ECOLOGICAL PSYCHOLOGY, *18*(4), 257–281

A Primer on the Use of Intentional Dynamics Measures and Methods in Applied Research

Ittai Flascher
Department of Neurobiology
Weizmann Institute of Science, Rehovot, Israel

Robert E. Shaw and Claire F. Michaels
Center for the Ecological Study of Perception and Action
University of Connecticut

Oded M. Flascher
Science Applications International Corporation
San Diego, California

Intentional dynamics is a modeling framework that provides methods and measures with which to evaluate and predict goal-directed behavior. The authors introduce the framework by showing how it can be used to solve two general problems in the design of human–machine interaction as they are encountered by designers of driving assistance systems. The first problem the authors treat is that of path planning on tasks in which multiple goals have to be satisfied, necessitating the evaluation of performance relative to conflicting criteria. The second problem is that of design selection in which designers need to assess the effectiveness of competing designs in improving performance relative to such criteria. The authors close with a small empirical study testing the predictive validity of the approach.

Optimization principles have been widely used by engineers to plan and improve the movements and transportation of mobile robots (e.g., Lee, Kim, & Kohout, 2004; Pradalier et al., 2005). Similarly, the trajectories produced by skilled human actors on many motor tasks, including pointing, writing, and drawing, have been

Correspondence should be addressed to Ittai Flascher, Department of Neurobiology, Weizmann Institute of Science, Rehovot, 76100, Israel. E-mail: itadita@actcom.net.il

shown to optimize performance measures such as task duration, spatial accuracy, and movement energy, among others (D. Elliott, Hansen, Mendoza, & Tremblay, 2004; Engelbrecht, 2001; Flash & Hogan, 1995; Tanaka, Tai, & Qian, 2004; Todorov, 2004; Todorov & Jordan, 2002; Vetter, Flash, & Wolpert, 2002; Viviani & Flash, 1995). However, the assessment and prediction of performance and improvement on tasks that require the achievement of multiple objectives (e.g., minimizing time, distance, energy, etc.) remain a challenge across disciplines (Knowles & Corne, 2004; Shimansky, Kang, & He, 2004; Todorov, 2004).

The approach we propose in this article to overcome these challenges is rooted in the modeling framework of *intentional dynamics* (e.g., Shaw & Alley, 1985; Shaw, Flascher, & Kadar, 1995; Shaw, Kadar, & Kinsella-Shaw, 1994; Shaw, Kadar, Sim, & Repperger, 1992; Shaw & Kinsella-Shaw, 1988; Shaw, Kugler, & Kinsella-Shaw, 1990). That work has sought to formalize Gibson's (1979, pp. 232–235) rules for the perceptual control of action as an optimization principle of a general performance measure capable of predicting goal-directed behavior across tasks, environments, objectives, and even organisms. However, most of the work done to date has focused on the theoretical formulation of such a measure, which restricted its use in empirical studies and applied research (Effken & Kadar, 2001). More recently, the focus has shifted to the development of quantitative measures and methods that are required in the analysis and prediction of performance data (Effken & Kadar, 2001; Flascher, 2004; Flascher, Shaw, Michaels, & Flascher, 2003; Kadar, Maxwell, Stins, & Costall, 2002). The aim of this article is to provide an accessible introduction to the use of such methods and measures of intentional dynamics in robotics, assessment, and prediction of human goal-directed behavior, and the design of human–machine interaction.

We begin by making explicit two general challenges in the assessment of goal-directed performance as they are encountered by designers of driving assistance systems (DASs). Following that, we will use that example to show how the proposed methods and measures can be applied to overcome those challenges in DAS design and other applications. We close with a small empirical study testing the predictive validity of the approach.

CONSEQUENCES OF THE PERFORMANCE MEASUREMENT PROBLEM TO THE DESIGN OF DAS MANAGERS

Drivers commonly act to achieve multiple goals such as avoiding collisions, minimizing travel time, saving fuel, and experiencing a comfortable ride (Dia, 2002; Prokop, 2001; Stanard, Hutton, Warwick, McIlwaine, & McDermott, 2001). To facilitate the achievement of drivers' various goals, several DASs such as collision-avoidance systems, adaptive cruise control (ACC), lane-keeping and changing systems, and navigation aids are either currently under development or already available in cars (e.g., Marchau, van der Heijden, & Molin, 2005; Stevens, 2002). Several additional information and entertainment systems, as well as autonomous driving and parking systems, are being developed by car manufacturers and in academia and are expected to reach the market in the next few years (e.g., Aufrère et al., 2003; Pradalier et al., 2005; Yang, Cronin, Meltzer, & Zirker, 2003). The introduction of several DASs, however, necessitates their coordination and management (i.e., integration). Left uncoordinated, these systems can demonstrably hurt driving performance by flooding the driver with partial or conflicting information and warnings, and by adding the operation of their various controls to the driver's already heavy workload (e.g., Piechulla, Mayser, Gehrke, & König, 2003; Stanton, Young, Walker, Turner, & Randle, 2001; Takada & Shimoyama, 2001). To guarantee that DASs remain part of the solution to the challenges of driving, several integration systems are currently being developed (Rombaut, 1995; QinetiQ, 2003; Wood, Leivian, Massey, Bieker, & Summers, 2001). These "DAS managers" collect information from the separate DASs and make automated decisions regarding when and which of their available assistance capabilities to deploy to best promote the achievement of the various driving goals. To make deployment decisions, however, DAS managers must evaluate the extent to which several possibly conflicting goals are satisfied by the observed driving performance. That has proven to be a difficult problem with adverse effects on the ability of designers to plan driving paths appropriate for drivers as well as to their ability to evaluate the effectiveness of different designs in improving driving performance.

Path Planning

The first challenge we address is the path planning problem in which the DAS manager has to decide which path on the road the vehicle should follow (e.g., Lauffenburger, Basset, Coffin, & Gissinger, 2003). A typical scenario in which the difficulty encountered arises when the vehicle approaches another car traveling at a lower speed in the same lane. Using its sensors, the ACC detects the shrinking headway and attempts to reduce the cruise velocity. The collision avoidance systems, however, may attempt to accelerate the vehicle to change lanes and overtake the slower leading car. The difficulty in deciding which path achieves the goals better, and therefore should be promoted, stems from the conflicting performance demands that different goals specify. A takeover maneuver facilitates the achievement of the shorter travel-time goal yet can do so only by subjecting the driver to larger accelerations (forces) that violate the goal of maximizing comfort. As a consequence of such tradeoffs among the physical variables to be controlled, no driving path can completely satisfy all the goals. The DAS manager is therefore generally left with several possible path alternatives, each of which violates different goals to different degrees. To choose a driving path, the DAS manager requires a comprehensive performance measure relative to which the different alternatives can be ranked and the “best” (i.e., preferable) path identified. There is currently little agreement among DAS manager designers, however, on how to construct such a measure, or how human drivers assess (i.e., monitor) or choose a driving path relative to their goals (Clarke, Goodman, Perel, & Kipling, 1994; Hansberger & Holt, 2001; Helander, 1997; Macadam, 2003; Prokop, 2001).

The challenge of establishing a comprehensive performance measure to carry out a multiobjective path planning in the physical world is by no means unique to the driving task or to the workings of automated DAS managers. It is, rather, a general and open problem in the study of human motor control, robotics, and human–machine interaction (e.g., Fernandez, Gonzalez, Mandow, & Pérez-de-la-Cruz, 1999; Todorov, 2004). A principled approach to tackling such problems, therefore, would be of considerable value to the scientific and engineering communities.

Design Selection

The second challenge we address is that of deciding which of two competing DAS manager designs is more effective in facilitating the achievement of the driving goals. During the design process, researchers run tests in which participants drive vehicles equipped with competing designs and measure their relative ability to improve the driving performance. However, given that several measures are needed to assess improvement in performance relative to several goals, each of the competing designs can be shown to be superior to its competitor by some of the measures and inferior by others. This indeterminacy, inherent in the assessment of performance by multiple (partial) measures, has proven to be a general obstacle to the identification and selection of effective designs in practice; it is termed the performance *Measurement Problem* in the literature of human factors engineering (Helander, 1997, p. 10; Prothero, 1994; Sanders, 1991; Vreuls & Obermayer, 1985). In a survey of measures used to evaluate human–machine interfaces, Prothero (1994) writes,

> A chronic problem in the field of interface design is difficulty in measuring how well an interface performs. In the absence of such measures, the development of good interfaces tends to be an art, with little assurance when one is done that the final product is the best possible interface for the given task. The evolution of interface design from an art to engineering would seem to depend on the improvement of the measures used to test interfaces. (p. 1)

Therefore, providing a measure by which to quantify multiple-goal achievement and improvement independent of task would have applications well beyond the DAS manager and enable the generalization of theories of human performance across tasks and environments.

MULTIOBJECTIVE PATH PLANNING

Current Approaches to Measuring Driving Performance Relative to Multiple Objectives

Drivers' goals may specify to a DAS manager two types of criteria for choosing a driving path (e.g., Anderson, 2000). The first type, as in the case of the goal of avoiding collisions or adhering to traffic laws, is an *exclusion* criterion (*constraint*) specifying which paths are unacceptable and must never be chosen by the DAS manager. The second type of driving goals, such as minimizing driving time, distance, and accelerations, is a *selection* criterion (*objective*) specifying which of all the acceptable alternative paths is preferable.

There are several effective computational methods that a DAS manager may use to identify and discard candidate paths that violate the driving constraints. Furthermore, when only a single objective is employed as the decision criterion (e.g., minimize time), several methods are also available for DAS managers to identify the best path (e.g., Jaszkiewicz, 2001). However, when multiple objectives must be satisfied, DAS managers have to select, for example, between a path reaching the destination in 5 min, subjecting the driver to 100,000 N, along a route 1.5 km in length, and another path taking 7 min, 70,000 N, and 2.0 km. Which is the better path is obscured by the tradeoffs between the competing objectives (i.e., quick vs. comfortable) all measured using different units and dimensions. Choosing the better path depends then on how the different measures are reconciled into a comprehensive performance measure of goal attainment. Next, we briefly review the three approaches taken in the literature to establish such a measure (Knowles & Corne, 2004).

The reductionist approach. One approach taken in the literature for the selection of a driving path has been to avoid the problem of multiple goals by choosing a single goal as the decision criterion. Many applications of DAS managers, as well as autonomous vehicles, are therefore designed to promote paths of either minimal duration (Hassoun & Laugier, 1995), shortest distance (Lázaro et al., 2001), minimal energy consumption (Garg & Kumar, 2002), or minimal accelerations (Lubashevsky, Kalenkov, & Mahnke, 2002).

As effective as this approach is in controlling autonomous vehicles, not taking into consideration and supporting known users' goals increases the risk of users' dissatisfaction and lowers the product's prospects of gaining widespread acceptance. In other words, drivers and passengers may not be impressed with the DAS manager's ability to get them to their destination quickly if the driving feels to them like a roller-coaster ride. From a more general perspective, this approach simply does not offer a solution to the original problem of path planning in multiple-objective human motor tasks or robotics.

The weighted sum approach. Another prevalent approach to evaluating candidate-path performance is to assess how well they achieve each of the objectives separately (minimum time, distance, forces, etc.) and then sum all the measures into an overall performance score (e.g., Prokop, 2001; Soltani, Tawfik, Goulermas, & Fernando, 2002). That approach attempts to bypass the tradeoffs between the different objectives in the combined path selection criterion by considering some goals as more important than others. Therefore, every performance measure is assigned a weight (i.e., multiplied by some factor) according to how important it is and consequently the summed score emphasizes the achievement of the more important objectives at the expense of the competing goals.

In spite of the apparent simplicity of the approach, designers have to overcome two challenges before the desired weighted sum can be computed. First, time, distance, and forces are all quantities of different natures (dimensions) yielding different value ranges and units and therefore they cannot be added together (e.g., Zhimin, Xu, & Xianyi, 2000). To remedy that, researchers must establish a common scale for all the measures by defining the best and worst values on each of the measures and then mapping all these ranges into a common range (e.g., between 0–1). For example, designers may determine that the shortest possible route to the destination under the prevailing driving conditions is 1 km and that a route taking 5 km is unacceptably excessive regardless of how quick or comfortable it may be. Therefore, in this case a path 2 km in length may be assigned the score of (2 km – 1 km)/5 km = 0.2. Other measures are assigned scores in the same manner, yielding pure numbers (i.e., dimensionless) and a common scale of goal satisfaction, allowing the summation of the different performance measures. However, whereas DAS managers can in many cases estimate the limit on best performance possible on each measure given the wealth of data on road and traffic conditions available to them, establishing a maximal tolerance level for the tradeoff (e.g., 5 km) for each of the measures involves designers' subjective judgment, which may or may not be acceptable to different drivers.

Second, the relative importance of objectives is commonly established by consulting "experts" or letting drivers themselves indicate their preferences (e.g., fast vs. safe driving; Itoh, Miura, & Shirai, 1999; Miura, Itoh, & Shirai, 2002; Youssef, Sait, & Adiche, 2001). Weights cannot, however, be set according to subjective preferences or other extrinsic criteria alone. As remarked earlier, quantities such as time, distance, forces, and energy are physically related in the driving dynamics, and it is unlikely that even experts can estimate their tradeoffs accurately. For instance, drivers may consider minimal travel time to be their most important goal, yet they might be hard-pressed to estimate what saving a minute of driving time would imply for their comfort. Depending on the prevailing road and traffic conditions, saving a minute may involve quite unacceptably erratic driving. Therefore, even in cases in which drivers explicitly specify their goal preferences, determining the precise values of the different weights has remained a challenge (Miura et al., 2002; Uchibe, Yanase, & Asada, 2002; Youssef et al., 2001).

The ranking approach. A similar approach has been to do away with weighing different goals by setting a goal hierarchy or dominance. In the hierarchical case, some goals are considered more important than others, either through reasoning or explicit user preference, and optimization proceeds according to the order of importance (e.g., Fernandez et al., 1999). For example, it may be assumed that conserving fuel is more important than shortening driving time in some application; accordingly, a set of paths is sought in which energy consumption is small, and from among those, the ones with short durations are selected for further consideration. Therefore, the order of optimization plays a major part in determining which paths will be seen as the solution to be promoted by the DAS manager. Consequently, in cases in which there is not a single dominant goal, or in which the dominance relationship among paths is debatable, the hierarchical approach may not be effective.

In an alternative and often used ranking method, the preferred paths are those that satisfy *most* of the goals. More precisely, the paths selected are those that are at least as good as their competitors on all the goals, and better on at least one goal (i.e., *Pareto optimality*, Tan, Chew, & Lee, in press; Yuhara & Tajima, 2001; Zitzler & Thiele, 1999). Unfortunately, this method generally yields a set of paths (i.e., a front) that are equally valid under the rule, and additional criteria or a weighting technique must be further applied to choose among them. Furthermore, in many cases encountered in practice, competing paths are better than the alternatives on some of the measures and worse on others, and, therefore, cannot be ranked under the rule.

A Proposed Measure for Multiobjective Path Planning

Let us assume that the threefold driving objectives are to keep the values of force, distance, and time to a minimum. Then the product of *force · distance · time* yields a performance measure that is minimal when all the driving objectives are satisfied by a driving path and increases for every violation of any goal. Therefore, this physical quantity, termed *action*, does not require subjective criteria for its computation, nor does it require estimations of the physical tradeoffs between competing objectives, because those tradeoffs are expressed in the overall measure (i.e., product). Action has been shown to be minimized (i.e., stationary) in the motions of almost all physical particles, and the *principle of least action* is consequently considered as possibly the most general principle in physics (e.g., Tiffoli, 2003; Yourgrou & Mandelstam, 1968). Furthermore, the principle has been used to explain and predict phenomena in biology as well as skilled human drawing movements (Behera, 1996; Lebedev, Tsui, & van Gelder, 2001). However, to use action to assess and predict goal-directed driving, a number of changes to the usual definition of action have to be made.

To see the issues of interest clearly, we neglect much of the complexity inherent in a realistic car model (e.g., Miano, Gobbi, & Mastinu, 2004) and consider a (collision-free and legal) driving path $[\mathbf{q}] = \{\mathbf{q}_t \mid t_0 \leq t \leq t_f\}$ as a sequence (i.e., series) of the car's position in two dimensions $\mathbf{q} = \{x, y\}$ (e.g., the car's center of mass) from the 0th time slice, at which the car is at rest at the origin, to the final time slice *f*, at which the vehicle arrives at the destination. To minimize forces (*mass · acceleration*) along a path, the DAS manager (or driver) has to control two components of acceleration. The longitudinal (*Tangent*) component $a_t^{\|}$ arises from changes to the car's speed caused by the pressing on the brake pedal or the accelerator.[1] The lateral (*Normal*) component of the acceleration, $a_t^{\perp}$, arises from moving the steering wheel and changing the car's *direction* of motion.[2] However, in the traditional definition for the product of *force · distance*, known as physical *work*,

$$W_t = \sum_{t=i-1}^{t=i} F_t^{\|} \cdot \left\| \Delta \mathbf{q}_t \right\|$$

only speed changes $F_t^{\|} = a_t^{\|} \cdot m$ are considered (e.g., Marsden & Tromba, 1996, pp. 402–404), where m is the mass of the car. Therefore, the action measure

$$S[\mathbf{q}] = \sum_{t=1}^{t=f} W_t \cdot \Delta t,$$

which should include all the measures to be minimized by the driving path, does not account for all the forces and the Normal component,

$$W_t^{\perp} = \sum_{t=i-1}^{t=i} F_t^{\perp} \cdot \left\| \Delta \mathbf{q}_t \right\|,$$

must be added to the usual action calculations (see also Lebedev, et al., 2001):

$$S_{\perp}[\mathbf{q}] = \sum_{t=1}^{t=f} F_t^{\perp} \cdot \left\| \Delta \mathbf{q}_t \right\| \cdot \Delta t. \tag{1}$$

Furthermore, the values of the traditional work W_t are negative whenever the car is decelerating and reduce the (total) action value for the path. However, driving comfort is affected by both acceleration and deceleration alike and action should increase for either violation of the goal. That is achieved by taking the absolute value of the change in speed:

---

[1]Explicitly, $a_t^{\|} = (\dot{q}_t - \dot{q}_{t-1}) / \Delta t$ is the (tangent) component of the acceleration quantifying the rate of change in speeds between two time slices $\Delta t = t_i - t_{i-1}$, where t is any time slice $t_0 < t_i \leq t_f$ and speed $\dot{q}_t = \left\| \Delta \mathbf{q}_t \right\| / \Delta t = \sqrt{\left(x_t - x_{t-1}\right)^2 + \left(y_t - y_{t-1}\right)^2} / \Delta t$.

[2] $a_t^{\perp} = \kappa_t \cdot \dot{q}_t^2$, is the Normal component of acceleration quantifying the rate of change in the direction of motion, where $\kappa_t = \left\| \dot{\mathbf{T}}_t \right\| / \dot{q}_t$. is the path's curvature, $\mathbf{T}_t = \dot{\mathbf{q}}_t / \dot{q}_t$. is the (unit) tangent vector specifying the momentary *direction* of movement along the path, $\dot{\mathbf{T}}_t = \Delta \mathbf{T}_t / \Delta t$ is the rate of change in direction, and $\| \ \|$ is the vector norm (i.e., length, magnitude; e.g., Stewart, 1994, p.725).

$$S_{\parallel}\left[\mathbf{q}\right]=\sum_{t=1}^{t=f}\left|F_t^{\parallel}\right|\cdot\left\|\Delta\mathbf{q}_t\right\|\cdot\Delta t \tag{2}$$

Summing the different components of the action, we get a measure that will take into account every violation of the goals mentioned:[3]

$$S_W\left[\mathbf{q}\right]=S_{\parallel}\left[\mathbf{q}\right]+S_{\perp}\left[\mathbf{q}\right] \tag{3}$$

So far we have assumed that drivers prefer driving paths that maximize speed to satisfy their goal of minimizing time. We have further assumed that the maximal driving speed is constrained by the posted speed limits. Research has shown, however, that drivers act so as to maintain a preferred speed that is a function of additional variables such as road conditions, traffic regulations, and the distance from other vehicles, among other factors (e.g., Brockfeld & Wagner, 2003; Lubashevsky et al., 2002; Prokop, 2001; Treiber & Helbing, 2000). Therefore, maintaining the preferred driving speed is a separate and concurrent driving objective with minimizing travel time, and may or may not be below the legal speed limit (e.g., Comte, 2000; M. A. Elliott, Armitage, & Baughan, 2004). Consequently, action should include an additional term that increases for violations of that objective. Although in practice DAS managers have sufficient data on the driver, road, and traffic conditions to estimate the driver's preferred speed value by one of the models given in the literature (e.g., Brockfeld & Wagner, 2003; Treiber & Helbing, 2003) we simplify the presentation by equating the preferred speed to the road's speed limit $(\dot{q}_{\max})$:

$$S_{\Delta KE}\left[\mathbf{q}\right]=\sum_{t=1}^{t=f}\left|\left(\frac{m}{2}\cdot\dot{q}_{\max}^2\right)-\left(\frac{m}{2}\cdot\dot{q}_t^2\right)\right|\cdot\Delta t \tag{4}$$

The deviation in speed from the preferred value at every time slice is evaluated in terms of the absolute difference between the car's *kinetic energy*, $m/2\cdot\dot{q}_t^2$, and the preferred one. Multiplying the energy deviation by time yields an action quantity that may be summed with $S_W\left[\mathbf{q}\right]$ for a comprehensive measure of goal performance, which we term *total-generalized-action* (TGA). The DAS manager may then use the minimization of TGA as the criterion for selecting the path that best satisfies all the goals:

$$\widehat{S}\left[\mathbf{q}\right]=S_{\Delta KE}\left[\mathbf{q}\right]+S_W\left[\mathbf{q}\right] \tag{5}$$

[3]For simplicity we treat the car as a point, and so neglect forces relevant to the car's rotation (i.e., *torque*). That component of the action is achieved by the (inner) product of the torque along the angular displacement of the car (multiplied by time). We also do not treat fuel consumption minimization in this introduction due to the modeling complexity involved. For a treatment of fuel consumption minimization and its relations to engine and limb torques see, for example, Garg & Kumar (2002) and Prokop (2001).

TABLE 1
Tradeoffs Between Maximizing Speed and Minimizing Accelerations

Time Slice	*x,y*	$S_{\Delta KE}$	$S_{\|\|}$	$S_{\perp}$	$\hat{S}$
0	0,0	—	—	—	
1	17,17	0	578	578	
2	34,34	0	0	0	
3	51,51	0	0	0	
4	68,68	0	0	0	
5	85,85	0	0	0	
6	102,102	0	0	0	
7	119,119	0	0	0	
8	136,136	0	0	0	
9	153,153	0	0	0	
10	170,170	0	0	0	
Sum	—	0	578	578	1,156

To illustrate how TGA will reflect driving violations of the objectives, let us assume that the vehicle starts from rest at the origin {0,0} and must travel within a wide single-lane road to a destination located at grid point {170,170} with a speed limit close to 90 km/hr.[4] For a path maximizing speed at every time slice ($\Delta t = 1$ sec), and therefore never violating the goal of time minimization, the kinetic difference $S_{\Delta KE}$ of TGA remains zero as can be seen in Table 1.[5] However, to achieve maximal speed in one time slice, a very large acceleration is required that increases TGA through S_{W}.

As a second example, a car may travel at a constant speed along a meandering (zigzagging) path but that will increase the action cost due to the repeated changes in direction and deviation from the preferred (maximal) speed, as well as increasing the distance and time of travel, as shown in Table 2.

Therefore, a path that minimizes TGA, the path described in Table 3, is the one that balances the action between minimizing time and increasing speed gradually (smoothly) until the car travels at the preferred maximal speed, along the shortest route, in the direction of the destination.

We introduced TGA in the context of a simplified driving task. Nevertheless, there is nothing in the measure's definition that makes it specific to a particular task, environment, or the actor (i.e., a human or a robot). Therefore, the measure

[4]Because whole numbers are easier to follow, we set the maximal speed at 86.6 km/hr instead of 90, which translates into 17 m/sec along each of the axes. To further simplify the presentation, we took the mass of the car to be 1 kg to avoid the large values involved in the computations of action for the motion of a realistic passenger car ($500 \leq m \leq 1,500$ kg).

[5]Measuring time in seconds, distance in meters, and the mass of the car in kilograms yields an action unit of [joule · second], where a joule is the amount of energy exerted when a force of 1 N is applied over a distance of 1 m. One Newton is the amount of force required to accelerate a mass of 1 kg at a rate of 1 m/sec squared.

TABLE 2
The Different Components of TGA for a Zigzagging Path

Time Slice	*x,y*	$S_{\Delta KE}$	$S_{\parallel}$	$S_{\perp}$	$\hat{S}$
0	0,0	—	—	—	
1	16,17	16.5	545.00	545.00	
2	33,33	16.5	0.00	33.02	
3	49,50	16.5	0.00	33.02	
4	66,66	16.5	0.00	33.02	
5	82,83	16.5	0.00	33.02	
6	99,99	16.5	0.00	33.02	
7	115,116	16.5	0.00	33.02	
8	132,132	16.5	0.00	33.02	
9	148,149	16.5	0.00	33.02	
10	165,164	32.0	15.27	47.64	
11	170,170	258.5	116.07	9.33	
Sum	—	439.0	676.34	866.13	1,981

TABLE 3
Minimization of TGA in Our Example Is Achieved by a Gradual Increase in Accelerations at the Cost of an Additional Time Slice

Time Slice	*x,y*	$S_{\Delta KE}$	$S_{\parallel}$	$S_{\perp}$	$\hat{S}$
0	0,0	—	—	—	
1	2,2	285	8	8	
2	17,17	64	390	0	
3	34,34	0	68	0	
4	51,51	0	0	0	
5	68,68	0	0	0	
6	85,85	0	0	0	
7	102,102	0	0	0	
8	119,119	0	0	0	
9	136,136	0	0	0	
10	153,153	0	0	0	
11	170,170	0	0	0	
Sum	—	349	466	8	823

may be applied to the evaluation of goal-directed performance in general. For example, TGA may be applied equally well to assess the task performance of a waiter acting to minimize delivery time while minimizing accelerations to avoid spilling the drinks he is carrying. Furthermore, TGA may be applied to path planning under different objectives. For example, car racers' objectives will normally not include driving comfort, and instead racers will act to maximize their accelerations to minimize time. Given those objectives, TGA will not include $S_{\parallel}$ in the total sum, and will be minimized when speed is increased as quickly as possible. In another case, a soccer player may intend to kick the ball with as much force as possible yet the direction of the ball's flight may be all important, requiring $S_{\perp}$ in the calculation of TGA. Therefore, *TGA is goal-dependent measure of performance where the task's objectives implicate the components to be included in the sum.*

It remains to be determined empirically, however, whether human actors in fact travel paths that minimize TGA as they become experienced in driving or performing any other goal-directed task. The DAS manager can only identify drivers' preferred (optimized) paths and correctly anticipate drivers' assistance needs if the improvement in human goal-directed performance is predicted by a principle of least TGA. We test this prediction later, after making explicit how we measure improvement in terms of TGA.

DESIGN SELECTION

As a human assistance system, a DAS manager design is evaluated primarily in terms of its effectiveness in facilitating *improvement* in drivers' performance. There is more than one way in which drivers' performance may be improved, however, and we need to clarify how the different forms of improvement are measured and used in design selection. Let us assume then that a driver is asked to drive several times a day to a given destination and that the test is repeated on several consecutive days. The onboard DAS manager computes and records the value of TGA for every driven path (trial) and the data are examined at the end of each day to determine whether driving performance has improved. Let us consider the data in Figure 1 as depicting the relative frequencies with which different values of TGA were observed on a given day.

Because TGA measures the violations of the goals, one way that improvement in driving performance can manifest itself is in decreasing TGA values over time. In other words, over test days, the distribution of TGA values or a measure of its central tendency (e.g., mean, median, mode, etc.) will displace toward the origin. DAS managers, however, cannot be selected according to such a measure of driving improvement for two reasons.

First, there is currently no known theory by which to determine how close to the origin any given driver's TGA distribution may reach on any given day. Therefore, when displacement is no longer detected, designers cannot tell whether it is due to the success of the DAS manager in bringing the driver to perform at his or her very best (beyond which the driver cannot improve), or whether the design fails to optimize performance. Second, driving is a life-critical task and therefore improving performance on the average drive is an inadequate criterion for a design's effectiveness; each drive is a new and separate challenge and the DAS manager is required to improve performance on most drives, if not all.

A second type (dimension) of improvement is therefore given in the change of the relative frequencies with which different TGA values are performed over time. In Figure 2 samples from two TGA distributions are compared.

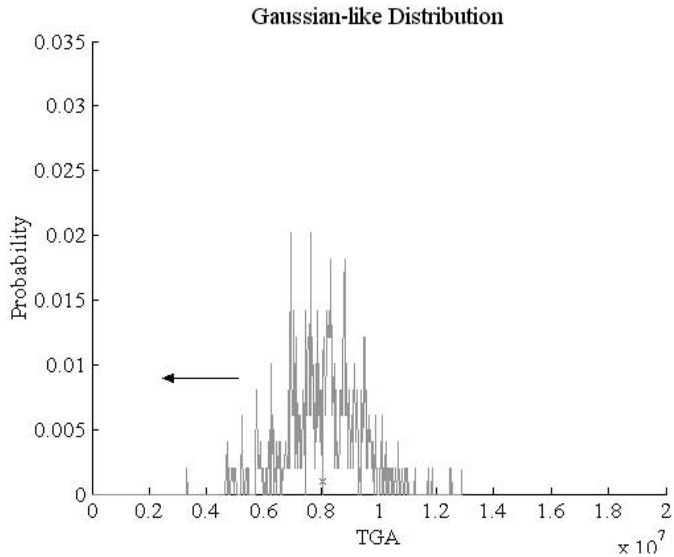

FIGURE 1 A sample of a Gaussian-like distribution.

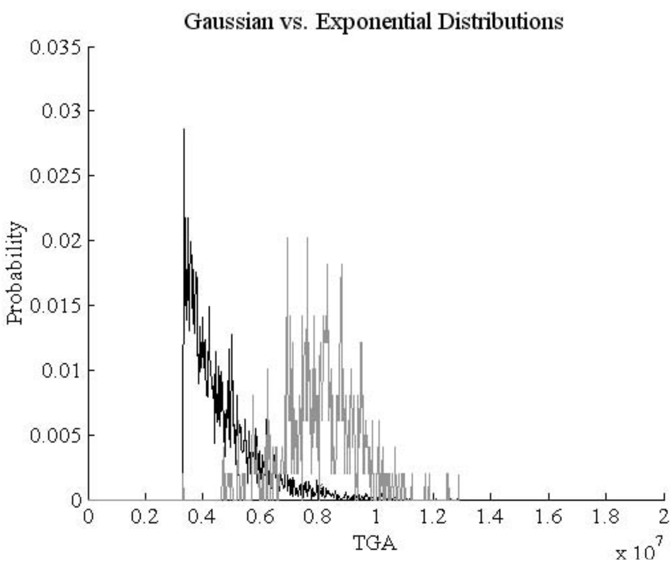

FIGURE 2 Comparing samples from a Gaussian (light) and Exponential (dark) distributions.

The Gaussian distribution shows driving performance that was rarely very good, mediocre on most drives, and seldom very bad. In contrast, the *Exponential* distribution indicates that the driver has performed close to his or her best most frequently and the occurrence of badly performed paths fell off rapidly. Therefore, the more effective DAS manager design is the one that brings most drivers to maximize the frequencies of minimal TGA drives in the shortest period of time, whatever their minimal value may be on any given day. More specifically, the effectiveness of the DAS manager can be measured by the rate of convergence from the TGA distribution observed on the first day to the exponential distribution, which minimizes TGA (i.e., the sum of TGA of all the paths performed in a day). However, as we remarked earlier, it remains to be determined whether the TGA distribution on which actors converge with experience is in fact the exponential one. We test that prediction next.

TESTING THE PREDICTIVE VALIDITY OF A LEAST TGA PRINCIPLE

We now test the feasibility of the approach and the predictive validity of the principle using a human motor-control task. In our small study, 4 participants used a joystick to move a graphically displayed sphere to a visible target (see Figure 3).[6] Participants were asked to displace the moveable sphere to the target sphere in minimum time on every trial, and to press a trigger to stop the timer and move on to the next trial. Each participant performed 10 sessions of 500 trials on different days.

To test the hypothesis that improvement in goal-directed performance follows (i.e., is predictable from) a least TGA principle, the distance between the observed and the exponential TGA distributions was computed for each session, and a statistical analysis was run to determine whether distances shrunk as the number of sessions performed increased. Pictorially, convergence will appear as a change over sessions in the shape of the observed distribution from a Gaussian-like distribution toward the exponential one, as shown in Figure 4.[7]

The details of each step in the analysis are made explicit in the following.

Computing a Path's TGA

The first step in the analysis is to compute TGA for each observed path (trial), where a path $[\mathbf{q}]$ is a sequence (series) of the sphere's position in two dimensions

---

[6]More details on the experimental methods can be found in Flascher (2004).

[7]We chose a smaller unit of action to quantify the dynamics of this fine motor control task than the $[J \cdot s]$ we used in the driving example. Given the sampling rate and spatial resolution of our measuring device, time was measured in units of 10 msec, distances in units of .03 mm, and mass in units of 1 g.

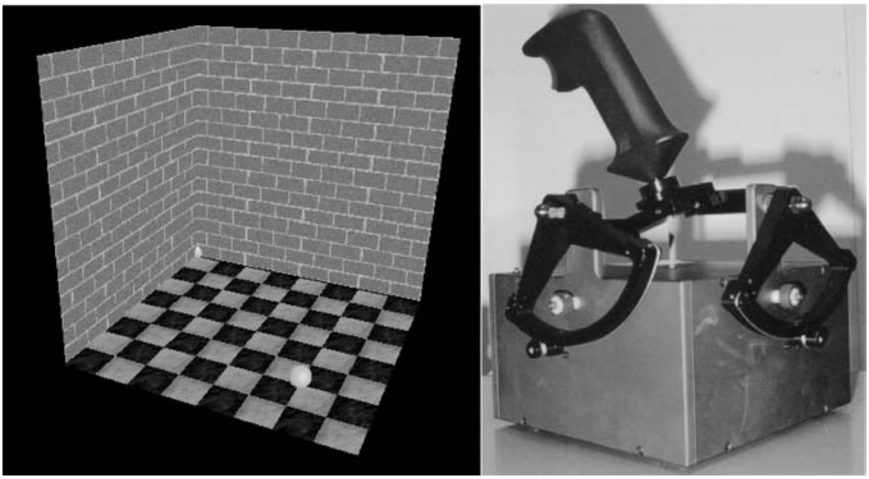

FIGURE 3 The experimental setup. A screenshot of the graphics display used in the experiment, and the force-feedback joystick (Immersion, 2001) used to control the motion of the sphere.

(determined by the joystick's handle position) sampled 100 times per second ($\Delta t =$ 10 msec). Trials (data collection) always started with the sphere (the handle pointing away from the participant) at the origin (at $t = t_0$) and terminated when the trigger was pressed (at $t = t_f$).

To compute a path's TGA we note that the goal of minimizing time specifies only two components of the measure that a path must minimize:

$$\widehat{S}[\mathbf{q}] = S_{\Delta KE}[\mathbf{q}] + S_{\perp}[\mathbf{q}] \tag{6}$$

In words, to reach as quickly as possible from the origin to the target, the participant must make the sphere travel along the shortest route, in the direction of the target, at the maximal speed possible, on every time slice.[8] Any slowdown would increase $S_{\Delta KE}$, and any unnecessary directional changes away from the shortest path would increase the $S_{\perp}$ component of TGA. In contrast, changes to speed are desirable under the given goal to achieve maximal speed as quickly as possible and therefore, $S_{\|}$, which increases with speed changes, is not included in the overall measure.

Unlike our driving example, however, in which drivers practically always reach their intended destination, in this fine motor control task participants may deviate

[8]The limiting maximal speed in the experimental setup was estimated from the results of an exploratory study in which participants attempted to move the handle as quickly as they could (without a display or any constraints). None of the participants in either the exploratory study or the experiment exceeded $\dot{q}_{\max} \approx 2.5$ m/sec on any time slice. Almost all trials lasted under half a second.

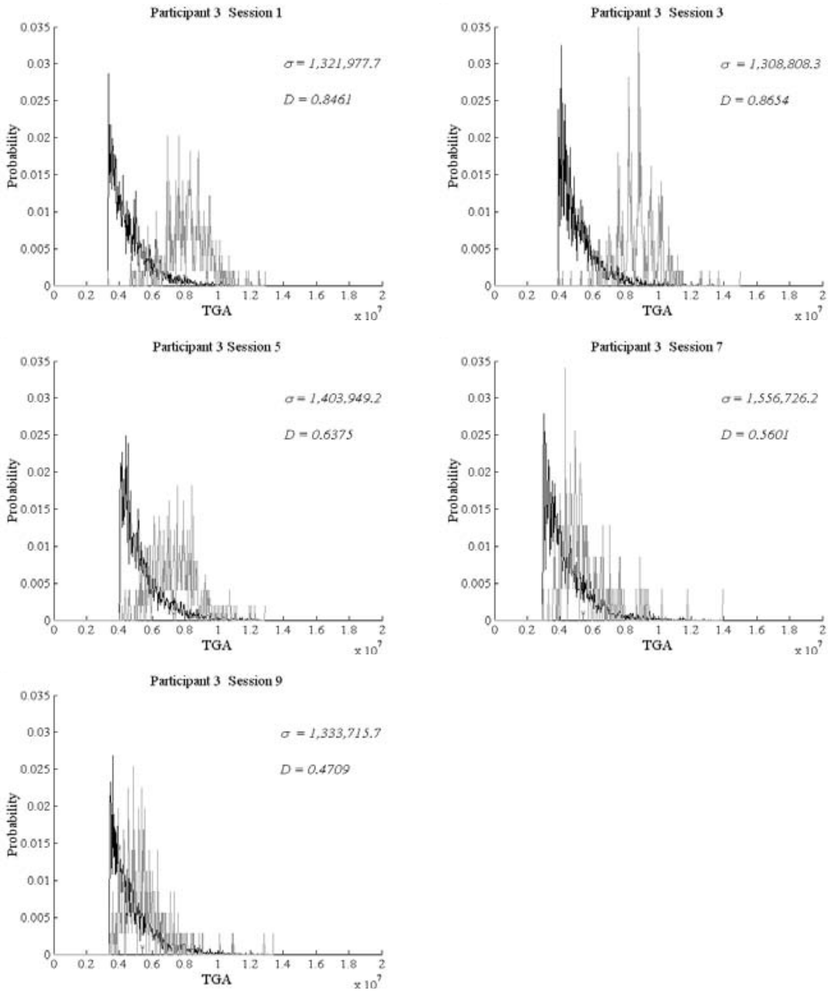

FIGURE 4 The distributions of TGA over a sequence of sessions performed by Participant 3. On Session 5 the Gaussian-like symmetry (around the central tendency) of the distribution is broken and the light (gray) observed distribution begins to show increasing convergence toward the optimal distribution (black).

from the target at t_f. Another objective of the task is, therefore, minimizing the deviation from the target. For such tasks, research has suggested that target deviations are a function of the energy and time costs of performing corrective movements to close that gap (D. Elliott et al., 2004; Sparrow & Newell, 1998). Therefore, to account for target deviation we add another action term in TGA by generating (simulating) a path $\left[\mathbf{q}^*\right] = \left\{\mathbf{q}_t^* \middle| t_f \le t \le t_{\text{Target}}\right\}$ starting from the last observed position $\mathbf{q}_f$ to the target as quickly as possible (according to the goal of the task) and computing $\hat{S}[\mathbf{q}^*]$ for that portion (see the Appendix for more details on how paths were simulated). Consequently, TGA for this multiobjective task is given by the sum of the observed (*retrospective*) and simulated (*prospective*) action components:

$$\hat{S}_j = \hat{S}[\mathbf{q}] + \hat{S}[\mathbf{q}^*] \tag{7}$$

Computing the Observed TGA Distribution

To determine the observed probability distribution we divide the relative frequencies of each TGA value observed in the session by the total number of trials performed. The joystick we used, however, is an extremely sensitive device capable of registering a .03-mm change in handle position. This spatial resolution is much finer than what humans can control, and therefore goal-directed behavior is not measurable at that precision level; if one attempts to hold the handle completely still, the path registered is nevertheless a random walk. Consequently, participants could never repeat the same precise TGA value twice in a session; therefore, TGA values not differing by more than 1% of the distribution's standard deviation were binned together for determining frequencies. Binning reduced the number of different TGA values observed in a session from 500 per session to between 300 and 400 bins.

Computing the Predicted (Optimal) TGA Distribution

The principle of least TGA predicts that as the number of trials N increases, the observed TGA distribution will converge on the exponential distribution:

$$N \to \infty : \Pr\left(\hat{S}_j\right) \to \pi_j(\sigma) = \frac{1}{Z(\sigma)} \exp\left(\frac{-\hat{S}_j}{\sigma}\right), \tag{8}$$

where,

$$Z(\sigma) = \sum_j \exp\left(\frac{-\hat{S}_j}{\sigma}\right). \tag{9}$$

In words, the probability of observing any possible value of TGA ($\hat{S}_j$) in a session will approach the *Boltzmann* optimal (equilibrium) distribution as more sessions are performed (see, e.g., Kirkpatrick, Gellat, & Vecchi, 1983; van Laarhoven & Aarts, 1987). Each path $[\mathbf{q}]_j$ that may be performed in a session is assigned a relative probability by the (negative) exponent of its measured TGA divided by the standard deviation σ of the observed TGA distribution. Due to the negative sign in the exponent, paths that minimize TGA have the highest relative probability of appearance. The probability that that TGA value will occur in the session is found by dividing the relative probability of such a path by the sum of *all* relative probabilities for all the paths that can possibly be performed $Z(\sigma)$.

There is a difficulty in computing $Z(\sigma)$, however, because there are more path possibilities to sum than we can hope to compute. For example, let us consider only paths lasting less than half a second (i.e., 50 time slices). At each time slice there are 600 · 600 spatial positions (Δq = 0.03 mm) that the sphere may occupy (at $\dot{q}_{\max} \approx 25 \text{ mm/10 ms}$) and therefore the number of different paths that could possibly be performed by a participant is $(600 \cdot 600)^{50} = 6.5332\text{e} = 277$. Even if we had a supercomputer of the future that was able to generate and compute TGA for a trillion paths each second, it would still take 2.0717e + 258 *years* to complete. Consequently, we have to compare the observed distribution to a sample from the Boltzmann distribution. We use a sampling method known as *simulated annealing* (e.g., Aarts, Korst, & van Laarhoven, 1988; Ingber, 1993) to generate a proper sample of the exponential distribution in the same value range (i.e., minimum and maximum values) and standard deviation observed at the given session (for details on the sampling method, see the Appendix). We therefore compare each participant's performance to the optimal distribution achievable at his or her current capabilities.

Testing Convergence to the Predicted TGA Distribution

The distance between distributions is commonly evaluated as the deviation between the distributions' means or other measures of central tendencies (e.g., D. Elliott et al., 2004). However, measures of central tendency reduce a distribution into a single point along the (TGA) value axis. As we saw, we need a measure of improvement along the frequency (probability) axis assessing the distance between the complete shapes of the distributions. To evaluate the distance between the observed and predicted probability distribution for a session we used the *variation distance* measure, also known as the *coefficient of ergodicity* (Denuit & van Bellegem, 2001; van Laarhoven & Aarts, 1987). That measure is used in the optimization literature to assess convergence to the optimal distribution:

$$D = \frac{1}{2}\sum_{i}\left|\pi_i(\sigma) - p_i^{obs}\right| \tag{10}$$

The measure is computed by summing the absolute difference between the observed and predicted probability values in all the bins along the TGA axis (indexed by i). When the two probability distributions are precisely congruent, and therefore the values of the observed and predicted distributions are identical in every bin, D = 0. When the probability distributions are completely independent and do not intersect (overlap) at any of the bins, the sum is 2; dividing by 2 then yields a probability measure (i.e., between 0 and 1.0).

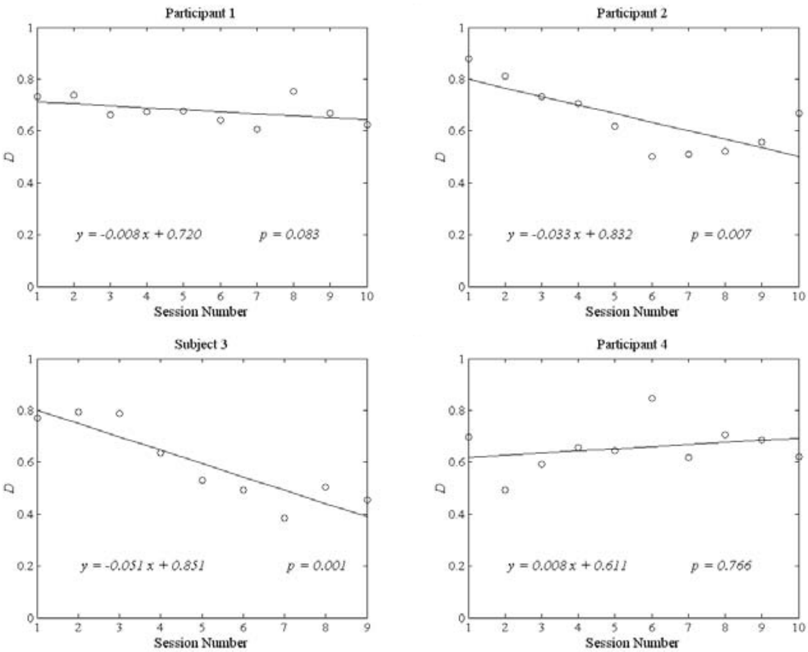

FIGURE 5 Convergence analysis results for the four participants.

To test for convergence we plotted the *D* values computed for each of the 10 sessions in the series. We then found the slope, θ, of the best fitting (least-squares) line, as shown in Figure 5. The research hypothesis we would like to test is then whether the slope of the fitted line is negative (decreasing *D* values). The null hypothesis for the experiment is, accordingly,

$$H_0 : \theta \geq 0 \tag{11}$$

In other words, the null hypothesis claims that the order of the *D* values in the experiment is in fact random and that any negative slope observed for any participant in the experiment is due to chance. To determine whether the observed slopes are (statistically) smaller than zero we ran a bootstrap test (e.g., Efron & Tibshirani, 1993; Lunneborg, 2000; Noreen, 1989). In that analysis the order of a participant's *D* values is randomly interchanged (permuted) to create a random sample (order) as the null hypothesis prescribes. θ_B is then computed for the reshuffled order and compared with the θ measured for the participant. A large number of reshuffles (*NS*) are performed and for every sequence (order) that θ_B is a number smaller or equal to θ a counter (*nse*) is increased by one. The significance level of the test is then given by the ratio $p = (nse + 1)/(NS + 1)$, which gives an estimate of how abundant slopes of the observed θ's magnitude are in a random distribution of sequences (i.e., the null distribution). Therefore, only when p values are small can we be confident that the observed decrease in the distance between the distributions is not a random occurrence.

The p values for three out of the four participants were small (0.083, 0.007, 0.001) and the null hypothesis can be rejected with some confidence ($NS = 1{,}000{,}000$). No convergence to the least-TGA distribution was found for Participant 4 ($p = 0.776$). Therefore, the results suggest that it is feasible that improvement in goal-directed performance converges toward the least-TGA distribution for this small sample.

SUMMARY AND CONCLUSIONS

In this article we described how the measures and methods of intentional dynamics can be used in the practice of applied research. In particular, we showed how a multiobjective measure of action (TGA) can be used to assess and assist human goal-directed performance as well as to plan paths for autonomous mobile robots. We have further shown how improvement in goal-directed performance can be measured, and how the distribution of TGA can be used in the identification and selection of assistance-system designs. However, the appropriateness (validity) of an approach using TGA to assess and assist human goal-directed performance largely depends on whether their performance tends to minimize TGA. The results of a feasibility study we ran indicate that three out of four participants showed the tendency to converge to the minimal TGA distribution when performing a simple motor control task. Nevertheless, there is still much need for further studies using larger samples and different tasks and objectives to test the applicability of the approach to the study of goal-directed human motor control and human assistance in general.

ACKNOWLEDGMENT

Ittai Flascher's research at the Weizmann Institute of Science was supported by the Ann Abrams Stone Postdoctoral Fellowship.

REFERENCES

Aarts, E. H. L., & Korst, J. H. M. (1989). *Simulated annealing and Boltzmann machines: A stochastic approach to combinatorial optimization and neural computing.* Chichester, NY: Wiley.

Aarts, E. H. L., Korst, J. H. M., & van Laarhoven, P. J. M. (1988). A quantitative analysis of the simulated annealing algorithm: A case study for the Traveling Salesman Problem. *Journal of Statistical Physics, 50,* 187–206.

Anderson, J. (2000). *A survey of multiobjective optimization in engineering design* (Tech. Rep. No. LiTH-IKP-R-1097). Linköping, Sweden: Linköping University. Retrieved October 1, 2003, from http://www.machine.ikp.liu.se/staff/johan/files/ikp_r_1097.pdf

Aufrère, R., Gowdy, J., Mertz, C., Thorpe, C., Wang, C-C., & Yata, T. (2003). Perception for collision avoidance and autonomous driving. *Mechatronics, 13,* 1149–1161.

Behera, N. (1996). Variational principles in evolution. *Bulletin of Mathematical Biology, 58*(1), 175–202.

Brockfeld, E., & Wagner, P. (2003). Calibration and validation of microscopic traffic flow models. In S. P. Hoogendoom, S. Luding, P. H. L. Bovy, M. Schreckenberg, & D. E. Wolf (Eds.), *Proceedings of Traffic and Granular Flow '03.* Berlin, Germany: Springer.

Clarke, R. M., Goodman, M. J., Perel, M., & Kipling, R. R. (1994). *Driver performance and IVHS collision avoidance systems: A search for design-relevant measurement protocols.* ITS America Annual Meeting, Washington, DC. Retrieved December 15, 2003, from http://www.itsdocs.fhwa.dot.gov//JPODOCS/REPTS_TE/2W601!.PDF

Comte, S. L. (2000). New systems: New behaviour? *Transportation Research Part F, 3,* 95–111.

Denuit, M., & van Bellegem, S. (2001). On the stop-loss and total variation distances between random sums. *Statistics & Probability Letters, 53,* 153–165.

Dia, H. (2002). An agent-based approach to modelling driver route choice behaviour under the influence of real-time information. *Transportation Research Part C, 10,* 331–349.

Effken, J., & Kadar, E. E. (2001). Learning as progressive biasing of a random walk process: Toward an ecological theory of exploratory learning. *Ecological Psychology, 13,* 1–30.

Efron, B., & Tibshirani, R. J. (1993). An introduction to the bootstrap. New York: Chapman & Hall.

Elliott, D., Hansen, S., Mendoza, J., & Tremblay, L. (2004). Learning to optimize speed, accuracy, and energy expenditure: A framework for understanding speed–accuracy relations in goal-directed aiming. *Journal of Motor Behavior, 36,* 339–351.

Elliott, M. A., Armitage, C. J., & Baughan, C. J. (2004). Drivers' compliance with speed limits: An application of the theory of planned behaviour. *Journal of Applied Psychology, 88,* 964–972.

Engelbrecht, S. E. (2001). Minimum principles in motor control. *Journal of Mathematical Psychology, 45,* 497–542.

Fernandez, J. A., Gonzalez, J., Mandow, L., & Pérez-de-la-Cruz, J. L. (1999). Mobile robot path planning: A multicriteria approach. *Engineering Applications of Artificial Intelligence, 12,* 543–554.

Flascher, I. (2004). *Goal-centered approach to the measurement of human-systems performance* (Doctoral dissertation, University of Connecticut, 2004). *Dissertation Abstracts International, 65*(1056), 02B.

Flascher, I., Shaw, R. E., Michaels, C. F., & Flascher, O. M. (2003). Methods for online management of AmI capabilities relative to users' goals. In E. Aarts, R. Collier, E. van Loenen, & B. de Ruyter (Eds.), *Proceedings of EUSAI 2003, First European Symposium on Ambient Intelligence* (pp. 334–348). Berlin, Germany: Springer.

Flash, T., & Hogan, N. (1995). Optimization principles in motor control. In M. A. Arbib (Ed.), *The handbook of brain theory and neural networks* (pp. 682–685). Cambridge, MA: MIT Press.

Garg, D. P., & Kumar, M. (2002). Optimization techniques applied to multiple manipulators for path planning and torque minimization. *Engineering Applications of Artificial Intelligence, 15,* 241–252.

Gibson, J. J. (1986). *The ecological approach to visual perception.* Hillsdale, NJ: Lawrence Erlbaum Associates, Inc. (Original work published 1979)

Hansberger, J. T., & Holt, R. W. (2001). The implications of goals and experience for driving expertise. In *Proceedings of the 45th Annual Meeting of the Human Factors & Ergonomics Society* (pp. 921–925), Minneapolis/St. Paul, MN: HFES.

Hassoun, M., & Laugier, C. (1995). An architecture for motion planning and motion control of a car-like vehicle. *Mathematical Computer Modelling, 22,* 329–341.

Helander, M. G. (1997). The human factors profession. In G. Salvendy (Ed.), *Handbook of human factors and ergonomics* (2nd ed., pp. 3–16). New York: Wiley.

Immersion. (2001). Impulse Engine 2000, Software Development Kit (Release 4.2) [Computer software]. Retrieved January 1, 2002, from http://www.immersion.com/industrial/docs/iesdk4_ 2.pdf

Ingber, L. (1993). Simulated annealing: Practice versus theory. *Journal of Mathematical Computer Modelling, 18*(11), 29–57. Retrieved from http://www.ingber.com/asa93_sapvt.pdf

Itoh, M., Miura, J., & Shirai, Y. (1999). Towards intelligent navigator that can provide timely advice on safe efficient driving. In *Proceedings of the 1999 IEEE/IEEJ/JSAI International Conference on Intelligent Transportation Systems* (pp. 981–986).

Jaszkiewicz, A. (2001). *A comparative study of multiple-objective metaheuristics on the bi-objective set covering problem and the Pareto memetic algorithm* (Working paper RA-003/01). Poznan, Poland: Poznan University of Technology. Available from http://citeseer.ist.psu.edu/522037.html

Kadar, E. E., Maxwell, J. P., Stins, J., & Costall, A. (2002). Drifting towards a diffuse control model of exploratory motor learning: A comparison of global and within-trial performance measures. *Biological Cybernetics, 87,* 1–9.

Kirkpatrick, S., Gellat, C. D., & Vecchi, M. P. (1983). Optimization by simulated annealing. *Science, 220*(4598), 671–680. Retrieved June 15, 2005, from http://citeseer.ist.psu.edu/kirkpatrick83 optimization.html

Knowles, J., & Corne, D. (2004). Memetic algorithms for multiobjective optimization: Issues, methods and prospects. In N. Krasnogor, J. E. Smith, & W. E. Hart (Eds.), *Recent advances in memetic algorithms* (pp. 313–352). Heidelberg, Germany: Springer.

Landau, D. P., & Binder, K. (2000). *A guide to Monte Carlo simulations in statistical physics.* Cambridge, England: Cambridge University Press.

Landau, R. H., & Páez, M. J. (1997). *Computational physics: Problem solving with computers.* New York: Wiley.

Lauffenburger, J. P., Basset, M., Coffin, F., & Gissinger, G. L. (2003). Driver-aid system using path-planning for lateral vehicle control. *Control Engineering Practice, 11,* 217–231.

Lázaro, J. L., García, J. C., Mazo, M., Gardel, A., Martín, P., Fernández, I., et al. (2001). Distributed architecture for control and path planning of autonomous vehicles. *Microprocessors and Microsystems, 25,* 159–166.

Lebedev, S., Tsui, W. H., & van Gelder, P. (2001). Drawing movements as an outcome of the principle of least action. *Journal of Mathematical Psychology, 45,* 43–52.

Lee, Y., Kim, Y., & Kohout, L. J. (2004). An intelligent collision avoidance system for AUVs using fuzzy relational products. *Information Sciences, 158,* 209–232.

Lubashevsky, I., Kalenkov, S., & Mahnke, R. (2002). Towards a variational principle for motivated vehicle motion. *Physical Review E, 65,* 036140–036145.

Lunneborg, C. E. (2000). *Data analysis by resampling: Concepts and applications.* Pacific Grove, CA: Duxbury.

Macadam, C. C. (2003). Understanding and modeling the human driver. *Vehicle System Dynamics, 40*(1–3), 101–134.

Marchau, V. A. W. J., van der Heijden, R. E. C. M., & Molin, E. J. E. (2005). Desirability of advanced driver assistance from road safety perspective: The case of ISA. *Safety Science, 43,* 11–27.

Marsden, J. E., & Tromba, A. J. (1996). *Vector calculus* (4th ed.). New York: Freeman.

Martínez-Alfaro, H., & Gómez-García, S. (1998). Mobile robot path planning and tracking using simulated annealing and fuzzy logic control. *Expert Systems With Applications, 15,* 421–429.

Metropolis, M., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. (1953). Equations of state calculations by fast computing machine. *Journal of Chemical Physics, 21,* 1087–1092.

Miano, C., Gobbi, M., & Mastinu, G. (2004). Multi-objective optimization of the handling performances of a road vehicle: A fundamental study on tire selection. *Journal of Mechanical Design, 126,* 687–702.

Miura, J., Itoh, M., & Shirai Y. (2002). Toward vision-based intelligent navigator: Its concepts and prototype. *IEEE Transactions on Intelligent Transportation Systems, 3*(2), 136–146.

Noreen, E. W. (1989). *Computer intensive methods for testing hypotheses: An introduction.* New York: Wiley.

Piechulla, W., Mayser, C., Gehrke, H., & König, W. (2003). Reducing drivers' mental workload by means of an adaptive man–machine interface. *Transportation Research Part F, 6,* 233–248.

Pradalier, C., Hermosillo, J., Koike, C., Braillon, C., Bessière, P., & Laugier, C. (2005). The *CyCab*: A car-like robot navigating autonomously and safely among pedestrians. *Robotics and Autonomous Systems, 50,* 51–67.

Prokop, G. (2001). Modeling human vehicle driving by model predictive online optimization. *Vehicle System Dynamics, 35*(1), 19–53.

Prothero, J. (1994). *A survey of interface goodness measures* (Human Interface Technology Laboratory Tech. Rep. R-94–1). University of Washington. Retrieved October 5, 2002, from http://www.hitl.washington.edu/publications/r-94–1/r-94–1.pdf

QinetiQ. (2003). *EVAID* [UK trademark No. 2299745]. London: UK Patent Office. Retrieved June 15, 2005, from http://webdb4.patent.gov.uk/tm/number?detailsrequested=C&trademark= 2299745

Rombaut, M. (1995). ProLab2: A driving assistance system. *Mathematical Computational Modelling, 22*(4–7), 103–118.

Sanders, A. F. (1991). Simulation as a tool in the measurement of human performance. *Ergonomics, 34,* 995–1025.

Shaw, R. E., & Alley, T. R. (1985). How to draw learning curves: Their use and justification. In T. D. Johnston & A. T. Pietrewicz (Eds.), *Issues in the ecological study of learning* (pp. 275–304). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Shaw, R. E., Flascher, O. M., & Kadar, E. E. (1995). Dimensionless invariants for intentional systems: Measuring the fit of vehicular activities to environmental layout. In J. Flach, P. Hancock, J. Caird, & K. Vicente (Eds.), *Global perspectives on the ecology of human–machine systems* (Vol. 1, pp. 293–358). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Shaw, R. E., Kadar, E. E., & Kinsella-Shaw, J. (1994). Modelling systems with intentional dynamics: A lesson from quantum mechanics. In Karl Pribram (Ed.), *Origins: Brains & self-organization* (pp. 53–101). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Shaw, R. E., Kadar, E. E., Sim, M., & Repperger, D. W. (1992). The intentional spring: A strategy for modeling systems that learn to perform intentional acts. *Journal of Motor Behavior, 24*(1), 3–28.

Shaw, R. E., & Kinsella-Shaw, J. (1988). Ecological mechanics: A physical geometry for intentional constraints. *Human Movement Science, 7,* 155–200.

Shaw, R. E., Kugler, P. N., & Kinsella-Shaw, J. M. (1990). Reciprocities of intentional systems. In R. Warren & A. Wertheim (Eds.), *Control of self-motion* (pp. 579–619). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Shimansky, Y. P., Kang, T., & He, J. (2004). A novel model of motor learning capable of developing an optimal movement control law online from scratch. *Biological Cybernetics, 90,* 133–145.

Soltani, A. R, Tawfik, H., Goulermas, J. Y., & Fernando, T. (2002). Path planning in construction sites: Performance evaluation of the Dijkstra, A*, and GA search algorithms. *Advanced Engineering Informatics, 16,* 291–303.

Sparrow, W. A., & Newell, K. M. (1998). Metabolic energy expenditure and the regulation of movement economy. *Psychonomic Bulletin & Reviews, 5,* 173–196.

Stanard, T., Hutton, R. J. B., Warwick, W., McIlwaine, S., & McDermott, P. L. (2001). A computational model of driver decision making at an intersection controlled by a traffic light. In *Proceedings of Driving Assessment 2001: International Driving Symposium on Human Factors in Driver Assessment, Training and Vehicle Design, Aspen, Colorado, August 14–17*. Retrieved from http://ppc.uiowa.edu/driving-assessment//2001/Summaries/Driving%20Assessment%20Papers/65_ stand.pdf

Stanton, N. A., Young, M. S., Walker, G. H., Turner, H., & Randle, S. (2001). Automating the driver's control tasks. *International Journal of Cognitive Ergonomics, 5,* 221–236.

Stewart, J. (1994). *Calculus—Early transcendentals* (3rd ed.). Boston: Brooks/Cole.

Stevens, A. (2002). *Area 2B: Human machine interfaces.* Retrieved December 1, 2003, from http://www.trg.soton.ac.uk/rosetta/workareas/2b_hmi/hmi_pr1_driver.htm

Tan, K. C., Chew, Y. H., & Lee, L. H. (in press). A hybrid multi-objective evolutionary algorithm for solving truck and trailer vehicle routing problems. *European Journal of Operational Research*.

Takada, Y., & Shimoyama, O. (2001). Evaluation of driving-assistance systems based on drivers' workload. In *Proceedings of Driving Assessment 2001: International Driving Symposium on Human Factors in Driver Assessment, Training and Vehicle Design, Aspen, Colorado, August 14–17*. Retrieved from http://ppc.uiowa.edu/driving-assessment/2001/Summaries/Driving%20Assessment%20Papers/41_takada_yuji2.htm

Tanaka, H., Tai, M., & Qian, N. (2004). Different predictions by the minimum variance and minimum torque-change models on the skewness of movement velocity profiles. *Neural Computation, 16,* 2021–2040.

Tiffoli, T. (2003). What is the Lagrangian counting? *International Journal of Theoretical Physics, 42,* 363–381.

Todorov, E. (2004). Optimality principles in sensorimotor control. *Nature Neuroscience, 7,* 907–915.

Todorov, E., & Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nature Neuroscience, 5,* 1226–1235.

Treiber, M., & Helbing, D. (2000). Congested traffic states in empirical observations and microscopic simulations. *Physical Review E, 62*(2), 1805–1824.

Treiber, M., & Helbing, D. (2003). Memory effects in microscopic traffic models and wide scattering in flow-density data. *Physical Review E, 68,* 046119–046127.

Uchibe, E., Yanase, M., & Asada, M. (2002). Behavior generation for a mobile robot based on the adaptive fitness function. *Robotics and Autonomous Systems, 40,* 69–77.

van Laarhoven, P. J. M. (1988). Theoretical and computational aspects of simulated annealing. (Doctoral dissertation, Erasmus University Rotterdam, 1988).

van Laarhoven, P. J. M., & Aarts, E. H. L. (1987). *Simulated annealing: Theory and applications.* Dordrecht, The Netherlands: Reidel.

Vetter, P., Flash, T., & Wolpert, D. M. (2002). Planning movements in a simple redundant task. *Current Biology, 12,* 488–491.

Viviani, P., & Flash, T. (1995). Minimum-jerk, two-thirds power law, and isochrony: Converging approaches to movement planning. *Journal of Experimental Psychology: Human Perception and Performance, 21*(1), 32–53.

Vreuls, D, & Obermayer, R. W. (1985). Human-system performance measurement in training simulators. *Human Factors, 27,* 241–250.

Wood, C., Leivian, R., Massey, N., Bieker, J., & Summers, J. (2001). Driver Advocate™ Tool. In *Proceeings of Driving Assessment 2001: International Driving Symposium on Human Factors in Driver Assessment, Training and Vehicle Design, Aspen, Colorado, August 14–17*. Retrieved February 3, 2003, from http://ppc.uiowa.edu/driving-assessment/2001/Summaries/Driving%20Assessment%20Papers/61_Wood_Chip.htm

Yang, D. C. Y., Cronin, B. P., Meltzer, N. R., & Zirker, M. E. (2003). *Status report on transit intelligent vehicle initiative studies* (Tech. Rep. No. FTA-TRI-11-2003.1). Retrieved June 12, 2004, from http://www.itsa.org/ITSNEWS.NSF/4e0650bef6193b3e852562350056a3a7/c42207272082f75385256d56006f34fa/$FILE/2003_Transit_IVI_Status_Report.pdf

Yourgrau, W., & Mandelstam, S. (1968). *Variational principles in dynamics and quantum theory.* New York: Dover.

Youssef, H., Sait, S. M., & Adiche, H. (2001). Evolutionary algorithms, simulated annealing and tabu search: A comparative study. *Engineering Applications of Artificial Intelligence, 14,* 167–181.

Yuhara, N., & Tajima, J. (2001). Advanced steering system adaptable to lateral control task and driver's intention. *Vehicle System Dynamics, 36*(2–3), 119–158.

Zhimin, Y., Xu, W., & Xianyi, Z. (2000). Multi-criteria control system parameters optimization based on genetic algorithm. In *Proceedings of the 3rd World Congress on intelligent control and automation, June 28–July 2, Hefei, P.R. China* (pp. 651–655). Piscataway, NJ: IEEE Standards Office.

Zitzler, E., & Thiele, L. (1999). Multiobjective evolutionary algorithms: A comparative case study and strength Pareto approach. *IEEE Transactions on Evolutionary Computation, 3,* 257–271.

APPENDIX

Here we outline the process by which paths are simulated and the predicted TGA distribution is generated.

The simulated annealing method constructs a sample of the predicted (Boltzmann) distribution in two stages. In the first stage—the *Generating* stage—the method constructs a random series of positions (a path) $[\mathbf{q}]_j$ that satisfies all the task constraints. Therefore, in our case the algorithm will generate only paths that do not violate the maximal speed in the task (or collision-free paths in the driving example). TGA $\hat{S}_j$ is then computed for that path, and only those paths with a value that is within the participant's observed range are used in further computations. The algorithm then constructs a new path, $[\mathbf{q}]_{j+1}$, by randomly changing the position of the (simulated) sphere on a randomly selected time slice. It then computes $\hat{S}_{j+1}$ for the new path.

In the second stage—the *Accepting (Selecting)* stage—the algorithm decides which of the paths is accepted as a member of the predicted distribution. The rule for selecting paths is known as the *Metropolis* acceptance criterion (Metropolis et al., 1953): If $\hat{S}_{j+1} \leq \hat{S}_j$, then the value of $\hat{S}_{j+1}$ is recorded in the frequency table (the counter for this value is increased by one). If $\hat{S}_{j+1} > \hat{S}_j$, $\hat{S}_{j+1}$ is not always rejected, as long as the deviation, $\Delta\hat{S} = \hat{S}_{j+1} - \hat{S}_j$, is not "large":

$$\text{if,} \quad \exp\left(\frac{-\Delta\hat{S}}{\sigma}\right) > rand\left[0,\ 1\right), \ \text{then accept} \tag{12}$$

In words, the probability of still accepting the new path (with larger TGA) decreases exponentially with increased TGA excess. Sigma (σ) is the standard deviation of the observed TGA distribution and therefore the variability of optimal distribution is appropriately scaled to the current capabilities (resolution) of the participant. If $\hat{S}_{j+1}$ is rejected, a counter in the frequency table is increased for the $\hat{S}_j$ value. After sampling hundreds of thousands of paths (properly) the frequencies in the table are divided by the number of samples yielding the predicted probability distribution $\pi_j(\sigma) = \exp\left(-\hat{S}_j/\sigma\right)/Z(\sigma)$.

For more information on the Metropolis algorithm and the simulated annealing method, see the references listed previously and, for example, Aarts and Korst (1989), van Laarhoven (1988), Landau and Binder (2000, chap. 4), Landau and Páez (1997), and Martínez-Alfaro and Gómez-García (1998).