

# USER FRONT-FOLLOWING BEHAVIOUR FOR A MOBILITY ASSISTANCE ROBOT: A KINEMATIC CONTROL APPROACH

George P. Moustris<sup>(a)</sup>, Athanasios Dometios<sup>(b)</sup>, Costas S. Tzafestas<sup>(c)</sup>

Intelligent Robotics & Automation Laboratory  
School of Electrical & Computer Engineering,  
National Technical University of Athens

<sup>(a)</sup>[groustri@mail.ntua.gr](mailto:groustri@mail.ntua.gr), <sup>(b)</sup>[athdom@mail.ntua.gr](mailto:athdom@mail.ntua.gr), <sup>(c)</sup>[ktzaf@cs.ntua.gr](mailto:ktzaf@cs.ntua.gr)

## ABSTRACT

We present a robot following behavior, which enables a mobility assistance robot to follow the user from the front. Although this behavior is desirable in various circumstances, it has received scarce attention. Our proposed solution consists of a kinematic control scheme, tied to a human position estimator based on a Laser Range Finder. Experiments have been performed in an indoor environment with ten users, and the results have been analyzed and presented. We show that the control is feasible but inserts a cognitive load on the users, who tend to “drive” the robot to the optimal paths they would take under normal conditions.

Keywords: kinematic control, front following, gait, assistance robot

## 1. INTRODUCTION

Human-robot interaction has received increasing attention in the past decade, especially for service robots. Human-aware navigation involves scenarios where the robot must navigate in public places e.g. libraries, hospitals, warehouses, and avoid collision with humans and obstacles (Kruse et al. 2013). Another typical task is human-following, where the robot must assist the user in various tasks by following him/her through the environment. This mode can be seen for example, in telepresence robots (Cosgun, Florencio, and Christensen 2013), assistance robots in hospitals (Gockley, Forlizzi, and Simmons 2007), companion robots (Ohya and Muneakata 2002) and other. The main assumption in human tracking is that the user resides always in front of the robot. However, in general we can discern three cases for human following (Ho, Hu, and Wang 2012):

1. “behind the leader”
2. “side-by-side”
3. “in front of the leader”

The vast majority of the research volume addresses only the first case; that is, following the user from behind. This eases the control problem as the user intention can be discerned from his/hers trajectory. However, it also presents shortcomings in the interaction. Jung, Yi, and Yuta (2012) have noticed that when the robot tracks the human from behind, the human looks backs to see where the robot is. This imposes a cognitive load and

causes the human to pay attention to the robot either from curiosity e.g. see where the robot has gone, or even from fear of the robot bumping onto the human. To alleviate this problem, recently the “side-by-side” and “following from the front” tasks has been addressed by some researchers. From a control-theoretic point of view, the “following from behind” task needs only to know the position of the human relative to the robot. As a control problem, the robot has but to retrace the human’s path in order to stay behind the user. In a more general case, the problem can be classified as a *pursuit* problem, which has received ample attention in the literature. However, in the “side-by-side” and “following from the front” tasks, the difficulty increases significantly. Specifically, in the first, the user orientation must also be taken into account in order for the robot to stay by the side of the human. In the second, prediction of the user’s intention must also be incorporated, while in some cases the problem might be undecidable, requiring user feedback to escape deadlocks.

Here we present a kinematic control strategy for enabling a mobility assistance robot to follow the user from the front. This behavior has been developed in the settings of the “MOBOT” project ([www.mobot-project.eu](http://www.mobot-project.eu)). MOBOT is an active mobility assistance robot for indoor environments that provides user-centered, context-adaptive and natural support (Papageorgiou et al. 2014).



Figure 1: The MOBOT assistance robot

The “front-following” behavior is intended to allow the robot to oversee the patient, walk along with him/her and provide assistance either on demand or autonomously.

## 2. PROBLEM CLASSIFICATION

The problem of following from the front can be divided into two general cases; following the human in free space i.e. in an obstacle-free space with no objects and, following the human in a structured environment e.g. in an office building, corridor etc. The two problems have different complexity with the former being substantially simpler than the latter. Specifically, in *free space following*, the problem can be cast as a control problem where the goal is to minimize some error measures e.g. minimize the distance and orientation errors between the human and the robot. This approach is singularly treated in the current literature. In the *structured environment* case, the task involves avoiding obstacles, either static or moving, as well as deciding where the human actually wants to go; a possibly undecidable problem. See for example Figure 2.

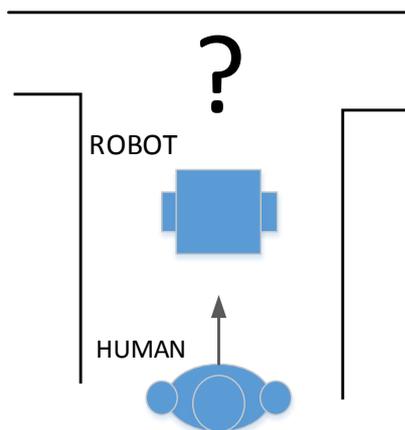


Figure 2: Undecidability of the front-following problem in structured environments

It is clear that the robot has no way of knowing where the human wants to turn by examining solely the human motion. This problem requires the addition of further information into the control loop by letting the human show the robot to turn left/right using some kind of feedback e.g. audio, posture, gestures etc. Thus, the human must also “steer” the robot and not just act as an observable for the robot. The control strategy for this problem is radically different from the *free space* following problem, and has received no attention in the literature.

## 3. RELATED WORK

As mentioned above, the front-following problem has received scarce attention from the research community. Our survey has produced only three papers dealing with subject. All three deal with the free-space following problem. In (Jung, Yi, and Yuta 2012) the authors use a Laser Range Finder (LRF) to scan the human torso, which serves as a more robust scanning target than the

legs. Using a particle filter employing a constant velocity model, they track the pose of the human during motion. The control algorithm uses a *virtual target* based on the human and robot poses. The aim is for the robot to track the target, which lays in the approximate direction of the human velocity vector. (Ho, Hu, and Wang 2012) use an RGBD sensor (Microsoft Kinect) to track the human position relative to the robot. Following, they use the nonholonomic human model (Arechavaleta et al. 2008; Arechavaleta et al. 2006) to calculate the human’s orientation, combined with an Unscented Kalman Filter to provide a smooth estimate of the human orientation, linear velocity and angular velocity. The controller is an ad-hoc solution aiming to align the human-robot poses while putting the robot in front. Simulation and experimental results are promising, although not very extensive.

(Cifuentes et al. 2014) use a different approach. They combine readings from a wearable IMU sensor on the human, along with LRF data of the legs in order to provide an estimate of the human pose and linear/angular velocities. They use an inverse kinematics controller to exponentially stabilize a position and orientation error between the human and the robot. In this setup, they perform experiments in straight line tracking, as well as in tracking the human along an 8-shaped path.

## 4. HUMAN POSITION ESTIMATION

The first step towards human following is the detection/estimation of the human pose. A basic assumption is that the human is detected by a LRF located on the robot, which scans the user legs. Furthermore, the kinematic controller only needs the position of the human, not the orientation and velocity. This simplifies the control and is more robust to estimation errors.

To filter out environment artefacts and obstacles, we borrow the idea of a Human Interaction Zone (HIZ) from (Cifuentes et al. 2014), which consists of a parallelogram of width 2m and length 2m, centered at the LRF (Figure 3).

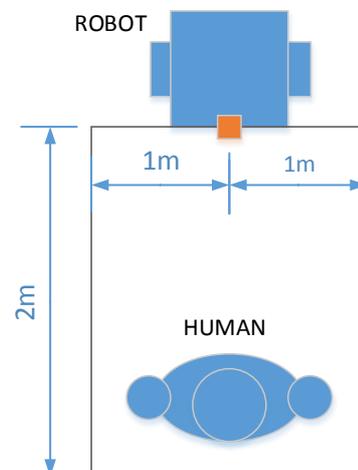


Figure 3: Definition of the Human Interaction Zone

Based on the laser scans inside the HIZ, a *centroid* is calculated by taking the average in each  $x, y$  coordinates. Thus, if  $k$  scans lay inside the HIZ, the centroid coordinates are,

$$\begin{bmatrix} x_H \\ y_H \end{bmatrix} = \begin{bmatrix} 1/k \sum_k x_L^i \\ 1/k \sum_k y_L^i \end{bmatrix} \quad (1)$$

To enable more valid detection results, in order to exclude false positives from walls, furniture etc. we have inserted an adaptive algorithm based on the previous valid centroid position. Specifically, in the beginning, the robot considers only scans inside an *initial window*, similar to the HIZ but with a width of 0.8m. This implies that the human who is intended to be followed, approaches the robot in a narrow region. Following, the algorithm estimates the centroid coordinates  $x_H^i, y_H^i$  at loop “ $i$ ”. In the next loop “ $i+1$ ”, the algorithm scans inside a small *leg window*, of width 0.3m and height 0.2m. Thus the detection area is the rectangle  $[x_H^i \pm 0.3, y_H^i \pm 0.2]$ . In this way, the algorithm tracks the human as he/she moves inside the HIZ, and discards other unrelated objects.

## 5. KINEMATIC CONTROLLER

The proposed solution for the front-following problem, is a *virtual pushing* approach through a kinematic controller. We define an equilibrium distance  $x_0$  where the system is at rest. If the human passes the equilibrium point and approaches the robot, then the robot starts to move depending on the human-robot distance.

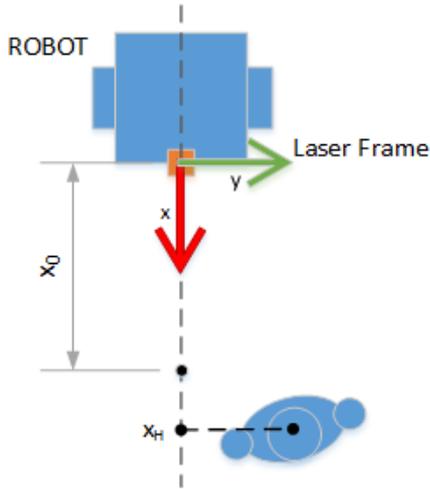


Figure 4: Depiction of the Laser Frame and the Equilibrium distance  $x_0$

The robot model used is the Unicycle robot (Figure 5), described by the equations,

$$\begin{bmatrix} \dot{x}_R \\ \dot{y}_R \\ \dot{\theta}_R \end{bmatrix} = \begin{bmatrix} \cos \theta_R \\ \sin \theta_R \\ 0 \end{bmatrix} v_R + \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \omega_R \quad (2)$$

where  $x_R, y_R$  are the coordinates of robot with respect to a *world frame*, and  $\theta_R$  its orientation. The inputs  $v_R, \omega_R$  are the linear and angular velocities respectively.

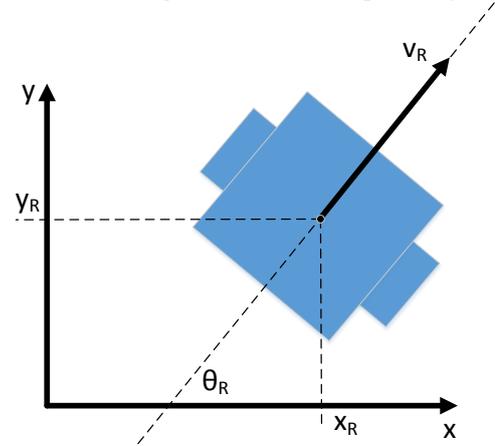


Figure 5. The unicycle robot model

Rigidly attached to the robot is the *laser frame*, in which the user centroid  $x_H, y_H$  is calculated. The robot's linear velocity is given by,

$$v_R = \lambda(y_H) v(x_H) \quad (3)$$

where,

$$v = \begin{cases} 0 & , x_H > x_0 \\ k_1(x_H - x_0) & , x_2 \leq x_H \leq x_0 \\ v_{\text{walk}} & , x_1 \leq x_H \leq x_2 \\ v_{\text{max}} - k_2 x_H & , 0 \leq x_H \leq x_1 \end{cases} \quad (4)$$

$$k_1 = \frac{v_{\text{walk}}}{x_2 - x_0}, \quad k_2 = \frac{v_{\text{max}} - v_{\text{walk}}}{x_1}$$

The term  $\lambda$  is a velocity modulating term (see next page for a more thorough analysis). Equation (4) defines a piece-wise linear velocity profile, consisting of three regions; the *approach region*, the *walking region* and the *collision region*.

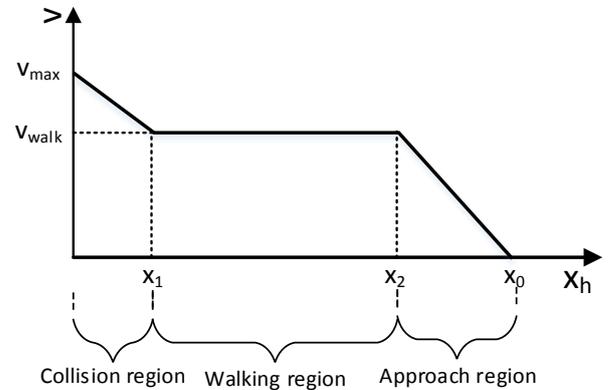


Figure 6: Profile of the linear velocity input

The *walking region* is the set on the  $x$ -axis of the LRF frame, within which the robot has a constant velocity, namely the *walking velocity*  $v_{walk}$ . In this region the robot moves synchronously with the user. If the human moves very close to the robot, he/she enters into the *collision region*, in which the robot accelerates up to a maximum velocity  $v_{max}$ . Conversely, if the human falls behind (or enters the HIZ from a distance greater than the Equilibrium distance  $x_0$ ), the *approach region* is considered, where the robot accelerates from halt up to the walking velocity. The second robot input, the angular velocity  $\omega_R$ , is described by the following equations,

$$\omega_R = \begin{cases} 0 & , |y_H| < \varepsilon \\ k_\omega \operatorname{sgn}(y_H)(|y_H| - \varepsilon_0) & , |y_H| > \varepsilon \end{cases} \quad (5)$$

$$k_\omega = \frac{\omega_{max}}{HIZ_w / 2 - \varepsilon}$$

Here  $\omega_{max}$  is the maximum angular velocity,  $HIZ_w$  is the width of the HIZ and  $\varepsilon$  is a deadband about the  $x$ -axis. The deadband is inserted in order to filter out natural gait oscillations during walking, as well as noise from the centroid estimator. In our experiments  $\varepsilon$  was set to 10cm.

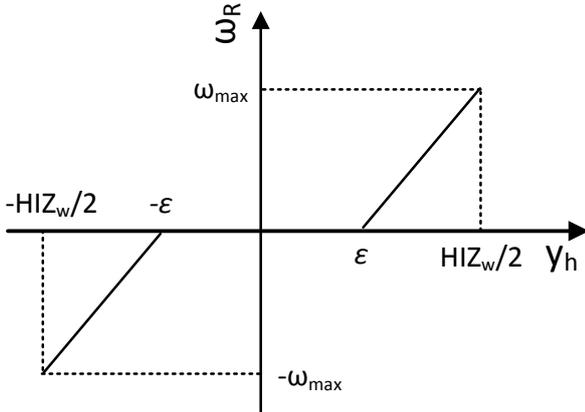


Figure 7: Profile of the angular velocity input

Using Eq.(5), the robot essentially turns in such a way as to always face the user. During experiments it was observed that in corners the users place themselves on the outer limits of the  $y$ -axis to make the robot turn fast enough. This oversteers the robot and in order to correct its heading, they must swiftly move on the other end of the axis. At the same time the robot is moving forward with a linear velocity, making the reaction time rather short and leading to unstable behaviors. To prevent this situation, we have inserted a velocity modulating term  $\lambda(y_h)$  in Eq.(3). The term is given by,

$$\lambda = \begin{cases} 1 & , |y_H| < y_a \\ \frac{y_b - |y_H|}{y_b - y_a} & , y_a \leq |y_H| \leq y_b \\ 0 & , y_b < |y_H| \end{cases} \quad (6)$$

$$y_b = \varepsilon + b(HIZ_w / 2 - \varepsilon)$$

$$y_a = \varepsilon + a(HIZ_w / 2 - \varepsilon)$$

The parameters  $0 < a < b < 1$  are percentages with respect to the deadband. A graphical depiction of  $\lambda$  can be seen in Figure 8.

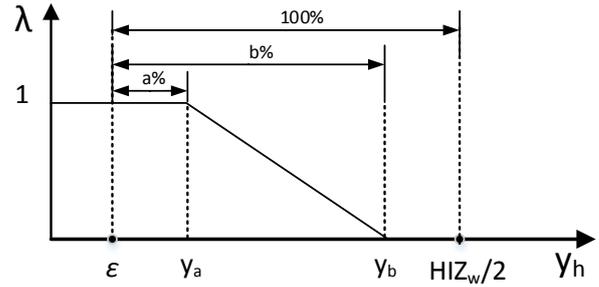


Figure 8: Illustration of the  $\lambda$  function

The  $\lambda$  term reduces the linear velocity as the user increases his/hers lateral displacement. On the outer regions, the robot halts and turns on the spot to face the human. For our experiments the parameters were set to  $a=0.3$  and  $b=0.6$ .

## 6. EXPERIMENTAL RESULTS

The control scheme presented in the previous, has been implemented on a Pioneer 3DX differential drive robot, with a Hokuyo UBG-04LX-F01 laser range finder. The experiments considered here, aim to assess the gait pattern of the users with and without the robot following them from the front.

Ten healthy subjects were asked to walk naturally from an initial predefined position, around a corner and stop at a designated target position. Each subject performed two runs, thus in total 20 paths were collected as a baseline. The subjects were tracked with the laser scanner on top of the robot, which in turn was placed statically at the head of the corner, overseeing the experimental field. In post processing, using the detection algorithm, the centroid traces were extracted, as seen in Figure 9.

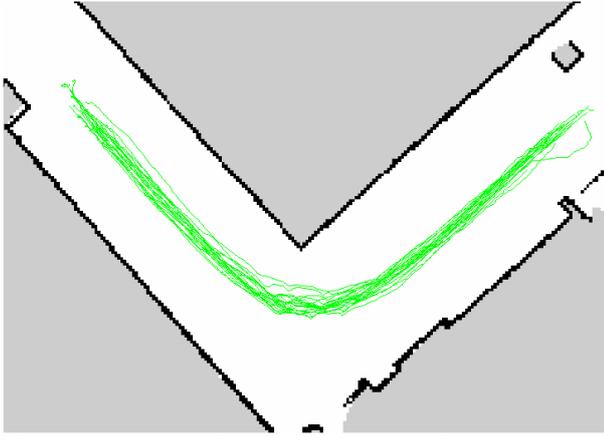


Figure 9: Traces of the baseline experiments (green). The subjects started on the right and progressed to the left.

Following, the subjects were asked to perform the experiment again, but with the robot following them from the front. Each subject did two test runs in order to get acquainted with the robot behavior. Then, they performed the experiment twice. The total collected paths are again 20.

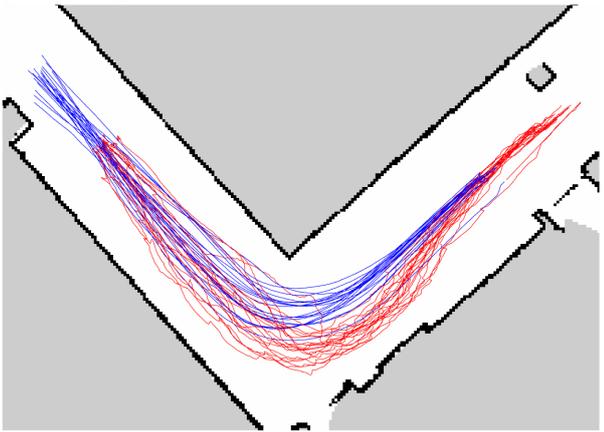


Figure 10: Traces of the following experiments (Human-red, Robot-blue). The subjects started on the right and progressed to the left

To analyze the paths, we have divided the plane into a grid of  $48 \times 26$  square cells with an edge of 20 cm each. Then, for each path we collected the binary mask consisting of those cells that the path has traversed. By counting the number of masks each cell appears in, we have produced a 2D histogram of those masks. Apparently, since we have 20 paths in each case, the count of each cell goes from zero up to twenty. The three histograms are,

$$\begin{aligned}
 H_B(i, j) &: \text{Baseline paths} \\
 H_U(i, j) &: \text{User paths} \\
 H_R(i, j) &: \text{Robot paths} \\
 i \in [1, 48] \quad , \quad j \in [1, 26]
 \end{aligned} \tag{7}$$

The histograms are presented in the following figures.

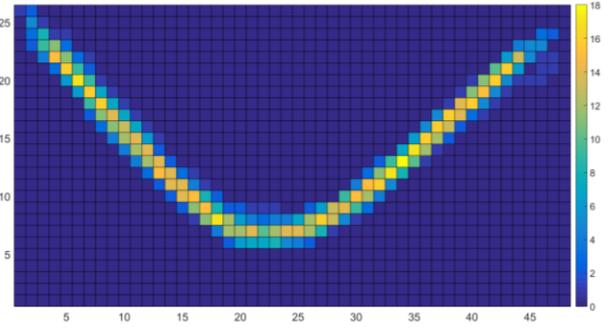


Figure 11: Histogram of the baseline paths

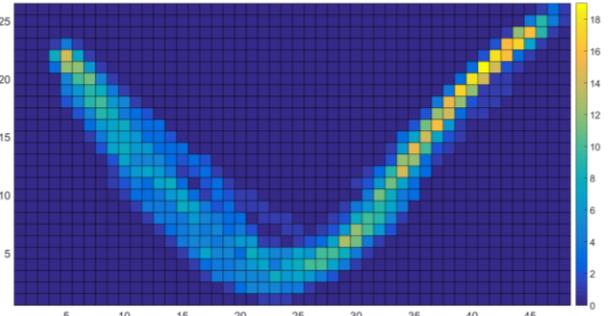


Figure 12: Histogram of the users' paths (following)

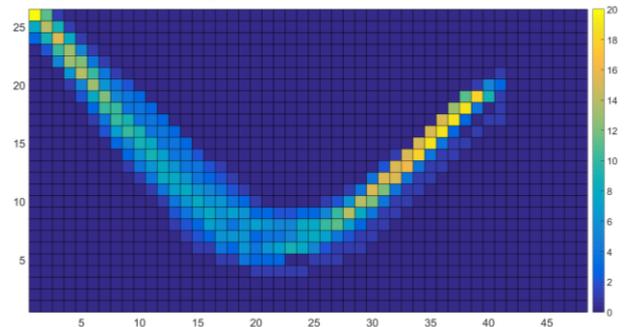


Figure 13: Histogram of the robot paths (following)

From the three histograms we can produce two new sets of distributions. By dividing the count of each cell with the total number of paths, we produce the probability of each cell being traversed by a path, viz.

$$\begin{aligned}
 T_B(i, j) &= H_B(i, j) / 20 \\
 T_U(i, j) &= H_U(i, j) / 20 \\
 T_R(i, j) &= H_R(i, j) / 20
 \end{aligned} \tag{8}$$

Thus a cell with high such a probability means that it is traversed by most of the paths. Note that these are not probability distributions as they don't sum up to one. Another set of distributions can be produced by dividing each cell with the total count of its respective histogram, i.e.

$$\begin{aligned}
 P_B(i, j) &= H_B(i, j) / \sum_{i,j} H_B(i, j) \\
 P_U(i, j) &= H_U(i, j) / \sum_{i,j} H_U(i, j) \\
 P_R(i, j) &= H_R(i, j) / \sum_{i,j} H_R(i, j)
 \end{aligned} \tag{9}$$

These express the probability of a user/robot being on a specific cell and are probability density functions. Equations (7)(8)(9) are similar up to a scaling factor (for each group “B”, “U”, “R”), thus all three have the same shape.

To compare the three groups, we resort to the *Hellinger distance* (Pollard 2002) which is a measure of statistical distance between two distributions  $P, Q$  given by.

$$H(p, q) = \frac{1}{\sqrt{2}} \sum_k (\sqrt{p_k} - \sqrt{q_k})^2 \quad (10)$$

The Hellinger distance ranges from zero to one, with zero being identical distributions and one completely disjoint. The distances of  $P_U$  to  $P_B$  and  $P_R$  to  $P_B$  are,

$$H(P_U, P_B) = 0.6265 \quad , \quad H(P_R, P_B) = 0.4907 \quad (11)$$

We see that the Robot path distribution is more similar to the Baseline distribution than the Users’ distribution. This means that the users actually tend to “drive” the robot to the path they consider “optimal” i.e. the one that *they* would take under normal conditions (the baseline paths). Doing so, they deviate from their normal gait patterns.

To further compare the three groups, we resort to Eq.(8) which describes the probability of each cell being traversed by a path. Taking the histogram of each  $T$  divided by the total number of cells *being traversed in the grid*, we get the probability of a cell being traversed by a specific fraction (or percentage if multiplied by 100%) of the paths, denoted by  $TP(a)$ . For example, if the total number of cells being traversed by paths in the grid is  $N$ , and there are  $M$  cells being traversed by  $K$  out of 20 paths, then  $TP(K/20)=M/N$ . Mathematically,

$$\begin{aligned} TP_B(a) &= \text{hist}(T_B) / \text{count}(T_B) \\ TP_U(a) &= \text{hist}(T_U) / \text{count}(T_U) \\ TP_R(a) &= \text{hist}(T_R) / \text{count}(T_R) \end{aligned} \quad (12)$$

where  $\text{count}()$  returns the number of non-zero element, and  $a$  is the fraction of paths with the given probability.

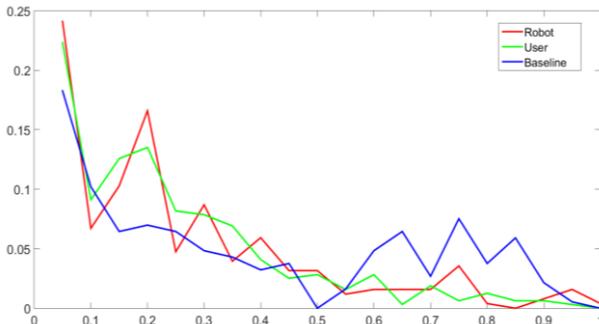


Figure 14: The distributions  $TP_B, TP_U$  and  $TP_R$

High probabilities for small fractions, means that most of the cells being traversed have been traversed by few paths, and so, the paths are “spread out” (large

variance). On the other hand, if we have high probabilities for high fractions, most of the cells that have been traversed, are done so by many paths and so the paths are localized.

We see that both the “user” and “robot” distributions are skewed towards small fractions whereas the “baseline” group is skewed towards large ones. The table with the standard deviations is given below:

Table 1: Standard Deviation of  $TP_B, TP_U$  and  $TP_R$

Standard Deviation		
$TP_B$	$TP_U$	$TP_R$
0.0412	0.0584	0.0608

As expected, we see that the “User” and “Robot” groups have higher standard deviation than the “Baseline” group. Another measure of dispersion is the relative differences between  $\text{count}(TP_R) - \text{count}(TP_B)$  and  $\text{count}(TP_U) - \text{count}(TP_B)$ , since the *count* function measures the number of cells a distribution contains. Thus the relative difference is a measure of the *extent* of a group with respect to the baseline group.

Table 2: Measure of the *extent* of the “User” and “Robot” groups with respect to the “Baseline” group

Results		
	<i>Count()</i>	<i>% rel.diff.</i>
$TP_B$	186	-
$TP_U$	318	70.96%
$TP_R$	253	36.02%

From Table 2 we see that the users cover almost 71% more cells trying to steer the robot, than when walking normally, which is almost twice the cells the robot covers. This can be regarded as a measure of cognitive load since it shows that the users walk through a wider area.

## 7. CONCLUSION

We have described a simple kinematic control strategy for enabling a non-holonomic mobile robot to follow a user from the front. The controller is simple and robust, based on a *virtual pushing* principle. This enables the following behavior to be incorporated into a structured environment, by including the human in the loop. Experimental results have shown the feasibility of our approach, albeit the users tend to deviate from their normal walking patterns trying to “steer” the robot around. Results have shown that the actual robot paths are closer to the optimal human paths (without the robot following them), than the path produced by the humans “steering” the robot. This implies that the controller inserts a cognitive load on the user by shifting a large control effort to him/her.

Future work will focus on refining the control strategy by inserting more intelligent features on the robot e.g. obstacle avoidance, undecidability detection, human intent recognition and prediction etc. These features will

aim to alleviate the control burden of the human and shift it to the robot. Our ultimate goal is to develop a successful “front-following” behavior in structured environment, in order to enable the mobility assistance robot follow the patient from the front, and provide help when needed.

## ACKNOWLEDGMENTS

This work is supported by 7th Framework Program of the European Union, ICT Challenge 2, Cognitive Systems and Robotics, contract "EU-FP7-ICT-2011-9 2.1 - 600796 - MOBOT: Intelligent Active MObility Assistance RoBOT Integrating Multimodal Sensory Processing, Proactive Autonomy and Adaptive Interaction".

## REFERENCES

- Evans W.A., 1994. Approaches to intelligent information retrieval. *Information Processing and Management*, 7 (2), 147–168.
- Arechavaleta G., Laumond J.-P., Hicheur H., and Berthoz A., 2006. “The Nonholonomic Nature of Human Locomotion: A Modeling Study.” In *The First IEEE/RAS-EMBS International Conference on Biomedical Robotics and Biomechanics, 2006. BioRob 2006*, 158–63.
- Arechavaleta G., Laumond J.-P., Hicheur H., and Berthoz A., 2008. “On the Nonholonomic Nature of Human Locomotion.” *Auton. Robots* 25 (1-2): 25–35.
- Cifuentes C. A., Frizera A., Carelli R., and Bastos T., 2014. “Human–robot Interaction Based on Wearable IMU Sensor and Laser Range Finder.” *Robotics and Autonomous Systems* 62 (10): 1425–39.
- Cosgun A, Florencio D. A., and Christensen H. I., 2013. “Autonomous Person Following for Telepresence Robots.” In *2013 IEEE International Conference on Robotics and Automation (ICRA)*, 4335–42.
- Gockley R., Forlizzi J., and Simmons R.. 2007. “Natural Person-Following Behavior for Social Robots.” In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction*, 17–24. HRI '07. New York, NY, USA.
- Ho D-M., Hu J-S., and Wang J-J. 2012. “Behavior Control of the Mobile Robot for Accompanying in Front of a Human.” In *2012 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM)*, 377–82.
- Jung E-J., Yi B-J., and Yuta S., 2012. “Control Algorithms for a Mobile Robot Tracking a Human in Front.” In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2411–16.
- Kruse T., Kumar A. P, Alami R., and Kirsch A., 2013. “Human-Aware Robot Navigation: A Survey.” *Robotics and Autonomous Systems* 61 (12): 1726–43.

Ohya A., and Munekata T., 2002. “Intelligent Escort Robot Moving Together with Human-Interaction in Accompanying Behavior.” In *Proceedings of the 2002 FIRA Robot Congress*, 31–35.

Papageorgiou X.S., Tzafestas C.S., Maragos P., Pavlakos G., Chalvatzaki G., Moustiris G.P., Kokkinos I., et al. 2014. “Advances in Intelligent Mobility Assistance Robot Integrating Multimodal Sensory Processing.” In *Universal Access in Human-Computer Interaction. Aging and Assistive Environments*, edited by C. Stephanidis and M. Antona, 692–703. Lecture Notes in Computer Science 8515. Springer International Publishing.

Pollard, D.E., 2002. “A user's guide to measure theoretic probability”. Cambridge, UK: Cambridge University Press

## AUTHORS BIOGRAPHY

**George P. Moustiris** received his M.Eng. in electrical & computer engineering (2003) from the ECE school at the Aristotle University of Thessaloniki (AUTH) and his Ph.D. in robotics and control at ECE/NTUA (2010). From 2011-2013 he carried out his Post-Doctoral research at the Intelligent Robotics & Automation Lab (IRAL) of ECE/NTUA in medical/surgical robotics. He is currently a Research Associate at the Institute of Computers & Communications Systems and member of IRA. His current research interests are focused on medical, surgical and assistive robots, aiming towards automating various functions of their operation. He has authored more than 15 publications and has been involved in many national and international research projects. He is a member of IEEE/RAS and the Technical Chamber of Greece.

**Athanasios Dometios** graduated from the School of Electrical and Computer Engineering (ECE) of the National Technical University of Athens (2013). He completed his Bachelor thesis in the Institute of Automatic Control Engineering (LSR) of the Technical University of Munich (TUM). From 2014 he has joined the Intelligent Robotics and Automation Laboratory of at ECE/NTUA as a Ph.D. Student. His research interests include 3D mapping, path-planning, sensor fusion, SLAM, control of mobile robots and robotic manipulators.

**Costas S. Tzafestas** holds an Electrical and Computer Engineering Degree (1993) from the National Technical University of Athens (NTUA, Greece), as well as a D.E.A. (1994) and Ph.D. (1998) degrees on robotics from the Université Pierre et Marie Curie (Paris 6, France). He is currently an Assistant Professor on Robotics at the School of Electrical and Computer Engineering of NTUA. He has previously served as a Research Fellow at the Institute of Informatics and Telecommunications of the Hellenic National Center for Scientific Research “Demokritos”. His main research interests include human/robot interaction, telerobotics, and haptics, also spanning robust, adaptive and

intelligent control methods with applications in dextrous robotic manipulation, as well as in walking and mobile robots. He has authored and co-authored more than 90 scientific publications and has been a principal investigator and/or technical manager in several national and international research programs. He currently serves as an Associate Editor of the Journal of Intelligent and Robotic Systems. Dr. Tzafestas is a member of the IEEE and of the Greek Technical Chamber.

