A Probabilistic Representation of User Intent for Assistive Robots

Matthew Derry[1] and Brenna Argall[2]

I. Introduction

Many situations exist in which the human control of a robot or machine can be difficult. For individuals with certain motor impairments—like spasticity or paralysis—driving powered wheelchairs can be challenging if not impossible [9], navigating telepresence robots is made all the more difficult by the reduced embodiment experienced by the user, and controlling machines that have more degrees of freedom (DoF) than the control interface, (e.g. a 6-DoF robot arm with a 3-DoF joystick) can be cumbersome. These situations can all benefit from introducing automation to assist the user in accomplishing difficult tasks. This assistance can come in many forms, from corrective inputs augmenting the user input to the complete assumption of control by the automated process. Understanding when and how much assistance to provide can be enhanced by an accurate estimate of the goal the the human might be trying to achieve. Research into user intent prediction and goal estimation aims to address this problem, which remains an open research area.

In order to produce an estimate of a user's intended goal, the goal hypotheses can come from analyzing the environment [6], modelling the user's control inputs to predict what controls they might issue to reach a goal [7], [8], using inverse optimal control to learn a cost function that captures the user's behavior and comparing plans to different goals using the learned cost function [5], or confidence metrics based on inverse models of robot state [2].

In our proposed framework, we combine the environmental structure via Voronoi graphs [1], which provide a global context by identifying homotopically distinct paths to global goals, with a general local driving model to identify the intended local subgoals represented by nodes in the Voronoi graph. In this preliminary work we have implemented the local model as a Gaussian Mixture Model (GMM), and use Gaussian Mixture Regression (GMR) to estimate the probability that a local goal is the target, given some input (e.g. user commands or robot state). The result is a probabilistic representation of the goals a user might be trying to achieve. Prior research in this area has both used graph representations of the environment as well as driving models of individual users [3]. One aspect in which our work differs is in the development of a general, user-agnostic driving model for local goals, combined with a graph representation of the environment for global context. The result is a framework for estimating the intended goal of a user without requiring training for specific individuals or an *a priori* map.

This work is developed specifically in support of research in the assistive control of powered wheelchairs. In this context, global goals might include doorways for traversing, tables for docking or ramps for boarding a vehicle in addition to many others. The perception of doorways and safe docking locations at tables has been developed as a part of the larger scope of our smart wheelchair research [4]. Our framework would be appropriate to extend to other domains such as telepresence robots or robot arm teleoperation.

II. Proposed Framework and Proof of Concept

Our framework consists of two major components: the global context provided by a Voronoi graph of the current map, and a local model for estimating the probability that a local subgoal (i.e. any Voronoi node adjacent to the Voronoi cell occupied by the robot) is the intended subgoal. The local subgoals then serve as the starting points of a search of the Voronoi graph in which the global goals are the end points.

We evaluated the algorithm in a simulated environment mirroring a lab setting (Figure 2). Four global goals were placed in the scene, two in doorways and two at desks. The user was asked to drive to a particular goal and the estimated intended goal was recorded during the task execution.

A. Global Context

Paths are homotopically distinct when there is no continuous transformation between them. This property allows for a more compact representation of meaningful paths and decision points (i.e. Voronoi nodes) without encoding the infinite number of paths that can occur between two points.

Voronoi graphs provide a convenient method for identifying homotopically distinct paths that exist in a map. The edges of Voronoi graphs are equidistant from the two nearest obstacles, and the nodes of Voronoi graphs are equidistant from the three nearest obstacles. Areas enclosed by Voronoi edges are known as Voronoi regions. Figure 2 inset shows an example of a Voronoi graph with the goals inserted in to the graph, shown as green dots.

B. Local Model

The local model in our framework is built as a Gaussian Mixture Model (GMM). A weighted mixture of Gaussian distributions is fit to a series of training data $\boldsymbol{\xi} = [\boldsymbol{\xi_i}; \boldsymbol{\xi_o}]$.

$$P(\boldsymbol{\xi}|\boldsymbol{\theta}) = \sum_{k=1}^{K} w_k \mathcal{N}(\boldsymbol{\xi_i}, \boldsymbol{\xi_o}; \boldsymbol{\theta^k}) \quad (1)$$

[1]M. Derry is with the Department of Electrical Engineering and Computer Science, Northwestern University, Evanston, IL 60208, USA mderry@u.northwestern.edu

[2]B. Argall is jointly appointed with the Faculty of Electrical Engineering and Computer Science, Northwestern University, Evanston, IL 60208, USA and the Rehabilitation Institute of Chicago, Chicago, IL 60611, USA brenna.argall@northwestern.edu

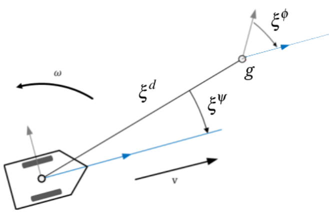

Fig. 1. Ego-centric coordinates $\boldsymbol{\xi}$ of goal g used in the proof-of-concept, local model formulation. Curvature is calculated using the commanded rotational velocity ω and commanded linear velocity v.

where $\boldsymbol{\xi_i}$ is the input, $\boldsymbol{\xi_o}$ is the output, K is the number of Gaussian components and $\boldsymbol{\theta}^k$ contains the prior w, mean μ, and covariance Σ for the kth Gaussian component. The parameters $\boldsymbol{\theta}$ are estimated with Expectation Maximization on a dataset of expert demonstrations driving to several goals.

With the parameters of the GMM estimated, GMR is used to compute the resulting conditional probability distribution, which determines the likelihood that the output data $\boldsymbol{\xi_o}$ was generated by the GMM, given the input data $\boldsymbol{\xi_i}$:

$$P(\boldsymbol{\xi_o}|\boldsymbol{\xi_i};\boldsymbol{\theta}) \sim \mathcal{N}(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}}) \tag{2}$$

In our proof-of-concept formulation for the local model, the input data $\boldsymbol{\xi_i}$ is the egocentric position of a subgoal (Figure 1), which contains the distance from the robot to the subgoal ξ_i^d, the difference between the robot heading and the heading to the subgoal ξ_i^ψ, and the difference between the robot heading and the subgoal orientation ξ_i^ϕ. The output data $\boldsymbol{\xi_o}$, is the current user-commanded curvature.

We assessed this proof-of-concept formulation in a pilot user study with 5 participants, each performing 12 trials. The results of this pilot study were encouraging, with the estimate of the user's goal frequently matching the true goal. However, the study also highlighted many areas for improvement. For example, when close to local subgoals, the user does not actually tend to drive to the subgoal with much precision since the subgoal, being a nearby Voronoi node, is an artificial construct not necessarily apparent to the user, whereas the global goal—such as a doorway—is a natural goal that could be perceived by both user and robot. This causes the local model to give the intended local subgoal a low probability due to the large swing in the heading *to* subgoal component of the input data, resulting in undesirable behavior near the Voronoi edges.

An alternative formulation which we are currently developing is to use the distance to the goal as the input data $\boldsymbol{\xi_i}$, and as the output data $\boldsymbol{\xi_o}$ the heading *to* goal and heading *of* goal components of the egocentric polar coordinate transformation of the local subgoal. The hypothesis is that the model will encode the high variance in the heading parameters seen when close to the subgoal, which can then be leveraged in the probability computation. An additional benefit is that by removing the user-commanded curvature from the formulation, the number of demonstrations required is reduced considerably. The trade-off is that the intent estimation will lag slightly since the probabilities will update only after the robot moves.

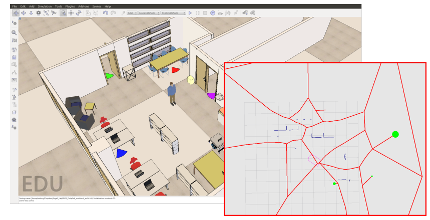

Fig. 2. A simulated lab environment, inset with a voronoi segmentation.

C. Combining the Local Model with Global Context

The last step is to apply the probability distribution over local subgoals to the global goals. In our proof-of-concept formulation, the shortest path from the current voronoi region to each global goal is found. Then each global goal is simply assigned the probability of the local subgoal that starts the shortest path to that global goal. Our pilot study highlighted the issue that this formulation suffered poor performance when the user's chosen path did not correspond to the shortest path along the Voronoi edges. To address this, an alternative formulation currently under development is to build a true probability distribution weighted by path length over all paths to a global goal.

III. Conclusion

Accurate estimation of user intent has the potential to improve the performance of assistive technologies that automate some portion of task execution. This work presents a framework that combines global context about the environment with a local probabilistic model to make estimates about user intent. This initial work has shown both promise and directions for improvement. The alternative formulations for the local model and application to the global goals, identified in Section II, are under active development, and their evaluation will be included in the workshop presentation.

References

[1] J. Canny and B. Donald. Simplified voronoi diagrams. *Discrete & Computational Geometry*, 3(1):219–236, 1988.

[2] T. Carlson and Y. Demiris. Human-wheelchair collaboration through prediction of intention and adaptive assistance. In *Proc. of IEEE International Conference on Robotics and Automation (ICRA'08)*, pages 3926–3931, 2008.

[3] E. Demeester, A. Hüntemann, D. Vanhooydonck, G. Vanacker, H. Van Brussel, and M. Nuttin. User-adapted plan recognition and user-adapted shared control: A bayesian approach to semi-autonomous wheelchair driving. *Autonomous Robots*, 24(2):193–211, 2008.

[4] M. Derry and B. Argall. Automated doorway detection for assistive shared-control wheelchairs. In *Proc. of IEEE International Conference on Robotics and Automation (ICRA'13)*, pages 1254–1259, 2013.

[5] A. Dragan, G. J. Gordon, and S. Srinivasa. Learning from experience in manipulation planning: Setting the right goals. 2011.

[6] R. Glinton, S. Owens, J. Giampapa, K. Sycara, M. Lewis, and C. Grindle. Intent inference using a potential field model of environmental influences. In *Proc. of the IEEE International Conference on Information Fusion (Fusion'05)*, volume 2, pages 8–pp, 2005.

[7] K. Hauser. Recognition, prediction, and planning for assisted teleoperation of freeform tasks. *Autonomous Robots*, 35(4):241–254, 2013.

[8] M. Kuderer, H. Kretzschmar, C. Sprunk, and W. Burgard. Feature-based prediction of trajectories for socially compliant navigation. In *Robotics: Science and Systems*, 2012.

[9] R. C. Simpson, E. F. LoPresti, and R. A. Cooper. How many people would benefit from a smart wheelchair? *Journal of Rehabilitation Research and Development*, 45(1):53–71, 2008.