Learning by Demonstration for Co-Operative Navigation with Assistive Mobility Devices

James Poon
University of Technology, Sydney, Australia
james.poon@student.uts.edu.au
Jaime Valls Miro
University of Technology, Sydney, Australia
jaime.vallsmiro@uts.edu.au

Abstract

This work proposes a learning by demonstration framework for intuitive navigational adaptation of human-robot interactive mobile systems. Co-navigation algorithms for mobile robots tend to be highly parameterised with variables that may be difficult to manually configure to an individual user's desired behaviours by non-technical personnel, e.g. carers and therapists overseing activities with would-be intelligent power mobility devices. The proposed framework automatically learns suitable joystick inputs for safe handling of the platform from a healthy user aware of a desired subset of behaviours (generic collision avoidance, wall-following and forward/reverse alignment manoeuvres) through performance of a small set of elementary training exercises, without the need and risk associated to trial-and-error variable tuning. The paper compares the semi-autonomous capability of the proposed learning scheme with the popular Vector Field Histogram local planner in a corridor navigation task, showing its ability to safely generalise to different environments despite the simplicity of the training demonstrations.

1 Motivation and Background

With predictions that the global population of people aged over 60 is set to approximately double by 2050 [C.B.O., 2005] there is a drive to improve the means by which services in the field of aged and disabled care are administered, especially towards those still leading active lives in their communities [Clarke and Colantonio, 2005]. In developed countries, people over the age of 65 are often reliant on power mobility devices (PMDs), with approximately 5 million in the United States alone [Kaye *et al.*, 2000]. These are often heavy and powerful machines operating in both indoor and outdoor spaces with proximity to hazards such as furnishings or pedestrians, with which collisions may incur dire consequences. To ensure that only those capable of competently handling a PMD have independent access, tests including the Wheelchair Skills Test [Kirby *et al.*, 2002] are manually conducted by qualified therapists to determine proficiency prior to prescription. As proficiency declines elderly or disabled individuals have to rely closely on carers to engage in everyday activities, resulting in a significant impact to their quality of life and an additional burden upon carers.

Many works have been conducted to introduce PMD systems capable of alleviating the burden of assisted navigation including [Demeester *et al.*, 2006; Taha *et al.*, 2007; Urdiales *et al.*, 2011]. These tend to rely on relatively static environments and goals which are quite clearly established, either as task objectives or likely points of interest such as doorways, furniture and light-switches. As such they generalise poorly to environments differing from those used during the learning phase. Moreover, platforms operating in any local space [Lankenau *et al.*, 1998; Oishi *et al.*, 2011] still require tuning of parameters that are likely alien to therapists/carers and the less technically-oriented. Learning from demonstration (LfD), where a human serves as a teacher for an intelligent learning system, can provide a viable foundation for collaborative human/robot navigation for unbounded objectives and changing environments. LfD is most commonly seen in the domain of robot manipulation; for example in [Amor *et al.*, 2014], where interactions such as high-fives are learned via motion capture and then transferred to the space of a robot arm as primitives. In [Edsinger and Kemp, 2007] a robot is trained to co-operate with a human in order to achieve collaborative tasks such as storing a box on a shelf. Within the scope of mobile robot motion control [Argall *et al.*, 2009] features the learning of three driving primitives (left, straight, right). These primitives can then be readily generalized to new tasks, provided features from an immediate short-term trajectory can be ascertained

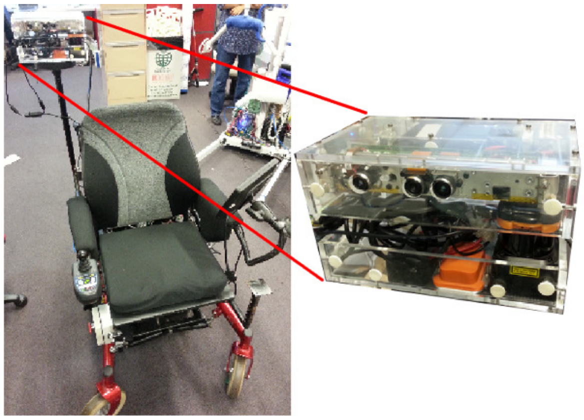

Figure 1: UTS CAS instrumented wheelchair with mounted sensor package

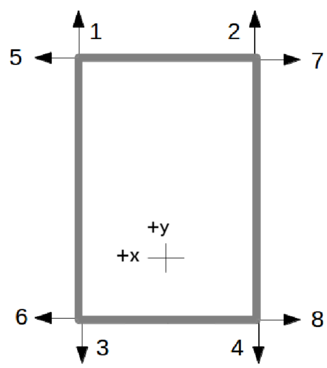

Figure 2: Laser scanner features for training. 1 and 2 face forward, with the cross indicating platform origin and joystick co-ordinate convention

in real-time or known apriori. [Soh and Demiris, 2013] introduces learning of appropriate 'when' and 'how' in assistive mobile navigation by training with a human teacher providing assistance to an autonomous navigation algorithm on a given test course, in order to improve lap times when the learned assistance scheme was integrated. Learning from demonstration thus allows intuitive configuration of intelligent systems to take place, and furthers the possibility of tuning or recalibration of human-interactive systems by non-specialist personnel. Accomplishing this in the space of assistive navigation is the objective of this paper.

2 Proposition

Given the relatively simple action space of a 2D non-holonomic mobile platform in comparison to, for example, a full humanoid robot, we can define a finite set of interaction primitives between the platform and its environment. Common actions include:

- Approaching objects from the front, e.g.moving towards tables and chairs
- Pulling over to objects on left/right sides, e.g. for wall following or docking alongside beds
- Approaching objects from the rear, e.g. for parking in an unobtrusive location

The robustness of a system only equipped with these limited primitives hence comes from the broader definition of what an "interaction" actually entails. In this operational domain it is defined as aligning the PMD appropriately as per one of the above tasks. A two-fold LfD system is thus proposed by this work aimed at simultaneously learning the above action events to assist in its successfull completion, as well as learning the appropriate safe input commands to do so depending on the PMD's surroundings. In effect, this work advocates for the fact that "assistance as needed" is further beneficial to those already familiar with PMD platforms by allowing performance as expected in situations where full control upon the user's part is deemed acceptable.

2.1 User Inputs and Environment Representation

As the most common human interface device (HID) on a PMD is a joystick, in this work this represents the means by which a demonstrator imparts knowledge as well as an intuitive shared medium for human/computer collaboration. Joysticks provide x and y co-ordinates; in the case of differential-drive control these values can be directly mapped into desired linear and angular velocities [1].

The definition of an object, such as a wall, furnishing or even a person, can be redefined in the PMD's 2D perception space to embody contiguous geometric segments along the ground plane; information readily available from sensors such as sonar arrays, depth/stereo imaging or LIDAR. This allows learning from demonstration without explicit attempts at identifying objects of interest. This is in contrast to propositions such as [Derry and Argall, 2013; Poon and Miro, 2014] where attempts have been made towards identifying objects such as doorways or narrow passages for navigational purposes, although these heuristic algorithms remain largely too brittle for real-world deployment. It is also presently intractable to attempt to identify the myriad of different objects and their countless variations that PMD users encounter throughout their everyday.

[1] While this work is presently restricted to simulation, in practice electronic low-pass filters are common features on PMD controllers in order to partially alleviate jerkiness as a result of improper grip or tremoring of the hand/arm. While these increase the level of comfort that a user experiences, it is also true that they introduce additional difficulty to the design and calibration of autonomous motor control schemes

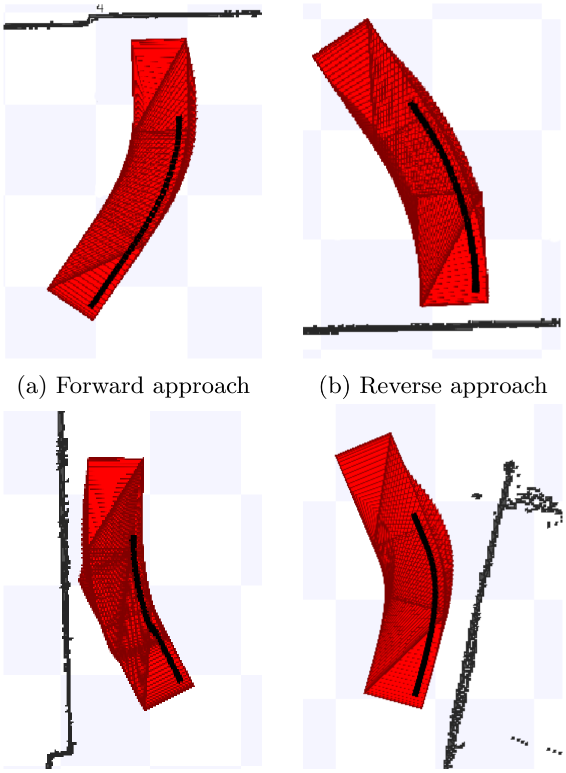

(a) Forward approach (b) Reverse approach

(c) Forward pull-over (d) Reverse pull-over

Figure 3: Demonstrated primitives

3 Learning from Healthy Demonstration Methodology

For learning to take place with minimal demonstration it is vital to select only features that are highly relevant to the task at hand as inputs. Collision avoidance for the front and back requires modulation of forward and reverse linear velocity (+ and − y axes), while collision from the sides requires modulation of angular velocity through the ±x axes. Treating each of the front, back and sides of the PMD as individual surfaces for collision likeliness is reasonable in this scenario; for example, what is immediately ahead of the platform should only affect the maximum safe magnitude of forward linear velocity for either safe bypass or docking. By only considering measurements strictly relevant to the modulation of each individual joystick axis, this approach carries the benefit of removing the need for elaborate demonstrative scenarios or time-consumingly extensive durations of driving. We thus decompose the space around the PMD into eight features from LIDAR as seen in Figure 2, representing distances to obstacles from each corner of the platform's footprint at 0 and 90 degrees with a measurement maximum of 2 metres.

A healthy user demonstrated several basic navigational exercises (Fig 3) representative of the chosen primitives and desired behaviours, with the eight range measurements logged alongside joystick input. These datasets are mirrored around the vertical axis to double the quantity of training data. Table 1 shows the features chosen for the modulation of each joystick axis corresponding to their cardinal sides of the PMD, as well as for modulation of angular velocity via the x axis for the three primitives. This minimalist approach to feature selection increases the robustness of each network when only a minimal amount of demonstrative training data is supplied. Each of the four joystick axes has its own network trained to yield the maximum magnitude for which the healthy user deemed to be safe when experiencing a particular observation set during demonstration, and each primitive also has its own network yielding an appropriate x axis value for action completion for a total of seven neural networks. Switching from axis modulation to primitive action completion is triggered when an action completion feature pair meets thresholds of mean and absolute difference α and β, as well as γ representing a magnitude threshold of joystick movement in the direction of the primitive's corresponding side of the PMD. These parameters are intuitively inferrable for individual users depending on their level of driving proficiency, embodying tolerances of distance and initial alignment suitable for a disabled user to indicate when autonomous alignment should commence. The user's requested linear

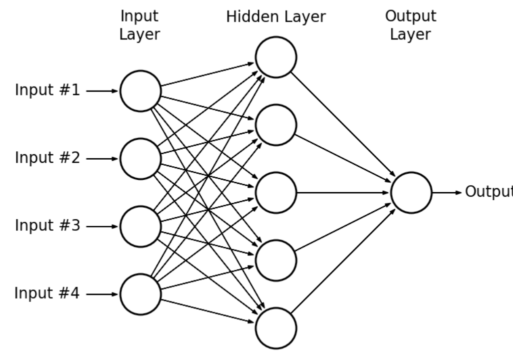

Figure 4: Example of a feed-forward Artificial Neural Network [Vanderplas, 2013]

Table 1: Feature Sets vs Modulated Axes for Neural Network Training

Features	Joystick Axis Modulated	Primitive
1,2	+y	fwd approach
3,4	−y	rev approach
5,8	+x	
6,7	−x	
5,6/7,8		wall follow

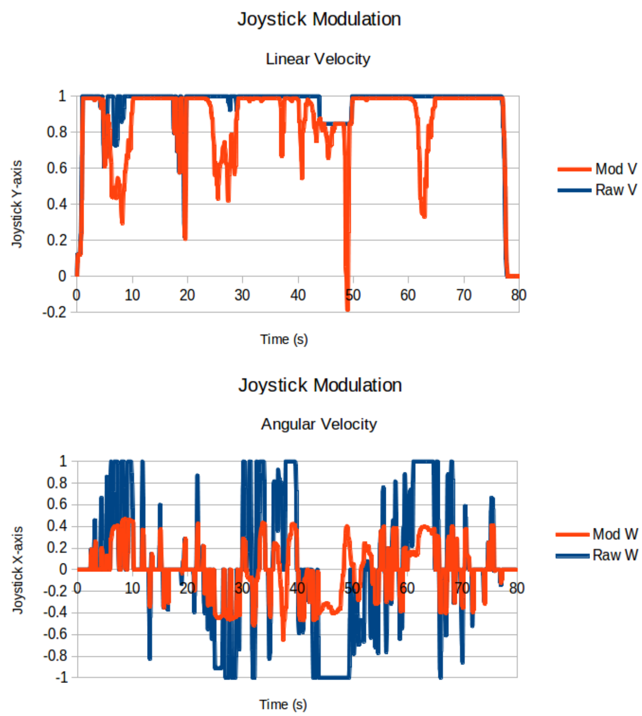

Figure 5: Joystick capping of user (blue) to safer orange

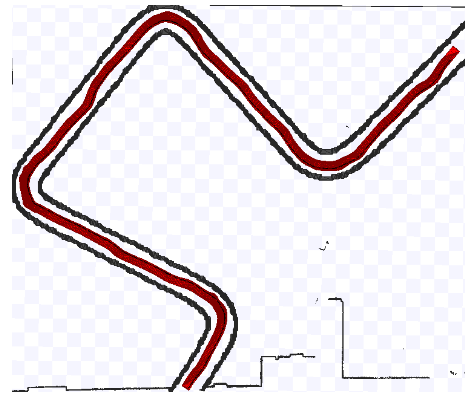

Figure 6: Path under input capping in Stage

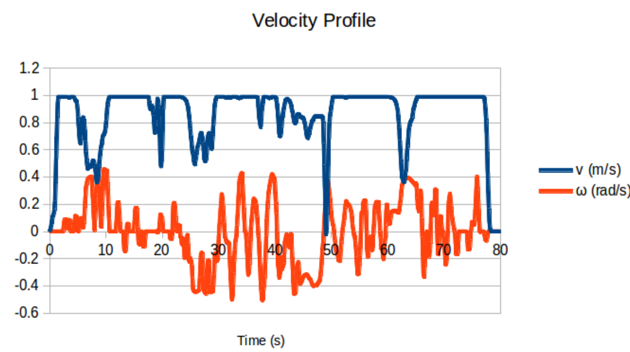

Figure 7: Velocity profiles for Figure 6

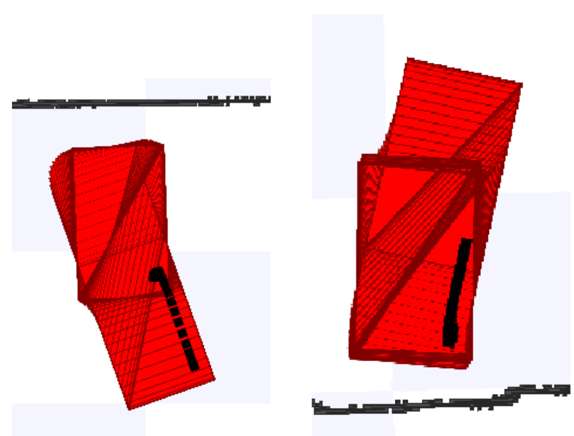

(a) Forward approach (b) Reverse approach

Figure 8: Learned primitive docking manoeuvres

velocity is thus always bounded by the y axes modulation limits, and their requested angular velocity either similarly bounded or substituted by an appropriate aligning velocity.

Neural networks [Maind and Wankar, 2013] are cognitive algorithms that attempt to replicate brain neuron behaviour in order to provide a machine learning mechanism for universal function approximation, which can take the form of regression or classification problems. Each neuron outputs a nonlinear transform of the weighted sum of the previous layer's outputs. A key factor of interest in these models is their ability to learn from input data, so that they can be re-trained in an efficient manner as more data becomes available. Figure 4 shows an example of a small neural network. For the experiments presented in this paper ANNs were trained using the iRProp+ heuristic update algorithm [Igel and Husken, 2000], chosen for its speed of convergence over traditional stochastic gradient descent.

4 Data Gathering

Preliminary experimentation was conducted within Stage [Vaughan, 2008] robot simulation software, configured to replicate the properties and behaviour of the UTS CAS intelligent wheelchair. The chair (Fig 1) is fitted with a sensor module housing a MS Kinect RGB-D camera, Hokuyo planar laser scanner and an Xsens inertial measurement unit. More information covering the details of the sensor module can be found in [Miro *et al.*, 2011]. A map of a basement level of UTS was built using Hector mapping [Kohlbrecher *et al.*, 2011] within the ROS middleware (`www.ros.org`) and served as our training environment.

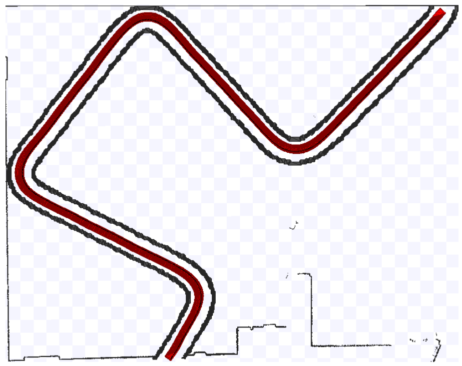

Figure 9: Tuned VFH trajectory under minimal user input

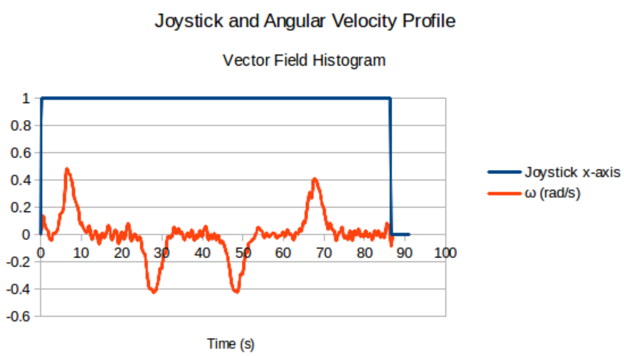

Figure 10: Profile corresponding to Figure 9

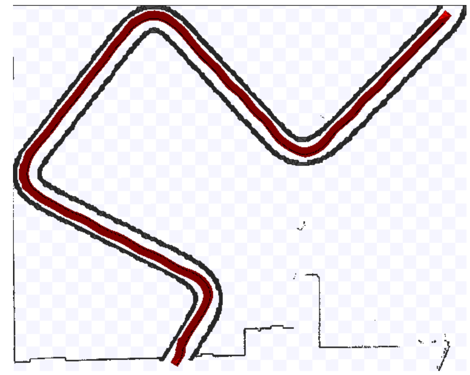

Figure 11: Learned capping trajectory under minimal user input

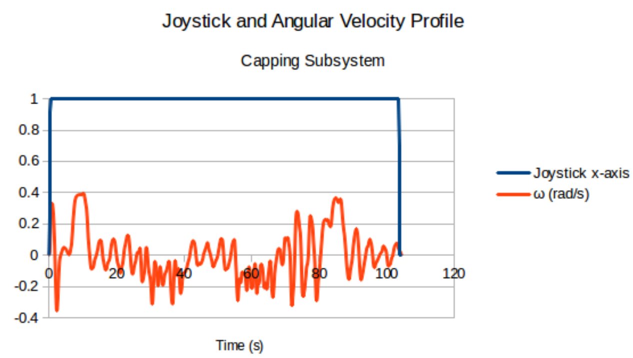

Figure 12: Profile corresponding to Figure 11

5 Results

The joystick capping was tested in a tight corridor measuring approximately twice the width of the wheelchair platform, adapted from the test course seen in [Argall *et al.*, 2009]. Figure 5 shows the joystick safety cappings that resulted in the trajectory and linear/angular velocity profile seen in Figures 6-7. Much of the able-bodied test user's deliberately poor driving was negated by the capping neural networks, and higher-frequency tremoring was largely diffused by the simulation's low-pass filtering for an overall collision-free, more comfortable drive. Figure 8 shows examples of the primitve action completion subsystem when the joystick was simply held in full forward and backward positions respectively, with α and β set to 0.5 and 0.4 metres for demonstration purposes.

5.1 Comparison with Vector Field Histogram

The Vector Field Histogram (VFH) [Borenstein and Koren, 1991] (Fig 13) is among the most well-known local mobile robot controllers. By constructing a histogram representation of immediate surrounding space, an angular target heading can be selected from a set of parameters describing physical platform characteristics and desired behaviour. Some such as size and the width of a histogram smoothing filter l are more intuitive than others, such as angular sector width α and histogram magnitude parameters a and b. Although readily tuneable for virtually any autonomous platform given a person familiar with VFH to obtain a powerful collision avoidance layer, it would be much more time-consuming to customize for a PMD user to take into account their individual preferences than to simply impart desirable behaviours via manual demonstration.

Figures 9,11 show the trajectories taken by both systems to navigate the same course along the left wall, at a constant desired linear velocity of 1.0 m/s. Both systems were able to complete the course with minimal user indications (Figs 10,12). Despite the forward approach neural network being disabled for this trial the VFH still finished faster due to taking full linear velocity around turns, as well as its generation of a smoother near-optimal trajectory owing to its calibration prior to the exercise. More contemporary control strategies such as Dynamic Window Approach [Fox *et al.*, 1997] and Nearness Diagram navigation [Minguez and Montano, 2004] unfortunately also involve significant amounts of parameters to tune, and thus suffer from the same drawbacks for personalization of mobility aids.

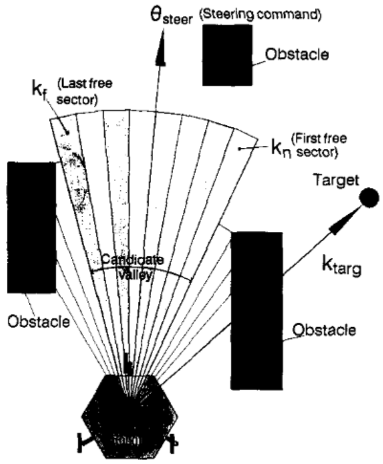

Figure 13: Vector Field Histogram local controller [Borenstein and Koren, 1991]

6 Conclusions and Future Work

This work demonstrates an approach for tuning desirable behaviours of assistive mobility devices without the need for manually tweaking parameters which would be risky/daunting in this context for non-technical personnel to attempt, and time-consuming for a professional appointment to calibrate a near-ideal parameter set by hand. By implementing a learning from demonstration framework taking the form of several small artificial neural networks, this parameter set can be automatically derived from several short driving exercises that can be easily conducted by a healthy person aware of platform handling and the behaviours best suited to the end-user. Enabling generic collision avoidance and elementary navigational manouevres will allow users of these devices to experience more independence and a greater quality of life, while simultaneously freeing up carer time from mobility supervision for other aspects of healthcare.

Future work will include a GUI for the demonstration process in order to streamline the manner by which the neural networks are trained, as well as potential additional navigational aids for handling doorways and other difficult situations for truly active navigational assistance. Further learning towards automated initiation of action completion primitives to replace the heuristic triggering parameters is additionally desirable. We also hope to further this work with a clinical trial utilizing a broad user base under the oversight of qualified care staff in order to gauge efficacy and ascertain insights into advancing this approach.

References

[C.B.O., 2005] Congressional Budget Office (C. B. O.), *Global population aging in the 21st century and its economic implications.* `http://www.cbo.gov/sites/default/files/cbofiles/ftpdocs/69xx/doc6952/12-12-global.pdf`. Accessed: Aug 23, 2015.

[Clarke and Colantonio, 2005] P. Clarke and A. Colantonio, *Wheelchair use among community dwelling older adults: prevalence and risk factors in a national sample.* Canadian Journal on Aging, 24:191–198, 2005.

[Kaye *et al.*, 2000] H. S. Kaye, T. Kang and M. P. LaPlante, *Mobility device use in the united states. Disability statistics report 14.* Disability Statistics Center, University of California, San Francisco, 2000.

[Kirby *et al.*, 2002] R. L. Kirby, J. Swuste, D. J. Dupuis, D. A. MacLeod, and R. Monroe, *The wheelchair skills test: A pilot study of a new outcome measure.* Archives of Physical Med. and Rehab., 83(1):10–18, 2002.

[Demeester *et al.*, 2006] Demeester, E. and Huntemann, A. and Vanhooydonck, D. and Vanacker, G. and Degeest, A. and Van Brussel, H. and Nuttin, M. *Bayesian Estimation of Wheelchair Driver Intents: Modeling Intents as Geometric Paths Tracked by the Driver.* Intelligent Robots and Systems, 2006 IEEE/RSJ International Conference on, 5775–5780, Oct 2006.

[Taha *et al.*, 2007] Taha, T. and Miro, J.V. and Dissanayake, G. *Wheelchair Driver Assistance and Intention Prediction using POMDPs.* Intelligent Sensors, Sensor Networks and Information, 2007. ISSNIP 2007. 3rd International Conference on, 449–454, Dec 2007.

[Urdiales *et al.*, 2011] Urdiales, C. and Peula, J.M. and Fdez-Carmona, M. and Barru, C. and Prez, E.J. and Snchez-Tato, I. and del Toro, J.C. and Galluppi, F. and Corts, U. and Annichiaricco, R. and Caltagirone, C. and Sandoval, F. *A new multi-criteria optimization strategy for shared control inwheelchair assisted navigation.* Autonomous Robots, 30(2):179–197, 2011.

[Lankenau *et al.*, 1998] Lankenau, A. and Meyer, O. and Krieg-Bruckner, B. *Safety in robotics: the Bremen Autonomous Wheelchair.* Advanced Motion Control, 1998. AMC '98-Coimbra., 5th International Workshop on, 524–529, Jun 1998.

[Oishi *et al.*, 2011] Oishi, M. and Cheng, A. and Bibalan, PT. and Mitchell, I. *Building a smart wheelchair on a flexible software platform* RESNA International Conference on Technology and Aging, 2011.

[Amor *et al.*, 2014] Ben Amor, H. and Neumann, G. and Kamthe, S. and Kroemer, O. and Peters, J. *Interaction primitives for human-robot cooperation tasks.* Robotics and Automation (ICRA), 2014 IEEE International Conference on, 2831–2837, 2014.

[Edsinger and Kemp, 2007] Edsinger, A. and Kemp, C.C. *Human-Robot Interaction for Cooperative Manipulation: Handing Objects to One Another*. Robot and Human interactive Communication, 2007. RO-MAN 2007. The 16th IEEE International Symposium on, 1167–1172, Aug 2007.

[Argall *et al.*, 2009] Argall, Brenna and Browning, Brett and Veloso, Manuela. *Learning Mobile Robot Motion Control from Demonstrated Primitives and Human Feedback*. Springer Tracts in Advanced Robotics, 70:417–432, 2011.

[Soh and Demiris, 2013] Soh, H. and Demiris, Y. *When and how to help: An iterative probabilistic model for learning assistance by demonstration*. Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on, 3230–3236, Nov 2013.

[Derry and Argall, 2013] Derry, M. and Argall, B. *Automated doorway detection for assistive shared-control wheelchairs*. Robotics and Automation (ICRA), 2013 IEEE International Conference on, 1254–1259, May 2013.

[Poon and Miro, 2014] Poon, J. and Miro, J.V. *A Multimodal Utility to Assist Powered Mobility Device Navigation Tasks*. Social Robotics, 8755:300–309, 2013.

[Vaughan, 2008] Vaughan, R. *Massively Multiple Robot Simulations in Stage*. Swarm Intelligence, 2(2–4):189–208, 2008.

[Miro *et al.*, 2011] Miro, J.V. and Black, R. and de Bruijn, F. and Dissanayake, G. *Semi-autonomous competency assessment of powered mobility device users*. IEEE Intl. Conference on Rehab. Robotics, 1–6, 2011.

[Kohlbrecher *et al.*, 2011] Kohlbrecher, S. and Meyer, J. and von Stryk, O. and Klingauf, U. *A Flexible and Scalable SLAM System with Full 3D Motion Estimation*. Proc. IEEE International Symposium on Safety, Security and Rescue Robotics, 2011.

[Vanderplas, 2013] Vandelplas, J., *Neural network diagram*. `http://www.astroml.org/book_figures/appendix/fig_neural_network.html`. Accessed: Aug 24, 2015.

[Maind and Wankar, 2013] Maind, S.B. and Wankar, P. *Research paper on basic of artificial neural network*. International Journal on Recent and Innovation Trends in Computing and Communication, 2(1):96–100, 2014.

[Igel and Husken, 2000] Igel, C. and Husken, M. *Improving the Rprop Learning Algorithm*. Second International Symposium on Neural Computation, 115–121, 2000.

[Borenstein and Koren, 1991] Borenstein, J. and Koren, Y. *The vector field histogram-fast obstacle avoidance for mobile robots*. Robotics and Automation, IEEE Transactions on, 7(3):278–288, 1991.

[Fox *et al.*, 1997] Fox, D. and Burgard, W. and Thrun, S. *The dynamic window approach to collision avoidance*. Robotics Automation Magazine, IEEE, 4(1):23-33, 1997.

[Minguez and Montano, 2004] Minguez, J. and Montano, L. *Nearness diagram (ND) navigation: collision avoidance in troublesome scenarios*. Robotics and Automation, IEEE Transactions on, 20(1):45–59, 2004.