INFERRING DRIVER INTENT: A CASE STUDY IN LANE-CHANGE DETECTION

Dario D. Salvucci
Drexel University
Philadelphia, PA

This paper introduces a robust, real-time system for detecting driver lane changes. Under the framework of a "mind-tracking architecture," the system simulates a set of possible driver intentions and their resulting behaviors using an approximation of a rigorous and validated model of driver behavior. The system compares these simulations with a driver's actual observed behavior, thus inferring the driver's unobservable intentions. The paper demonstrates how this system can detect a driver's intention to change lanes, achieving an accuracy of 85% with a false alarm rate of 4%; detecting 80% of lane changes within 1/2 second and 90% within 1 second; and detecting 90% before the vehicle moves 1/4 of the lane width laterally — that is, approximately when the vehicle first touches the destination lane line.

INTRODUCTION

Intelligent transportation systems (ITS) represent a core part of research and development of the next generation of vehicles. While the many flavors of ITS span a broad scope of devices, infrastructure, and technology, one major goal of ITS in general arises in helping drivers to do what they intend to do more safely and effectively. Embedded in this goal is a simple idea: for systems to help drivers do what they intend, they must somehow infer drivers' intentions from their observable behavior. Humans infer one another's intentions continually — for instance when a passenger warns the driver of a car in the blind spot, the passenger has likely first inferred the driver's intention to change lanes, otherwise the warning would be needless and probably distracting. Similarly, intelligent machine systems need some way of inferring driver intent to provide safe assistance in both mundane and critical driving situations.

My research team and I are currently developing a general computational architecture for inferring human intent which we call "mind tracking." In essence, the architecture provides a computational framework for representing and tracking a person's intentions based on their observable behavior. As one test of this architecture, this paper describes an application of mind tracking to the problem of inferring driver intent, with the ultimate goal of being included in a variety of ITS systems. In particular, the paper addresses one very common aspect of driver behavior, namely lane changing, and explores how mind tracking can be used to detect a driver's intention to change lanes. A first attempt at lane-change detection (Salvucci & Siedlecki, 2003) yielded reasonable detection rates but, in practice, led to sometimes inconsistent and misleading predictions. This paper represents a different and improved approach to model representation and mind tracking, and offers significant promise as a robust, real-time algorithm ready for incorporation into intelligent vehicle systems.

The paper begins with an overview of the mind-tracking architecture as well as a detailed description of its application to lane-change detection. It then discusses an evaluation study in which the system was applied to human driver data taken from a realistic highway simulation, demonstrating that the system indeed can successfully detect driver lane changes with high accuracy and low false-alarm rates.

LANE-CHANGE DETECTION BY MIND TRACKING

The problem of detecting lane changes has proven deceptively challenging for existing methods, including some of the most powerful methods of statistical pattern recognition. Pentland and Liu (1999) and Oliver and Pentland (2000) employed hidden Markov models to recognize lane changes, but used a batch algorithm that detected whole instances of a maneuver rather than continuous recognition of streaming, real-time data as would be needed in a real vehicle. Kuge et al. (2000) developed a recognition system for this continuous case; however, their system used no information about the surrounding environment, but rather focused only on steering-based features. For true real-time recognition, a system must utilize contextual information to provide robust, consistent detections as each new data point is collected from the driver, vehicle, and environment.

This paper attempts to overcome the limitations of previous work by detecting lane changes using the mind-tracking architecture. As mentioned, the mind-tracking architecture is a computational framework for mapping a person's observable actions to their unobservable intentions. At its core, the architecture uses a computational model that is capable of predicting possible expected behavior(s) given a particular intention — for instance, predicting the driver behavior resulting from the intention to change lanes, turn, or simply stay in the current lane. The architecture can be best described as an iterating cycle of four steps, illustrated in Figure 1: (1) data collection, (2) model simulation, (3) action tracking, and (4) thought inference. The following sections describe each of these steps and how they are instantiated for the particular goal of detecting the intention to change lanes.

Downloaded from pro.sagepub.com at PENNSYLVANIA STATE UNIV on May 17, 2016

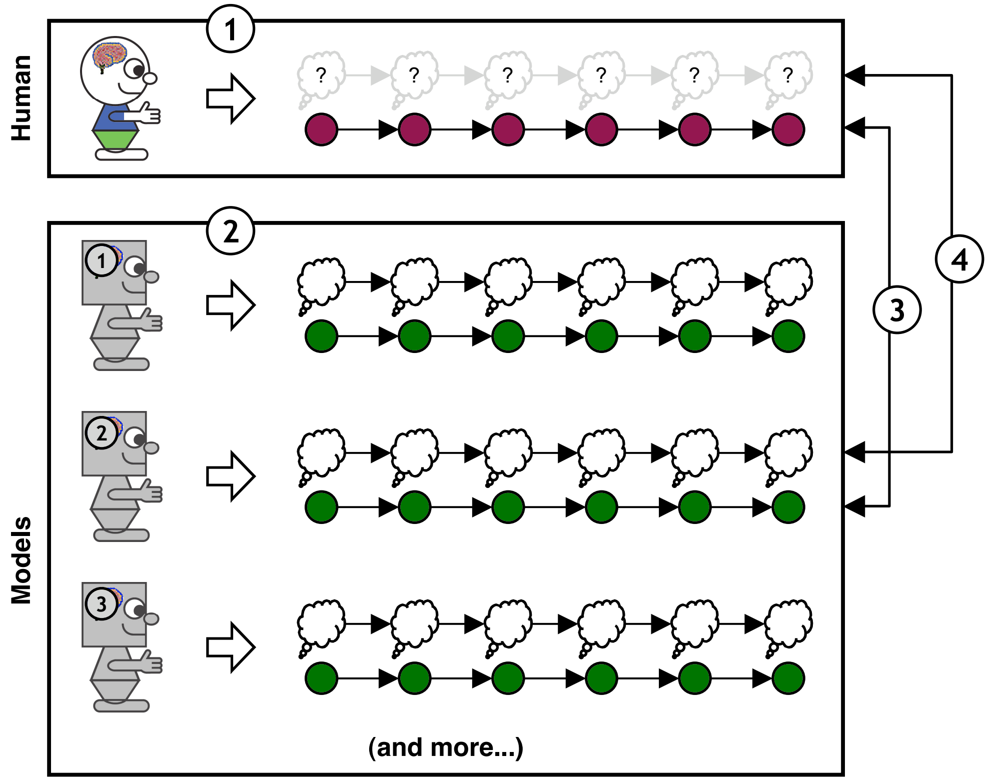

Figure 1: Schematic of the mind-tracking architecture.

Data Collection

The first step of mind tracking involves collecting a person's observable behavior and recording the behavior as a time-ordered vector of multimodal data. The data are typically sampled at a constant rate set in consideration of both the temporal density of the data's information and the density of predictions from the computational model. In addition to any observable data, mind tracking also records current environmental data in order to enable association of environment factors with resulting behavior. For the lane-change application, the system records steering-wheel angle and accelerator depression as well as environmental data (lateral position and time headway) as described next.

Model Simulation

The second step of the process involves running several versions of the cognitive model in parallel (conceptually if not actually), each representing a particular stream of possible intentions and actions. The cognitive model itself is a computational representation of a person's intentions and actions. The model used here is based on a cognitive model of driver behavior (Salvucci, Boer, & Liu, 2001) implemented in the ACT-R cognitive architecture (Anderson & Lebiere, 1998). The ACT-R driver model includes a cognitively plausible model of lateral and longitudinal control and has been validated to behave like human drivers in many aspects of common driving scenarios (e.g., curve negotiation and lane changing on a multi-lane highway). While the full driver model would suit our purposes here, a far simpler model based on this one suffices for tracking intentions.

The driver model used here is structured as follows. For lateral control, we assume the model has access to two salient visual features, namely the orthogonal lateral distance x_{near} (in meters) of the road 10 m ahead to the vehicle's current heading, and the analogous quantity x_{far} calculated at 30 m ahead. Using this information the model calculates a desired steering angle φ as

$$\varphi = k_{near}\left(x_{near} + x_{lc}\right) + k_{far}\left(x_{far} + x_{lc}\right)$$

where x_{lc} is zero during lane keeping and non-zero when lane changing, representing the desired displacement of the vehicle during the maneuver (roughly equivalent to desired lateral speed) with a sign dependent on the desired lane-change direction (left or right). The model also sets the accelerator position α based on another environmental variable, namely the minimum time headway thw to either the lead vehicle or, if changing lanes, the lead vehicle in the destination lane:

$$\alpha = k_{acc}\left(thw - thw_{follow}\right)$$

In this formulation, thw_{follow} is the desired following time headway, and the resulting value α is limited to the range [-1,1] indicating zero-to-maximum depression of the throttle

Downloaded from pro.sagepub.com at PENNSYLVANIA STATE UNIV on May 17, 2016

for positive numbers and zero-to-maximum depression of the brake for negative numbers.

Admittedly this model of the driver is grossly simplified — for instance, the steering angle does not take into account the vehicle's current speed. Nevertheless, we have found that this simple model is quite sufficient in producing the desired effect — effective tracking of driver intent — and is also computationally straightforward, making possible the real-time version of the full system presented here.

During simulation, the system runs simultaneous simulations of several versions of the model. Specifically, it maintains a set of models and spawns off new models for the next time step as follows: lane keeping (LK) stays LK and spawns two lane changing models, left and right (LC-L and LC-R); LC-L stays LC-L and LC-R stays LC-R until they reach their destination lane, then terminate and return to LK (thus, the set remains finite and reaches steady state with respect to number of models). A moving time window of 2 s is maintained for each simulation along with the human data.

The model parameters, including constants *k*, were approximated from the original ACT-R driver model and adjusted informally to produce better tracking performance during evaluation. The final parameter values were as follows: k_{near}=2, k_{far}=20, k_{acc}=1, x_{lc} = 1.75 m (approximately one-half the lane width), and thw_{follow}=1.0 s.

Action Tracking

The third step of mind tracking involves matching the observed behavior of the human driver with the predicted behaviors of the model simulations. Because each model generates an action sequence analogous to the human driver, we can compare the sequences directly and determine which model sequence best matches the human sequence. This requires a similarity metric between a model *M*'s simulation and the observed human data, computed as

$$S(M) = \prod_i G(\varphi_i^M, \hat{\varphi}_i, \sigma_\varphi) \cdot G(\alpha_i^M, \hat{\alpha}_i, \sigma_\alpha)$$

as the product over all sample indices *i* in the moving window. In the equation, φ_i^M and α_i^M are the steering angle and accelerator position (respectively) for the model at sample *i*; $\hat{\varphi}_i$ and $\hat{\alpha}_i$ are the analogous quantities observed from the human driver; and *G* is a Gaussian distribution for the given value, mean, and standard deviation (estimated along with the model parameters). Finally, a lane-change score is computed using the most probable models *LK* for lane keeping and *LC* for lane changing as

$$Score = \frac{\log S(LK)}{\log S(LC) + \log S(LK)}$$

where a score > .5 indicates a lane change.

Thought Inference

In the final step, mind tracking finds the inferred driver intentions simply by examining the "thoughts" of the best-matching model — that is, the intentions that produced this model's action sequence. Thus, the end result of mind tracking is inferred sequence of intentions over the length of the moving window. The process then repeats, shifting the window by one sample and iterating the four-step process.

EVALUATION STUDY

We evaluated the proposed lane-change detection system by running the system on a set of 11 human-driver protocols collected during free-form driving on a multi-lane highway (Salvucci, Boer, & Liu, 2001). To evaluate the algorithm, we required a rigorous definition of a lane change against which to compare the tracker's predictions. To this end, we classified a lane change as a segment in which the vehicle starts moving toward another lane and continues, without reversal, through to that lane.

Figure 2 graphs the true-positive versus false-positive (or false alarm) rate for the system on these protocols (a so-called ROC curve), where perfect recognition would pull the curve through (0,1). We see that the mind-tracking system performed very well for both true and false positives. In particular, when using the 0.5 score threshold, the system achieves 85% with 4% false alarms.

Figure 3 shows the ratio of lane changes detected over time from the start of the maneuver. The tracker already detects 65% of lane changes at the very start, indicating that to some extent, the tracker can detect the initial behavior that leads into the maneuver. After a half second, the tracker reaches 80% detection, and reaches over 90% after one second. As another way to view detection over time, Figure 4 shows the detection ratio as a function of lateral movement from the vehicle's position at the start of the lane change. Again the tracker detects lane changes rapidly; most notably, it achieves over 90% detection by the time the vehicle has moved 1/4 of the lane width — that is, roughly when the vehicle first touches the adjacent lane line.

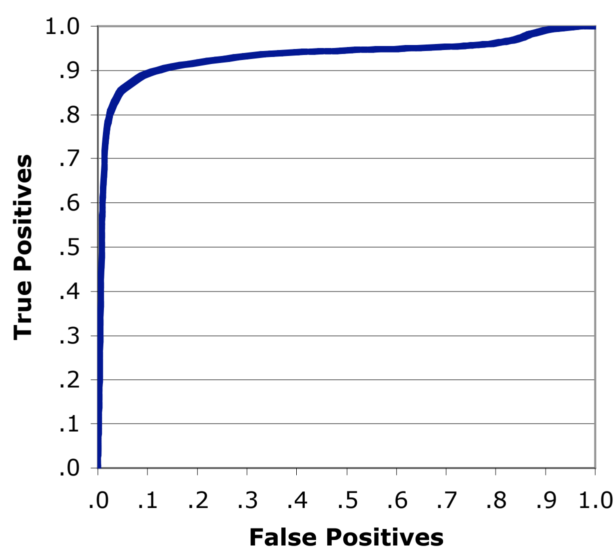

Figure 2: Detection true vs. false positives.

Downloaded from pro.sagepub.com at PENNSYLVANIA STATE UNIV on May 17, 2016

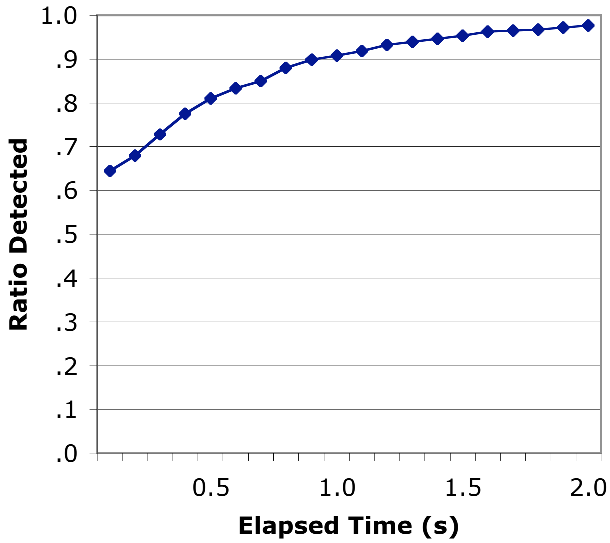

Figure 3: Detection ratio over time.

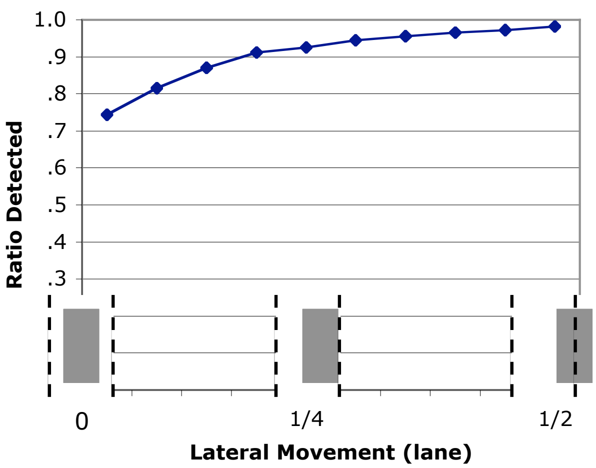

Figure 4: Detection over lateral movement.

DISCUSSION

As demonstrated, the mind-tracking architecture can quickly and accurately detect a drivers' intention to make a lane change. The architecture is by no means limited to lane-change detection, however; in fact, the basic ideas in the architecture generalize well to other intentions such as turning, stopping and starting, etc. We are now exploring the application of the architecture to other intentions, including the development of further computational models for expected behavior during these intentions and integrating these models into the mind-tracking process to enable real-time driver intent inference.

ACKNOWLEDGMENTS

This work was supported by Nissan Motor Company and National Science Foundation Grant #IIS-0133083. Many thanks to Nobuyuki Kuge, Tomohiro Yamamura, David Wipf, Erwin Boer, and the entire Nissan team for their valuable guidance and suggestions, and to Hiren Mandalia and Mark Zuber for their help with coding and analysis.

REFERENCES

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

Kuge, N., Yamamura, T., Shimoyama, O., & Liu, A. (2000). A driver behavior recognition method based on a drive model framework. In *Proceedings of the Society of Automotive Engineers World Congress 2000.*

Oliver, N., & Pentland, A. P. (2000). Graphical models for driver behavior recognition in a SmartCar. In *Proceedings of the IEEE Intelligent Vehicles Symposium 2000*.

Pentland, A., & Liu, A. (1999). Modeling and prediction of human behavior. *Neural Computation*, *11*, 229-242.

Salvucci, D. D., Boer, E. R., & Liu, A. (2001). Toward an integrated model of driver behavior in a cognitive architecture. *Transportation Research Record*, 1779.

Salvucci, D.D., & Siedlecki, T. (2003). Toward a unified framework for tracking cognitive processes. In *Proceedings of the 25th Annual Conference of the Cognitive Science Society*. Mahwah, NJ: Erlbaum.

Downloaded from pro.sagepub.com at PENNSYLVANIA STATE UNIV on May 17, 2016