

Individualized Performance Quantification of Power Wheelchair Driving

Tim D. Yang Ameya Patil Yih-Kuen Jan

(Honorable Mention Award)

Abstract—Power wheelchair users may present secondary impairments that inhibit effective joystick control. Accordingly, driving assessments are needed to ensure the safety of both the drivers and their surroundings. We introduce a quantitative method to evaluate power wheelchair driving performance using the Markov decision process. Our dynamic, stochastic approach supplements previous quantitative evaluations by accounting for individual differences between drivers. For example, a driver who frequently pulls to the left on accident might stay far from the left wall and make extremely wide left turns, but these inefficiencies may actually be appropriate given the context. We include preliminary results to demonstrate this assessment method.

I. INTRODUCTION

Power wheelchairs provide autonomy and promote self-reliance in people with severe mobility impairments. In 1995, powered mobility devices were used by an estimated 291,000 Americans and exhibited growth trends of 5% per year [1]. Based on these projections, there could be well over a million powered mobility devices in use by 2020. Since power wheelchair users may also present accompanying cognitive, sensory, and motor impairments, driving assessments are needed to ensure the safety of both the drivers and their surroundings.


Current clinical evaluations are generally qualitative in nature. For example, pediatric wheelchairs are prescribed based on clinical intuition and “trial and error” [2] or based on assessment batteries, such as the Pediatric Powered Wheelchair Screening Test [3]. For general power wheelchair evaluation, the Wheelchair Skills Test includes a series of common driving tasks that are rated by clinicians [4]. To evaluate driving performance more objectively, some studies have begun quantifying joystick metrics with descriptive statistics: reaction time, total driving time, and target accuracy [5]; movement error, movement variability, average speed, and average acceleration [6]; and joystick activations, total driving time, and joystick directional variability [7].

We thank Permobil, Inc. for providing the power wheelchair and the S. N. Bose Scholars Program for facilitating scientific collaboration between the U. S. and India.


T. D. Yang is with the Department of Kinesiology & Community Health and Program in Computational Science & Engineering, University of Illinois at Urbana-Champaign, Champaign, IL 61820, USA (tyang15 at illinois.edu).

A. Patil is with the Department of Electrical Engineering, Indian Institute of Technology, Hyderabad, India (ee10b030 at iith.ac.in).

Y. K. Jan is with the Department of Kinesiology & Community Health, Program in Computational Science & Engineering, and Center on Health, Aging, & Disability, University of Illinois at Urbana-Champaign, Champaign, IL 61820, USA (217-300-7253; yjan at illinois.edu).



(a) Deterministic



(b) Nondeterministic

Fig. 1. A comparison of deterministic and nondeterministic decisions for reaching the positive block while avoiding the negative block. (a) The deterministic agent always chooses the shortest path. (b) The nondeterministic agent chooses a longer path to lower the risk hitting the negative block by avoiding the narrow corridor.

However, as a dynamic, stochastic, and nonlinear task [8], human driving contains complexities that are difficult to express with descriptive statistics and ratings. Because driving involves real-time adaptation based on both personal skill and environmental factors, optimal driving decisions are contingent upon each specific individual and the current surroundings. Individual differences (e.g., in fine motor skill) inject nondeterminism into the driving process.

To observe how this stochastic factor affects optimal control decisions, consider a simplified grid environment in which the goal at any location is to reach the positive block while avoiding the negative block (Fig. 1). When a purely deterministic agent makes decisions, those decisions always yield the expected outcome. For example, if it attempts to move forward, it does not have to worry about accidentally moving orthogonally to its left or right. This allows the deterministic agent to use the shortest path distance to its goal. However, when a

TABLE I
MDP FORMULATION

MDP Component	Corresponding Driving Component
State s	Wheelchair position
Action a	Wheelchair heading
Reward $R(s)$	Environmental context
Transition $P(s' s, a)$	Uncertainty

nondeterministic agent makes decisions, those decisions do not always yield the expected outcome. In this example, the agent has a high probability of moving orthogonally against its intended direction, so it avoids the narrow corridor rather than taking the risk of accidentally falling into the orange block. Although this results in a longer path distance, the nondeterministic agent is still making optimal decisions within the context of its personal skill and situational environment. In stochastic scenarios such as this, the optimal decision always involves a careful balance between the movement probabilities and movement rewards. This type of stochastic decision optimization problem can be formalized as a first-order Markov decision process (MDP).

II. METHODS

The MDP framework [9] consists of a set of states (e.g., wheelchair positions), set of actions (e.g., wheelchair heading vectors), probability distribution of movement transitions (e.g., movement uncertainty), and reward function (e.g., environmental cues). These components allow us to model the task of driving as a sequence of MDP decisions, accounting for differences in driving skill and environmental contexts (Table I). We then develop an individualized performance measure by comparing the driver's actual actions with those computed from the MDP.

A. Mapping Wheelchair Headings into MDP Actions

Unlike cars which turn via differential gearing, power wheelchairs turn via differential steering (Fig. 2), such that each wheel's velocity is explicitly governed. For example, the power wheelchair turns to the right by making the left wheel rotate faster than the right wheel. To estimate the driver's actions, joystick positions are first mapped into wheel velocities (Fig. 3) and then converted into heading vectors using rigid body mechanics.

B. Mapping Uncertainty into MDP Transition Probabilities

Driving is a stochastic process in large part because the result of each decision is always uncertain. Intended decisions do not perfectly manifest themselves in the physical domain, especially for people with impaired fine motor control. The transition probabilities represent an individualized model for a person's driving tendencies. For this preliminary report, we use a simulated transition model, such that the driver is assumed to have an 80% probability of completing the intended action (Fig. 4).




Fig. 2. In a modern car, differential gears allow its wheels to freely adapt to the needed speeds when negotiating turns. In power wheelchairs, the wheels turn via differential steering, such that each wheel's velocity is explicitly governed.




Fig. 3. An overhead diagram of the joystick. The joystick's Cartesian coordinates were mapped to pairs of wheel velocities based on the output voltages. The values for v_1 , v_2 , v_3 , θ_1 , and θ_2 were obtained empirically.

C. Mapping Environmental Context into MDP Rewards

We use the Kinect sensor (Microsoft Corporation, Redmond, WA) to sense the environment. The Kinect captures color and depth images within a 57° horizontal field of view and 4,000 mm radial distance. Based on the joystick actions and movement transitions, we model the MDPs states as a 3-branch, 6-depth graph (Fig. 5). Each branch represents one of the three joystick actions (left, middle, right) that can be chosen at a given state.

D. Optimizing the MDP Policy

In the MDP framework, decisions can be optimized into a policy of actions that maximizes the agent's expected utility using a simplified Bellman equation,

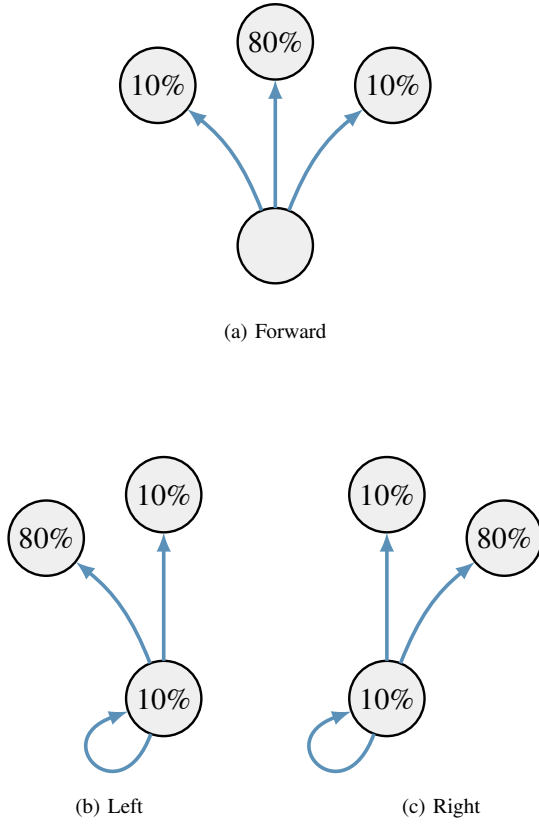


Fig. 4. The transition values represent movement probabilities when attempting to move (a) forward, (b) left, and (c) right.

$$U(s) = R(s) + \max_a \sum_{s'} P(s'|s, a) U(s'), \quad (1)$$

where $U(s)$ is the utility value for a state s , $R(s)$ is the reward for a state s , and $P(s'|s, a)$ is the transition probability of reaching state s' given the current state s and action a . Optimization of the utility values is performed using backward induction [9].

E. Assessing the Driving Performance

Although the backward induction algorithm computes a utility-optimal policy, there is nothing inherently wrong with suboptimal decisions. Mathematically, it simply means that the driver can no longer achieve the highest possible utility. Practically, the impact could range anywhere from negligible to catastrophic. The severity of the deviations can be quantified in terms of the resulting expected utility.

At the root node of each graph, the computed and observed actions are compared based on their resulting expected utilities. A risk index is calculated by normalizing the utility difference against the optimal utility,

$$I(t) = \frac{U_{\text{obs}}(t) - U_{\text{mdp}}(t)}{U_{\text{mdp}}(t)}, \quad (2)$$




Fig. 5. An overhead depth image of two obstacles. The graph is overlaid on the depth image. All graph nodes of obstacles and their children are assigned negative values, inversely weighted by the distance from the sensor.

where $I(t)$ is the risk index at time t , $U_{\text{obs}}(t)$ is the utility value of the observed action at the root node at time t , and $U_{\text{mdp}}(t)$ is the utility value of the MDP action at the root node at time t . Thus, when the computed and observed actions match, the risk index evaluates to zero; otherwise, riskier mismatches evaluate to higher values.

III. RESULTS

We performed preliminary experiments on a C400 power wheelchair (Permobil, Inc., Lebanon, TN). Fig. 6 presents a selected time series of the calculated risk indices. Sample camera images, which have been superimposed with computed and observed heading vectors, are shown for both high-risk (Fig. 7) and low-risk (Fig. 8) situations.

IV. DISCUSSION

We used a simulated transition model that assumed the driver to possess a relatively high probability (20%) of mistakenly moving to the right or left of each intended direction. Fig. 7 shows an example of a high-risk mismatch between the computed MDP action (move left) and the observed action (move forward). For an able-bodied driver, the observed action might not be considered risky. However, the MDP was provided a simulated model of a relatively unpredictable driver, causing the optimization algorithm to compute a left action. For such a driver, attempting to move straight might accidentally cause the driver to move to the right and collide with the extremely proximate door. Fig. 8 illustrates another mismatch. Again, the MDP accounted for the model's driving uncertainty and computed a conservative optimal policy to perform a right turn away from the wall. However, the risk index was negligible. In this case, the wall was more distant than in Fig. 7, so the expected utility difference between the




Fig. 6. Risk indices from one driving trial.




Fig. 8. Example of a low risk index, $I(31.0 \text{ s}) = 0.01$.




Fig. 7. Example of a high risk index, $I(12.5 \text{ s}) = 0.59$.

mismatched actions was lower. These examples illustrate the model's ability to account for both driving tendencies and environmental cues.

Although our reported examples focused on driving decisions, the value of individualized driving evaluation extends beyond the explicit path decisions. For example, a driver may actually be executing optimal driving decisions, even if the joystick activation counts or driving durations are not absolutely minimal. As depicted earlier in the simplified grid example, optimal paths may be altered significantly based on a driver's tendencies. Perhaps a driver, caregiver, or clinician has noticed that the driver frequently pulls to the left accidentally. Such a driver might instinctively stay farther from the left wall and/or make extremely wide left turns, but these inefficiencies may actually be desirable.

Our MDP quantification method provides awareness of situational context (i.e., reward function via the Kinect data) and individual differences (i.e., transition model). In this study, we used a simulated transition model as a demonstration of the method. In future work, we plan to empirically estimate the transition model of each driver's joystick control, allowing the MDP to compute individualized optimal policies as a reference to evaluate each driver's performance.

REFERENCES

- [1] M. P. LaPlante and H. S. Kaye, "Demographics and trends in wheeled mobility equipment use and accessibility in the community," *Assistive Technology*, vol. 22, no. 1, pp. 3–17, Spring 2010. doi: 10.1080/10400430903501413
- [2] P. Guerette, D. Tefft, and J. Furumasu, "Pediatric powered wheelchairs: Results of a national survey of providers," *Assistive Technology*, vol. 17, no. 2, pp. 144–158, Fall 2005.
- [3] J. Furumasu, P. Guerette, and D. Tefft, "Relevance of the Pediatric Powered Wheelchair Screening Test for children with cerebral palsy," *Developmental Medicine and Child Neurology*, vol. 46, pp. 468–474, 2004.
- [4] 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 Medicine and Rehabilitation*, vol. 83, no. 1, pp. 10–18, 2002. doi: 10.1053/apmr.2002.26823
- [5] B. E. Dicianno, D. M. Spaeth, R. A. Cooper, S. G. Fitzgerald, and M. L. Boninger, "Advancements in power wheelchair joystick technology: Effects of isometric joysticks and signal conditioning on driving performance," *American Journal of Physical Medicine and Rehabilitation*, vol. 85, pp. 631–639, Aug 2006. doi: 10.1097/01.phm.0000228519.54763.b4
- [6] B. E. Dicianno, H. Mahajan, A. S. Guirand, and R. A. Cooper, "Virtual electric power wheelchair driving performance of individuals with spastic cerebral palsy," *American Journal of Physical Medicine and Rehabilitation*, vol. 91, pp. 823–30, Oct 2012. doi: 10.1097/PHM.0b013e31825a1497
- [7] G. U. Sorrento, P. S. Archambault, F. Routhier, D. Dessureault, and P. Boissy, "Assessment of joystick control during the performance of powered wheelchair driving tasks," *Journal of NeuroEngineering and Rehabilitation*, vol. 8, p. 31, 2011.
- [8] M. C. Nechyba and Y. Xu, "Human control strategy: Abstraction, verification, and replication," *IEEE Control Systems*, vol. 17, no. 5, pp. 48–61, 1997.
- [9] S. J. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 3 ed.: Prentice Hall, 2010.