

The Exploratory Advanced Research Program



Automated Video Feature Extraction

WORKSHOP SUMMARY REPORT • October 10-11 2012



U.S. Department
of Transportation
**Federal Highway
Administration**

Notice

This document is disseminated in the interest of information exchange under the sponsorship of the Department of Transportation. The United States Government assumes no liability for its contents or use thereof. This report does not constitute a standard, specification, or regulation.

The United States Government does not endorse products or manufacturers. Trade and manufacturers' names appear in this report only because they are essential to the object of the document.

Quality Assurance Statement


The Federal Highway Administration (FHWA) provides high-quality information to serve government, industry, and the public in a manner that promotes public understanding. Standards and policies are used to ensure and maximize the quality, objectivity, utility, and integrity of its information. FHWA periodically reviews quality issues and adjusts its programs and processes to ensure continuous quality improvement.

Cover image: © shutterstock.com/Sergei Nivens.

TECHNICAL REPORT DOCUMENTATION PAGE

1. Report No. FHWA-HRT-13-037	2. Government Accession No.	3. Recipient's Catalog No.	
4. Title and Subtitle Automated Video Feature Extraction Workshop Summary Report		5. Report Date December 2012	
		6. Performing Organization Code:	
7. Author(s) Tom Morton		8. Performing Organization Report No.	
9. Performing Organization Name and Address Woodward Communications, Inc. 1420 N Street, NW, Suite 102 Washington, D.C. 20005		10. Work Unit No.	
		11. Contract or Grant No. Contract DTFH61-09-F-00027	
12. Sponsoring Agency Name and Address Office of Safety Research and Development and Office of Corporate Research, Technology, and Innovation Management Federal Highway Administration 6300 Georgetown Pike McLean, Virginia 22101		13. Type of Report and Period Covered Summary Report October 2012	
		14. Sponsoring Agency Code HRDS-2 and HRTM-30	
15. Supplementary Notes FHWA's Contracting Officer's Task Manager (COTM): Zachary Ellis, HRTM-30 Technical Contact: Lincoln Cobb, HRDS-2			
16. Abstract This report summarizes a 2-day workshop on automated video feature extraction. Discussion focused on the Naturalistic Driving Study, funded by the second Strategic Highway Research Program, and also involved the companion roadway inventory dataset. The specific objectives of the workshop were to begin a discussion on how Government, academia, and the private sector can cooperate to advance the state of the practice in the automated analysis of video data from naturalistic driving studies. A panel of expert speakers presented the state of knowledge in video feature extraction and demonstrated and described a range of analytical capabilities that could be automated. Following the presentations, the participants discussed what could be learned from the data, identified naturalistic data challenges, examined near- and long-term technical approaches, and reviewed organizational approaches for advancing the practice of automated feature extraction.			
17. Key Words Automated Video Feature Extraction, Naturalistic Driving Study, Video Analytics, Automated Analysis, Video Data, Real-Time Analysis, Computer Vision, Big Data, Data Sets, Driver Behavior, Human Factors, Driver Distraction.		18. Distribution Statement No restrictions. This document is available to the public through the National Technical Information Service, Springfield, VA 22161.	
19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages TBD	22. Price N/A

EXECUTIVE SUMMARY

 On October 10–11, 2012, at the Turner-Fairbank Highway Research Center in McLean, VA, the Federal Highway Administration’s Office of Safety Research and Development and Exploratory Advanced Research Program convened a 2-day workshop on automated video feature extraction. The objective of the workshop was to begin to answer the question: how can Government, academia, and the private sector cooperate to advance the state of the practice in the automated analysis of data from naturalistic driving studies?

With new, smaller, and less obtrusive sensor technology, researchers are able to engage for the first time in gathering massive amounts of data about driving behavior. Naturalistic driving studies—including the one undertaken by the Strategic Highway Research Program 2, the largest ever naturalistic driving study—provide detailed information about driver behavior, vehicle state, and roadway using video cameras and other types of sensors. Naturalistic driving data provide an opportunity to improve understanding of vehicle crashes, particularly providing useful information on driver distraction and driver behavior leading up to a crash.

The workshop began with an introduction from U.S. Department of Transportation representatives and a brief background to the workshop. A panel of expert speakers

then presented on state of knowledge in video feature extraction and demonstrated and described a range of analytical capabilities that could be automated.

Virginia Tech Transportation Institute’s Jon Hankey provided video examples of the data collection system in action. Hankey also provided information on privacy issues and ways to improve data access. Participants were told about remote secure enclaves and a smaller driver study offering easier access. John Lee of the University of Wisconsin–Madison’s Department of Industrial and Systems Engineering explained the importance of putting driving into context and the unique and valuable role naturalistic driving data can play in that process. Lee also covered some of the challenges that come with big data before highlighting the importance of fully understanding driver distraction and avoiding disciplinary myopia.

Human–computer interaction was the focus of Margrit Betke, from Boston University’s Computer Science Department. Betke addressed some of the research challenges—including privacy issues—and then made several recommendations, including the possibility of creating a three-dimensional (3D) facial reconstruction to assist with data privacy. Qiang Ji of the Rensselaer Polytechnic Institute’s Department of Electrical, Computer, and Systems Engineering then explained his research

into real-time driver-state monitoring and recognition. Video demonstrations showed face detection and tracking, and Ji highlighted the ability to characterize driver state using facial expression analysis.

Next, Yuanqing Lin of NEC Laboratories America's Department of Media Analytics explained the concept of low-level sensing and high-level understanding, also highlighting the importance of 3D reconstruction. Finally, Mohan Trivedi of the Laboratory for Intelligent and Safe Automobiles (LISA) at the University of California at San Diego discussed some key remarks regarding computer vision, driving context, and adopting a holistic approach to understand driving. Trivedi addressed the importance of focused research and provided an overview of LISA's research into a complete driving context capture system.

The second part of the workshop included an opportunity for workshop participants—including safety regulators, driver behavior researchers, and human factors experts—to review some of the key points raised during

part one before moving on to address data needs, identify challenges, and collaborate on potential solutions.

Throughout the workshop, contextual data were consistently discussed as essential to the behavioral analysis of video data. It was noted the parties to the conversation have complementary needs: researchers in computer science, image processing, and machine learning need access to data and some degree of financial support to do their work; researchers in highway safety and driver behavior need tools to process the existing and forthcoming data.

Several possible approaches were identified for further investigation, including the concept of “remote secure enclaves” and creating reduced datasets to bypass some of the existing privacy management issues. Generating more manageable datasets was a key discussion point and several methods were discussed to achieve this goal, including developing tools to help researchers identify data of interest from thousands of hours of footage.

TABLE OF CONTENTS

INTRODUCTION	1
PART ONE: SPEAKER PRESENTATIONS	2
SHRP2 Safety: Making a Significant Improvement in Highway Safety	3
Behavioral Safety and Driver Distraction	6
Data Collection	8
Video Analytics: Putting Driving in Context	10
Challenges of Video-Based Human-Computer Interfaces and Experiences in Analyzing Videos of Driver Faces	13
Automatic Video-Based Driver State Analysis and Recognition	15
Our Roadmap of Automatic Feature Extraction for Driving Safety	18
Vision for Driver Assistance: Looking at People in a Vehicle	20
PART TWO: DISCUSSION	22
APPENDIX A: WORKSHOP PARTICIPANTS	29
REFERENCES	32

LIST OF FIGURES

Figure 1. Diagram. The data acquisition system.

Figure 2. Photo. Footage is placed into a single grid.

Figure 3. Photo. An intentionally blurred still image of the inside of the car.

Figure 4. Photo. An instrumented van.

Figure 5. Photo. The van produces a sequence of photographs for a photo log of its journey.

Figure 6. Photo. Light “blooming” can impede facial recognition.

Figure 7. Screen Capture. A facial recognition system from earlier research.


Figure 8. Chart. 3D trajectory analysis.

Figure 9. Diagram. System approach.

Figure 10. Photo. Facial feature detection automatically superimposes 28 points onto a face.

Figure 11. Photo. Infrared is used to detect cornea reflection.

INTRODUCTION

 On October 10-11, 2012, at the Turner-Fairbank Highway Research Center (TFHRC) in McLean, VA, the Federal Highway Administration's (FHWA) Office of Safety Research and Development (R&D) and Exploratory Advanced Research (EAR) Program convened a 2-day workshop on automated video feature extraction.

Workshop discussion focused on the Naturalistic Driving Study (NDS), funded by the second Strategic Highway Research Program (SHRP2), and also involved the companion roadway inventory database. The largest study of its kind ever undertaken, once completed researchers will have access to almost 3,700 driver years and a database containing an estimated 2.5 million trip files.

Feature extraction is designed to simplify the resources required to analyze a large dataset accurately; however, large datasets make conventional manual coding techniques unfeasible. Therefore, to maximize the use of NDS data there is an urgent need to develop automated feature extraction and efficiently reduce the data to more manageable elements of interest. It is essential to identify short-term solutions and incremental improvements to help analyze data more efficiently for application in the safety community and beyond.

The specific objectives of the workshop were to begin a discussion on how Government, academia, and the private sector can cooperate to advance the state of the practice in the automated analysis of video data from naturalistic driving studies. A panel of expert speakers presented the state of knowledge in video feature extraction and demonstrated and described a range of real-time analytical capabilities.

Following the presentations, the participants—including safety regulators, driver behavior researchers, and human factors experts—discussed what could be learned from the data, identified naturalistic data challenges, examined near- and long-term technical approaches, and reviewed organizational approaches for advancing the practice of automated feature extraction.

The Government's goals are both short and long term. In the short term, the Government wants to begin to extract value from the NDS data, and welcomes immediately deployable and partial solutions toward that goal. In the long term, the Government wants to ensure that the data being collected will improve transportation safety to the maximum possible extent.

**PART ONE:
SPEAKER PRESENTATIONS**

SHRP2 Safety: Making a Significant Improvement in Highway Safety

Ken Campbell, *Chief Program Officer*

National Academy of Sciences, Transportation Research Board, SHRP2

SHRP2 Background

Ken Campbell, Chief Program Officer at the National Academy of Sciences, began with a background to the second Strategic Highway Research Program (SHRP2). SHRP2 was authorized by Congress to address some of the most pressing needs related to the Nation's highway system. SHRP2 is administered by the Transportation Research Board (TRB) of the National Academies under a Memorandum of Understanding with FHWA and the America Association of State Highway and Transportation Officials (AASHTO). Contractors for the safety area include Virginia Tech Transportation Institute (VTTI), the Iowa State University Center for Transportation Research and Education (CTRE), Fugro, CUBRC, Battelle, the University of South Florida, Westat, Penn State, and Indiana University.

Why Study Naturalistic Driving?

Campbell explained that the safety area of SHRP2 is conducting the largest ever naturalistic driving study (NDS) to better understand the interaction among various factors involved in highway crashes—driver, vehicle, and infrastructure—so that better safety countermeasures can be developed and applied to save lives. The study is a focus on driver behavior and addresses the notion that it is possible to obtain more and better information on what people do when they drive—not just in the moments before they get into a collision but on a day-to-day basis.

NDS helps provide the solid baseline reference that is required to provide context to assess what is high-risk behavior and what is considered normal. For example, when analyzing what a person does every day, it may become apparent that a perceived risk factor just forms part of the subject's normal driving habits.

NDS Design

There are two major databases forming the SHRP2 data. Campbell informed workshop participants that the NDS ultimately will involve 2,800 primary drivers of all age and gender groups in passenger cars, minivans, SUVs, and pickup trucks. By the end of the study, this will add up to a database containing an estimated 5 million trip files that will prove invaluable to the next generation of researchers.

The data being collected by instrumentation within the participating vehicles include information from multiple videos; lane trackers; accelerometers; global positioning systems (GPS); radar; cell phone records; alcohol sensors; turn-signal status; and specific vehicle data on accelerator and brake pedal inputs, gear selection, steering angle, speed, seat belt information, and more, as illustrated in figure 1.⁽¹⁾

The Instrumentation Package

Campbell explained that a video stream is recorded at 15 Hz and 640 x 320 pixels. This is compressed before it is stored, so only the compressed vehicle data are available for analysis. The images are placed into a single

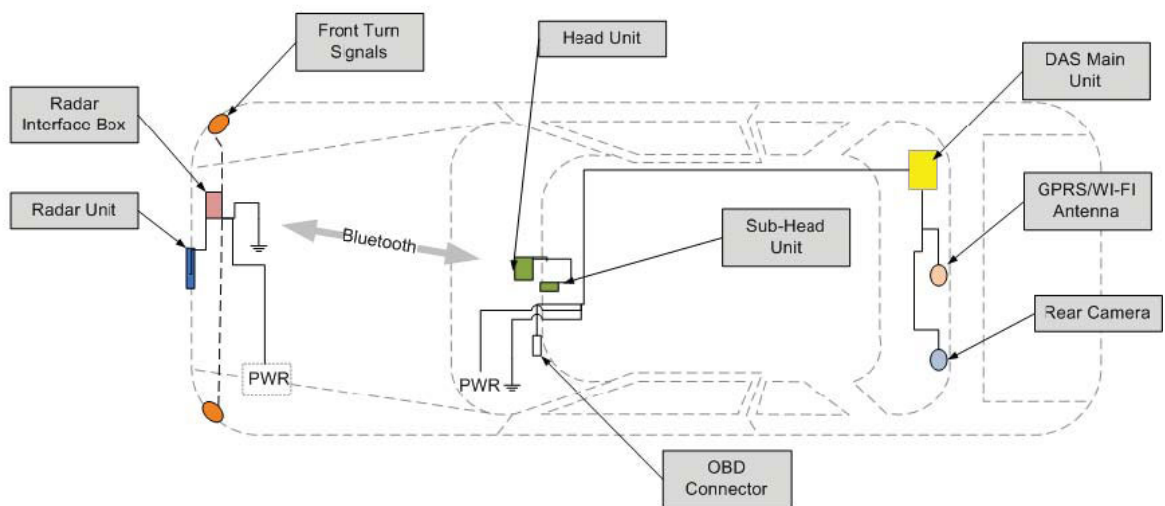


Figure 1. Chart. Full range of automation and cooperation alternatives.

grid, with most of the pixels allocated to the forward camera view for maximum clarity, as shown in figure 2.⁽¹⁾ In addition to multiple video camera angles, an intentionally blurred still image of the inside of the car is taken once every several seconds to show if there are other occupants in the vehicle, as shown in figure 3.⁽¹⁾



Figure 2. Photo. Footage is placed into a single grid.

In addition to vehicle data, roadway data are formed from a combination of three data sources. Researchers sent vans, as shown in figure 4, to measure 12,500 centerline miles (20,100 km) (the length of a highway regardless of the pavement width or the number of lanes) of roads across the test sites with a focus on data needed for lane departure and intersections. These include

variables such as curvature location, grade, cross slope, lane, shoulder type, speed-limit signs, medians, rumble strips, lighting, intersection location, and number of approaches. The researchers also added supplemental data from State inventory and data from each of the sites—including crash history, information from work zones, weather and traffic information, and other specific topics.

The roadway data collection also includes a traditional photo log where still pictures are taken every 21 ft (6 m) to produce a sequence of photographs of where the vans are driven, as shown in figure 5.



Figure 3. Photo. An intentionally blurred still image of the inside of the car alternatives.



© Fugro Roadware

Figure 4. Photo. An instrumented van.

Next Steps

Campbell stated that four separate analysis projects have started based on the collected data. These projects focus on testing analysis methods on rural two-lane curves; offset left-turn bays; driver inattention; and crashes on congested freeways. Future applications could lead to more cost-effective roadway measures to prevent crashes, cost-effective intersection design, new vehicle technology to track driver attention, and effective methods to warn drivers of congestion ahead.

In the long term, there will be many applications to address significant safety issues using the NDS and roadway data, which will help gain insights into driver behavior that cannot be obtained any other way. Some short-term goals include assisting users in defining research questions and accessing the databases and identifying who or what the user will be.

The research team will produce specific user tools and data files to support the analysis activities. These will include trip summary

files based around trips of interest for easy access—containing information on trip, roadway, vehicle, and driver variables. Specific event files could also help identify particular areas of interest (e.g., collision data) and suggestions for events and triggers are being solicited. The possibility of reduced datasets for individual users is also under development.

Questions

Following the presentation, workshop participants had an opportunity to ask questions. One question focused on how the vehicle instrumentation and roadside image data could be linked together. Campbell explained that this is one of the challenging problems of the study. Although the vehicle instrumentation does include inexpensive GPS sensors (recording location and heading once per second) and the location of the roadway data is stored in a spatial database, how to actually link them together is yet to be resolved.



© Fugro Roadware

Figure 5. Photo. The van produces a sequence of photographs for a photo log of its journey.

Behavioral Safety and Driver Distraction

Richard Compton, *Director, Office of Behavioral Safety Research*
National Highway Traffic Safety Administration

Analyzing Crash Data

Richard Compton, Director of the National Highway Traffic Safety Administration (NHTSA) Office of Behavioral Safety Research, began by informing workshop participants that studying naturalistic driving data offers NHTSA a valuable opportunity to greatly improve the understanding of how crashes occur. He also noted that there are many potential uses for the data in the future. Participants were told that most of the knowledge of the past 50 years comes from post-crash analysis; although information gathered from the scene of a crash is very useful, there is a lot of information leading up to a crash that cannot be gathered using this method.

Many hours were spent performing manual coding of video to allow researchers to see exactly what drivers were doing in their cars on a normal basis and leading up to a crash.

Compton stated that most crashes are not due to roadway or vehicle defects. 85 to 95 percent of crashes are due to operator behavior, but it is very hard to understand how the person was behaving leading up to the crash after the event. Measurements like pre-crash speed are extremely difficult for a police officer to calculate at the scene after a crash has occurred. Visual indicators, such as skid marks, used to provided valuable information but most

cars today have anti-lock braking systems which take such clues away.

It was noted that modern cars typically contain an electronic data recorder (EDR); however, these recorders operate on a first-in first-out policy and are usually triggered by an airbag deployment, saving just a couple of seconds of extremely limited data. The EDR systems do not perform like the more sophisticated black boxes installed in airplanes and trains that provide detailed information on the last hour of travel.

Assessing Driver Distraction

Compton informed participants that this naturalistic behavior study will supply that critical information and also provide an opportunity to study and analyze what normal behaviors are and assess abnormal behavior.

Previously, a 100-car study operated as a pilot study for the NDS and part of the preliminary study aimed to get a sense of driver distraction factors. Many hours were spent performing manual coding of video to allow researchers to see exactly what drivers were doing in their cars on a normal basis and leading up to a crash. One of the key factors to emerge from this study was the importance of keeping eyes on the forward roadway. Eyes that strayed off the forward roadway for more than 2 seconds showed a clear link to the crash data. For the first time this provided objective evidence and a clear correlation between an event happening when eyes were off the forward roadway.

Compton added that NHTSA has a great interest in getting involved in different types of driver distraction issues—whether they are manual, visual, or cognitive distractions. For example, many drivers now use electronic navigation systems in their vehicles but there is no clear understanding of exactly how people are using these devices. This type of naturalistic driving data gives insight into those uses.

Assessing a Driver’s Perception of Safety

Compton highlighted that naturalistic behavior data are also extremely useful for studying the use of seat belts. He noted that, although NHTSA performs an annual survey to assess how many people use their seat belts and conducts self-report surveys, it is difficult to get a real and accurate understanding of seat belt use without having objective data to work from. Researchers for the original 100-car study recorded seat belt use of participants over a year of driving and were able to classify people into three categories:

- » Consistent seat belt users (95 percent of trips or more).
- » Occasional seat belt users (40 to 85 percent of all trips).
- » Infrequent seat belt users (less than 35 percent of all trips).

The occasional group was of particular interest to researchers because it provided an opportunity to assess when and why driver behavior changed. One hypothesis for the variation in seat belt use between journeys included a change in use when driving was considered more risky (e.g., at night, or

in adverse weather conditions); however, people’s perceptions are not always accurate. For example, a driver is much more likely to get into a crash on a low speed road as opposed to the interstate, but the occasional group’s belt use was significantly higher on high speed, limited-access roads versus low speed roads. It becomes much easier for the safety community to do a better job to improve countermeasures to encourage seat belt use when it is known how and why drivers behave the way they do.

New Insights Into Driver Behavior

Compton stated that naturalistic driving data offer tremendous potential to provide new information and insights into how driver behavior and inattention are contributing to crashes. NHTSA wants to know what people do when they are driving and when and why this behavior changes. For example, the behaviors associated with speeding, aggressive driving, and drowsy driving are specific areas of interest that could tremendously benefit from the insight offered by naturalistic driving data. There is also an invaluable opportunity to help refine the metrics for vehicle safety systems. Additionally, safety researchers are able to gain new information on how people understand and interact with automated vehicles and collision avoidance systems.

Questions

One question addressed whether the study assesses how drivers respond to advanced avoidance warning systems. Compton confirmed that there was an initial intent to examine this in the original study design, but the study had trouble recruiting people with those systems within their vehicles.

Data Collection

*Jon Hankey, Senior Associate Director for Research and Development
Virginia Tech Transportation Institute*

Video Examples

Jon Hankey, Senior Associate Director for Research and Development at VTTI, provided several video examples of the VTTI data collection system. The video footage included various in-vehicle camera angles, including driving during the day and at night (using infrared cameras). One video example showed a forward-facing video of a distracted driver leaving a lane and striking a tire on the curb. Examples of these videos can be viewed at <http://forums.shrp2nds.us/>.

Securing Approval

Due to Institutional Research Board (IRB) regulations concerning personally identifiable information (PII), a researcher needs to obtain IRB approval to access the SHRP2 data. Study participants signed a consent form for future use of the data; however, the form states that data will always be used under IRB approval and a data-sharing agreement. To date nobody has had to undergo a full review. In addition to the video of the driver, GPS information has also been deemed to be personally identifiable information under certain situations. Data access differs according to how identifiable it is, and VTTI is therefore trying to make as much of it as unidentifiable as possible. It was noted that part of the workshop's mission is to move toward enabling a computer to perform the automated analysis of the data and bypass human involvement altogether.

Improved Data Access

Workshop participants were informed that it is hoped workshop discussion will help

gain a better idea of research needs and figure out a way to make the data available but still meet all the privacy requirements. Currently, the use of the sensitive SHRP2 NDS data is only permissible by physically traveling to a secure facility at VTTI; however, FHWA and SHRP2 are exploring the potential of "remote secure enclaves," which offer all of the essential privacy protections offered by the current situation. VTTI is also conducting a 24-driver experiment in which each driver is behind the wheel for 45 minutes using the same equipment and installation procedures as the larger study. These drivers will sign especially broad releases, providing easier access to the data than is permissible with the SHRP2 data. It should be noted that although this is a small sample relative to the SHRP2 data, it is substantially larger than the samples generally used in the publication of academic journal articles. Although there was substantial interest among the academics present in accessing the data, questions arose about the scalability of models based on small datasets to the analysis of thousands, or millions, of hours of video data. It was highlighted that appropriate validation will be essential.

Final Video Examples

Another demonstration video showed a rear-end collision. Researchers asked how to provide an indication of how fast the car was travelling when it ran into the back of the research vehicle. The ability to calculate that sort of information

would be very useful. Finally, the blurred photos of the vehicle interior taken every 10 seconds were shown to attendees. It was noted that it would be useful if there was an automated way to go through the blurred photos and count the number of passengers, identify

which seats were being used, whether seat belts were used, if a child seat was installed, and if occupants were male or female. Hankey also noted that the images were blurred intentionally in real time because IRB permission was obtained for drivers not passengers.

Video Analytics: Putting Driving in Context

John D. Lee, Professor, Department of Industrial and Systems Engineering
University of Wisconsin-Madison

Video Analytics

John Lee, a Professor at the University of Wisconsin-Madison's Department of Industrial and Systems Engineering, began by highlighting that video analytics has potential to be an amazing tool within the realm of naturalistic data. Importantly, it could help put driving into context, more so than any other analysis technique such as simulators. Lee noted that naturalistic driving data are interesting but it is also hugely challenging to make it scientifically

The data are there but extracting it manually would take centuries. This is, therefore, a coding challenge for researchers to develop ways to extract what is meaningful.

meaningful and useful. It was also noted that although video analytics can code driver "state," its bigger value may be about driving "context"—something that would provide great value to the safety community.

Overcoming a Data Deluge

Lee stated that naturalistic data are creating a deluge of data that is challenging to deal with. This issue is not specific to the driving community and is being witnessed in all domains with access to lots of data. Accordingly, there is much to be learned from others involved in these "big data" projects. For example, many of the big

data techniques that have been developed for the Large Hadron Collider project in Switzerland would be equally applicable to driving data analysis.

Lee also noted that there is a need to get around the "event centric" approach to obtaining data. Naturalistic data acquisition has mostly been examined from the perspective of event-triggered analysis but this is limiting because it provides a very limited context of what is going on around an event such as a rear-end collision. The data are there but extracting it manually would take centuries. This is, therefore, a coding challenge for researchers to develop ways to extract what is meaningful.

An example of where this has already been performed successfully can be seen in the work of Rob Radwin at the University of Wisconsin-Madison. Radwin developed video analysis capable of carrying out automatic measurement of hand activity in the workspace and producing, in real time, the risk posed to people doing manual labor tasks. This represents a prime example of a human factors expert collaborating with a video analytics expert to develop an application that works very well—a successful partnership that could also be applied to driving analysis.

Driver Distraction

Driver distraction is an issue of increasing concern as technology continually moves into cars and understanding distraction is a huge challenge with major implications.

Lee explained that driver distraction has previously been defined as “the diversion of attention away from activities critical to safe driving toward a competing activity.”

Looking at the road environment and the driver state are equally important, and video analytics provides a lot of potential for coding the driver state. Lee provided a list of requirements (a “wish list”) that would help characterize driver state and lead to better understanding of distraction:

- » Are there other passengers in the car and what is their age and gender?
- » What is the driver’s body pose and reach status?
- » What is the driver’s head pose?
- » Where is the driver looking?
- » What is the driver’s facial expression?
- » What is the driver’s pupil response?

Lee highlighted that video data could offer a lot of value to researchers but to understand distraction it is necessary to look at roadway demands in context of the competing demands of the task.

Distraction Analysis

Lee stated that the SHRP2 distraction analysis looks at continuous severity measure where severity is defined by safety margin and injury risk, defined by the field of safe travel. Video analytics can look at the driving environment from the driver’s perspective and recreate some of these elements to better understand what the driving environment and context are that might combine with the distractions to reduce

safety. Lee explained that understanding the road context can include the following:

- » What is the speed limit?
- » Is the car ahead braking or not?
- » What is the traffic signal state?
- » What is the road type and level of service?
- » What is the traffic level surrounding the driver?
- » How frequently does the driver have to look at the road in front to maintain safety?
- » What is the attention field?

Computational Models

Lee noted that computational models of road context, in terms of visual attention, are capable of predicting where people are going to look based on the characteristics of various scenes. It is possible to predict where the eyes are likely to fall, what the driver sees, and where the driver is likely to look in a scene. This sort of analysis could be applied to video data that are coming out of SHRP2.

In addition to analyzing the facial expression of drivers, Lee noted that cars also give clues to their driver’s behavior through external “facial” expression. Drivers communicate through their cars using turn signals, lights, and horns. If video analytics can tap into that communication it may be possible to gain a deeper understanding of what is going on, as opposed to just looking at the driver’s face.

Disciplinary Myopia

Lee informed workshop participants that there are thousands of potential projects to undertake using video analytics, some offering huge safety benefits, some with very small safety benefits. Some projects may be interesting from a purely video analytics perspective, some may be of interest only from a safety perspective. He noted that because the really interesting driving safety problems may be uninteresting from a video analytics perspective, this is a challenge that needs to be addressed. Video analytics needs to be considered as one of several approaches to deal with the naturalistic data deluge. For example, an exclusive focus on video analytics to identify driver state may miss the bigger opportunity to look at the driver context. Therefore, efforts to avoid disciplinary myopia need multidisciplinary teams to include driver behavior, unstructured data design experts, and video analytics experts.

Questions

During questions that followed the presentation, the importance of context was raised with the example of how driver demand on a familiar road will be completely different to someone tackling a road for the first time. It was also noted that there is a hierarchy of needs required from video analytics, from the simplest to the most complex. The simplest need is establishing where the person is looking at any given time. This can be extended to whether they are happy or sad—things that are not definable from an engineering perspective. Manufacturers and safety experts also have needs in different places than video analytics people, whose interest in serving those needs may be very small. Effective matchmaking is therefore critical.

It was discussed that the possibility of developing an avatar that could be placed on faces, to retain head pose and eye direction, would solve a lot of problems, and access to data and resulting privacy issues could be minimized.

Challenges of Video-Based Human-Computer Interfaces and Experiences in Analyzing Videos of Driver Faces

Margrit Betke, Professor, Computer Science Department
Boston University

Video Analytics Experience

Margrit Betke, a Professor at Boston University's Computer Science Department, provided a brief background on human-computer interaction (HCI) and tracking experiences, specifically in the fields of multiobject tracking and video-based HCI. The concept of the camera mouse was then explained. This is a video-based interface for non-verbal people with severe motion impairments. The system tracks body features and movements and converts them into mouse pointer movements. It is then possible to use this to interact with on-screen keyboards and other communication interfaces. Development of this technology involved a lot of facial analysis but it has proved highly successful. Since 2007, there have been over 1 million downloads of the free mouse software from www.cameramouse.org.

Developing Computer Vision

In 2000, Betke originally examined active computer vision in vehicles, specifically in terms of studying drivers and traffic in realistic conditions. The goal at that time was to develop a system that locates the face and eyes of a driver in realistic conditions and in real time. Betke made driving videos with students but several challenges emerged. These included when operating at night, adjusting cameras after leaving a garage, and handling light blooming effects, as shown in figure 6. There were also considerable difficulties producing results in real time with the technology available in 2000, as shown in figure 7.

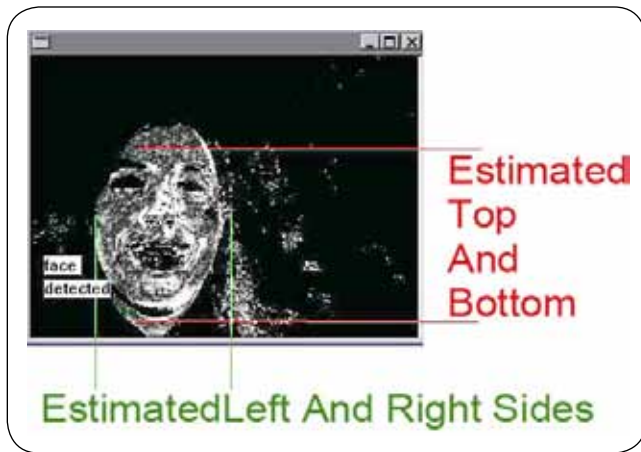


Figure 6. Photo. Light "blooming" can impede facial recognition.

Research Challenges

Betke informed workshop participants that there was a long pause in this research from 2000 to 2011. This was due to several reasons, in particular a lack of funding opportunities that support basic science approaches and also problems reproducing previous papers due, in part, to privacy. It was also difficult to obtain the large amounts of data required for such research.

The benefit of adopting multiple cameras that have views in different spatial domains was noted by Betke for creating three-dimensional (3D) gesture analysis. This involves using stereoscopy (also known as 3D imaging) to see what users are doing. Betke highlighted that a lot of motion can be lost using a single camera; but with 3D analysis, test subjects were able to move a pointer around the screen with their facial gestures using 3D-trajectory analysis, as illustrated in figure 8.



© Margrit Betke

Figure 7. Screen Capture. A facial recognition system from earlier research.

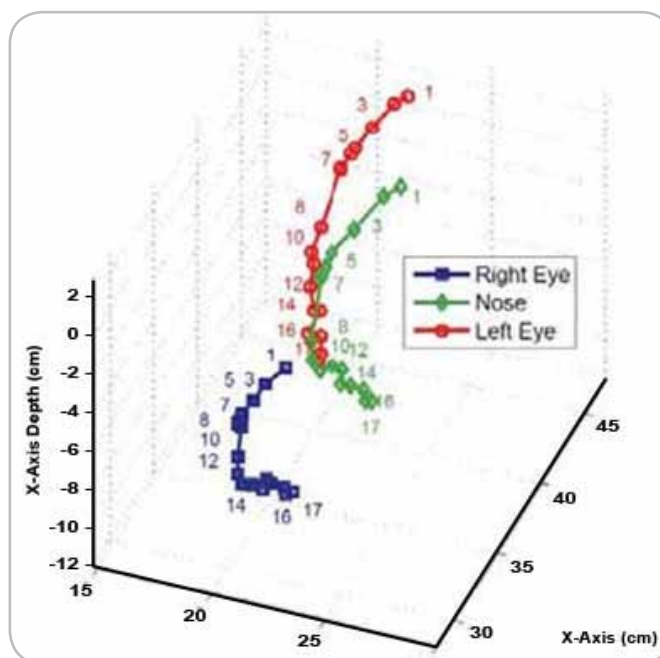
Recommendations

Betke concluded with a number of recommendations. Recording the driver's face with two or three calibrated cameras was suggested as a potentially good idea for the future naturalistic driving studies. This would enable researchers to create a 3D reconstruction and also make the implementation of an avatar much easier. Such a multi-camera system would also offer a level of redundancy in case of occlusion of the driver's facial features.

Another recommendation was to ensure that all cameras used in a vehicle are calibrated spatially and synchronized temporally.

This includes the cameras with the fields of view of the road behind and in front of the vehicle, and the cameras pointing to the steering wheel, the driver's face, and the passenger seats. Spatial and temporal camera calibration will enable research of correlations of events inside and outside the vehicle.

Potential database characteristics were also discussed. Although including basic factors such as age, gender, and hair type are important, other features such as personality should also be covered. For example, some people are very still behind the wheel, others move a lot, so tracking these varying parameters will be very different and should be catered for. Other suggested factors to consider include what the car interior is like, the lighting and weather conditions outside the vehicle, and various passenger configurations within the vehicle.



© Margrit Betke

Figure 8. Chart. 3D trajectory analysis.

Automatic Video-Based Driver State Analysis and Recognition

Qiang Ji, Professor, Department of Electrical, Computer, and Systems Engineering
Rensselaer Polytechnic Institute

System Overview

Qiang Ji, a Professor at Rensselaer Polytechnic Institute's Department of Electrical, Computer, and Systems Engineering, said that his research has been underway for 10 years, with a focus on real-time driver-state monitoring and recognition systems. The current system uses multiple cameras to monitor the driver and acquire video from different angles. The system uses narrow-view cameras to focus on the eyes and wide-view cameras to focus on the face and upper body. This video data are then processed to recognize different driver behaviors, particularly focusing on facial and body behavior. These behaviors are categorized and then fed into a probabilistic model which integrates the parameters with contextual information to provide a comprehensive and robust characterization of the driver's state, as illustrated in figure 9. The system is able to continually monitor the driver and decide what is the best information to provide to keep them safe and productive.

Monitoring Multiple Behaviors

Ji noted that it is important to monitor multiple behaviors. Accordingly, several visual behaviors are monitored—including eyelid movement, head movement, eye gaze, facial expression, and upper body movement. Computer vision methods include multiview face detection and tracking, eye detection and tracking, facial feature point tracking, head-pose estimation and tracking, facial expression analysis and recognition, eye-gaze tracking, and upper-body tracking and gesture recognition. Multiview face detection and tracking is considered to be very important because the face of the driver is not always front facing and therefore needs to be tracked from different angles.

Detecting and Tracking a Range of Movements

Video demonstrations of face detection and tracking using different face angles showed the system working in real time and successfully detecting and tracking faces with significant movement and expression

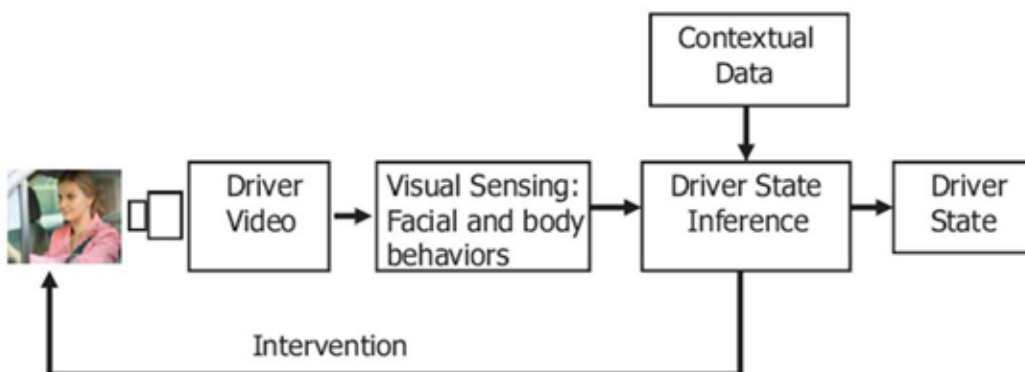


Figure 9. Diagram. System approach.

© Qiang Ji

changes. Another video showed that the analysis can also be applied outdoors to detect faces and facial orientation as people walk by.

Eye detection and tracking forms a key part of this research, and two specific techniques are employed by the system. One technique detects and tracks eyes during the day under normal sunlight conditions, another is able to detect and track eyes under poor illumination at night using infrared cameras. The system has proved capable of automatically detecting above 99 percent of eyes from the detected face, day or night.


Facial feature point detection and tracking are another key component of this research. Using a system of 28 points located around each major facial component—near the eyes, mouth, and eye brows—the system is able to detect and then track specific movements. It automatically locates the

face, normalizes the image, and then superimposes facial feature locations onto the live face image, as shown in figure 10. The 28 points also enable the system to estimate head pose, and a technique has been developed to work out the 3D angle in real time.

Characterizing Driver State

Another component of the system involves facial expression analysis—enabling the system to specifically characterize the state of the driver. Ji explained that much work has been done in this area, specifically focusing on two areas for analysis: (1) global facial expression analysis—recognizing six facial expressions (e.g., happy or sad); and (2) local facial recognition—capturing each muscle movement around the eyebrows, mouth, and cheek to indicate the state of the person.

Ji noted that it is a challenge to perform this subtle local recognition; however,



© Qiang Ji

Figure 10. Photo. Facial feature detection automatically superimposes 28 points onto a face.

to mitigate this challenge a probabilistic model has been developed for inference. As a result, the system is able to automatically detect a person's expression in real time and indicate the orientation of the face.

After visual expression detection, the next step is to perform eye detection. A driver's eye gaze can reveal the intent of a driver and indicates the focus of his attention. The eye-gaze tracking is designed to be able to operate under natural head movement with minimum personal calibration. The goal is to determine the visual axis of the person and use that to determine the line of sight. To do this, a narrow-angle camera focuses on the eyes and infrared detects cornea reflection to reveal the center of the cornea, as shown in figure 11. This can then

be connected with the center of the pupil to detect the optical axis. The system then intersects the visual axis with the optic axis to produce the gaze point.

Monitoring Driving Behavior

In addition to monitoring facial expressions, Ji explained that by using cameras to track arm and hand movement the system also could recognize upper body gestures. It can use this information to recognize whether the driver is performing normal driving behaviors or eating, using the phone, applying makeup, texting, or adjusting the radio. So far the system has been successfully applied to several different real-world experiments, including those investigating fatigue, stress, workload, and distraction factors.

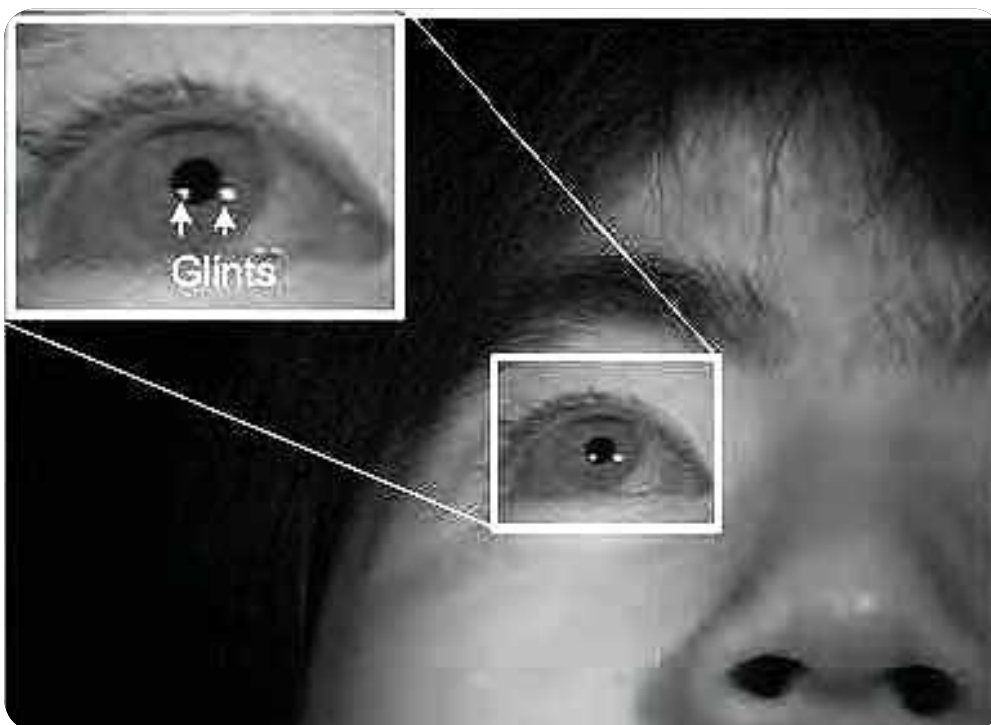


Figure 11. Photo. Infrared is used to detect cornea reflection.

Our Roadmap of Automatic Feature Extraction for Driving Safety

Yuanqing Lin, *Department Head, Department of Media Analytics
NEC Laboratories America*

Background

Yuanqing Lin, Department Head of Media Analytics at NEC Laboratories America, began by explaining that the NEC Laboratories America research team is relatively new to the field of driver safety, having been initially inspired by the Google automated car project. Lin explained there are currently nine full-time researchers in the media analytics lab with a mission to solve fundamental problems in computer vision and develop state-of-the-art systems. A strong collaboration with universities also exists, and the major research direction is investigating recognition and 3D reconstruction. Lin noted that driving safety is considered an interesting project because it specifically integrates lots of technologies from both of these research areas.

Sensing Levels

Lin said that sensing is a two-stage problem that requires low-level sensing and high-level understanding, both inside and outside the car. For the low-level sensing, there are many different sources of data that can be established. These can be from onboard sensors monitoring inside elements such as the gas or brake pedals, or from monitoring driver features such as eyes, head movement, or foot pose. Another factor involves looking at the outside surrounding environment and using 3D reconstruction to detect factors such as the lane, road, nearby pedestrians, bicycles, or cars.

Lin explained that with the low-level sensing complete it is then possible to

move to high-level understanding to establish inside factors such as a driver's mental state and behavior, and also understand outside factors such as the road scene. Although low-level sensing can be performed without driver video data, the data could still prove to be very helpful. Video from naturalistic driver study data can play a critical role in high-level understanding of behaviors.

Automatic Feature Extraction

Lin said that there is a need to determine which low-level features are useful for understanding certain types of behavior. Lin said that not a lot of work has so far been conducted by NEC Laboratories America on high-level understanding—partly because of a lack of access to the data. Accordingly, the current focus is on the ongoing work in low-level sensing, specifically focused on event detection, image recognition, and object detection.

NEC Laboratories America has conducted research on using automated detection to pick out people performing tasks, such as operating a cell phone or pointing. A key feature of the technology behind NEC Laboratories America's success is feature extraction, which takes place unsupervised using local coordinate coding or super-vector coding.

The existing state-of-the-art method for detecting objects such as cars, bikes, and pedestrians is to use a deformable part-

based model—a class of detection algorithm in which various parts of an image are used separately to determine if and where an object of interest exists. This method of detection has had a low object detection success rate, lower than the NEC Laboratories America’s object-centric pooling method.

Once objects have been successfully detected, Lin noted that the SHRP2 naturalistic driving study data will be very important to achieve the high-level of understanding. It may then be possible to figure out a real danger from the low-level sensing data. For example, high-level understanding could potentially establish that an accident or near accident happened in a similar case and automatically extract a danger scene from the available information.

3D Reconstruction

Lin highlighted that one particularly important ingredient required to put detected objects into 3D worlds is the concept of 3D reconstruction. Participants were shown a single camera video demonstration of real-time structure-from-motion being used for 3D reconstruction. The technique uses a challenging real-world dataset, the KITTI Vision Benchmark Suite (an open-access software suite from the Karlsruhe Institute of Technology), to evaluate interference from pedestrians, other cars, large illumination changes, large speed variations, and other environmental factors. This enables comprehensive evaluation for rotation and translation errors.

Vision for Driver Assistance: Looking at People in a Vehicle

Mohan M. Trivedi, *Professor, Laboratory for Intelligent and Safe Automobiles
University of California at San Diego*

Opening Remarks

Mohan Trivedi, a Professor at the Laboratory for Intelligent and Safe Automobiles (LISA) at the University of California at San Diego, began with a few key remarks. These included that successful computer vision is difficult to realize and that researchers need to recognize that vision is purposeful, and it is important to understand a picture in context.

A Holistic Approach to Driving

Trivedi highlighted that driving is very complex and the most dangerous act the average person does every day. Capturing naturalistic driving is a very different task to capturing driving behavior in laboratory or simulator conditions. Trivedi noted that driving is not just made up of one act at one time but is made up of three different types of tasks:

- » Strategic tasks
- » Tactical tasks
- » Operational tasks

Robustness and reliability are considered the hardest part of such a system because it needs to perform critical tasks day in and day out. Addressing this is a research problem, and metrics are needed to establish reliability measures and prove exactly how good something is. This raises the question of how to define the metrics that are associated with the cameras and sensors put in a car. For example, what are the performance parameters for face recognition?

Research Focus

Trivedi noted that very specific experiment design and data requirements need to be established. Clarifying the experiment objectives is an important first step in the design. For example, driving consists of a lot more than just lane changes but this is still a critical issue and driving component that needs to be looked into in detail. It illustrates the point that researchers must sometimes narrow their focus and hone in on one element.

According to Trivedi, a simple search of papers published in the Institute of Electrical and Electronics Engineers Transactions on Intelligent Transportation Systems revealed that papers and citations referencing “driver, eyes, and cameras” produced 11 papers per year from 2000 to 2004; 35 per year from 2005 to 2009; and 80 per year from 2010 onwards. This represents a clear trend and growing interest in the area.

LISA’s Research Agenda

Trivedi explained that the LISA research agenda adopts a multidisciplinary focus on development of a complete driving context capture system. Research includes robust computational algorithms for context and intent analysis, detailed behavioral analysis of driver and driving tasks, mental models for attention and multitasking, and multimodal interfaces for driver attention management.

LISA’s first project, in 2000, was to examine automobiles as a context-aware

space. Researchers placed a camera and a cellphone in the vehicle system to inform a caller if it is the right time to call or not. This experiment addressed how to use a camera to tell what the surrounding information is indicating. The second project was with Volkswagen and required the development of smart airbags. A robust real-time vision system was created for sensing occupant body posture in vehicles and providing safe airbag deployment.

Trivedi stated that a lot of previous studies about attention and drowsiness were looking at eyes. This is generally considered a hard problem; however, an alternative approach suggested by Trivedi is to try to solve the easy problems first. For example, eyes are part of the head so the head is a good starting point for cameras to focus on. If focusing on the head cannot provide everything that is needed to detect drowsiness, for example, then it may be necessary to move on and then look at the eyes.

PART TWO: DISCUSSION

For the second part of the workshop, attendees were given an opportunity to review items discussed during part one and then focus on data needs and requirements. Attendees then participated in general group discussion to identify challenges and put forward possible solutions.

Developing Tools

Lincoln Cobb, Technology Facilitator at FHWA's Office of Safety Research and Development, began by highlighting that in addition to the valuable NDS data, there is also the equally useful roadway information database running in parallel and providing helpful context. Cobb stated that this other rich resource is useful both on its own and for integrating with the NDS data.

Cobb explained that it is exciting to learn more about the breadth of applications for these types of data and that there are many research questions. As the data improve, more will be collected and the need for automated extraction will only get bigger. Cobb stated that looking at roadway safety applications is the most relevant use for the data; however, automated feature extraction applications under discussion cross all modes—wherever human beings are controlling the system these approaches will be valid (e.g., in marine and freight train areas).

SHRP2 is a year away from completing NDS data collection, and there are ongoing discussions between TRB, FHWA, NHSTA, and AASHTO looking at ways to make the data more useful and easily accessible. Specific items being looked at include reduced datasets that would be easy to access and answer limited questions with; using richer trip headers; and basic linkage between trip files and road segments is also being discussed.

Cobb informed workshop participants that the workshop is part of the effort to support the development and deployment of tools to make the extraction of information

economical. There are two groups of users to cater to: (1) researchers who want to take the data and apply them to a specific project; and (2) tool developers. Researchers need easy access to clear data so there are several privacy and security details to overcome to reach this point given the strong protections in place for the subjects. For example, Cobb highlighted that the current situation is a long way from being able to send 2 petabytes of data to a researcher via the internet. He confirmed that USDOT is looking for guidance in this regard and would like to identify what makes the data useful. Questions to be addressed include:

- » What reduced datasets would be of particular use to researchers?
- » What user tool will be helpful?
- » How can access be improved?

Cobb confirmed these items are being discussed, highlighting that the more economical, timely, and automated the process of data extraction the better for all. Any progress that researchers can make with tools that look at the data being collected now will be of particular value in the medium term. Cobb confirmed that tools for the short term are just as important as the long term. He noted that partial solutions that are deployable can be extremely helpful—reiterating that a comprehensive solution for automated data extraction is not essential at this time.

Cobb concluded with a reminder of some of the key observations to emerge from researchers at the workshop. These included

the theory that context is crucial—it is important not to just focus on any one thing or center of your interest. Cobb also noted that a great deal of research is underway and there are many opportunities for collaboration. He highlighted that progress in automation seems to be moving ahead in many fields and research focus areas but the basic research that will support the development of deployable tools is also interesting.

Advanced Research Agenda

Next, David Kuehn provided some concluding remarks from the perspective of the FHWA's EAR Program. Kuehn noted the value of the multidisciplinary and multidomain solutions discussed during the workshop and confirmed that this is an approach that fits well within the context of EAR Program-sponsored projects. He noted that automated video analytics research is important for saving lives and reducing injuries and hopefully will also be something of interest to the larger research community.

Kuehn highlighted that it would be useful to capture why the recent advances in machine learning mean it is possible to solve issues now that were not possible 5 to 10 years ago. He also noted that it would be useful to know some of the approaches that would be most effective moving forward and what else needs to be known to advance.

Kuehn stated that there appears to be some uncertainty about what the ideal features are for extracting. He noted that it would be useful to identify the salient features for driver safety that could easily be extracted

for short-term solutions. Finally, the issue of context was highlighted—specifically, addressing the balance of how much supporting context is needed to go with extracted features so data can be easily sent out to investigators as soon as possible.

State of the Practice

Wider group discussion initially focused on the state of the practice in feature extraction and how it can be applied to feature extraction problems. It was noted that although some features are easy to obtain, it is more challenging if a precise and exact orientation is required. The level of detail depends on the quality of image, desired features, and the purpose of those features. It was also noted that it would be useful to start out with a hierarchy of needs, even though currently not all of those needs can be met—for example, it would be useful to know exactly where a person is looking but eye trackers do not currently exist to obtain those data (although obtaining head pose is possible).

Another question addressed how much information can be obtained from feature extraction. For example, knowing if hands are on or off the steering wheel would be particularly useful; however, it was also highlighted that information must already be available within available naturalistic driving study data—if another camera is needed to collect those data, it is impossible to proceed in that direction.

One workshop participant explained that many hours are spent on crude video coding—for example, using head pose to calculate what the driver is doing. Despite this, coders cannot tell exactly where the

driver's eyes are directed because of issues such as resolution, wearing glasses, and glare. It was suggested that dividing the cockpit into regions and identifying where the driver is looking at that moment in time would be good enough to perform some analysis. It was then explained that it currently takes approximately 1 minute to code 10 minutes of video, so to code hours of video takes considerable time.

Group discussion suggested that having a tool that could do that operation as well, or better, than coders would be a big advantage. It was agreed that obtaining head pose from videos is currently possible and hand tracking can be performed very well with algorithms, so the technique would potentially be possible and reliably repeatable with such a computer-based approach. One final comment highlighted that a lot of open-source software already exists that can perform this automated processing but there is still a lot of research to be done depending on application.

Data Sources

Discussion also focused on the idea that the NDS dataset is rich with not just videos but other variables. It was suggested that when performing the video coding, if visual information is ambiguous about where a driver is looking it could be possible to go into other types of data to attempt to confirm where they are looking and disambiguate the video using the other data sources. Other data sources to help identify where a person is looking could include current speed, braking status, and weather conditions. An automated system capable of performing that task would be more in line with what a human does.

One approach could be to use an integrated dataset, comprised of video and extracted features, that takes the video and reduces it to a simplified dataset that can be combined with other data. It was noted that computer vision researchers prefer data that are not reduced, preferring to go through the original data and perform the reduction that would provide the necessary analysis to pass on for behavior analysis.

Discussion then moved to the requirement for multidisciplinary teams to connect with the needs of the analysis. A two-step process was put forward: (1) produce features and analysis out of videos; (2) datamine these for behavioral analysis. It was noted that a team effort would be required to pull these two stages together.

Integrating Data

One workshop participant stated that a lot of the current focus is about what happens inside a vehicle; however, researchers are also interested in events outside the vehicle. In particular, the roadway features are crucial so, using the front view camera, it would be useful to link roadway data and onboard vehicle data acquisition system data.

One possible application could be to examine electronic billboards that may be distracting people with motion. It was noted that it would be useful if there was a feature that could be pulled out of the video to show the difference between a static and animated billboard and objectively assess how long a driver stares at those. A question on this subject focused on how hard it would be to automatically decode something that is moving versus a static

sign. For example, would it be possible to automatically go through the video and find every sign, or to find someone doing certain maneuvers. It was explained that these are called gesture events, and there has been a lot of work in this field recently; however, event recognition is currently not as good as static object recognition.

Another suggestion was to develop the ability to pick out where a bicyclist is on the road to study how a driver interacts with a bicyclist, or bicycle lane, without trawling through trips manually. It was noted that these tasks are not something that have to be performed in real time and can all be processed offline after the data have been collected.

Data Reduction

One workshop participant stated that partial solutions are important and researchers should not attempt to process too much data on the first step. Features that are too high level should not be extracted at this time but should be saved for later. It was suggested that the first step should be to reduce the huge size of the video and remove the personally identifiable data and then let a researcher deal with low-level features. For example, just to focus on and study the eyes would be a huge reduction in the size of the data. Extracting some simple features would make the data more easily available.

Another workshop participant said that computer vision researchers want access to the original data. For example, with 1.5 million trip files, computer vision could automatically distinguish from a bike or a jogger using existing examples; however,

examples of where bikes show up in the data would be needed to train the computers to recognize that.

A potential application could be to study the difference between having a conversation with a passenger and using a hands-free device. It was suggested that it would be useful to objectively investigate if this changes where people look; however, to achieve this it will be necessary to identify when drivers are conversing within millions of trip files and then analyze the effect on driving. It was also noted that when using the term “reduced dataset,” it does not mean picking and choosing which elements are important but instead means picking entire segments in which there appear to be instances of something that is useful.

Advances In Machine Learning

Workshop participants proceeded to discuss some of the recent advances in machine learning over the past 5 years. For example, if researchers wanted to find vehicles that had specific hand-held devices in them, is there something that is common among them that machine learning could recognize? It was suggested that instead of attempting to recognize what type of device a driver is holding, it would be better to look for the action of picking up a device. It would then be possible to go to the database and retrieve all similar events and look at when this happens and at what time. An example was given that if people do not turn their phones off, there is a tendency to visually check to see who is calling even if the phone does not get answered. If glancing at the phone supports the

supposition that Bluetooth makes driving safer, it would be invaluable information for safety researchers and automakers.

Participants agreed that information on whether the driver is communicating, sitting, or stretching is all of use to researchers. It was noted that head pose and hand position can provide a good indication of what a driver is doing and a face camera can also say a lot about a driver. For example, an algorithm for facial expression could potentially identify and index each video sequence to state if the driver is happy or angry.

An important issue raised by workshop participants was the computational challenges involved in processing these data. The computational cost of running algorithms on such a large dataset requires an understanding of the machines and hardware required to process them. It was noted that although it takes time to perform offline learning to learn the features, once a feature is learned the classification does not take much time and can even be performed in real time by some systems. Even running classification algorithms in real time would take over a million hours of video, so simply attempting to index a selection of five features would take a long time. Extracting features will require large

computing infrastructure (such as multiple parallel processing) to process these big data, which relates to the wider issue of data reduction and segmenting the data into interesting pieces.

Next Steps

FHWA's Office of Safety R&D and SHRP2 will continue to pursue incremental and potentially breakthrough approaches to making the NDS data more accessible and useful. On the incremental end of the spectrum, possibilities include the concept of "remote secure enclaves;" reduced datasets for which privacy management issues could be mitigated; richer trip headers to help researchers identify data of interest; and the development of a database of linkages between trip data files and the roadway segments on which the trips occurred.

The EAR Program is considering plans for continuing to engage the community assembled at the workshop. This is anticipated to include some follow-up conversations and outreach in the upcoming year. In addition, the EAR Program is considering the release of a Broad Agency Announcement soliciting proposals for automated video feature extraction sometime in 2013.

Appendix

APPENDIX A WORKSHOP PARTICIPANTS

Randolph Atkins
Social Science Researcher
National Highway Traffic Safety
Administration

Margrit Betke
Professor
Boston University

David Bolme
Oak Ridge National Laboratory

Ann Brach
Director
National Academy of Science,
Transportation Research Board, SHRP2

Ken Campbell
Chief Program Officer
National Academy of Science,
Transportation Research Board, SHRP2

Zhaofu Chen
Research Assistant
Northwestern University

Anil Cheriyyadat
Research and Development Staff
Oak Ridge National Laboratory

Lincoln Cobb
Technology Facilitator
Office of Research and Development
Federal Highway Administration

Richard Compton
Director
Office of Behavioral Safety Research
National Highway Traffic Safety
Administration

Eric Cosatto
Senior Researcher
NEC Laboratories America

Matt Cuddy
Operations Research Analyst
Volpe Center

Monique Evans
Director
Office of Research and Development
Federal Highway Administration

Charles Fay
Senior Program Officer, Safety
National Academy of Science,
Transportation Research Board, SHRP2

Marie Flanigan
Principal Scientist
CUBRC

Lisandra Garay-Vega
Human Factors Engineer
National Highway Traffic Safety
Administration

Jurek Grabowski
Research Director
AAA Foundation for Traffic Safety

Pujitha Gunaratne
Senior Scientist
Toyota

Terry Halkyard
Transportation Specialist
Office of Research and Development
Federal Highway Administration

Jon Hankey
Senior Director
Center for Automotive Safety Research
Virginia Tech Transportation Institute

Qiang Ji
Professor
Rensselaer Polytechnic Institute

Thomas Karnowski
Research and Development Staff Member
Oak Ridge National Laboratory

Dev Kochar
Technical Expert
Ford

David Kuehn
Team Director/EAR Program Manager
Office of Research and Development
Federal Highway Administration

John Lee
Professor
University of Wisconsin-Madison

Yuanqing Lin
Department Head
Media Analytics
NEC Laboratories America

Brian Philips
Senior Research Psychologist
Office of Research and Development
Federal Highway Administration

Venkat Rajagopalan
Senior Robotics Engineer
Carnegie Mellon University

Terry Regan
Community Planner
Volpe Center

Jim Shurbutt
Research Psychologist
Office of Research and Development
Federal Highway Administration

John Sullivan
Assistant Research Scientist
University of Michigan Transportation
Research Institute

Brian Tefft
Senior Research Associate
AAA Foundation for Traffic Safety

Mohan Trivedi
Professor and Director
University of California, San Diego

Jing-Shiarn Wang
Mathematical Statistician
National Highway Traffic Safety
Administration

Jingshu Wu
Mathematical Statistician
National Highway Traffic Safety
Administration

Ken Wu
NRC Research Fellow
National Research Council
Federal Highway Administration

David Yang
Research Engineer
Office of Research and Development
Federal Highway Administration

References

1. Campbell, K. The SHRP2 Naturalistic Driving Study, Addressing Driver Performance and Behavior in Traffic Safety. In *TR News 282* (pp. 32-33). Transportation Research Board of the National Academies, Washington, DC, 2012. Retrieved December 29, 2012, from <http://onlinepubs.trb.org/onlinepubs/trnews/trnews282SHRP2nds.pdf>.
2. Lee, J., Regan, M., & Young, K. *Driver Distraction: Theory, Effects, and Mitigation*. CRC Press, 2008.

Office of Safety Research and Development
Turner-Fairbank Highway Research Center
6300 Georgetown Pike
McLean, VA 22101-2296

Publication No. FHWA-HRT-13-037
HRTM-30/12-12(WEB)E