Hanxiao Fu for their contributions to this work.

References

[1] H.-M. Gross, A. Scheidig, S. Müller, B. Schütz, C. Fricke, and S. Meyer, "Living with a mobile companion robot in your own apartment-final implementation and results of a 20-weeks field study with 20 seniors," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 2253–2259.

[2] G. Wilson, C. Pereyda, N. Raghunath, G. de la Cruz, S. Goel, S. Nesaei, B. Minor, M. Schmitter-Edgecombe, M. E. Taylor, and D. J. Cook, "Robot-enabled support of daily activities in smart home environments," *Cognitive Systems Research*, vol. 54, pp. 258–272, 2019.

[3] H. M. Do, M. Pham, W. Sheng, D. Yang, and M. Liu, "Rish: A robot-integrated smart home for elderly care," *Robotics and Autonomous Systems*, vol. 101, pp. 74–92, 2018.

[4] Z. Fu, T. Z. Zhao, and C. Finn, "Mobile aloha: Learning bimanual mobile manipulation with low-cost whole-body teleoperation," *arXiv preprint arXiv:2401.02117*, 2024.

[5] iRobot Corporation. (2024) Roomba® robot vacuum cleaners. [Online]. Available: https://www.irobot.com/en_US/roomba.html

[6] Mayfield Robotics. (2017) Kuri. [Online]. Available: https://robotsguide.com/robots/kuri

[7] Amazon.com, Inc. (2023) Amazon astro, household robot for home monitoring. [Online]. Available: https://www.amazon.com/astro

[8] D. Bryant, T. Etiene, A. Howard, W. D. Smart, and D. F. Glas, "Teaching a robot where to park: A scalable crowdsourcing approach," in *2023 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, 2023, pp. 226–233.

[9] D. F. Glas, "System for determining constraint regions for an autonomous mobile device," U.S. Patent US-11 853 077-B1, Dec 26, 2023.

[10] S. Chik, C. Yeong, E. Su, T. Lim, Y. Subramaniam, and P. Chin, "A review of social-aware navigation frameworks for service robot in dynamic human environments," *Journal of Telecommunication, Electronic and Computer Engineering (JTEC)*, vol. 8, no. 11, pp. 41–50, 2016.

[11] K. Charalampous, I. Kostavelis, and A. Gasteratos, "Recent trends in social aware robot navigation: A survey," *Robotics and Autonomous Systems*, vol. 93, pp. 85–104, 2017.

[12] R. Möller, A. Furnari, S. Battiato, A. Härmä, and G. M. Farinella, "A survey on human-aware robot navigation," *Robotics and Autonomous Systems*, vol. 145, p. 103837, 2021.

[13] A. Rudenko, L. Palmieri, M. Herman, K. M. Kitani, D. M. Gavrila, and K. O. Arras, "Human motion trajectory prediction: a survey," *The International Journal of Robotics Research*, vol. 39, no. 8, pp. 895–935, 2020.

[14] T. Kanda, D. F. Glas, M. Shiomi, and N. Hagita, "Abstracting people's trajectories for social robots to proactively approach customers," *IEEE Transactions on Robotics*, vol. 25, no. 6, pp. 1382–1396, 2009.

[15] H. Noguchi, T. Yamada, T. Mori, and T. Sato, "Mobile robot path planning using human prediction model based on massive trajectories," in *2012 Ninth International Conference on Networked Sensing (INSS)*. IEEE, 2012, pp. 1–7.

[16] H. Kidokoro, T. Kanda, D. Brščic, and M. Shiomi, "Will i bother here? - a robot anticipating its influence on pedestrian walking comfort," in *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2013, pp. 259–266.

[17] P. Liu, D. F. Glas, T. Kanda, and H. Ishiguro, "Data-driven hri: Learning social behaviors by example from human-human interaction," *IEEE Transactions on Robotics*, vol. 32, no. 4, pp. 988–1008, 2016.

[18] T. Kitade, S. Satake, T. Kanda, and M. Imai, "Understanding suitable locations for waiting," in *Proceedings of the 8th ACM/IEEE International Conference on Human-Robot Interaction*, ser. HRI '13. IEEE Press, 2013, p. 57–64.

[19] P. Patompak, S. Jeong, I. Nilkhamhang, and N. Y. Chong, "Learning proxemics for personalized human–robot social interaction," *International Journal of Social Robotics*, vol. 12, pp. 267–280, 2020.

[20] P. Ratsamee, Y. Mae, K. Ohara, T. Takubo, and T. Arai, "Human–robot collision avoidance using a modified social force model with body pose and face orientation," *International Journal of Humanoid Robotics*, vol. 10, no. 01, p. 1350008, 2013.

[21] Y. Hiroi and A. Ito, "A pedestrian avoidance method considering personal space for a guide robot," *Robotics*, vol. 8, no. 4, 2019.

[22] S. Satake, T. Kanda, D. F. Glas, M. Imai, H. Ishiguro, and N. Hagita, "A robot that approaches pedestrians," *IEEE Transactions on Robotics*, vol. 29, no. 2, pp. 508–524, 2012.

[23] M. Vázquez, E. J. Carter, B. McDorman, J. Forlizzi, A. Steinfeld, and S. E. Hudson, "Towards robot autonomy in group conversations: Understanding the effects of body orientation and gaze," in *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*, 2017, pp. 42–52.

[24] M. Fu, R. Liu, B. Qi, and R. R. Issa, "Generating straight skeleton-based navigation networks with industry foundation classes for indoor way-finding," *Automation in Construction*, vol. 112, p. 103057, 2020.

[25] J. Lee, "A spatial access-oriented implementation of a 3-d gis topological data model for urban entities," *GeoInformatica*, vol. 8, pp. 237–264, 2004.

[26] L. Perkins, "Terrain analysis in real-time strategy games: An integrated approach to choke point detection and region decomposition," in *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, vol. 6, no. 1, 2010, pp. 168–173.

[27] R. Fisher, S. Perkins, A. Walker, and E. Wolfart. (2000) Hypermedia image processing reference. [Online]. Available: https://homepages.inf.ed.ac.uk/rbf/HIPR2/skeleton.htm

[28] H. Blum, "A transformation for extracting new descriptions of shape," *Models for the perception of speech and visual form*, pp. 362–380, 1967.

[29] D.-T. Lee, "Medial axis transformation of a planar shape," *IEEE Transactions on pattern analysis and machine intelligence*, no. 4, pp. 363–369, 1982.

[30] R. Tam and W. Heidrich, "Shape simplification based on the medial axis transform," in *IEEE Visualization, 2003. VIS 2003.* IEEE, 2003, pp. 481–488.

[31] J. Hu, B. Wang, L. Qian, Y. Pan, X. Guo, L. Liu, and W. Wang, "Mat-net: Medial axis transform network for 3d object recognition." in *IJCAI*, 2019, pp. 774–781.

[32] G. Bradski, "The OpenCV Library," *Dr. Dobb's Journal of Software Tools*, 2000.

[33] S. van der Walt, J. L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, T. Yu, and the scikit-image contributors, "scikit-image: image processing in Python," *PeerJ*, vol. 2, p. e453, 6 2014. [Online]. Available: https://doi.org/10.7717/peerj.453

[34] E. T. Hall, *The hidden dimension*. Garden City, NY: Doubleday, 1966.