Modeling Driver Behavior in a Cognitive Architecture

Dario D. Salvucci, Drexel University, Philadelphia, Pennsylvania

Objective: This paper explores the development of a rigorous computational model of driver behavior in a cognitive architecture – a computational framework with underlying psychological theories that incorporate basic properties and limitations of the human system. **Background:** Computational modeling has emerged as a powerful tool for studying the complex task of driving, allowing researchers to simulate driver behavior and explore the parameters and constraints of this behavior. **Method:** An integrated driver model developed in the ACT-R (Adaptive Control of Thought-Rational) cognitive architecture is described that focuses on the component processes of control, monitoring, and decision making in a multilane highway environment. **Results:** This model accounts for the steering profiles, lateral position profiles, and gaze distributions of human drivers during lane keeping, curve negotiation, and lane changing. **Conclusion:** The model demonstrates how cognitive architectures facilitate understanding of driver behavior in the context of general human abilities and constraints and how the driving domain benefits cognitive architectures by pushing model development toward more complex, realistic tasks. **Application:** The model can also serve as a core computational engine for practical applications that predict and recognize driver behavior and distraction.

INTRODUCTION

Driving is a very common yet highly complex task that involves dynamic interleaving and execution of multiple critical subtasks. To explore how people perform this complex task, researchers have developed a variety of models to account for and simulate driver behavior. Some of these models are primarily conceptual models that help one to understand the representational and procedural components of the driving task (e.g., Boer & Hoedemaeker, 1998). Others are computational models that compute, simulate, and predict various aspects of driver behavior (e.g., Donges, 1978; Godthelp, 1986; Hildreth, Beusmans, Boer, & Royden, 2000). These computational models have emerged as powerful tools for both theoretical study of driver behavior (e.g., study of the perception-action aspects of steering) and practical development of real-world intelligent vehicle systems (e.g., intelligent lane guidance and warning systems). In particular, the research community has recently witnessed a growing push for integrated driver models – models that unify the many aspects of driving into a single, larger scale computational model of behavior. Past and ongoing efforts toward integrated driver models (e.g., Levison & Cramer, 1995; Tsimhoni & Liu, 2003) have shown great promise, accounting for aspects of behavior during normal driving and even (albeit to a much more limited extent) performance when driving while performing secondary tasks.

Driving and Integrated Driver Modeling

Even with these successes, the community has a great deal more work to do in achieving a truly integrated driver model that simulates and predicts real-world driver behavior. To better understand the power and limitations of existing models, it is useful to view driving and driver modeling in the context of the embodied cognition, task, and artifact (ETA) framework (Byrne, 2001; Gray, 2000; Gray & Boehm-Davis, 2000). As the name suggests, this framework emphasizes three components of an integrated modeling effort: the task that a person attempts to perform, the artifact (or

Address correspondence to Dario D. Salvucci, Department of Computer Science, Drexel University, 3141 Chestnut St., Philadelphia, PA 19104; salvucci@cs.drexel.edu. ***HUMAN FACTORS*****,** Vol. 48, No. 2, Summer 2006, pp. 362–380.

instrument) by which the person performs the task, and the embodied cognition by which the person perceives, thinks, and acts in the world through the artifact. A sound understanding of each component is critical to developing rigorous integrated models of driver behavior.

The task of driving is in fact an ever-changing set of basic tasks that must be integrated and interleaved. Michon (1985) identified three classes of task processes for driving: operational processes that involve manipulating control inputs for stable driving, tactical processes that govern safe interactions with the environment and other vehicles, and strategic processes for higher level reasoning and planning. Driving typically involves all three types of processes working together to achieve safe, stable navigation – for instance, monitoring a traffic light and controlling the vehicle to stop and start, or deciding to make a turn and controlling the vehicle through the turn. Some tasks are not continual but intermittent, arising in specific situations – for instance, parking a vehicle at a final destination. In addition, driving may include secondary tasks, perhaps related to the primary driving task (e.g., using a navigation device) or perhaps mostly or entirely unrelated (e.g., tuning a radio or dialing a cellular phone).

The "artifact" for driving is the vehicle itself and the interface between the human and the vehicle. Most recognizably, this interface includes the steering wheel, the accelerator (throttle) and brake pedals, and possibly the clutch pedal (on a manual transmission); it also includes related controls such as turn signals, headlights, and windshield wipers. These components of the vehicle interface are, in large part, standardized among different vehicles. For secondary tasks, the artifact also includes any interface to the secondary device – typically, knobs, buttons, and other inputs, along with small displays and other outputs; unlike the control-related components, secondary task components are much less standardized among today's vehicles. Of course, these interface components can be incorporated into any number of vehicle types, from compact to midsize to luxury vehicles, to sports utility vehicles and light trucks, to large trucks, military vehicles, and so forth, producing in effect a seemingly endless variety of possible artifacts for the driving task.

Embodied cognition is the integrated cognitive, perceptual, and motor processes that manipulate the vehicle and execute the desired tasks. Of course, driving requires cognition during even routine driving, most obviously for higher level decision making, more subtly for lower level vehicle control and situation awareness (see, e.g., Groeger, 2000). Between cognition and the vehicle lies the embodiment of the driver, namely the perceptual processes (visual, aural, vestibular, etc.) and motor processes (hands, feet) that provide the input from and output to the external world. Not surprisingly, there can be parallelism in this integrated system – for instance, moving the hand while visually encoding the lead car – but there are also capacity constraints and/or bottlenecks that sometimes result in degraded performance.

The goal of integrated driver modeling is to rigorously address all three of these components: handling as many driving-related tasks as possible, incorporating realistic controls and vehicle dynamics, and performing the tasks through cognitive processes that interact through realistic perceptual and motor processes. Many existing models of driver behavior and performance emphasize only one or two of the ETA-triad components. Perception-and-action models of control (e.g., Fajen & Warren, 2003; Rushton, Harris, Lloyd, & Wann, 1998; Salvucci & Gray, 2004; Wilkie & Wann, 2003) provide a firm theoretical basis for how perception and action interact in basic tasks such as lateral and longitudinal control (for driving or related tasks such as locomotion); however, these models are often formulated as continuous-valued functions without realistic vehicle dynamics and without explicit constraints on the perceptual and/or motor systems (e.g., delays in perceiving objects or turning a wheel). More control-theoretic models of driving (e.g., Donges, 1978; Godthelp, 1986; Hess & Modjtahedzadeh, 1990) often incorporate rigorous vehicle dynamic models for basic control, but they sometimes ignore issues of whether or how model inputs can be readily perceived from the external environment and how drivers interleave multiple tasks to produce satisficing rather than optimal performance (see Boer, 1999, for an excellent discussion). Machine-learning models of vehicle control (e.g., Pomerleau & Jochem, 1996) focus on the task of lane keeping and implicitly account for vehicle dynamics in their trained systems; however, because of their emphasis on application, these models reveal little about driver cognition and behavior and thus do not generalize well to phenomena such as cognitive bottlenecks or distractions from secondary tasks. Nevertheless, these categories of models are not diametrically opposed to integrated driver models; in fact, the many successes of these models have demonstrated the importance of rigorous modeling efforts for both theoretical understanding of driver behavior and practical application of these theories in real-world system development.

Integrated Driver Modeling in the ACT-R Cognitive Architecture

The approach to integrated driver modeling explored here centers on the development of driver models in the framework of a *cognitive architecture.* A cognitive architecture is a general framework for specifying computational behavioral models of human cognitive performance (e.g., Anderson & Lebiere, 1998; Just & Carpenter, 1992; Laird, Newell, & Rosenbloom, 1987; Liu, 1996; Meyer & Kieras, 1997; Newell, 1990). The architecture embodies both the abilities and constraints of the human system – for instance, abilities such as memory storage and recall, learning, perception, and motor action; and constraints such as memory decay, foveal versus peripheral visual encoding, and limited motor performance. As such, a cognitive architecture helps to ensure that cognitive models developed in the framework are rigorous and psychologically valid, thus abiding by all the limitations of the human system. The chosen framework for this driver model is the ACT-R (Adaptive Control of Thought-Rational) cognitive architecture (Anderson et al., 2004; see also Anderson & Lebiere, 1998), a hybrid architecture based on chunks of declarative knowledge and condition-action production rules that operate on these chunks. The particular advantages of the ACT-R architecture will be discussed in the course of the paper. However, it is important to note that many of the central themes and ideas remain the same for models developed in any cognitive architecture – for instance, Aasman's (1995) driver model developed in the Soar architecture (Laird et al., 1987; Newell, 1990).

Integrated driver models developed in a cognitive architecture such as ACT-R are especially well suited to addressing all three components of the ETA triad. Cognitive architectures have demonstrated the ability to model tasks ranging from basic laboratory tasks (e.g., serial recall: Anderson & Matessa, 1997) to higher level cognition and decision making in complex dynamic tasks (e.g., fighter piloting: Jones et al., 1999; air traffic control: F. J. Lee & Anderson, 2001; human-computer interfaces: Ritter, Baxter, Jones, & Young, 2000). Architectural models typically interact with a simulated environment identical to, or almost identical to, the environment used by human participants, and thus the models must abide by the same input/output limitations and environment dynamics as human participants. In doing so, architectural models represent and account for both the internal workings of human cognition and the external manifestations of cognition through perceptual processes and motor behavior. All these features make cognitive architectures extremely amenable to modeling many of the most important aspects of driver behavior.

This paper represents the culmination of several years of work focused on developing an integrated driver model in the ACT-R cognitive architecture. The initial prototype model (Salvucci, Boer, & Liu, 2001) served as the first proof-of-concept of the feasibility and power of the cognitive architecture approach. Since that time, several significant developments have helped to shape and formalize the initial model. First, the basic elements of its two-level control model have been validated outside the context of the architecture as a simple, stand-alone computational model, and it has been demonstrated how the model can account for curve-negotiation and lane-changing behavior results across several empirical studies (Salvucci & Gray, 2004). Second, a number of lessons have been learned from applications of the initial model to predicting driver distraction both from cell phone dialing (Salvucci, 2001b; Salvucci & Macuga, 2002) and from cognitive tasks (Salvucci, 2002). Third, the ACT-R cognitive architecture itself has undergone a major evolution, resulting in a novel "buffer"-centered architecture that fits nicely with many of the evolving concepts in the driver model (see Anderson et al., 2004); the new architecture posits buffers of information through which production rules can communicate with both declarative memory and the external environment (through perceptual and motor modules). (The initial model was developed in Version 4.0 of the ACT-R architecture [Anderson & Lebiere, 1998]; the model described here represents a complete reworking, conceptually and practically, in Version 5.0 of the architecture [Anderson et al., 2004], taking full advantage of the new version's buffer-centered architecture.) All these developments have combined to produce the new ACT-R integrated driver model presented here.

THE ACT-R INTEGRATED DRIVER MODEL

The ACT-R driver model is a computational model of driver behavior implemented in the ACT-R cognitive architecture (Anderson et al., 2004; see also Anderson & Lebiere, 1998). The current driver model is intended to represent behavior for a particular task and artifact, namely that of driving a standard midsize vehicle on a multilane highway with moderate traffic. Highway driving is a common driving scenario that accounts for a large percentage of vehicle miles – for instance, highway driving accounted for 72% of vehicle miles in the United States in 2000 (Federal Highway Administration, 2001). It is hoped that the model will be extended to other common contexts (e.g., city driving) and possibly other vehicles (e.g., trucks and buses) in the near future, but the common highway driving context has enabled exploration of a number of interesting specific processes (e.g., lane changing) and general processes (e.g., monitoring for situation awareness) that should generalize well to other tasks and artifacts.

The driver model includes three main components that derive from and emphasize specific aspects of Michon's (1985) hierarchical control structure. The *control* component, analogous to Michon's (1985) operational level, manages all aspects of perception of the external world and mapping of specific perceptual variables to manipulation of vehicle controls (i.e., steering, acceleration, braking). The *monitoring* component, part of Michon's (1985) tactical level, maintains awareness of the current situation by periodically perceiving and encoding the surrounding environment. The *decision making* component, also part of Michon's (1985) tactical level, handles tactical decisions for individual maneuvers (e.g., lane changes) based on knowledge of the current environment. These three components are implemented in the ACT-R architecture to take advantage of the architecture's built-in features and human-like limitations that result in a more psychologically plausible model of driver behavior. I will first provide a brief description of the ACT-R architecture and then describe each of the core components of the driver model as implemented in this architecture.

The ACT-R Cognitive Architecture

The ACT-R cognitive architecture, like all cognitive architectures, is simultaneously a rigorous theory of human cognition and a working framework in which to build computational models of human behavior. ACT-R posits two separate but interacting knowledge stores. The first type of knowledge, *declarative knowledge,* is made up of *chunks,* or small logical units, of symbolic information. Declarative chunks can encode simple facts (e.g., Philadelphia is in Pennsylvania), current goals (e.g., try to change lanes), and even ephemeral situational information (e.g., there is a car to my left). Chunks are also associated with "subsymbolic" parameters that encode continuous-valued properties of each chunk – for instance, chunk *activation* as a representation of the relative ease with which the chunk can be recalled. In addition, learning mechanisms can affect these subsymbolic parameters – for instance, chunk activation decays over time to mimic forgetting but also increases as the chunk is practiced and recalled more often.

The second type of knowledge, *procedural knowledge,* is made up of *production rules* representing procedural skills that manipulate declarative knowledge as well as the environment. Each production rule is essentially a condition-action rule that generates the specified actions if the specified conditions are satisfied. Rule conditions typically specify the current goal and match other current values in the architectural buffers (e.g., in the visual or retrieval buffers). When all conditions match and the rule "fires," rule actions can add to or alter declarative memory, set a new current goal, and/or issue perceptual or motor commands (e.g., find a particular visual object or type a key). Like declarative chunks, production rules are associated with subsymbolic parameters that affect their behavior – for example, "conflict resolution" parameters that determine which rule fires given several possible matches. Also like declarative parameters, parameters for procedural knowledge can adapt over time given different choices and situations.

One critical advantage of a cognitive architecture for driving (or any other complex task, for that matter) is the incorporation of built-in features that mimic human-like abilities. ACT-R in particular has built-in perceptual and motor mechanisms (Byrne, 2001; Byrne & Anderson, 2001) that allow ACT-R models to interact with external simulations. These mechanisms can perceive changes in the environment and can manipulate the environment using programmed modules that interface with the actual simulation. Also, ACT-R has the ability to perform some processes in parallel such that, for example, the perceptual module can look at a new item while the motor module performs a physical movement. Of note, this allows ACT-R models to run in simulation – often in real time, or at least generating data on the same time scale as that of human participants – and predict exactly those data and measures that are collected from drivers (steering wheel, throttle, and brake positions; turn signals; eye movements; etc.). The prediction of real-world measures greatly facilitates model validation through direct comparison with human data.

At the same time, ACT-R places certain limitations and constraints on models that mimic the constraints of the human system. One of the most important constraints for the driver model is that although perceptual and motor processes can run in parallel with cognition, the cognitive processor itself is serial and, in essence, can "think" only one thing at a time. The cognitive processor is responsible for collecting all information from perceptual modules and issuing all motor commands, and thus it serves as the central bottleneck for behavior. This fact is critical for applications such as predicting driver distraction: When the driver model attempts to perform a secondary task such as dialing a phone, the cognitive processor must interleave the secondary task with the primary driving task, thus potentially leading to reduced performance. In fact, the model can account not only for interference from primarily perception- and motor-oriented secondary tasks but also primarily cognitive secondary tasks, because the latter would interfere with the cognitive processor's ability to manage the primary task.

Model Specification

As mentioned, the ACT-R driver model has three primary components: control, monitoring, and decision making. The three components are integrated to run in ACT-R's serial cognitive processor as a tight loop of small cognitive (and related) operations. The entire model is implemented as an ACT-R production system including relevant procedural and declarative knowledge. This section describes each component, the integration of the components into a working implementation, and finally estimation of model parameters and integration with the simulated driving environment.

Control. The control component of the driver model manages all perception of lower level visual cues and manipulation of vehicle controls for lateral control (i.e., steering) and longitudinal control (i.e., acceleration and braking). Lateral control centers on a new steering model (Salvucci & Gray, 2004) that utilizes "two-level" control based on the perception of two salient visual points (see Donges, 1978; Land & Horwood, 1995). First, the *near point* represents the vehicle's current lane position, used to judge how close the vehicle is to the center of the roadway. The near point is characterized as a point in the center of the near lane visible in front of the vehicle, set at a distance of 10 m from the vehicle's center. Second, the *far point* indicates the curvature of the upcoming roadway, used to judge what the driver should execute to anticipate the upcoming curvature and remain in the current lane. The far point is characterized as one of three targets: (a) the vanishing point of a straight roadway, up to a maximum distance equivalent to 2 s of time headway; (b) the tangent point of an upcoming curve (Land & Lee, 1994); or (c) the lead vehicle – that is, the vehicle immediately in front of the driver's vehicle. Together, the near and far points provide complementary pieces of information that allow for adjustment to the lane center at the current position (using the near point) and for predictive compensation at a near-future position (using the far point). Figure 1 illustrates the near and far points for a straight road segment (with vanishing point), a curved road segment (with tangent point), and a road segment with a lead vehicle.

In the model, lateral control requires perception of these salient points and subsequent motor execution of control. The model first moves its visual attention to the near point, then to the far point, noting the visual angles θ_{near} and θ_{far} of the two points, respectively. Also, assuming the far

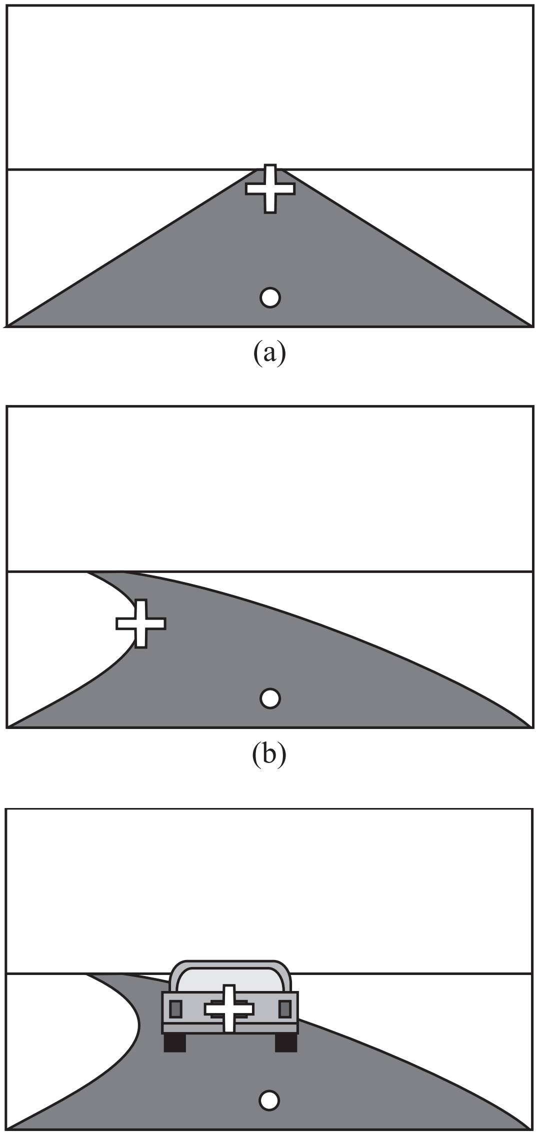

Figure 1. Illustration of near point (circle) and far point (cross) for (a) straight road segments, (b) curved road segments, and (c) any road with lead car. (From Salvucci & Gray, 2004. Reproduced with permission from Pion Limited, London.)

point has remained the same from the last cycle of the main control loop (e.g., it hasn't changed from a vanishing point to a tangent point), the model calculates differences from the last cycle, namely $\Delta\theta_{near}$, $\Delta\theta_{far}$, and Δt (given the time of the last cycle). Finally, the model uses these quantities to adjust the vehicle's steering angle by some incremental value. The model strives for the most simple but effective method of control and, to this end, utilizes a simple steering control law that relies on perceived visual direction to the near and far points, as described by Salvucci and Gray (2004). The control law for steering angle φ can be expressed in its discrete form as

$$\Delta\varphi = k_{far}\Delta\theta_{far} + k_{near}\Delta\theta_{near} + k_I\theta_{near}\Delta t \qquad (1)$$

(See Salvucci & Gray, 2004, for derivation from the continuous form.) The control law essentially attempts to impose three constraints: a steady far point ($\Delta\theta_{far} = 0$), a steady near point ($\Delta\theta_{near} = 0$), and a near point at the center of the lane ($\theta_{near} = 0$). The three constants vary the weights associated with each of these constraints. (An enhanced version of the model could adapt these constants as a function of speed to compensate for corresponding changes in vehicle response; however, the constants in the current model have been found to work well for typical highway speeds and thus suffice for this paper's primary focus on highway driving.) Salvucci and Gray (2004) have shown how this control law, as a stand-alone model outside the context of a cognitive architecture, accounts for driver steering profiles during curve negotiation (e.g., Land & Horwood, 1995) as well as corrective steering (Hildreth et al., 2000).

Longitudinal control (i.e., speed control) embodies a very similar process. The model encodes the position of the lead car and derives the time headway thw_{car} to this vehicle. Again, it computes differences from the last instantiation of control, deriving Δthw_{car} along with the previously mentioned Δt. These two values then result in an updated value for acceleration ψ. For simplicity, the model utilizes a longitudinal control law very similar to the lateral control law, namely,

$$\Delta\psi = k_{car}\Delta thw_{car} + k_{follow}(thw_{car} - thw_{follow})\Delta t. \qquad (2)$$

The acceleration equation attempts to impose two constraints: a steady time headway ($\Delta thw_{car} = 0$) and a time headway approximately equal to a desired time headway for following a lead vehicle ($thw_{car} = thw_{follow}$). Again, the two constants determine the weights of the two constraints. The acceleration value ψ actually manipulates two controls:

A positive value translates to depression of the accelerator (throttle), and a negative value translates to depression of the brake pedal, with values from 0 to 1 representing *no depression* to *full depression,* respectively.

The ACT-R architecture includes several constraints that limit how such control can be employed. Most significantly, the serial cognitive processor cannot instantiate a continuous function and instead must update control inputs discretely, as indicated previously. Also, the architecture's 50-ms "cycle time" – the time needed to fire a production rule – dictates the minimum time between updates of control. In the production-system implementation, the minimum time for execution of this control loop is 150 ms, representing the firing of three production rules (in essence, one rule to encode the near point, one to encode the far point, and one to issue the motor commands, as detailed in an upcoming section). Perception in the architecture happens through shifts of visual attention and resulting eye movements predicted by ACT-R's EMMA (Eye Movements and Movement of Attention) module (Salvucci, 2001a); thus, perception of a salient visual point requires not only time to fire the relevant production rule or rules but also time to encode the visual object into declarative memory. For motor processes, ACT-R has no built-in motor modules for steering and pedal movements, so these have been added to the architecture for the driving domain. For most motor movements, drivers do not need to operate near peak performance (e.g., turning the steering wheel), and thus the model imposes only one significant constraint: A foot movement between the accelerator and brake requires 200 ms of preparation time (based on preparation of four motor parameters at 50 ms each; see Byrne & Anderson, 2001) and 500 ms of execution time (an approximation based on a range of 420–630 ms found by J. D. Lee, McGehee, Brown, & Reyes, 2002, for successful braking situations).

In a highway environment, the majority of lower level control arises in the context of lane keeping and curve negotiation. However, one other common subtask that frequently occurs in highway driving is that of lane changing. Interestingly, the basic equations stated previously, particularly that for steering control, generalize in a completely straightforward manner to lane changing: When the driver enacts a lane change, he or she simply begins to use the near and far points of the destination lane rather than the current lane (Salvucci & Liu, 2002). This change has the immediate effect of creating a large θ_{near}, thus initiating a large steering motion in the direction of the destination lane. However, this effect is tempered by the compensation of the $\Delta\theta_{near}$ and $\Delta\theta_{far}$ terms in the equation, which attempt to maintain a steady transition and prevent the vehicle from swinging wildly into the other lane. This compensation limits the maximum rate at which the driver will steer. Salvucci and Gray's (2004) analysis showed how this idea nicely accounts for steering profiles during lane changing (Salvucci & Liu, 2002). In addition, some drivers may wish to steer even less rigorously for safety precautions; thus, the model uses a revised lateral control law,

$$\Delta\varphi = k_{far}\Delta\theta_{far} + k_{near}\Delta\theta_{near} + k_I \min(\theta_{near}, \theta_{nmax})\Delta t, \quad (3)$$

incorporating a parameter θ_{nmax} that limits the contribution of θ_{near} to changes in steering angle.

Monitoring. The monitoring component of the driver model handles the continual maintenance of situation awareness. For this model in the highway driving environment, situation awareness centers critically on awareness of the position of other vehicles around the driver's vehicle. Monitoring is currently based on a random-sampling model that checks, with some probability $p_{monitor}$, one of four areas – namely, either the left or right lane, and either forward or backward (i.e., in the mirror) – with equal likelihood. When the model decides to monitor a particular lane and direction, it moves visual attention to that area and determines whether there is any vehicle present. If so, the model notes the vehicle's current lane, direction, and distance in ACT-R's declarative memory. Thus, declarative knowledge continually maintains the awareness of surrounding vehicles. The model could, of course, be extended in a straightforward way to note other aspects of the surrounding environment (on- and off-ramps, signs, billboards, etc.), but these are not currently included.

The use of ACT-R's declarative memory for encoding of the current environment provides immediate predictions about potential driver errors: Because of memory decay mechanisms built into the architecture, the chunks that encode vehicle position and distance decay rapidly and may be forgotten when not refreshed often, thus providing predictions (albeit yet-untested predictions) about mistaken estimation of surroundings and potentially hazardous lane changes. For instance, if the model were to use a modified strategy relying on memory instead of visual checking before a lane change, it could forget about a vehicle observed in the "blind spot" or even mistakenly retrieve an obsolete memory of a clear blind spot. The model could thus begin to address drivers' differential use of knowledge "in the world" versus "in the head" (Gray & Fu, 2004).

Decision making. The decision-making component of the driver model uses the information gathered during control and monitoring to determine whether any tactical decisions must be made. In the highway environment, the most common decision-making opportunity arises in the determination of whether and when to execute a lane change. The decision of *whether* to change lanes depends on the driver's current lane, given that drivers (in the United States) attempt to stay in the right lane during normal driving and use the left lane for passing only (ideally). If the driver's vehicle is in the right lane, the model checks the current time headway to the lead vehicle (if any); if the lead car time headway drops below a desired time headway thw_{pass}, the model decides to change lanes to pass the vehicle. If the driver vehicle is in the left lane, the model checks instead simply for the presence of a lead vehicle. If there is a lead vehicle, the model remains in the left lane (because this vehicle is also passing other vehicles); otherwise it decides to change lanes to return to the right lane.

Even after the model decides to change lanes, it must still determine *when* to initiate the maneuver at an appropriate, safe time. First, the model attempts to recall from declarative knowledge the nearest vehicle to the driver's vehicle in the other lane; if such a vehicle can be recalled and that vehicle is within a safe distance d_{safe} of the driver's vehicle, the lane change is aborted. In this way, the model can avoid vehicles in its blind spot even when it cannot directly view the vehicle; however, as mentioned, the model may also forget about this vehicle if it remains in the blind spot for a longer period of time. Then, the model monitors the destination lane both front and back to check for the presence of another vehicle and updates its mental model of the environment. If this monitoring does not observe a vehicle nearer than the safe distance d_{safe}, the model initiates the execution of the lane change; otherwise, the model immediately aborts the maneuver. As described earlier, the execution of the lane change corresponds simply to the use of the near and far points of the destination lane rather than the current lane. The lane-change maneuver continues until the vehicle reaches the center of the destination lane.

Component integration and multitasking. The integration of the three core model components of control, monitoring, and decision making requires some method of multitasking in performing each of the three respective subtasks. Because of its implementation in the ACT-R cognitive architecture, the model is constrained to a serial cognitive processor that cannot perform all three tasks at once but rather must interleave the tasks serially. To this end, the model contains a single, tight main loop that performs an incremental step of each subtask: For control, this incremental step involves a single instantiation of the lateral and longitudinal control laws; for monitoring, the step involves (with probability $p_{monitor}$) the check of a single vehicle in a chosen lane and direction; and for decision making, the step involves a single decision as to whether or not to initiate a lane change while lane keeping or to end the lane change while lane changing. In addition, work with driver distraction (e.g., Salvucci, 2001b) suggests that drivers deviate from the primary control task only when they judge themselves to be "safe" given the current conditions, including current lane position and lateral velocity (i.e., velocity side to side). Thus, the model includes the requirement that before switching from control to monitoring or decision making, the vehicle must be within lane-position and lateral-velocity bounds; specifically, the model checks lateral position with the constraint that $\theta_{near} < \theta_{stable}$ and checks lateral velocity with the constraint that $\dot{\theta}_{near} < \dot{\theta}_{stable}$ (where $\dot{\theta}$ becomes $\Delta\theta_{near}/\Delta t$ in the discrete case).

The serialization of the three model subtasks is critical to the model's ability to predict realistic driver performance that incorporates human-like limitations. Because even this tight loop requires some time to execute, the driver model does not produce perfect performance even on a straight roadway with no other secondary tasks, just as human drivers would also not produce perfect performance. In addition, if the driver attempts to perform a secondary task along with the primary driving task, the secondary task will further interrupt the main loop and may potentially worsen performance; this observation forms the basic premise of my approach to predicting driver distraction, mentioned in the final discussion. Ideally, arbitrary secondary tasks would be added to the primary component tasks to allow a rigorous theory of multitasking to arbitrate the execution of the tasks using some notion of prioritized task scheduling. However, currently, the default ACT-R cognitive architecture does not incorporate such a theory, although initial efforts have begun to specify such a theory (Salvucci, 2005). Nevertheless, the driver model's somewhat naive account of task scheduling has proven adequate for the purposes of validating the basic model and demonstrating its feasibility for integration into practical applications such as predicting driver distraction.

Production-system implementation. A schematic of the driver model's implementation as an ACT-R production system is shown in Figure 2. In the figure, the title of each large box indicates the goal (e.g., drive), the names below indicate production rules for that goal (e.g., control-attend-near), the arrows indicate flow of control between goals, and the asterisks (*) indicate where a goal completes and returns to its parent goal. The top-level drive goal includes production rules to initiate a control update (control-attend-near), to initiate lane monitoring (monitor-lane), and to initiate or complete a lane change (try-lane-change and end-lane-change). These rules create subgoals to perform the associated subtasks of control, monitoring, and decision making.

For control, after the drive goal attends the near point (control-attend-near), the control goal first processes (i.e., notes the position of) the near point and attends the far point (process-near-attend-far). Then, it processes the far point using all the stored information to update the steering and acceleration inputs (process-car in the presence of a lead car, process-far otherwise). As noted earlier, this central loop for updating control thus requires the firing of three production rules, or roughly 150 ms. Because control updates require the "old" control goal to calculate changes from the last update, a special rule fires in the absence of the old goal (process-without-old) and immediately performs a second update. A final rule (control-failure) handles all other (anomalous) cases and also immediately restarts the update.

For monitoring, the model fires one of four drive productions (monitor-lane-{l,r,lm,rm}) with probability p_{monitor} to monitor one of the four possible visible lane areas, namely the left or right lane ahead or the left or right lane behind (in the mirror). The monitor goal simply processes the presence

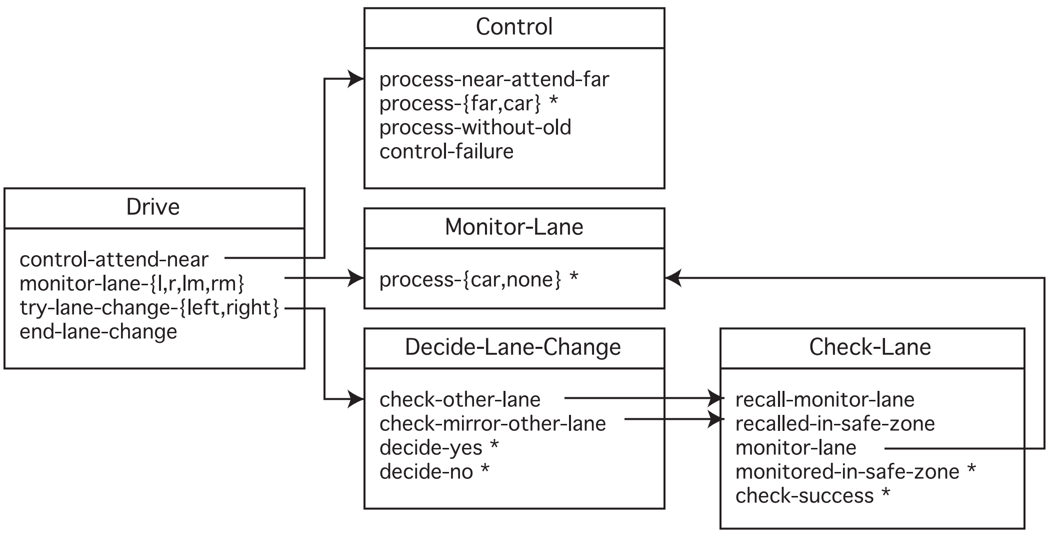

Figure 2. Schematic of the ACT-R production system. Box titles indicate goal types; names below indicate production rules; arrows indicate flow of control; and asterisks mark rules that return control to the parent goal.

(process-car) or absence (process-none) of a car in this area. Notably, the goal chunk remains in declarative memory and thus can be later retrieved for the purposes of decision making and/or general situation awareness.

For decision making, the model fires one of two drive productions (try-lane-change-{left,right}) when a lane change is desired, each of which initiates the decide-lane-change goal. This goal in turn initiates a check-lane subgoal to check for vehicles in the other lane forward (check-other-lane) and backward (check-mirror-other-lane). As mentioned earlier, check-lane first attempts to recall a vehicle in that lane (i.e., recall a monitor-lane chunk for that area) and then explicitly monitors the lane with a call to monitor-lane. If any vehicles are found within the safe distance d_{safe} (recalled-in-safe-zone and monitored-in-safe-zone), check-lane returns a failure and decide-lane-change chooses not to change lanes (decide-no); otherwise decide-lane-change chooses to change lanes (decide-yes) and updates this fact in the parent drive goal. In this latter case, the lane change continues until the vehicle reaches the destination lane and the model reverts back to normal lane keeping (end-lane-change).

Parameter estimation. Given the complexity of the driving task, it is reasonable that any model of driver behavior involves a number of domain-specific parameters, some of which may vary among individual drivers. At this stage of development for the ACT-R driver model, not yet having tackled the problem of individual differences in a rigorous manner, a single set of estimated values is used for all model parameters, with an eye for future individual differences in these parameters. The parameter estimation process involved two stages for two separate sets of parameters. First, the bulk of the parameters were simply set to reasonable values based on informal observation of the model driving as well as approximations derived from available empirical literature. These parameters and their values are listed in Table 1 as "informal." The remaining parameters represented those most critical to the resulting driver behavior for the various detailed measures that will be described in the Model Validation section. These parameters were estimated by setting them to reasonable values, observing the resulting qualitative and quantitative fits given these values, and revising the values accordingly; most importantly, the constant weights in the steering control required rigorous estimation because these weights most drastically affected the resulting behavior. These parameters are listed in Table 1 as "estimated."

Along with the domain-specific parameters, the model also inherits through the ACT-R architecture a set of domain-independent parameters. Fortunately, the ACT-R community has converged on default settings for these parameters through previous modeling efforts, and the vast majority of these defaults are utilized in the driver model. The only changed parameters were modified to represent the fact that the procedural skills and declarative chunks that constitute the model are well learned: Because the model assumes an experienced driver well past the learning stage, it includes parameter values that reflect this assumption – namely, declarative chunks are given 100 "references" and a creation time of –1000 s, which indicates that their activation is stable and relatively fixed (i.e., will not be easily forgotten). (ACT-R's base-level learning was activated with optimized learning enabled and a default decay rate of 0.5.) Further explanation of these parameters can be found in Anderson and Lebiere (1998).

Integration with simulation environment. One of the primary goals of the ACT-R driver model is to facilitate rigorous evaluation and validation by having the model drive in the same environment as human drivers. The environment used for the driver model and the validation participants is a

TABLE 1: ACT-R Driver Model Parameters, Values, and Method of Determining Values (Informal Observation or Estimation for Best Model Fit)

Parameter	Value	Method
k_{far}	16.0	Estimated
k_{near}	4.0	Estimated
k_I	3.0	Estimated
θ_{nmax}	0.07 rad	Estimated
k_{car}	3.0	Informal
k_{follow}	1.0	Informal
thw_{follow}	1.0 s	Informal
thw_{pass}	2.0 s	Informal
$p_{monitor}$	.20	Informal
d_{safe}	40 m	Informal
θ_{stable}	0.07 rad (≈1/4 lane)	Informal
$\dot{\theta}_{stable}$	0.035 rad/s (≈1/8 lane/s)	Informal

multilane highway environment integrated with a fixed-base driving simulator. (The original driving simulator was located at Nissan Cambridge Basic Research in Cambridge, Massachusetts; the simulator software for the current model has evolved from this simulator and is now being developed further at Drexel University.) The environment incorporates a model of vehicle dynamics as well as automated vehicles that simulate realistic traffic around the driver. When a human driver navigates this environment in the driving simulator, the system generates a lengthy protocol that includes vehicle control data (e.g., steering wheel and throttle position), driver gaze (eye movement) data, and environment information (e.g., location of all vehicles). The ACT-R driver model has its own fully equivalent simulation that differs only in that the human-viewable graphics have been removed. When the driver model navigates this environment, the system generates exactly the same protocol as the original simulation system – for instance, the model turns a virtual steering wheel to control lateral position and also focuses its virtual eyes on various components of the visual field (including the control points and other vehicles). In fact, the model protocol can be replayed in the original simulator with no modification. Thus, not only can the qualitative nature of the model's behavior be examined, but also direct quantitative comparisons of the model's and humans' behavior can be generated, greatly facilitating the quantitative evaluation and validation of the model with empirical data.

Comparison with Related Driver Models

The ACT-R driver model is most closely related to three categories of previously developed driver models. First, many early models of driver steering and lane keeping focused primarily on control-theoretic descriptions of steering control (e.g., Carson & Wierwille, 1978; Donges, 1978; Godthelp, 1986; Hess & Modjtahedzadeh, 1990; McRuer, Allen, Weir, & Klein, 1977; van Winsum & Godthelp, 1996; Weir & McRuer, 1973). Although they have captured some aspects of driver lane keeping and curve negotiation, these models, as pointed out by Boer (1999), may be unreasonable in the sense that they rely on inputs not readily perceivable from the external environment – for instance, vehicle yaw, or time to lane crossing. In addition, some models incorporate additional parameters (e.g., Weir & McRuer, 1973) and/or smoothing techniques (Carson & Wierwille, 1978) to account for delays in driver perception and response. In contrast, the ACT-R driver model utilizes readily perceivable inputs – namely the visual direction to the near and far points – and integrates these into a fuller theory of cognition, perception, and motor action to provide a more psychologically plausible model that incorporates the constraints of the human system.

A related class of models arises in recent perception-action models of control (e.g., Fajen & Warren, 2003; Rushton et al., 1998; Salvucci & Gray, 2004; Wilkie & Wann, 2003). Whether by the use of visual direction (e.g., Salvucci & Gray, 2004) or optic flow (e.g., Wilkie & Wann, 2003), these models follow more closely the types of perceptual constraints inherent in the human system. However, unlike the earlier control-theoretic models mentioned previously, these models have not typically incorporated rigorous models or delays of vehicle dynamics and human physical movement, which certainly contribute to the dynamic process of steering and control.

The models most closely related to the proposed model are integrated driver models that attempt to unify many aspects of the driving task. The model of Levison and Cramer (1995) integrates driver and vehicle models to predict performance measures for typical driving scenarios. The models of Aasman (1995) and Tsimhoni and Liu (2003) instantiate models of driver behavior within the context of cognitive architectures similar to ACT-R, namely Soar (Laird et al., 1987; Newell, 1990) and Queueing Network-Model Human Processor (QN-MHP; Liu, 1996), respectively. These last two models in particular were developed in very much the same spirit as was the ACT-R model, with complementary goals of exploring driver behavior and, simultaneously, exploring the generality and applicability of the cognitive architecture. Currently, the ACT-R driver model has been applied in a wider range of domains than have the Soar and QN-MHP models (see, e.g., Salvucci, 2001b, 2002, 2005; Salvucci, Chavez, & Lee, 2004; Salvucci & Macuga, 2002), which have focused on intersection approach and lane keeping, respectively. Nevertheless, because of their implementation in general cognitive architectures, both models show promise similar to that of the ACT-R model in being applied to a variety of driving-related phenomena.

MODEL VALIDATION

Given that the goal of the ACT-R driver model is to accurately represent driver behavior, the model's behavior requires validation and comparison with real human behavior. This is no small endeavor: Just as no single method, measure, or metric will suffice for understanding human driver behavior, no single one will suffice to validate that the model indeed corresponds well to human drivers. Nevertheless, one can validate the most critical parts of a driver model by focusing on key scenarios and analyzing the most important observable data involved in these scenarios. To this end, how the ACT-R model fits several aspects of driver data will now be examined from two common scenarios in normal highway driving: lane keeping/curve negotiation and lane changing. For these scenarios, the examination focuses on three important measures of behavior – steering angle, lateral position, and eye movements (as a surrogate for the locus of visual attention) – in the form of aggregate results and time-course profiles.

Human and Model Data

The computational nature of the ACT-R driver model, combined with its ability to interact with the same simulation environment that human drivers use, greatly facilitates the collection and comparison of human and driver data. Human data from 11 drivers were collected in the original study (Salvucci et al., 2001) conducted in the previously mentioned driving simulator. Model data were collected by running five 10-min model simulations in the same conditions and same environment as the original experiment; note that the model, like a human driver, produces variability in behavior, and thus several simulation runs are desirable to achieve more stable results. The following analysis includes a total of 311 min (548 km) of driving data for human participants and 50 min (94.9 km) of driving data for the model simulations. Because the human and model simulation protocols are identical in form, each set is analyzed in the same manner so as to generate directly comparable measures of driver behavior and performance.

Lane Keeping and Curve Negotiation

The most common component of highway driving (and most driving) is lane keeping, or simply steering down the center (or near the center) of a lane. Under informal observation, the ACT-R driver model can clearly negotiate down the current lane and maintain a reasonably central lane position. However, drivers' lane-keeping behavior can be quantified more formally in several ways. First, perhaps the most interesting aspect of the lane-keeping task involves how drivers negotiate curves – that is, how they steer into and through a curved section of roadway. Figure 3 shows time-course plots of drivers' steering profiles and lateral positions while negotiating right and left curves. The construction of these plots requires some explanation. First, the sections of the protocols that corresponded with entering and exiting the curve, as noted by the roadway at the vehicle's current position, were extracted; the roadway used in data collection had curved segments of different lengths and curvatures, but all segments had a constant curvature throughout the segment. Next, each segment was divided into 10 equal-sized units, and these units were extended out to before and after the actual curved road segment. Finally, all protocols were aggregated by averaging together all the values within each segment.

Figure 3 includes the time-course plots for the (a) human data and (b) model simulations for both steering wheel angle ($R^2 = .98$, $RMSE = .06$, comparing all human data points with all model data points in the graph) and lane position ($R^2 = .55$, $RMSE = .06$). The human drivers' steering profiles show that they began to change steering angle before approaching the curve, flattened out to a fairly constant angle, and then began to steer back to center before the end of the curve and gradually returned to center. The model showed this same trend, particularly in that it began steering off center before the curve started and again before it ended. The model reached the peak steering deflection slightly more sharply than did the human drivers at the start of the curve, but it exhibited the same smooth steering toward the end of and after the curve. The lateral position profiles show that both human drivers and the model tended slightly to steer toward the inner part of the road during a curve (i.e., toward the left for left curves and toward the right for right curves). This result

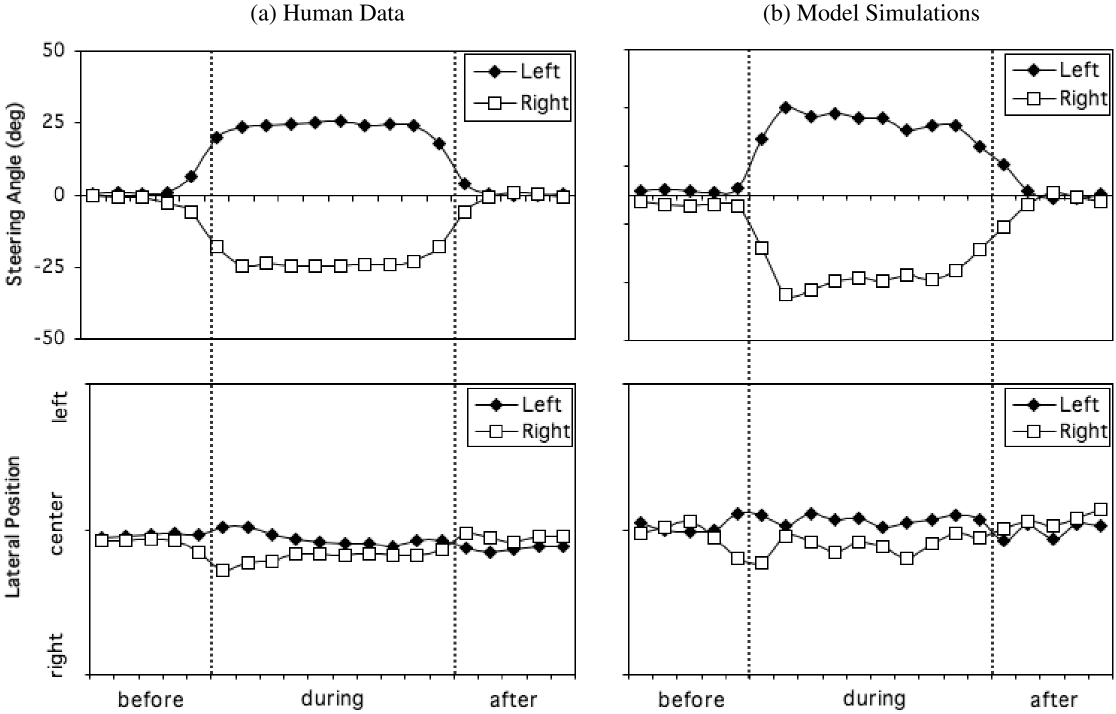

Figure 3. Curve negotiation profiles for steering angle and lateral position for (a) human data and (b) model simulations. The vertical dotted lines indicate the start (leftmost line) and end (rightmost line) of the curve.

was especially interesting given that the model had not been explicitly designed to approach the inner part of the curve. This behavior arises because the model observes the tangent point in the distance and attempts to keep it stable as the vehicle approaches, thus drawing the vehicle slightly toward the tangent point as it approaches the curve; thus, the model's tendency toward the inner part of the curve is an emergent prediction of the model's simple control law that nicely reflects human driver performance.

Another measure by which lane-keeping behavior can be examined is the distribution of driver gaze to various parts of the visual environment – that is, the distribution of where drivers look as they drive. Although drivers are generally expected to maintain gaze in front of their vehicle at a far point or lead car, occasional gazes are also expected to other areas of the environment for purposes of monitoring and situation awareness (and potentially other reasons, even boredom). Figure 4 shows the proportion gaze time for the human and model data (R^2 = .93, *RMSE* = .03), expressed as the proportion of time spent looking at one of several salient visual areas: the current lane's near point, vanishing point, tangent point, lead car, or other cars (in front); the same areas for the other lane; and finally the rear-view mirror, oncoming vehicles, and unclassified gazes ("none").

Perhaps not surprisingly, the human drivers spent the most time looking at the far area in their own lane, namely the lead car and the two types of far points (vanishing point and tangent point). The model also predicted this majority of gaze time on far objects, albeit with a slight overprediction for the far areas in the current lane (attributable in part, I believe, to humans' significant "none" gazes, which could not be classified). The model produced an overall nice fit with the other data points as well, particularly in that (a) both humans and model looked at the lead car in the other lane for the same amount of time, (b) both humans and model looked at the mirror roughly 5% of the time, and (c) both humans and model rarely looked at the near point for either lane. This last point is another emergent prediction of the driver model, arising from the fact that the model gathers very little information from the near point (only the visual angle on the x axis) and thus can typically acquire this information peripherally without

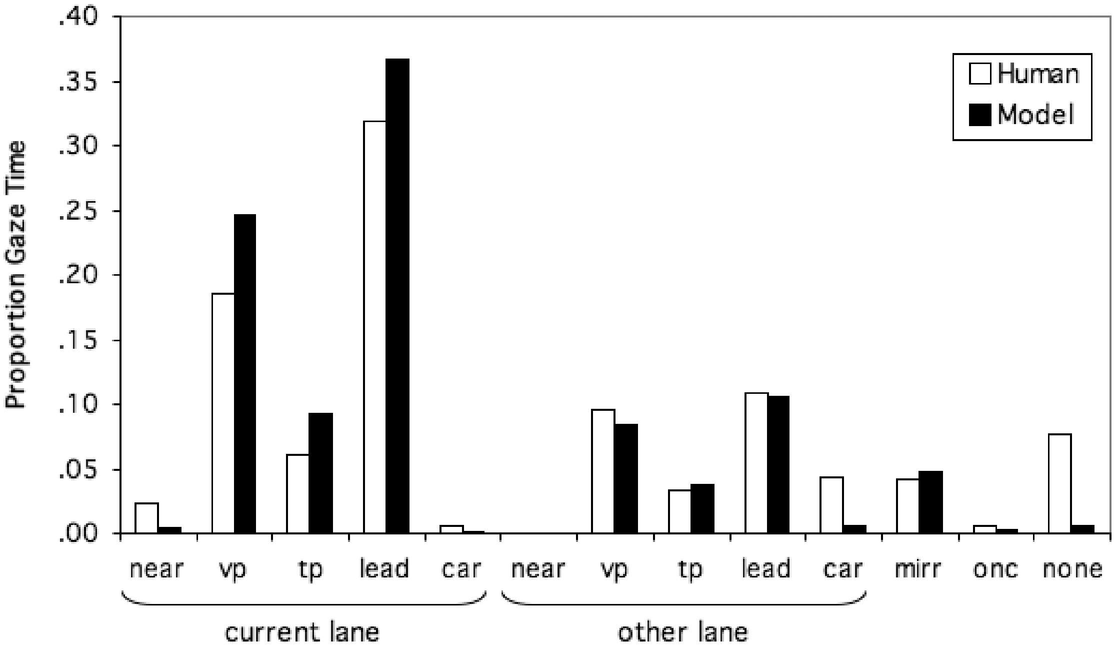

Figure 4. Lane-keeping proportion gaze time for human and model data. Near = lane near point, vp = vanishing point, tp = tangent point, lead = lead vehicle, car = other vehicle, mirr = rear-view mirror, onc = oncoming vehicle, none = none of the above.

actually fixating the near point directly (see Salvucci, 2001a, for further information about how eye movements are associated with shifts in visual attention).

Lane Changing

The other common scenario in highway driving, complementary to that of lane keeping, is that of lane changing. As was done for lane keeping, the three basic measures of lane position, steering wheel angle, and proportion gaze time can be examined to elucidate and compare the behavior of the driver model and human drivers. Figure 5 shows aggregated time-course profiles for all lane changes in the human and model protocols. These plots were generated in the same manner as those for curve negotiation, except that the boundaries of the protocol segments were dictated not by roadway curvatures but by verbal protocols: both human drivers and the model produced a verbal utterance when (a) they formed the intention to change lanes and (b) they completed this goal and reverted back to lane keeping. As in the plots for lane keeping and curve negotiation, the protocol segments were broken into 10 units, extended before and after the actual lane-change segment, and averaged together in each segment to form aggregate plots; to avoid averaging in the large steering angles that occur during curved road segments, the data occurring during curves were omitted. One unit in the profile is approximately equivalent to one-half second of real time (0.51 s for the human drivers, 0.64 s for the model).

Figure 5 shows the steering and lateral position profiles for lane changes in both directions. The steering profiles show that human drivers executed the lane change maneuver by turning the wheel in the direction of the destination lane, then back through center to the opposite direction, and finally settling back in the center position; the second peak has slightly lower overall magnitude, albeit not by a large amount. The model executed a very similar profile (R^2 = .79, *RMSE* = .02, again comparing all human with all model data points); the model's steering was slightly less smooth and flattened out more around the center 0° steering angle but, overall, reproduced both the pattern and the magnitude of the drivers' profile. Both human and model steering profiles were in large part symmetrical for left versus right lane changes (R^2 = .97 comparing human left and right profiles, R^2 = .93 comparing model profiles). The lateral position

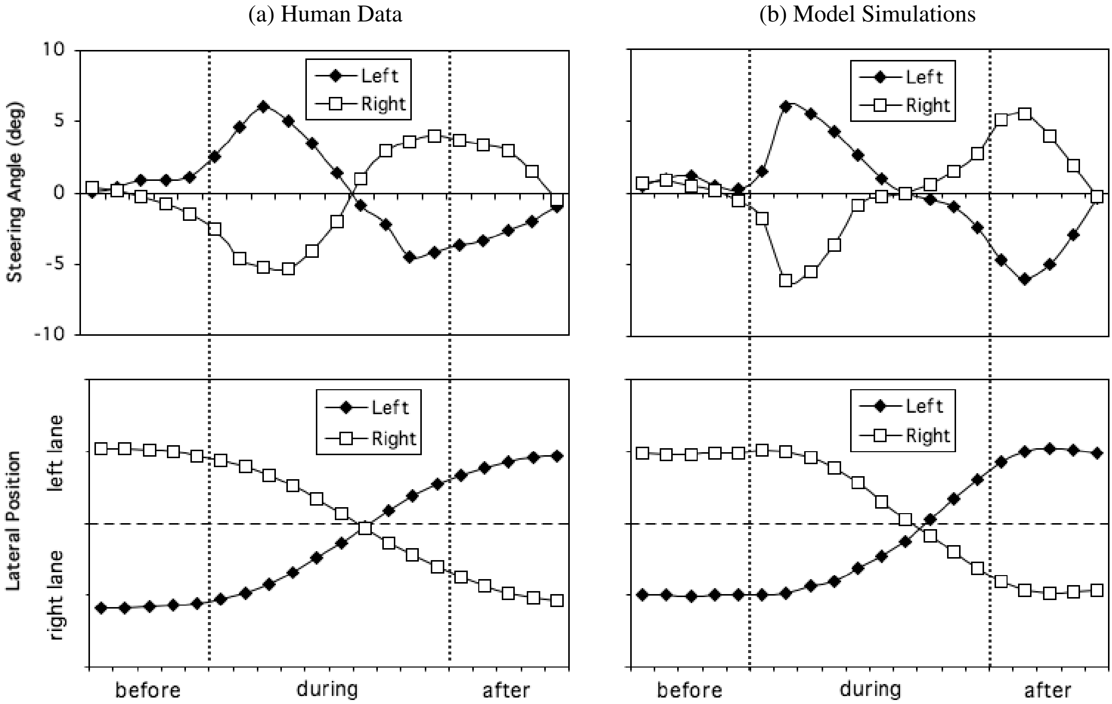

Figure 5. Lane-changing profiles for steering angle and lateral position for (a) human data and (b) model simulations. The vertical dotted lines indicate the start (leftmost line) and end (rightmost line) of the lane change.

profiles provide further illustration of the smooth execution of their lane changes, again with a close correspondence between model and human data (R^2 = .98, *RMSE* = .07).

In addition to steering and lateral position profiles, proportion gaze time can also be examined as a time-course profile over time during the lane change. The profiles for model and human data are shown in Figure 6, where each data point represents the equivalent of five units in Figure 5 to alleviate the large variability in these eye movement data. The figure presents results for three groupings of visual areas: the start lane of the lane change, the end lane, and the rear-view mirror. In examining the human driver profiles, one can see a noticeable shift from the start lane to the end lane; perhaps surprisingly, this shift happened not in the middle of the lane change, as the vehicle actually crossed lanes, but rather at the very start (or even before the start) of the lane change (see also Salvucci & Liu, 2002). The model predicted this same shift because of the manner in which it changes lanes – namely, simply switching to the control points of the end lane. However, the model's shift was much more severe: Although the model still gazed occasionally at the less frequent lane, it clearly directed most of its focused attention on the current control lane. Also, both the human drivers and model exhibited a larger proportion of gazes to the mirror before the lane change, steadily dropping through the progression of the maneuver; for both humans and model, the early gazes represent searching for surrounding vehicles and checking behind for sufficient clearance to execute the lane change. Although the quantitative fits could admittedly be better, the model does manage to capture the critical effects and shifts of attention exhibited by human drivers (R^2 = .65, *RMSE* = .26).

GENERAL DISCUSSION

Computational cognitive modeling is quickly maturing to address increasingly complex phenomena at an increasingly high level of rigor. More specifically, cognitive architectures have proven very successful at capturing both lower level performance and higher level decision making in complex dynamic tasks. The ACT-R driver model represents a contribution toward this effort with

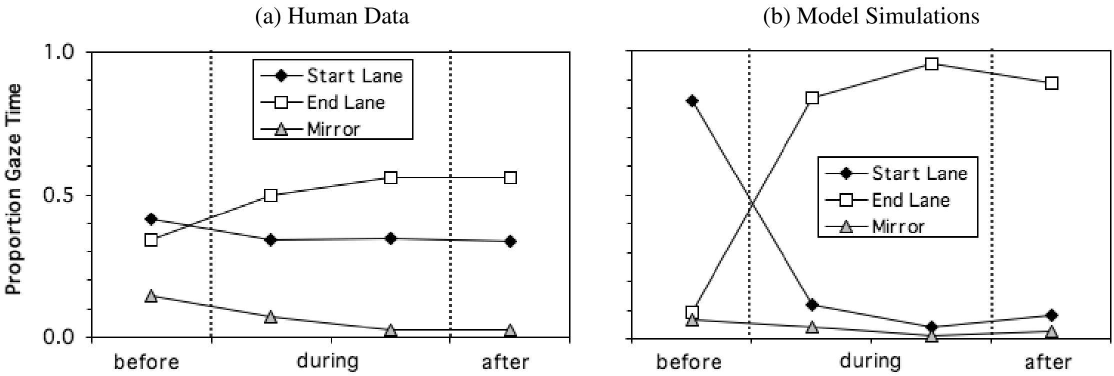

Figure 6. Lane-changing profiles for proportion gaze time on the start lane, end lane, and mirror for (a) human data and (b) model data. The vertical dotted lines indicate the start (leftmost line) and end (rightmost line) of the lane change.

a novel approach to integrating the lower level (i.e., operational) and higher level (i.e., tactical) aspects of driver behavior in the framework of the ACT-R cognitive architecture. Of course, the ACT-R driver model does not yet provide a complete picture of driver behavior – further work extending the task, artifact, and/or embodied cognition addressed by the model could take any number of directions. Nevertheless, I am confident that both model and architecture can evolve significantly from the current state of the art to capture a broader and deeper range of the phenomena surrounding driver behavior.

Benefits of Cognitive Architectures

As mentioned earlier, cognitive architectures offer many benefits for developing integrated models of driver behavior. The ACT-R architecture in particular has a number of features that relate directly to driver modeling and either facilitate modeling in the current model or will facilitate modeling in future versions of the model. A few of the most important benefits will now be considered and, as before, one should note that many of these benefits are not particular to ACT-R but would generalize to other cognitive architectures as well.

One of the most fundamental aspects of the ACT-R architecture is the treatment of seriality and parallelism in its various processes. ACT-R includes several processes that run in parallel, namely the cognitive, visual, and motor processes. However, each of these processes is itself a serial stream: The cognitive processor can run only one "thought" (i.e., production rule firing) at a time, and the visual and motor processors can execute only one visual/motor operation at a time (although they can prepare one while executing another). Cognition serves as the central bottleneck in the system: All visual and motor operations must be initiated by the cognitive processor, although while they are in operation, cognition may, if desired, process other rule firings. (In highly optimized cases, this can lead to extreme interleaving, as can be seen in models of "psychological refractory period" effects: e.g., Byrne & Anderson, 2001.) The immediate consequence of this limited parallelism for the driver model is that all driving tasks must share the serial cognitive processor, leading to contention for this precious resource.

In the current model, the three core components of control, monitoring, and decision making all share cognition – for instance, when the model is monitoring its environment, it cannot update vehicle control. When integrating this model with models of secondary tasks, the resource contention results in driver distraction or inattention and can adversely affect driver performance (e.g., Salvucci, 2001b). As noted earlier, when and how the model should schedule and interleave the various processes is a crucial concern for proper multitasking. The model currently incorporates a domain-specific model of multitasking (or a "customized executive": see Kieras, Meyer, Ballas, & Lauber, 2000) in which the interleaving is tuned for the defined set of tasks; a domain-independent theory of multitasking (or "general executive," e.g., Salvucci, 2005) would help to fold the driver model's method of interleaving into a more general cognitive framework.

Another key aspect of the architecture and model is the notion of limited attention: Drivers simply cannot attend – visually, cognitively, or otherwise – to everything at once. Constraints on "cognitive attention" come primarily from the limited parallelism and cognitive bottleneck described earlier. For visual attention, ACT-R incorporates a rigorous perceptual module (Byrne, 2001) that quantifies the constraints on visual attention in two important ways. First, the architecture can attend to only a single visual object at one time, and thus to attend to many objects it must shift attention between them. In attending to the near point, far point, and other vehicles, the driver model shifts visual attention between the various points and stores information about each (e.g., the visual angle to the near point) in the current goal chunk. Second, when using the EMMA module extension (Salvucci, 2001b), ACT-R generates emergent predictions about when and where the eyes move when following visual attention; in essence, every shift of visual attention also initiates an eye movement to that location, although this eye movement may not actually occur if attention shifts again before the eye movement starts executing. Also, a model can more easily encode objects near the current foveal location, and thus farther objects are more likely to be fixated with an eye movement. Thus, ACT-R incorporates several realistic constraints on visual attention, and all ACT-R models, like the driver model, inherit and must abide by these constraints. In addition, ACT-R's visual system produces observable predictions of driver behavior, namely driver eye movements, which can be compared directly with human drivers.

A third aspect of the model and architecture involves accounting for the many individual differences among drivers. Although individual differences have only begun to be addressed with the model, two initial studies have shown promising results in this area. First, in developing and testing the two-level control model, Salvucci and Gray (2004) found that scaling the model's parameter values allowed it to capture individual drivers' steering profiles for corrective maneuvers and lane changes. Second, Salvucci et al. (2004) explored the effects of age on driver performance and distraction: In comparing younger versus older drivers, we successfully modeled the effects of age-related cognitive speed differences on performance during normal driving and while executing a secondary task. This work addressed only the "hardware" differences (Meyer, Glass, Mueller, Seymour, & Kieras, 2001) among individuals, namely the changes in core processes such as memory and motor movements; the work did not address "software" differences in declarative knowledge and/or procedural strategies, perhaps arising from differences in background knowledge and experience. Because these software differences are typically domain specific, they are more elusive than hardware differences in terms of finding a comprehensive, general theoretical solution. Nevertheless, the modeling of strategic differences among drivers remains a long-term if not a short-term goal.

Practical and Theoretical Implications

The ACT-R driver model has several practical and more general theoretical implications for both driving and cognitive architectures. Just as the driver model aims to provide a rigorous theoretical account of driver behavior, it simultaneously strives to be a useful practical tool for real-world applications. In the past few years, several possible applications have been explored that generally fall into two categories: systems that attempt to recognize and infer driver intentions from actions (e.g., Salvucci, 2004) and systems that attempt to predict driver behavior given current situations (e.g., Salvucci, Zuber, Beregovaia, & Markley, 2005).

One application that has shown particularly good promise involves using the model as a tool to predict the effects of driver distraction for evaluating in-vehicle devices. Using an "integrated model approach," developers and/or designers can create cognitive models of behavior for their new devices and integrate them with the driver model production system. Because of the constraints imposed by the cognitive architecture, the behavior for the secondary task device interacts with the behavior for the primary driving task, potentially producing effects of the secondary task on driving (and also effects of driving on the secondary task). To date, the original prototype model (Salvucci et al., 2001) has been successfully used to account for effects of driver distraction in several studies. The initial study of the integrated model approach showed how the model could predict effects of cell phone dialing in different modalities on lateral control (Salvucci, 2001b). A follow-up study using existing cellular phone dialing methods and a challenging car-following task demonstrated similar predictive power for both lateral and longitudinal control (Salvucci & Macuga, 2002). In addition, a study of "cognitive distraction" showed successful predictions of driver performance during a high-load sentence-span task (Salvucci, 2002). As the ACT-R driver model continues to evolve to capture additional aspects of task (e.g., nonhighway driving), artifact (e.g., nonmidsize vehicles), and embodied cognition (e.g., haptic perception), the greater predictive power of the theoretical description immediately benefits real-world applications and, it is hoped, increases the impact and benefit of such applications for practical design and development.

In more general theoretical terms, the ACT-R cognitive architecture and the domain of driving have enjoyed a symbiotic relationship in which each benefits from interactions with the other. The driving domain challenges ACT-R to expand beyond the boundaries of basic laboratory tasks to the full complexity of real-world complex tasks. In doing so, driving and related complex domains have pushed the ACT-R research community to more rigorously address larger issues relevant in real-world task modeling, such as navigating in a three-dimensional world and coordinating low-level perception and action with higher level decision making. At the same time, ACT-R benefits the driving community by enabling researchers to view driver behavior through the eyes of the architecture, thus explaining or elucidating interesting aspects of behavior; for example, one can derive better understanding of driver visual processing and action, integration of low-level control and higher level decision making, and multitasking within driving and with secondary tasks all by placing driving in the context of the cognitive architecture. The ACT-R architecture is thus helping to shape scientific understanding of driving and, in turn, helping to provide a sound theoretical basis for practical applications that address real-world issues such as predicting driver distraction and performance.

ACKNOWLEDGMENTS

This work is supported in part by Office of Naval Research Grant #N00014-03-1-0036 and National Science Foundation Grant #IIS-0133083. I sincerely thank Wayne Gray and three anonymous reviewers for many excellent comments and helpful suggestions.

REFERENCES

Aasman, J. (1995). *Modelling driver behaviour in Soar.* Leidschendam, Netherlands: KPN Research.

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review, 111,* 1036–1060.

Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought.* Mahwah, NJ: Erlbaum.

Anderson, J. R., & Matessa, M. P. (1997). A production system theory of serial memory. *Psychological Review, 104,* 728–748.

Boer, E. R. (1999). Car following from the driver's perspective. *Transportation Research – Part F, 2,* 201–206.

Boer, E. R., & Hoedemaeker, M. (1998). *Modeling driver behavior with different degrees of automation: A hierarchical decision framework of interacting mental models.* Paper presented at the 17th European Annual Conference on Human Decision Making and Manual Control held in Valenciennes, France, December 14–16, 1998.

Byrne, M. D. (2001). ACT-R/PM and menu selection: Applying a cognitive architecture to HCI. *International Journal of Human-Computer Studies, 55,* 41–84.

Byrne, M. D., & Anderson, J. R. (2001). Serial modules in parallel: The psychological refractory period and perfect time-sharing. *Psychological Review, 108,* 847–869.

Carson, J. M., & Wierwille, W. W. (1978). Development of a strategy model of the driver in lane keeping. *Vehicle System Dynamics, 7,* 233–253.

Donges, E. (1978). A two-level model of driver steering behavior. *Human Factors, 20,* 691–707.

Fajen, B. R., & Warren, W. H. (2003). Behavioral dynamics of steering, obstacle avoidance, and route selection. *Journal of Experimental Psychology: Human Perception and Performance, 29,* 343–362.

Federal Highway Administration. (2001). *Our nation's highways 2000* (Tech. Rep. No. FHWA-PL-01-1012). Washington, DC: U.S. Department of Transportation.

Godthelp, H. (1986). Vehicle control during curve driving. *Human Factors, 28,* 211–221.

Gray, W. D. (2000). The nature and processing of errors in interactive behavior. *Cognitive Science, 24,* 205–248.

Gray, W. D., & Boehm-Davis, D. A. (2000). Milliseconds matter: An introduction to microstrategies and to their use in describing and predicting interactive behavior. *Journal of Experiment Psychology: Applied, 6,* 322–335.

Gray, W. D., & Fu, W.-T. (2004). Soft constraints in interactive behavior: The case of ignoring perfect knowledge in-the-world for imperfect knowledge in-the-head. *Cognitive Science, 28,* 359–382.

Groeger, J. A. (2000). *Understanding driving: Applying cognitive psychology to a complex everyday task.* Philadelphia: Psychology Press.

Hess, R. A., & Modjtahedzadeh, A. (1990). A control theoretic model of driver steering behavior. *IEEE Control Systems Magazine, 8,* 3–8.

Hildreth, E. C., Beusmans, J. M. H., Boer, E. R., & Royden, C. S. (2000). From vision to action: Experiments and models of steering control during driving. *Journal of Experimental Psychology: Human Perception and Performance, 26,* 1106–1132.

Jones, R. M., Laird, J. E., Nielsen P. E., Coulter, K., Kenny, P., & Koss, F. (1999). Automated intelligent pilots for combat flight simulation. *AI Magazine, 20,* 27–42.

Just, M. A., & Carpenter, P. A. (1992). A capacity theory of comprehension: Individual differences in working memory. *Psychological Review, 99,* 122–149.

Kieras, D. E., Meyer, D. E., Ballas, J. A., & Lauber, E. J. (2000). Modern computational perspectives on executive mental processes and cognitive control: Where to from here? In S. Monsell & J. Driver (Eds.), *Control of cognitive processes: Attention and performance XVIII* (pp. 681–712). Cambridge, MA: MIT Press.

Laird, J. E., Newell, A., & Rosenbloom, P. S. (1987). Soar: An architecture for general intelligence. *Artificial Intelligence, 33,* 1–64.
Land, M. F., & Horwood, J. (1995). Which parts of the road guide steering? *Nature, 377,* 339–340.
Land, M. F., & Lee, D. N. (1994). Where we look when we steer. *Nature, 369,* 742–744.
Lee, F. J., & Anderson, J. R. (2001). Does learning of a complex task have to be complex? A study in learning decomposition. *Cognitive Psychology, 42,* 267–316.
Lee, J. D., McGehee D. V., Brown T. L., & Reyes, M. L. (2002). Collision warning timing, driver distraction, and driver response to imminent rear-end collisions in a high-fidelity driving simulator. *Human Factors, 44,* 314–334.
Levison, W. H., & Cramer, N. L. (1995). *Description of the integrated driver model* (Tech. Rep. No. FHWA-RD-94-092). McLean, VA: Federal Highway Administration.
Liu, Y. (1996). Queuing network modeling of elementary mental processes. *Psychological Review, 103,* 116–136.
McRuer, D. T., Allen, R. W., Weir, D. H., & Klein, R. H. (1977). New results in driver steering control models. *Human Factors, 19,* 381–397.
Meyer, D. E., Glass, J. M., Mueller, S. T., Seymour, T. L., & Kieras, D. E. (2001). Executive-process interactive control: A unified computational theory for answering twenty questions (and more) about cognitive ageing. *European Journal of Cognitive Psychology, 13,* 123–164.
Meyer, D. E., & Kieras, D. E. (1997). A computational theory of executive cognitive processes and multiple-task performance: Part 1. Basic mechanisms. *Psychological Review, 104,* 3–65.
Michon, J. A. (1985). A critical view of driver behavior models: What do we know, what should we do? In L. Evans & R. C. Schwing (Eds.), *Human behavior and traffic safety* (pp. 485–520). New York: Plenum Press.
Newell, A. (1990). *Unified theories of cognition.* Cambridge, MA: Harvard University Press.
Pomerleau, D., & Jochem, T. (1996). Rapidly adapting machine vision for automated vehicle steering. *IEEE Expert, 112,* 19–27.
Ritter, F. E., Baxter, G. D., Jones, G., & Young, R. M. (2000). Supporting cognitive models as users. *ACM Transactions on Computer-Human Interaction, 7,* 141–173.
Rushton, S. K., Harris, J. M., Lloyd, M. R., & Wann, J. P. (1998). Guidance of locomotion on foot uses perceived target location rather than optic flow. *Current Biology, 8,* 1191–1194.
Salvucci, D. D. (2001a). An integrated model of eye movements and visual encoding. *Cognitive Systems Research, 1,* 201–220.
Salvucci, D. D. (2001b). Predicting the effects of in-car interface use on driver performance: An integrated model approach. *International Journal of Human-Computer Studies, 55,* 85–107.
Salvucci, D. D. (2002). Modeling driver distraction from cognitive tasks. In *Proceedings of the 24th Annual Conference of the Cognitive Science Society* (pp. 792–797). Mahwah, NJ: Erlbaum.
Salvucci, D. D. (2004). Inferring driver intent: A case study in lane-change detection. In *Proceedings of the Human Factors and Ergonomics Society 48th Annual Meeting* (pp. 2228–2231). Santa Monica, CA: Human Factors and Ergonomics Society.
Salvucci, D. D. (2005). A multitasking general executive for compound continuous tasks. *Cognitive Science, 29,* 457–492.
Salvucci, D. D., Boer, E. R., & Liu, A. (2001). Toward an integrated model of driver behavior in a cognitive architecture. *Transportation Research Record, 1779,* 9–16.
Salvucci, D. D., Chavez, A. K., & Lee, F. J. (2004). Modeling effects of age in complex tasks: A case study in driving. In *Proceedings of the 26th Annual Conference of the Cognitive Science Society* (pp. 1197–1202). Mahwah, NJ: Lawrence Erlbaum Associates.
Salvucci, D. D., & Gray, R. (2004). A two-point visual control model of steering. *Perception, 33,* 1233–1248.
Salvucci, D. D., & Liu, A. (2002). The time course of a lane change: Driver control and eye-movement behavior. *Transportation Research Part F, 5,* 123–132.
Salvucci, D. D., & Macuga, K. L. (2002). Predicting the effects of cellular-phone dialing on driver performance. *Cognitive Systems Research, 3,* 95–102.
Salvucci, D. D., Zuber, M., Beregovaia, E., & Markley, D. (2005). Distract-R: Rapid prototyping and evaluation of in-vehicle interfaces. In *Human Factors in Computing Systems: CHI 2005 Conference Proceedings* (pp. 581–589). New York, NY: ACM Press.
Tsimhoni, O. & Liu, Y. (2003). Modeling steering using the Queuing Network-Model Human Processor (QN-MHP). In *Proceedings of the Human Factors and Ergonomics Society 47th Annual Meeting* (pp. 1875–1879). Santa Monica, CA: Human Factors and Ergonomics Society.
van Winsum, W., & Godthelp, H. (1996). Speed choice and steering behavior in curve driving. *Human Factors, 38,* 434–441.
Weir, D. H., & McRuer, D. T. (1973). Measurement and interpretation of driver/vehicle system dynamic response. *Human Factors, 15,* 367–378.
Wilkie, R. M., & Wann, J. P. (2003). Controlling steering and judging heading: Retinal flow, visual direction and extra-retinal information. *Journal of Experimental Psychology: Human Perception and Performance, 29,* 363–378.

Dario D. Salvucci is an assistant professor at Drexel University. He received his Ph.D. in computer science from Carnegie Mellon University in 1999.

Date received: March 26, 2004
Date accepted: February 16, 2005