Assessing the Feasibility of Vehicle-Based Sensors to Detect Alcohol Impairment

by

John D. Lee, Dary Fiorentino, Michelle L. Reyes, Timothy L. Brown, Omar Ahmad, James Fell, Nic Ward, Robert Dufour

National Advanced Driving Simulator, The University of Iowa

Executive Summary
Prepared for the

National Highway Traffic Safety Administration
1200 New Jersey Avenue, SE, West Building
Washington, DC 20590Sponsor

August 2010

EXECUTIVE SUMMARY

EXECUTIVE SUMMARY

The most notable findings from this study include:

• The National Advanced Driving Simulator (NADS-1) and a low-workload scenario are sensitive to alcohol. Alcohol effects were apparent in a simulation scenario representing a typical drive home at night from an urban bar.

• These effects were evident in drivers' control of vehicle lane position and speed.

Standard deviation of lane position and average speed differentiated BAC conditions most precisely.

• The most sensitive indicators of impairment are associated with continuous performance

(e.g., lane keeping) rather than discrete events (e.g., response to a traffic signal or use of

turn signals).

• The three algorithms detected impairment at and above the legal limit about as well as the Standardized Field Sobriety Test (SFST), with sensitivity increasing with BAC level.

• This project demonstrates the feasibility of a driving-behavior-based approach to detecting alcohol impairment in real time.

Background

Despite persistent efforts at the local, state, and federal levels, alcohol-impaired driving crashes still contribute to approximately 31% of all traffic fatalities. Although regulatory and educational approaches have helped reduce alcohol-related fatalities, other approaches merit investigation. One such approach detects alcohol impairment in real time using the increasingly sophisticated sensor and computational platform that is available on many production vehicles. It may be possible to detect impairment using driver state (e.g., eye movements), drivers' control inputs (e.g., steering and accelerator movements), and vehicle state (e.g., speed or lane position). Once detected, this information can support interventions that discourage drivers from driving while impaired and prevent alcohol-related crashes. This study assessed how well algorithms could detect impairment in a widely applicable and timely manner.

Objectives

The long-term research objective is to use algorithms that detect impairment as feedback to drivers to discourage or prevent drinking and driving. This report describes how, individually and in combination, driver actions reveal signatures of alcohol impairment, and how well algorithms built on these signatures detect drivers with BAC levels that are over the legal limit. Specific objectives include:

• Understand how driving-related metrics reflect the impairment associated with BACs at 0.05% and 0.10% (currently, the legal limit in the United States is 0.08%)

- Determine how well these metrics apply to different roadway situations and to different drivers (i.e., determine robustness)
- Develop algorithms to detect alcohol-related impairment
- Compare robustness of metrics and algorithms

Method

Data were collected in the National Advanced Driving Simulator (NADS) from 108 moderate to heavy drinkers dosed at placebo (0.00%), 0.05%, and 0.10% BAC on three separate visits. Drivers were divided into equal groups by age (21-34, 38-51, and 35-68) and gender. The participants drove a scenario representative of a drive home from an urban bar: a nighttime trip that involved a stretch of freeway and ended on a rural gravel road. The drives started with an urban segment composed of a two-lane roadway through a city with posted speed limits of 25 to 45 mph with signal-controlled and uncontrolled intersections. An interstate segment followed that consisted of a four-lane divided expressway with a posted speed limit of 70 mph. The drives concluded with a rural segment composed of a two-lane undivided road with curves. A portion of the rural segment was gravel. Drivers' steering, accelerator, and brake inputs, vehicle lane position and speed, the driving context (whether the vehicle was in a urban, interstate, or rural environment), and driver eye and eyelid movements were captured in representative driving situations, with precise control and in great detail.

Results

The objectives were addressed with two broad sets of analyses. The first focused on whether BAC affected performance. The second focused on detection of impairment. These analyses show the simulator and scenario to be sensitive to alcohol, and that algorithms can detect alcohol-related impairment.

Driving performance measures (i.e., mean speed, standard deviation of speed, and standard deviation of lane position) indicated systematic differences between BAC conditions. No statistically reliable effects of age and gender were found for lane deviation, but BAC affected lane deviation. Normalized lane deviation for the entire drive was 46.77 at 0.00% BAC, 49.79 at 0.05% BAC, and 54.31 at 0.10% BAC. Age reliably affected average speed, with average speed increasing with increasing age. BAC also affected average speed, with a higher BAC, in general, leading to lower average speed. Neither age nor BAC reliably affected speed deviation. Surprisingly, gender reliably affected speed deviation, with speed deviation greater for males than for females. These results are notable because the alcohol effects are apparent even though all participants were moderate to heavy drinkers and the driving situation was representative of daily driving, placing relatively low demands on the driver.

Taken together, the results from the impairment analyses indicate that alcohol affected performance, and that the NADS is sensitive to those changes. The next set of analyses focused on whether it is possible to classify BAC status (BAC $< 0.08\%$ v. BAC $\geq 0.08\%$) on the basis of those changes. Three algorithms were developed to predict BAC status based on logistic regression, decision tree modeling, and support vector machines. The algorithms were assessed in terms of sensitivity, robustness, timeliness, and bias.

Sensitivity is the degree to which the algorithm can differentiate between drivers above and below the legal limit. These algorithms show sensitivity comparable to that of the SFST.

Accuracy of the algorithms ranged from 73 to 86%. The logistic regression algorithm achieved an accuracy of 74.4% by combining information across the entire drive, achieving maximum sensitivity after approximately 25 minutes of driving. Decision tree and Support Vector Machine algorithms are much more sensitive and timely, identifying impairment in the situations tested with greater precision after only approximately eight minutes of driving. Greatest sensitivity was achieved by the decision tree, which combined driving performance indicators tailored to individual drivers. The most sensitive indicators of impairment were associated with continuous performance (e.g., lane keeping) rather than discrete events (e.g., response to a traffic signal or use of turn signals). Performance of the algorithms showed substantial differences in the degree to which they and their constituent measures provide robust and timely indications of impairment.

Robustness is insensitivity to confounding factors, such as different driving environments, and it applies to many factors affecting algorithm performance. An important element of robustness addressed in this study concerns the dependence of the algorithm on differences between drivers and the driving environment (i.e., urban, freeway, rural). Consistent with previous research, algorithms tailored to individuals outperformed generic algorithms by approximately 13%. Algorithm performance also depends on the driving context: different driving situations provide different measures, and these measures differ in their sensitivity. Current vehicle technology makes it quite feasible to capitalize on the benefits of tailoring algorithms to individuals by comparing a driver’s performance to his or her past performance in similar roadway situations.

Timeliness is the speed with which an algorithm is able to detect impairment. Timeliness is a critical consideration for real-time algorithms because some interventions rely on impairment detection well before the end of the drive. Timely impairment detection depends critically on the driving context with some events and variables being more sensitive than others. The most sensitive indicators of impairment involve continuous measures cumulated over time, such as the standard deviation of lane position. In addition, important signatures of alcohol impairment (straddling, weaving, and gaze concentration) are defined by behavior that evolves over a relatively long time horizon, requiring samples of driving behavior that extends over 30 seconds to several minutes. Even with such constraints, sensitivity comparable to the SFST was obtained over approximately eight minutes of driving.

Bias refers to the tendency of the algorithm to favor correctly detecting impairment at the expense of incorrectly identifying impairment when there is none. The fundamental differences in the optimization approaches between decision trees and SVM lead to differences in bias. These complementary differences can be leveraged to minimize false detection of impairment and to maximize detection of impaired drivers. Algorithms can be combined according to the benefits of detecting impairment and the costs of failing to detect impairment, so that one algorithm is used to maximize impairment detection and another is used to avoid false detection. Support Vector Machines show a tendency to outperform decision trees in maximizing detection.

Recommendations and conclusions

This study demonstrates the feasibility of vehicle-based sensors to detect alcohol-related impairment in real time: sensitivity is comparable to the SFST. This sensitivity is likely a very conservative estimate relative to sensitivity in detecting higher BAC levels. Because 66% of alcohol-related fatalities occur with BAC levels above 0.15% (compared with impaired drivers at 0.08 BAC or greater), the greatest value of a vehicle-based countermeasure may lie in detecting high BAC levels, where algorithms are likely to be very sensitive. These results suggest substantial promise in detecting other impairments, such as drowsiness, distraction, and even age-related cognitive decline.

The ultimate aim of impairment-detection algorithms is to support interventions that guide the driver to safer behavior. The desirability and feasibility of any particular algorithm depends on how it meets the particular needs of an intervention. This study demonstrates that algorithms differ substantially on these dimensions and that design must consider the inevitable tradeoffs. Algorithms become more sensitive, but less timely, as measures are integrated over time. The ultimate feasibility of impairment-detection algorithms depends on matching the performance profile of an algorithm to an intervention.

This project demonstrated the feasibility of a behavior-based approach to detecting alcohol impairment. It also identified many issues that merit further investigation. This study focused on moderate levels of alcohol (0.05% and 0.10% BAC) in people who reported to be moderate to heavy drinkers, but not problem drinkers. Given the assumption that moderate and heavy drinkers show less obvious indications of impaired driving, algorithms that are able to detect impaired driving from this population are likely to be much more effective for people who are light drinkers or for higher BAC levels, whereas the effectiveness of the algorithms for problem drinker is unknown. It would be useful to empirically assess how algorithm sensitivity differs at higher BACs, and for light and chronic drinkers.

More generally, this study identified a huge design space of sensors, measures, signatures and algorithms, algorithm parameter combinations, and meta algorithms. Assuming there are at least 10 sensors, 10 metrics for each sensor, 4 time scales, 10 algorithms, 10 implementations of each algorithm, and 4 meta algorithms, a total of more than 160,000 potential algorithms exist. This project sampled only a small region of that space. A specific challenge facing deployment of the algorithms developed in this study is the reliable measurement of lane position. Current lanetracking technology is vulnerable to sun glare, adverse weather, and poorly maintained lane markings. Steering wheel position is not subject to these limits and further analysis of its ability to detect alcohol-impaired driving is warranted. More generally, a promising direction for algorithm development is to identify classes of drivers and classes of driving situations, and an important direction for algorithm assessment is to develop metrics that relate to the interventions the algorithm intends to support. These results also suggest that further exploration is warranted, not just for alcohol impairment detection, but also for detecting impairment associated with drowsiness, distraction, and even age-related cognitive decline.