Predictive Haptic Guidance: Intelligent User Assistance for the Control of Dynamic Tasks

Benjamin A.C. Forsyth and Karon E. MacLean, *Member*, *IEEE*

Abstract—Intelligent systems are increasingly able to offer real-time information relevant to a user's manual control of an interactive system, such as dynamic system control space constraints for animation control and driving. However, it is difficult to present this information in a usable manner and other approaches which have employed haptic cues for manual control in "slow" systems often lead to instabilities in highly dynamic tasks. We present a predictive haptic guidance method based on a look-ahead algorithm, along with a user evaluation which compares it with other approaches (no guidance and a standard potential-field method) in a 1-DoF steered path-following scenario. Look-ahead guidance outperformed the other methods in both quantitative performance and subjective preference across a range of path complexity and visibility and a force analysis demonstrated that it applied smaller and fewer forces to users. These results (which appear to derive from the predictive guidance's supporting users in taking earlier and more subtle corrective action) suggest the potential of predictive methods in aiding manual control of dynamic interactive tasks where intelligent support is available.

Index Terms—Human factors, evaluation/methodology, haptic I/O, user-centered design.

1 Introduction

Intelligent systems now appear in highly interactive applications as diverse as automobile driving support, surgical simulation for training, animation design aids, and tools that teach skills based on physical gestures. A current challenge for interface designers is to transfer relevant information the "last 5 inches" from the computational element to the user, in a digestible form. This might include cues derived from an intelligent system's knowledge of the environment and/or from its assessment of the user's current, even momentary, capabilities, intentions, or needs.

Haptic force feedback can be utilized to this end. The many possible approaches to devising intuitive haptic control-sharing cues can be differentiated by the degree of control retained by the user. At one extreme, the system behaves autonomously, but allows the user limited intervention when desired; at the other, the user is completely responsible for interface control, but the intelligent system offers supplementary force suggestions. The latter space is most relevant to applications that require tightly coupled user interaction with highly dynamic content. Likewise, haptic cues offer a unique but risky opportunity, in that guidance cues can be overlaid directly onto the physical control channel or, alternatively, they may be supplied through another channel—e.g., delivered as a tactile stimulus to spatially separate site. The former potentially affords very immediate and easily integrated feedback to the user—but, if not well designed, it can also obscure the user's perception of the system, capture his attention, and disrupt his intended control actions.

- *The authors are with the Department of Computer Science, University of British Columbia, 2366 Main Mall, Vancouver, BC, Canada V6T 1Z4. E-mail: bacforsyth@gmail.com, maclean@cs.ubc.ca.*

Manuscript received 8 Dec. 2004; revised 16 Mar. 2005; accepted 27 May 2005; published online 9 Nov. 2005.
For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org, and reference IEEECS Log Number TVCGSI-0239-1204.

Our own initial efforts, as well as those of others have led us to believe that, for overlaid haptic guidance to be both usable and helpful, force cues must be introduced gradually rather than abruptly, particularly for tightly coupled, low-reaction-time applications, and oscillation-prone structures must be avoided. In short, there is a need for a guidance approach that will support *transparent* communication between the user and the intelligent system: It must offer motion suggestions without demanding attention or cognitive effort, while still allowing the user to maintain absolute control.

The objective of the work reported here is a better understanding of the basis of successful haptic cuing of user motions, with attention to both performance and user preference. We focus on the question of how overlaid haptic feedback can best be utilized to *suggest*, as opposed to *dictate* dynamic actions derived from an intelligent system's information. Specifically, we describe a predictive haptic guidance method devised as a response to our observation of usability problems with other approaches and compare it with a commonly used potential field/spring-like guidance method and with a baseline "no guidance" condition in a driving-type task.

In the remainder of this paper, we comment on related work and how our approach has built upon it, then define our new predictive method and the methods we compared it with. We present and describe the objective and subjective results of a user evaluation of these methods in a representative dynamic user-control task, where subjects use a 1-DoF haptic interface to navigate a vehicle along a path. Finally, we conclude with a summary of insights into overlaid haptic guidance.

2 Related Work

Our review of past work in conjunction with our own initial efforts suggests that, while effective in some situations, the potential field and related spring-damper methods used in the majority of the previous work on haptic guidance can be a source of usability problems, with a detrimental affect on performance, in higher-bandwidth and more complex control applications. Furthermore, usability issues (user preferences as well as performance) have not been considered nearly as carefully as algorithm development has, with the exception of very simple tasks.

Historically, force feedback has long been used to help users perform interactive tasks; for example, Adelstein and Rosen designed a haptic joystick in the early 1990s to support volitional movements in users with physiological tremor through tuned dynamic coupling [1]. Shortly thereafter, Rosenberg defined "virtual fixtures" as forces superimposed on a rendered environment to help guide a user's motion [2]. Since then, haptic feedback has been used to augment interaction in many tasks, including surgery, the learning of physical gestures, driving, and animation.

2.1 Successes

In several cases, researchers have found clear performance and usability benefits to overlaid haptic guidance. Teo et al. used a 6-DoF haptic interface to teach Chinese handwriting [3]. They virtually attached (damped spring) the interface's tip to either a recorded teacher's trajectory (spatial-temporal constraint) or to just the recorded path (spatial-only constraint). They reported that spatial path constraints *without* a temporal constraint were "agreeable to users" and resulted in a performance increase, especially for beginners.

Feygin et al. used a Phantom[1] to train users in an abstract sensorimotor skill: tracing and then recalling the spatio-temporal motion of a point on a complex 3D trajectory [5]. They utilized three presentation methods (haptic, visual, and haptic+visual), with the haptic content produced in a similar manner to Teo et al., and two recall methods: purely kinesthetic versus kinesthetic plus visual guidance (no guidance forces were presented). They conclude that haptic guidance can benefit performance, especially when training temporal aspects of a task, but their results also suggest a potential interference when sensory modes are changed between training and recall and they comment on the danger of abruptly withdrawing haptic guidance when a user has become dependent on it during training.

Haptic path guidance is used successfully in work by Okamura et al., with the goal of providing assistance in microsurgical applications [6], [7]. Their mechanism guides a user along a path by making movement parallel to the path easier than in the perpendicular direction and has good usability properties in microsurgical tasks. However, because it requires that the force display have at least as many DoFs as does the constraint, it is inappropriate for providing guidance in "underactuated" systems, which have more DoFs than are available or humanly manageable for control—e.g., animating an articulated figure or using a 1-DoF wheel to control position in 2-DoF through control of heading.

Steele and Gillespie studied haptic guidance in the shared control of a vehicle and its effect on visual and cognitive load in an apparently low-bandwidth task [8]. They required users to follow a straight, obstacle-ridden path while providing haptic feedback guiding the vehicle toward the center of the path. Users were instructed to avoid obstacles and stay on the center of the path; path and obstacles were displayed in response to user demand. Haptic guidance provided a significant decrease in both visual demand and lateral deviation as compared to no guidance. In a second experiment, subjects were additionally subjected to cognitive load; uniform performance in the load task suggested that the haptic guidance did not affect subjects' overall cognitive effort. Thus, for solely objective measures of a very simple and probably low-bandwidth path-following task, these experiments suggest that haptic guidance helps users follow a path without impinging on cognitive load.

All of the preceding examples focus on spatial movements that are inherently slow. Additionally, the training examples are repetitive and have a specific desired spatial outcome as opposed to a set of possible trajectories that a user must steer among. In these situations, a simple damped-spring model seems to be beneficial.

Finally, Donald and Henle use haptic guidance to interact with motion capture data of an articulated human figure in real-time [9]. They created a bidirectional transfer function between a 3-DoF haptic workspace and the 57-DoF configuration of the motion capture data. The result is a 3-DoF "force river" along which the interface's tip is drawn with a virtual damped spring. By providing forces to the end effector, a user can alter the configuration of the articulated figure. The authors did not report any evaluation of performance or usability. This work moves into the realm of haptic interaction with complex (here, high-dimensional) systems using an interesting *bidirectional* guidance metaphor. This is one of the areas where we see our guidance method ultimately offering a benefit, but it is not the focus of the current work.

2.2 Lessons from Nonhaptic Guidance

Some important contributions to the body of real-time guidance have occurred in contexts that are either non-interactive (e.g., autonomous) or which employ other sensor modalities for user feedback.

Rossetter et al. have extensively studied control of vehicle steering aimed at autonomous applications [10], [11], [12]. Their focus is on lane-keeping guidance using a potential field method with a look-ahead predictor and the safety and stability concerns of such a guidance method; it includes an intricate vehicle model and a formal stability analysis. This work is relevant here because it addresses higher-bandwidth control issues, but it does not consider user input. Rossetter does acknowledge that user interaction with the guidance system is an important issue that requires attention [12].

Reynolds presents a predictive guidance method for autonomous vehicles that was the inspiration for our own look-ahead guidance [13]. He used a simple linear predictor based on the vehicle's velocity to predict if the vehicle will be on or off of the path in the future; if off, the vehicle's course is adjusted in an attempt to stay on the path. As a test bed for his steering behaviors, Reynolds also developed a software toolkit, OpenSteer [14], which we used as the basis

1. The Phantom is a commercially available haptic interface with three actuated and six sensed degrees of freedom (DoF) [4].

of our simulation software. His results suggest that, for an autonomous system, this look-ahead approach improves system performance for the same reasons that we hypothesize it will help people, when properly cued: The system has more time to prepare for a control action and can adjust its course smoothly and without unstable oscillations triggered by overly abrupt commands.

Finally, Feng et al. developed a nonhaptic but visually interactive path guidance system, also using a look-ahead algorithm [15]. They use a more complex algorithm than does Reynolds to predict the future vehicle position and display it visually; however, they do not evaluate this aid's impact on user driving performance.

To summarize, none of these examples provide experimental evidence of a look-ahead algorithm's benefit, in any modality, but they indicate its promise.

2.3 Haptic Guidance for Fast, Complex Systems

While we are not aware of other efforts to overlay intelligent cues on systems involving complex, underactuated, and/or high bandwidth motion, some early forays by our group illustrate some inherent challenges.

In an initial unpublished effort, we used force feedback to display constraints on the steering angle of a vehicle as computed by an intelligent system developed by Kalisiak and van de Panne [16]. The most straightforward solution, which generated forces that pushed the Phantom's handle away from the constraint, resulted in poor usability: Oscillations were induced by the guidance forces when the abrupt impulse applied to avoid one constraint would launch the end effector toward the other constraint. This situation was not improved by introducing the forces earlier as a spring element attached to the wall and was mitigated only slightly by light damping of the spring. Further damping adversely affected the feel of the environment. We also found that users tended to reflexively fight abrupt guidance forces, reducing their effectiveness.

Conversely, users were able to substantially benefit from a subsequent scheme which incorporated simple damped-spring haptic feedback into a longitudinal high-speed driving control (the accelerator pedal) to convey following distance from a lead car [17]. We believe that this was in part because the cue algorithm is based on more gradual, sometimes even unnoticeable force introduction. Further, because the constraint was unidirectional, it did not tend to introduce oscillations.

These experiments tell us that, for sufficiently simple problems (e.g., longitudinal vehicle control), a straightforward damped spring can be helpful—and is thus almost certainly the preferred solution. But, when the context is highly dynamic, has multiple constraints, and/or is underactuated, more sophisticated tools are needed.

3 Guidance System Design

In this section, we describe our general design approach, and the vehicle model used in guidance method development and evaluation. We devised two guidance algorithms to be compared in a user study with a control:

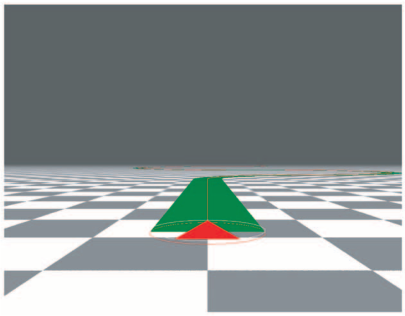

Fig. 1. View provided by simulation software for high visibility condition (viewed by subjects in color). The red triangle represents the vehicle; the path is green.

1. A potential field method (Section 3.3.2), to represent state-of-the-art as cited in Section 2.
2. A new predictive method (Section 3.3.3) targeting the problem of high-bandwidth usability by introducing guidance forces earlier and thus (we theorized) more smoothly.

3.1 Approach to System Design and Choice of Model

We needed a guidance system that would support transparent, bidirectional communication between the user and an intelligent system. Our premise was that the crafting of user interactions with the guidance forces must be an integral part of the design process.

Because our target was the basic usability challenges inherent in providing haptic guidance for highly interactive tasks, we chose a simple model and a 1-DoF controller combined with a common, low-DoF, high-dynamic-range control task of adjustable challenge: In a simple driving scenario, a user attempts to keep a vehicle inside of a path. This testbed avoided confounds between the guidance algorithm and details of an elaborate system model or the difficulty of learning a complex control. Further, steering (1-DoF, underactuated) tasks comprise an important subset of the possible space where we foresee predictive method applicability. While further study will be required to determine how insights gathered here will apply to higher-DoF systems, this seemed the best place to start.

3.2 Simulation Environment and Vehicle Model

Our experiment context required a simulator engine, a graphic display of the vehicle and path, an implementation of the guidance model, and low-level haptic interface control. Our simulation, based on OpenSteer [14], updated at 60 Hz and displayed an oblique, overhead view of the vehicle and a one-unit-wide path (Fig. 1).

OpenSteer supports only autonomous steering behaviors, so we developed our own user-controllable vehicle model (Fig. 2). The angle of the steering wheel, Δ, steers a tricycle's front wheel, θ, by $\theta = 0.7\Delta$. With the wheelbase, ℓ, this defines the vehicle turning radius, r:

$$r = \begin{cases} \frac{\ell}{\tan\theta} & : \theta \neq 0 \\ \infty & : \theta = 0. \end{cases} \tag{1}$$

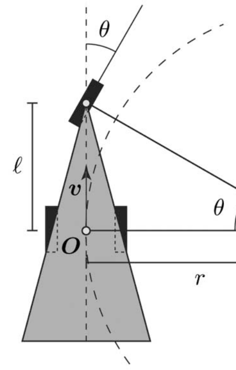

Fig. 2. Schematic of vehicle dynamics: the shaded area represents the vehicle. Tires are shown here for clarity, but are not rendered during the simulation. θ is the current vehicle steering angle.

To better simulate an automobile, we clamp the steering angle to $[-\theta_{\max}, \theta_{\max}]$, where, for our experiments, $\theta_{\max} = 15°$. We found that this also helps prevent users from getting into overly tight turns, with ensuing disorientation.

3.3 Haptic Guidance Methods

We implemented and compared three haptic guidance conditions, described below: 1) a set of underlying orienting forces common to all conditions, called "No Guidance" (NG), 2) a Potential Field Guidance (PFG) method, resembling those used by others, and 3) a predictive Look-Ahead Guidance (LAG) method of our own design, inspired by Reynold's autonomous algorithm [13].

Both guidance methods described below compute a desired heading change, ϕ, given the vehicle's current orientation relative to the path (ϕ_{pfg} in Fig. 3 and ϕ_{lag} in Fig. 4). A clamped, linear transfer function is used to map ϕ to a desired steering angle, θ_{desired} ((2), with $\phi_{\max} = 60°$):

$$\theta_{\text{desired}}(\phi) = \begin{cases} -\theta_{\max} & : \quad \phi \leq -\phi_{\max} \\ \frac{\phi}{\phi_{\max}} \theta_{\max} & : \quad -\phi_{\max} < \phi < \phi_{\max} \\ \theta_{\max} & : \quad \phi \geq \phi_{\max}, \end{cases} \tag{2}$$

where θ_{desired} is the setpoint for the haptic interface's PD controller (Section 3.4).

3.3.1 Baseline Forces

We implemented two baseline forces that are always present, a centering force and viscous damping. The centering force attempts to maintain a zero steering angle via a virtual spring, recreating the force that is felt on the steering wheel of a real car at speed when its tires tend to point straight ahead. We found that, when people interacted with a prototype lacking this, regardless of guidance method, they often became disoriented in tight turns. The viscous damping force was added to stabilize the centering force in nongrasped conditions. We used values of $K = 10.0$ and $B = 3.0$ and centering force output was clipped to ± 0.25 out of a total command range of ± 1.0.

3.3.2 Potential Field Guidance (PFG)

Based on its use in previous work, we considered potential field and spring-damper guidance methods to be the de facto standard and, therefore, a necessary component of a comparative evaluation. Whereas others have used higher-DoF Cartesian displays, our version of a PFG method was influenced by our 1-DoF controller; this issue is discussed in Section 6.2.

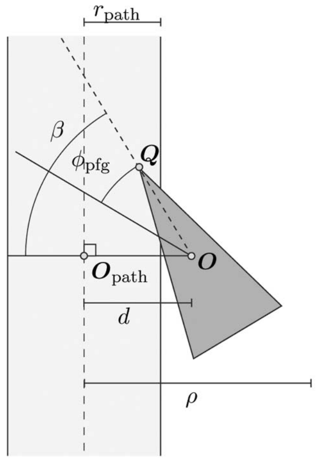

Fig. 3. Components of the potential field guidance method. ϕ is the desired vehicle heading.

As shown in Fig. 3, we apply a force proportional to the distance d between the vehicle's center (O) and the point on the path closest to the vehicle (O_{path}). This is accomplished by generating a desired heading change ϕ_{pfg} (and, thence, a desired knob angle from (2)) via (3) and (4):

$$\phi_{\text{raw}}(d) = \begin{cases} \frac{d}{\rho} \phi_{\max} & : \quad 0 \leq d < \rho \\ \phi_{\max} & : \quad d \geq \rho, \end{cases} \tag{3}$$

$$\phi_{\text{pfg}}(d, \beta) = \begin{cases} \frac{\beta}{|\beta|} \phi_{\text{raw}}(d) & : \quad \phi_{\text{raw}}(d) < |\beta| \\ \beta & : \quad \phi_{\text{raw}}(d) \geq |\beta|, \end{cases} \tag{4}$$

where ρ is the distance from the path at which the maximum guidance force is applied and ϕ_{raw} is the desired heading offset based only on the distance from the path. It is possible for the raw desired heading offset to steer *behind* the line straight back to the path and, thus, we limit ϕ_{raw} by the angle β

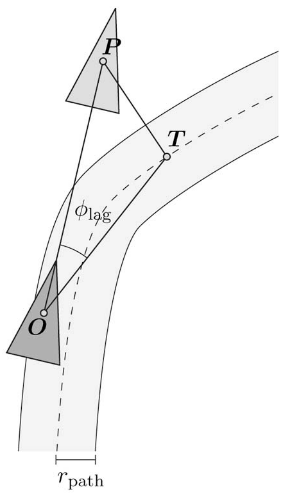

Fig. 4. Components of the look-ahead guidance method. ϕ is the desired vehicle heading.

Fig. 5. The haptic interface knob, motor, encoder, and stand.

(the angle between the current heading and the shortest line to the path) in (4). The sign of β indicates which side of the path the vehicle is on and is used to set steering direction.

ρ is the only tunable parameter for the PFG algorithm and was explored over a large range in pilot studies before being set to the level which gave best overall performance ($\rho = 1$ unit, as is the path width). In addition, the values of the PD controller (Section 3.4) impacted PFG performance (high proportional gains caused oscillation, while low gains left the constraint too soft). However, the same PD gains gave the best performance for both methods.

3.3.3 Look-Ahead Guidance (LAG)

We developed our predictive algorithm to avoid the strong and sudden guidance forces produced by reactive guidance methods such as PFG. With LAG, we first estimate the vehicle's position, P, t seconds into the future by $P = \mathbf{v}t$, where $\mathbf{v}$ is the vehicle's current velocity.

When this predicted point is outside of the path, we display a guidance force based on the desired vehicle heading change. ϕ_{lag} is the angle between the current vehicle heading and the line from the vehicle to the "target" T, the point on the path closest to P (Fig. 4): $\phi_{\text{lag}} = \angle POT$.

LAG's corrective forces tend to increase gently as a corner or obstacle is approached because, by responding to the predictive guidance with earlier corrective actions, the user is typically able to maintain the distance between P and T at a small value. He would generally first experience a gradually escalating guidance force, rather than "bump into" a constraint. We anticipate that this property will be of the greatest value for curves approached at relatively high speed—a situation not well addressed by other guidance methods.

3.4 Haptic Interface and Low-Level Control

Our haptic interface consisted of a 20 W Maxon motor and 4,000 cpr encoder, with a 9 cm beveled acrylic knob mounted directly on its shaft (Fig. 5). It was connected via a custom amplifier and PCI interface board to a Windows 2000 PC with a 2 GHz Pentium 4 Xeon processor (faster than required) and 512 MB RAM.

Low-level control was accomplished by a PD controller in a high priority thread running at 1,000 Hz, minimizing the difference between the actual and desired knob positions. To smooth the difference signal's derivative (used by the PD's derivative term), we used an adaptive windowing technique [18]. We found that gains of $K_P = 3.0$, $K_D = 0.2$ resulted in good performance for both methods.

The physical characteristics of our control knob afforded a somewhat different interaction style than does a full-sized steering wheel. Interaction with the knob is performed with one arm and primarily involves wrist and finger motions, with some small elbow rotations. A steering wheel is typically controlled with shoulder and elbow rotations from both arms. An advantage of a full-sized steering wheel is that small movements of the wheel engage similar motor units as do those of large wheel movements. In contrast, the typical five-fingered grasp of our device facilitated small knob movements and large movements involved significantly different motor units. Since we were not aiming for face validity with respect to driving, it was acceptable that our control knob elicited different motor patterns than a steering wheel, but, as will be discussed later, it did have ambiguous benefits in terms of task transfer.

4 Evaluation

4.1 Objectives and Approach

We sought an indication of the effectiveness and acceptability of a predictive haptic guidance method (as represented by LAG) relative to other methods currently in use and to no guidance at all, in the performance of high-bandwidth or complex tasks. To this end, we hypothesized that LAG would have measurable performance benefits over both NG and PFG, particularly in the more difficult versions of the tasks considered here, and that users would subjectively prefer LAG.

We tested these hypotheses using an MSE performance metric and a repeated-measures three-factor ANOVA, alongside a subjective evaluation. We designed our experiment for 18 subjects and, to minimize subject fatigue, restricted experiment duration to one hour.

4.2 Design

4.2.1 Experiment Task

Subjects steered along a path from beginning to end while attempting to keep the vehicle within the extent of the path (primary directive) and close to the center of the path (secondary directive). To retain experiment control over task difficulty, velocity was fixed to a level that resulted in good performance for "easy" factor settings during pilots, while still begin challenging for "hard" factor settings ($\mathbf{v} = 5.0$ units/sec and lookahead interval $t = 1.0$ sec). These criteria were chosen to exercise our hypothesis that LAG will outperform potential force methods principally for higher bandwidth or more complex tasks, i.e., it was less critical to consider slower velocities.

The basic task was repeated under different combinations of the experiment factors. The task was purposely designed to be similar to driving in an attempt to minimize the learning curve for interaction with our system; driving is not its only potential target application.

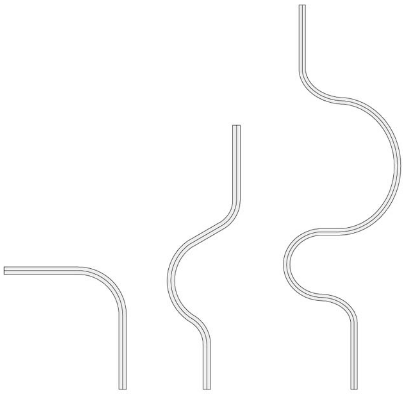

Fig. 6. Examples of the three path complexities used in the evaluation (from left to right: bump, curve, and zigzag).

4.2.2 Independent Variables and Levels

In addition to guidance method, candidates for our evaluation included vehicle velocity, look-ahead time, path width, visibility level, viewpoint, path complexity, and look-ahead predictor algorithm. Experiment duration limited us to two. We chose *path complexity* and *level of visibility* on the premise that they would contribute the most broadly applicable insights into the benefits and appropriate use of the guidance methods: Some methods are likely to be more or less helpful for challenging steering situations and, because LAG is a predictive method, we speculated that its strength might be in low visibility conditions.

Note that a range of subjective velocities (also likely to influence the moment-to-moment performance of either guidance method) are represented by the variation in path complexity and visibility: "fast" for driving on a straight road is different than for a curvy road in high fog conditions. Modulating too many sources of task difficulty would have made it hard to interpret results and we felt it less productive to study lower velocities based on past evidence that PFG performance there is adequate (e.g., [7]).

The levels per condition were likewise constrained by experiment duration to 2-3 per factor, including three difficulty levels for guidance method (NG, LAG, and PFG), and three path difficulty levels chosen based on the number and radii of corners (curve, bump, and zigzag, shown in Fig. 6). Finally, we used two visibility levels: High allowed viewing of the full extent of the path (Fig. 1), whereas low introduced a fog which restricted visibility to about three units (squares) ahead of the vehicle.

4.2.3 Repetition, Blocking, Randomization, and Learning

Five repetitions (of 18 factor combinations) were enough to address the individual variability observed in pilots. To ensure that subjects could become familiar with each guidance method, we blocked trials by guidance method while continuing to vary path complexity and visibility (which seemed both less subject to such a concern and more likely to occur variably in a realistic context).

The experiment task was chosen to minimize learning required, but we still anticipated that some would occur. We addressed this by two means, illustrated in Fig. 7. The subject began the session with a Familiarization Phase, repeating an 18-factor block (all factor combinations) until attaining sufficient skill to continue to a Learning Phase, followed by the final Test Phase. Both consisted of three blocks of 30 trials each such that each phase covered all three guidance methods. Test Phase blocks were used in our analysis; both phases were used to verify that learning had indeed stabilized. Block presentation order was counterbalanced across subjects.

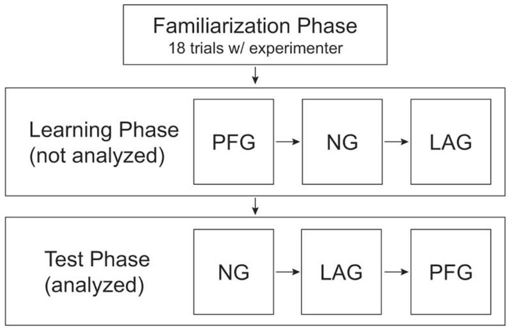

Fig. 7. An example protocol for an experiment session: The actual ordering of blocks was counterbalanced by subject. Each block (NG, LAG, or PFG) contained $3 \times 2 \times 5 = 30$ trials.

4.3 Metrics

4.3.1 Objective Performance Metric

We measured inverse path following performance for each trial by integrating the Mean Square Error (MSE) of the deviation of the vehicle from the center of the path, as sampled at 60 Hz:

$$MSE = \frac{\sum_{n=0}^{N-1} |O_n - O_{\mathrm{path}_n}|^2}{N}, \tag{5}$$

where O_n is the location of the vehicle at time n and O_{path_n} is the point on the centerline of the path closest to point O_n at time n.

For simplicity and clarity, we chose to use a single reasonable "all-around" quantitative performance metric rather than either a multivariate analysis or multiple single-metric analyses of several metrics. For both of the latter, the metrics' relative strengths and biases can be unclear and interpretation unwieldy. We considered many other rejected candidates, including trial duration, mean error, trajectory length, smoothness, shape, and frequency of leaving path. Each metric has its own drawback, often a result of subject strategy: For example, both trial duration and trajectory length reward taking a straight line from start to finish. Leaving the path early in a trial and staying outside for the duration of the trial results in a low path-leaving frequency, but not in good path following. The shape and smoothness of a trajectory do not capture a consistent offset from the path center. Neither ME nor MSE is susceptible to shortcuts or constant offsets. Of these, we chose MSE because it heavily penalizes large path departures, which seemed to better capture the notion of good path following.

Our MSE metric is not perfect: Computed from the path center, it does not account for the path width and does not

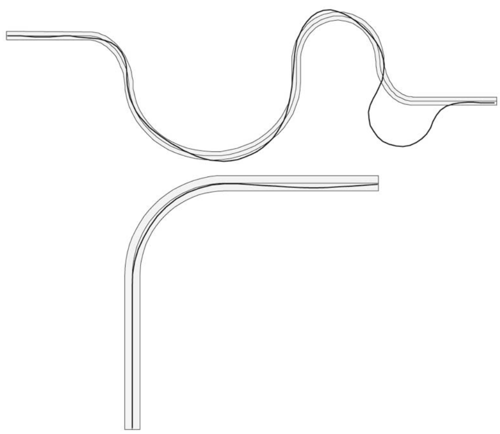

Fig. 8. Trajectories with the worst (top) and best (bottom) MSE scores.

heavily penalize minor corner-cutting. The secondary task of following the *center* of the path as closely as possible was designed to indirectly address this. It also can have a strategic influence: Subjects are likely to spatially shorten their focus of attention when required to increase the precision of their control, in turn limiting their perception of upcoming path features. However, the centerline requirement is just one of several precision-related components of control difficulty [19]; others are vehicle speed and path width. This shortcoming therefore seemed minor.

4.3.2 Subjective Measures

We measured user's reactions to the guidance methods via two mechanisms. After each block, we asked subjects about their experiences with the trials in that block, specifically perceived degree of vehicle control, helpfulness of the provided guidance, and the pleasurability of the haptic feedback, using a 5-point Likert scale. At the end of the session in a debriefing interview, we asked which of the last three blocks the subject liked the best and if they felt that their performance was improved by the guidance method in any of those blocks.

4.4 Procedure

The one-hour experiment sessions were conducted in an experimentation room with controlled acoustics and lighting. During trials, the full system state, including the position of the vehicle and distance from the path, was sampled and recorded at 60 Hz. Subjects were instructed to grasp the knob in a manner most natural to them.

5 Results

5.1 Subjects

Eighteeen subjects completed the experiment. All were university students aged 19-33, 12 male and six female. Seventeen held a driving license; eight stated they drove daily or weekly and nine infrequently. Fifteen reported previous experience with a haptic interface and, of these, five said their haptic experience was advanced. Upon initial analysis, it was clear that one subject did not follow the instructions; rejection of his data left 17 in our final analysis.

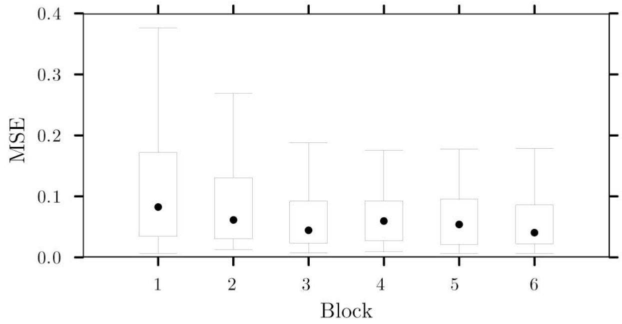

Fig. 9. Boxplot of performance across all subjects in order of block performance, showing stabilization of learning by Block 3 (end of the Learning Phase).

Fig. 8 illustrates the range of path following performance exhibited throughout the experiment.

Learning effects visualized in Fig. 9 suggest that performance has stabilized in the final three blocks used for analysis.

5.2 Results of Quantitative Performance Analysis

We used R [20] to perform a $3 \times 3 \times 2$ within-subject, repeated measures ANOVA with one restriction on randomization (blocks) and 18 data points per subject. Each data point is the mean of five repetitions. All main effects were significant (Guidance Method at $F_{2,32} = 4.860, p = 0.014$, Path at $F_{2,32} = 8.984, p = 0.001$, and Visibility at $F_{1,16} = 4.887$, $p = 0.042$). None of the tested interactions were significant.

We performed post-hoc, pairwise comparisons between the factor levels (using the Holm adjustment for multiple comparisons), which show a significant difference between LAG and both PFG ($p = 0.012$) and NG ($p = 0.004$), but not between PFG and NG. Likewise, there is a significant difference between Low and both Medium ($p = 0.018$) and High ($p = 0.010$) path complexity, but not between Medium and High. Fig. 10 shows the mean MSE values for each main effect. LAG guidance, Low path complexity, and High visibility resulted in the lowest average errors for the three conditions.

5.3 Comparison of Forces Generated by Each Method

To gain insight into both the quantitative performance and subjective response results, we examined the forces delivered to the user for each guidance method. The most useful and objective view is presented in Fig. 11, which displays these forces as histograms. Key observations are that 1) NG forces tend to peg at the 0.25 cut-off, implying heavy use of the centering force when it is the only guidance (this peak appears, diminished, for LAG as well); 2) LAG forces, summed, are no larger than NG (baseline forces only), whereas average PFG per-trial force is approximately double; and 3) LAG forces tend to be small, whereas PFG forces are distributed across the available range up to saturation. That is, under PFG, users experience more large-magnitude corrective guidance forces than under LAG.

Frequency power spectrums for both supplied force and knob position indicated similar frequency distributions for

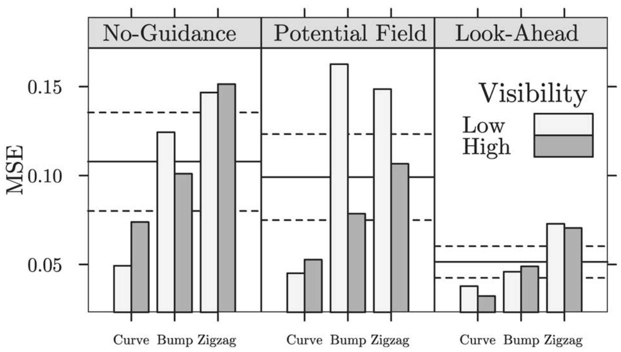

Fig. 10. Mean MSE for all combinations of independent variables (lower scores are better). Each bar represents 85 trials (five repetitions $\times$ 17 subjects). The horizontal solid and dashed lines represent overall mean and standard error of the mean, respectively, for that guidance method.

NG, PFG, and LAG (rolloff around 0.5 Hz); however, for all levels of path difficulty and particularly for force, PFG spectrums contained substantially more power than the others. This is consistent with the generally higher degree of PFG force application, and suggests that PFG users experience more variation in force levels as well as the generally more widespread and higher-magnitude application of force shown in Fig. 11.

5.4 Results of Subjective Response Analysis

We measured subjects' feelings about their interaction with the system after each block and when debriefing. Postblock questions were:

Q1. Did you feel any force feedback? (Yes/No).
Q2. What level of control did you feel you had over the vehicle? 1-5 (No Control-Complete Control). If the subject reported feeling force feedback:
Q3. How helpful did you find the force feedback? 1-5 (Very Unhelpful-Very Helpful).
Q4. How much did you like the force feedback? 1-5 (Strongly Dislike-Strongly Like).

Table 1 and Figs. 12, 13, and 14 show the results of questions Q1-Q4. Debriefing question Q5, "Which block did you like the most?" generated responses of NG (1), PFG (3), and LAG (13).

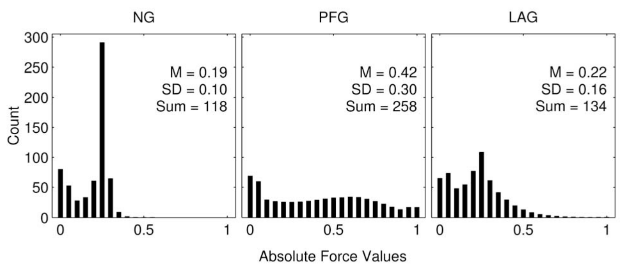

Fig. 11. Average absolute force values delivered to the user for each guidance method. Each subplot represents histogrammed forces (recorded at 60 Hz) for 510 trials (5 x 17 x 6 pathtype/visibility combinations); each element is divided by 510. M and SD are the mean and standard deviation of forces for that method, and Sum is the total of forces delivered, divided by ($60s^{-1} \times$ 510 trials). Note that trial duration also affects Sum.

TABLE 1
Response to Q1: *Force Feedback Felt*

Guidance Method	Yes	No
No-Guidance	3	14
Potential Field	15	2
Look-Ahead	17	0

In summary, the postblock subjective results suggest that LAG made subjects feel more in control of the vehicle and was perceived as more helpful and was better liked than the other guidance methods. When given the opportunity to express an overall preference, 13/17 subjects chose the block with LAG.

6 Discussion

We begin our discussion with an observation on experiment validity. The significance of path type and visibility confirmed that these supporting experiment manipulations were, in general, effective in exercising a range of responses to the guidance methods. The lack of a significant difference between the Medium and High path complexity levels suggests we could have been more extreme in terms of task challenge.

The results reported here are with respect to our experiment task and interface, the other "difficulty" factors employed, and to the MSE performance metric; however, these were selected to fairly represent the general task class (high bandwidth and/or complexity) where we believe predictive methods are most likely to have relevance. With value in this domain demonstrated, we can proceed to wider validation.

6.1 Quantitative Performance of Look-Ahead Guidance

The guidance method main effect and related post-hoc tests indicate that Look-Ahead Guidance exhibits significant performance benefits compared to both the baseline NG and the "standard" PFG, with a mean path deviation score of about half that of the other methods (Fig. 10). PFG did not significantly improve performance relative to NG.

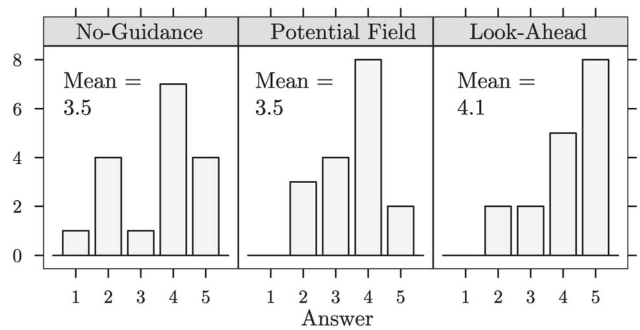

Fig. 12. Responses to Q2: *felt in control*, for each guidance method. 1 = No Control, 5 = Complete Control.

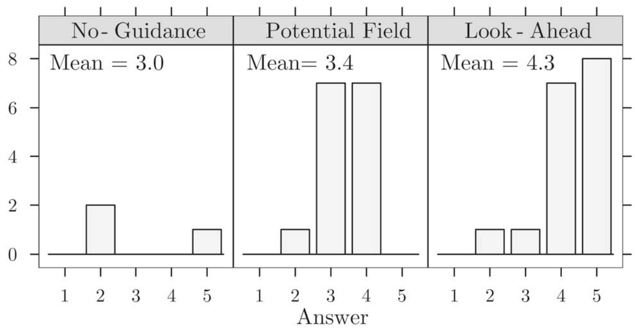

Fig. 13. Responses to Q3: *force feedback helpful*, for each guidance method; asked only of those who answered "Yes" to Q1 (3/15/17 respondents, respectively). 1 = Very Unhelpful, 5 = Very Helpful.

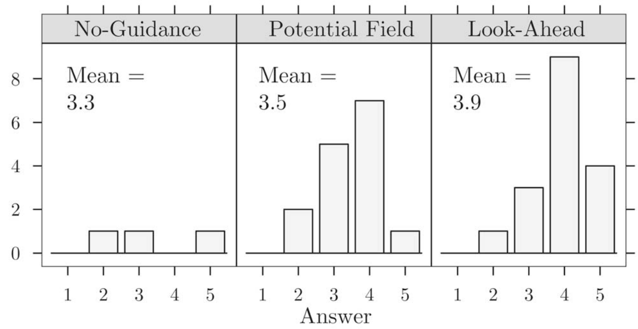

Fig. 14. Response to Q4: *liked force feedback*, for each guidance method; asked only of those who answered "Yes" to Q1. 1 = Strongly Dislike, 5 = Strongly Like.

6.1.1 Source of Performance Benefit

The force results of Fig. 11 and reported frequency analysis support our reasoning that a predictive method should generate smoother, lower magnitude, and generally *less* force. LAG warns users of impending problems and allows them to alter their course before PFG would indicate the same problem and we posit that this avoids the abrupt, oscillation-inducing forces that are problematic with PFG-style methods. Furthermore, while it may be possible to tune the PFG method to avoid generating abrupt forces (e.g., by reducing the position controller's gain), these forces still come too late to be of use.

6.1.2 Scope of Benefit across Task Difficulty

LAG did not prove *especially* useful for more complex paths compared to less complex paths: The improvement over other methods is relatively consistent across path types. We also hypothesized that LAG might offer more than the other methods in low-visibility conditions because the system can "see" further than the user. Fig. 10 shows that LAG does indeed perform much more consistently across visibility conditions than did the other methods: Low-visibility performance is brought up to the level of LAG high-visibility performance.

Thus, these results suggest that the LAG performance benefit applies across a range of path complexities and visibility levels.

6.2 Implementation of Guidance Methods

There are many possible approaches to designing both potential field and predictive guidance algorithms and optimality is difficult to guarantee. The versions used here were the best of many we tried during an extensive prototyping process. We were specifically interested in underactuated tasks because of their potential efficiency, but they are less straightforward to implement. Ours is intended as a 1-DoF analog of the 2-DoF path-tracing potential field method (Section 3.3.2), but it is atypical in using a desired steering angle, similarly to LAG (Section 3.3), and is thus more accurately a spring-damper method. We feel that this implementation nevertheless captures the spirit of potential field methods used elsewhere.

6.3 Performance Metric

For our MSE performance metric to function, we had to require subjects both to stay inside the path and, secondarily, to follow its centerline. It is possible that demanding fine control (by giving the centerline directive) may have contributed to LAG's observed consistent performance benefit: If it caused subjects to focus their attention closer in on the path than they otherwise would have, these subjects might also have been delayed in incorporating upcoming path features into their own path following strategy. These individuals would likely find more benefit in the predictive component of LAG. If accurate, this observation would support a premise that LAG is good for tasks requiring fine control, which is simply a different form of workload than the one we set out to control by varying path complexity and visibility.

In a different manner, our MSE metric tends to heavily weight large departures from the path: For example, a "bump" counts more than an oscillation. For some purposes, this might be undesirable and a different metric would be more appropriate. In our case, we find that the MSE metric gives one of many possible views on performance effects of the methods. A multivariate approach might supply a more general view, but will need to be weighted appropriately for the given task.

6.4 Subjective Performance of the Guidance Methods

Our subjective evaluation methods sufficiently addressed the issues in which we were interested, providing consistent results for perceived level of control and the preferred guidance method. There seemed to be no discrepancy between perceived helpfulness and aesthetic preference. Look-Ahead Guidance gave subjects a better sense of control over the vehicle compared to the other guidance methods, an important characteristic for usability and acceptance. When subjects reported feeling force feedback, LAG was reported to be more helpful and better liked than the other methods; 13/17 subjects reported preferring the block corresponding to LAG during the debriefing interview.

Reasons given for this preference were consistent. In interviews, we heard that some subjects preferred LAG because they could tell that it was improving their ability to follow the path; others reported that LAG felt the most natural. Some found PFG overwhelming at times.

6.5 Learnability and Task Transfer from Real World

We purposely made the experiment task similar to driving in an attempt to minimize learning. However, key differences between our setup and real driving (most notably, knob size and deliberate control over vehicle velocity) may have impacted task transfer and, hence, increased learning times and/or individual variability.

The size and shape of our interface's knob led some users to "scroll" by releasing it and repositioning the fingers. When released, the centering force sometimes made it move in a counterproductive direction; in a real steering wheel, a user would counter this with a two-handed grasp. It was actually possible to steer the knob without releasing, but a few subjects did not identify this without input from the experimenter.

Learning had settled by the end of the learning period, so we do not view data instability as a serious concern with respect to our result validity; likewise, any extraneous variability thus induced is unlikely to be systematic and did not prevent significant results. Thus, using a more ecologically valid driving setup in terms of knob size and velocity control would probably not have changed the evaluation's outcome.

6.6 Observations on When Haptic Guidance Is Useful

We observed that, when users were able to follow the path closely, the guidance force feedback was helpful; but, once they made significant deviations from the path and became lost, the guidance feedback was less useful and may actually have made the problem worse. We hypothesize that, once off the path, the user's objectives are no longer predictable by the system and, thus, the system does not display forces that correspond to the direction he wants to go.

This illustrates a need for an intelligent system to be able to assess the user's goals, e.g., his recovery strategy, and also to quantify the certainty it has about user goals. Likewise, the system must have a given level of certainty about the task constraints or other cues in order to evaluate whether its computed guidance forces will correspond to the user goals.

7 Conclusions and Future Work

This paper presents a haptic Look-Ahead Guidance method designed to test the hypothesis that predictive algorithms may solve problems with other guidance mechanisms currently in use, in haptically communicating stable, helpful motion cues.

For a simple but challenging steered path-following scenario, we found that LAG performed significantly better than either No-Guidance and the standard Potential Field Guidance, in terms of an MSE trajectory-following metric. Force analysis shows that LAG-produced forces were smaller and fewer than PFG, bearing out our premise that a predictive guidance method could be both more effective and more subtle by providing force cues early and helping users to react before errors build up. In terms of usability and probably for the same reasons, subjects overall preferred LAG: They felt more in control and found the force feedback more helpful and likeable.

We believe that the advantages of predictive guidance will particularly apply to highly dynamic tasks such as those made possible through intelligent-system support of animation control or driving and represented by our evaluation context. The results reported here suggest that LAG benefits may apply across a broader range of task dynamics than we had expected.

In future work, we plan to further explore this notion with a more extensive study of tasks that require a range of dynamic responsiveness, as well as an examination of guidance method success as a function of parameters such as look-ahead time and distance and, in tasks like the one used here, velocity and path curvature. We anticipate that, while LAG will be of the greatest benefit at a higher speed/trajectory frequency ratio, with an adjustable look-ahead time it should deliver consistent performance across a range of velocities. Haptic guidance cues need to be compared directly with visual guidance in order to both generalize on its strength and address a potential confound between force feedback and guidance, although we expect such a study to have a large task-type dependency. Finally, we would like to consider different kinds of predictors applied to tasks of higher dimensionality and to learn if LAG methods will be beneficial in low or nonvision tasks through their potential to "see" for the user.

References

[1] B.D. Adelstein and M.J. Rosen, "Design and Implementation of a Force Reflecting Manipulandum for Manual Control Research," *Advances in Robotics*, H. Kazerooni, ed., pp. 1-12, 1992.

[2] L. Rosenberg, "Virtual Fixtures: Perceptual Tools for Telerobotic Manipulation," *Proc. IEEE Virtual Reality Int'l Symp.*, pp. 76-82, 1993.

[3] C. Teo, E. Burdet, and H. Lim, "A Robotic Teacher of Chinese Handwriting," *Proc. Symp. Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pp. 335-341, Mar. 2002.

[4] T.H. Massie and J.K. Salisbury, "The Phantom Haptic Interface: A Device for Probing Virtual Objects," *Proc. ASME Winter Ann. Meeting, Symp. Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pp. 295-302, Nov. 1994.

[5] D. Feygin, M. Keehner, and F. Tendick, "Haptic Guidance: Experimental Evaluation of a Haptic Training Method for a Perceptual Motor Skill," *Proc. Symp. Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pp. 40-47, Mar. 2002.

[6] A. Bettini, S. Lang, A. Okamura, and G. Hager, "Vision Assisted Control for Manipulation Using Virtual Fixtures: Experiments at Macro and Micro Scales," *Proc. IEEE Int'l Conf. Robotics and Automation*, pp. 3354-3361, May 2002.

[7] P. Marayong and A. Okamura, "Speed-Accuracy Characteristics of Human-Machine Cooperative Manipulation Using Virtual Fixtures with Variable Admittance," *Human Factors*, vol. 46, no. 3, pp. 953-966, 2004.

[8] M. Steele and B. Gillespie, "Shared Control between Human and Machine," *Proc. Human Factors and Ergonomics Soc. 45th Ann. Meeting*, pp. 358-366, Oct. 2001.

[9] B.R. Donald and F. Henle, "Using Haptic Vector Fields for Animation Motion Control," *Proc. IEEE Int'l Conf. Robotics and Automation*, pp. 3435-3442, Apr. 2000.

[10] E.J. Rossetter and J.C. Gerdes, "A Study of Lateral Vehicle Control under a 'Virtual' Force Framework," *Proc. Int'l Symp. Advanced Vehicle Control*, Jan. 2002.

[11] E.J. Rossetter, J.P. Switkes, and J.C. Gerdes, "A Gentle Nudge towards Safety: Experimental Validation of the Potential Field Driver Assistance System," *Proc. Am. Control Conf.*, pp. 3744-3749, June 2003.

[12] E.J. Rossetter, "A Potential Field Framework for Active Vehicle Lanekeeping Assistance," PhD dissertation, Stanford Univ., Palo Alto, Calif., Aug. 2003, http://www-cdr.stanford.edu/dynamic/PF/papers/eric_diss.pdf.

[13] C.W. Reynolds, "Steering Behaviors for Autonomous Characters," *Proc. Game Developers Conf.,* pp. 763-782, Mar. 1999. http://www.red3d.com/cwr/papers/1999/gdc99steer.html.

[14] C.W. Reynolds, "OpenSteer: Steering Behaviors for Autonomous Characters," 2003, http://opensteer.sourceforge.net/.

[15] K.-T. Feng, H.-S. Tan, M. Tomizuka, and W.-B. Zhang, "Look-Ahead Human-Machine Interface for Assistance of Manual Vehicle Steering," *Proc. Am. Control Conf.,* pp. 1228-1232, June 1999.

[16] M. Kalisiak and M. van de Panne, "Approximate Safety Enforcement Using Computed Viability Envelopes," *Proc. IEEE Int'l Conf. Robotics and Automation,* 2004.

[17] M.J. Enriquez and K.E. MacLean, "Impact of Haptic Warning Signal Reliability in a Time-and-Safety-Critical Task," *Proc. Symp. Haptic Interfaces for Virtual Environments and Teleoperator Systems,* pp. 407-415, 2004.

[18] F. Janabi-Sharifi, V. Hayward, and C.-S.J. Chen, "Discrete-Time Adaptive Windowing for Velocity Estimation," *IEEE Trans. Control System Technology,* vol. 8, no. 6, pp. 1003-1009, 2000.

[19] J. Accot and S. Zhai, "Beyond Fitt's Law: Models for Trajectory-Based hci Tasks," *Proc. Conf. Human Factors in Computing Systems (CHI),* pp. 295-302, 1997.

[20] R Development Core Team, *R: A Language and Environment for Statistical Computing,* R Foundation for Statistical Computing, Vienna, Austria, 2003, http://www.R-project.org.

Benjamin A.C. Forsyth received the BSc (Hon.) and MSc degrees, both in computer science, from the University of British Columbia located in Vancouver, Canada, in 2001 and 2004. respectively. He currently works for Apple Computer, Inc. in Vancouver, Canada.

Karon E. MacLean received the BS degree in biology and mechanical engineering from Stanford University in 1986, and the MSc and PhD degrees from the Massachusetts Institute of Technology in mechanical engineering in 1988 and 1996, respectively. She worked for two years for the Center for Engineering Design at the University of Utah and, from 1996-2000, as a scientist at Interval Research in Palo Alto, California. She is currently an associate professor in the Department of Computer Science at the University of British Columbia, Vancouver, Canada, where her research focuses on haptic communication of information between people and intelligent systems. She is a member of the IEEE.