

In Press,  
IEEE Pervasive Systems

# Facial Expression Analysis for Predicting Unsafe Driving Behavior

Maria E. Jabon, Jeremy N. Bailenson, Emmanuel Pontikakis, Leila Takayama, and Clifford Nass

***Abstract***— Pervasive computing provides an ideal framework for active driver support systems in that context-aware systems are embedded in the car to support an ongoing human task. In the current study, we investigate how and with what success tracking driver facial features can add to the predictive accuracy of driver assistance systems. Using web cameras and a driving simulator, we captured facial expressions and driving behaviors of 49 participants while they drove a scripted 40 minute course. We extracted key facial features of the drivers using a facial recognition software library and trained machine learning classifiers on the movements of these facial features and the outputs from the car. We identified key facial features associated with driving accidents and evaluated their predictive accuracy at varying pre-accident intervals, uncovering important temporal trends. We also discuss implications for real life driver assistance systems.

***Index Terms***—Computer vision, Face and gesture recognition, Real-time systems, Pervasive computing, Human safety

---

Traffic related accidents are recognized as a serious and growing problem around the world. Every year in the U.S. alone, more than 42,000 Americans die as a result of 6.8 million automobile accidents [1]. Consequently, driver safety technology has become an active area of research in both industry and academia. Pervasive computing environments, with integrated sensors and networking may provide an ideal platform for the development of such technology. Taking effective counter measures to enhance safe operation of a vehicle requires merging information from many diverse layers of the system. As a first step, an *active driver safety system* (a system designed to prevent accidents from occurring) must monitor vehicle state and/or vehicle surroundings. (See “Related Work in Driver Safety” pullout.) However, in order to fully transform the vehicle into a smart environment [2], the driver must also be monitored. Human factors researchers have long studied the driver’s role in causing and preventing accidents and have found that the driver’s physical and emotional state, including driver fatigue [3] and stress levels [4], play a role in a significant number of traffic accidents. Thus, many researchers have turned to developing active driver safety systems that monitor not only the vehicle, but the driver as well.

In the current work we investigate the performance of an active driver safety framework that captures both vehicle dynamics and the driver’s face, and merges the two levels of data to produce an accident prediction. We distinguish the current study from previous work in active driver safety in four ways: (1) we use a bottom-up approach, analyzing the movement of a comprehensive set of 22 raw facial features, rather than simply analyzing eye gaze or head orientation measures, (2) we evaluate a wide range of time and frequency domain statistics in order to determine the most valuable statistics for driving accident prediction, (3) we predict major and minor accidents directly, not intermediate driver states such as fatigue, and (4) we explore the predictive accuracy of the face and car outputs at varying pre-accident intervals, uncovering important temporal trends in predictive accuracy for each feature subset.

## EXPERIMENTAL TESTBED

For our study, 49 undergraduate students were recruited to drive through a 40-minute simulated course in a STISIM driving simulator. The driving simulator, developed by Systems Technology, Inc. [5], was set up to run on a single PC and project the simulated image of the roadway onto a white wall in the lab (Fig. 1).



Fig. 1. Participant being monitored in STISIM driving simulator

During the driving course, we projected the sounds of the car, the road, and events happening around the car into the simulator room via four PC speakers. The course simulated driving in a suburban environment with conditions varying from light to intense traffic. We included many challenging situations, including busy intersections, unsafe drivers, construction zones, sharp turns, and jaywalking pedestrians in an effort to increase the complexity of the drive (Table 1).

Miles	Environment	Speed Limit	Intersections	Traffic	Challenges
0-1	Suburban	35 mph	6	Heavy	Many pedestrians
2-6	Highway	65 mph	0	Moderate	Narrow roads, tight curves
7-9	Suburban	35 mph	15	Heavy	Many pedestrians
10-11	Highway	55 mph	1	Light	Construction zone, obstacles
12-18	Suburban	35 mph	24	Heavy	Many pedestrians
19-20	Rural	35 mph	0	Light	Dirt roads, obstacles
21-22	Rural	35 mph	1	Moderate	Narrow roads, tight curves
23-32	Urban	55 mph	6	Moderate	Tight curves

Table 1. Simulation parameters

We opted for the context of a virtual driving simulator instead of real cars in order to safely collect a large sample size of accidents to use in our analyses. This allowed us to generate separate models for major accidents (e.g., hitting pedestrians, other vehicles, or off-road objects) and minor accidents (e.g., unwarranted change of lanes, driving off the road, or running a stoplight).

During the experimental sessions we recorded participants' faces with two Logitech web cameras at a rate of 15 frames per second. We compressed the AVI format videos in real-time using DirectX and DivX technology. Although many technologies exist for capturing face movements, we opted for image-based capture because it does not require special markers or user intervention. This makes our system less intrusive, increasing system transparency. We also recorded the output of the simulator during the driving sessions, which was a text-based log-file that contained the road conditions, steering wheel angle, lane tracking information, car speed, longitudinal acceleration (feet/second<sup>2</sup>) due to pedal presses and braking, braking information, number of accidents, and type of accident.

## ANALYSIS PROCEDURE

Using the collected face videos and the driving simulator data, we constructed a dataset to build our computational models. Figure 2 summarizes the phases of data analysis.

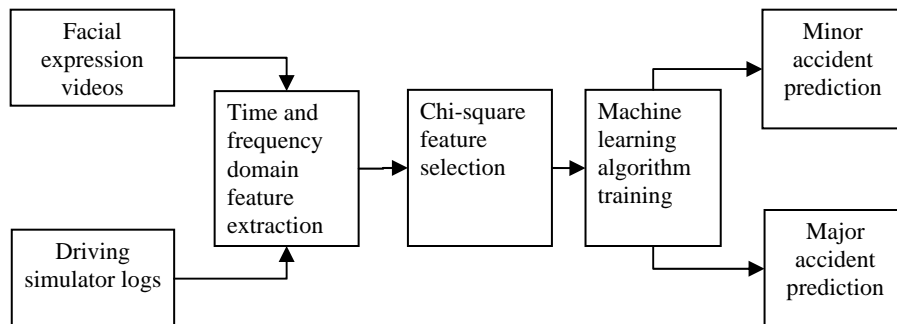


Fig. 2. Data analysis procedure

### A. Facial Feature Extraction

The first step in the construction of our datasets included post-processing the videos we collected to extract key facial features and head movements. For this processing we used the Neven Vision library [6]. For each frame of the video the Neven Vision library automatically detects, without any preset markers worn on the face, the  $x$  and  $y$  coordinates of 22 points on the face, eye and mouth openness levels, and head movements (e.g., yaw, pitch, and roll). This is done at a rate of 30 frames per second. In Figure 3, we

present a screenshot of the Neven face tracking software.




Fig. 3. Neven Vision tracking points on subject's face

### *B. Data Synchronization*

In the next phase of our analysis we synchronized the video outputs with the driving simulator outputs so that we knew when accidents occurred within the video and could extract the pre-accident facial geometry and vehicle information. Our goal was to determine the optimal way to combine the facial movements and vehicle outputs to predict when accidents would occur. Thus, we sampled a number of different pre-accident time intervals, including intervals beginning between one and four seconds before the accident and ranging from one to ten seconds in length. For each interval, we extracted the data for every major and minor accident in our dataset, along with a random number of non-accident intervals to use in our analyses.

### *C. Time Series Statistics Calculation*

After data synchronization, we computed a series of time-domain statistics on the coordinates in each interval to use as inputs to our classifiers. We calculated averages, velocities, maximums, minimums, standard deviations, and ranges. These values were computed for each of the Neven outputs and for the vehicle outputs such as speed, wheel angle, throttle, and braking outputs. For some important facial

characteristics, such as eye and mouth openness levels, we also created five-bin histograms from 0% to 100% to capture the distribution of eye and mouth states over the time interval.

#### *D. Frequency Domain Statistics Calculation*

We then calculated frequency domain statistics on each facial coordinate as well as for each car output. We used the MATLAB Wavelet toolbox to perform the discrete Wavelet transform, in particular the Daubechies Wavelet family with orders one, two, and four. For each order, we performed a level three decomposition of the input signal and collected statistics over the detail coefficients of each level including averages, ranges, histograms, and variances. We calculated these additional statistics because facial signals are dynamic, and we expected that their micro-momentary movements can leak information about the internal state of the person making the expression.

#### *E. Final Dataset Creation*

Finally, we combined all of the pre-accident and non-accident intervals for each of the 49 participants. To improve the reliability of the face measurements, we discarded intervals where the average face-tracking confidence (i.e., the measure of how confident the face tracking software was in its measurement) was lower than 60%. Our resulting dataset contained a total of 179 minor accident instances, 131 major accident instances, and 627 non-accident instances.

#### *F. Chi-Square Feature Extraction*

Initially, our datasets consisted of 7402 facial features and 1162 vehicle features for each pre-accident and non-accident vector. To speed up the training of our algorithms and to identify which facial features were the most important indicators of unsafe driving behavior, we performed a chi-square feature selection for each dataset. The top 20 car features and top 20 face features for our major and minor accident predictions are presented in Tables 2 and 3.

	<b>Feature</b>	<b>Statistic</b>	<b>Value</b>
--	----------------	------------------	--------------

<b>Face</b>	Lower Lip Center X	Wavelet 62	33.3
	Upper Lip Center Y	Wavelet 62	26.6
	Right Lower Lip X	Wavelet 62	25.8
	Right Pupil Y	Wavelet 9	25.3
	Right Pupil Y	Wavelet 39	25.3
	Right Upper Lip X	Wavelet 62	23.2
	Left Eye Aspect Ratio	Wavelet 88	23.0
	Left Inner Eye Brow Y	Wavelet 55	20.9
	Lower Lip Center Y	Wavelet 51	20.2
	Right Inner Eye Brow Y	Wavelet 55	20.2
	Left Nostril Y	Wavelet 5	18.4
	Left Mouth Corner X	Wavelet 72	17.7
	Right Eye Aspect Ratio	Wavelet 78	17.1
	Face X	Wavelet 42	16.6
	Left Eye Aspect Ratio	Wavelet 53	15.8
	Head Euler X	Wavelet 89	15.3
	Right Pupil Y	Wavelet 84	15.1
	Left Lower Lip Y	Wavelet 51	15.1
	Face Scale	Wavelet 55	15.1
	Right Lower Lip Y	Wavelet 80	14.8
<b>Car</b>	Velocity	Average	161.0
	Velocity	Minimum	154.7
	Velocity	Maximum	124.5
	Steering Wheel Angle	Wavelet 40	88.8
	Braking Input	Wavelet 65	86.9
	Steering Input	Wavelet 71	83.8
	Steering Wheel Angle	Wavelet 3	82.7
	Steering Wheel Angle	Velocity	78.9
	Steering Wheel Angle	Wavelet 43	78.1
	Steering Wheel Angle	Wavelet 72	76.2
	Steering Wheel Angle	Wavelet 35	76.1
	Steering Wheel Angle	Wavelet 41	76.1
	Steering Input	Wavelet 44	74.2
	Steering Wheel Angle	Wavelet 2	74.1
	Braking Input	Wavelet 42	73.9
	Steering Input	Wavelet 40	73.5
	Steering Wheel Angle	Wavelet 42	73.3
	Steering Input	Wavelet 37	72.5
	Steering Wheel Angle	Wavelet 11	72.5
	Steering Wheel Angle	Wavelet 37	73.2

Table 2. Top 20 facial and vehicle features for major accident prediction

	<b>Feature</b>	<b>Statistic</b>	<b>Value</b>
<b>F</b>	Right Outer Eye Corner Y	Minimum	41.6
	Mouth Aspect Ratio	Wavelet 57	33.6



	Left Outer Eye Corner Y	Maximum	32.6	
	Left Pupil Y	Average	31.5	
	Left Outer Eye Corner Y	Average	28.0	
	Left Pupil Y	Maximum	28.0	
	Left Inner Eye Brow X	Minimum	27.1	
	Left Pupil X	Minimum	26.4	
	Nose Tip Y	Minimum	26.2	
	Right Nostril X	Maximum	25.9	
	Nose Root X	Minimum	25.3	
	Nose Tip X	Maximum	25.3	
	Right Nostril X	Average	25.0	
	Face X	Velocity	24.7	
	Left Eye Brow Center X	Minimum	23.9	
	Left Outer Eye Corner X	Minimum	23.5	
	Left Inner Eye Corner X	Wavelet 28	23.3	
	Right Nostril Y	Minimum	23.1	
	Left Inner Eye Brow X	Average	23.0	
	Left Inner Eye Corner X	Minimum	22.4	
	<b>Car</b>	Steering Wheel Angle	Maximum	143.9
		Steering Wheel Angle	Velocity	136.7
Steering Input		Velocity	128.4	
Steering Wheel Angle		Wavelet 3	122.0	
Steering Input		Maximum	119.7	
Steering Input		Wavelet 3	110.3	
Steering Wheel Angle		Variance	102.2	
Steering Wheel Angle		Wavelet 4	98.4	
Steering Wheel Angle		Wavelet 69	98.3	
Steering Wheel Angle		Wavelet 70	96.1	
Steering Wheel Angle		Average	96.0	
Steering wheel Angle		Wavelet 2	93.5	
Steering Wheel Angle		Wavelet 72	93.4	
Steering Input		Variance	93.4	
Steering Input		Wavelet 4	89.9	
Steering Wheel Angle		Wavelet 34	89.4	
Steering Wheel Angle		Wavelet 44	89.4	
Steering Input		Range	87.8	
Steering Wheel Angle		Wavelet 42	87.3	
Steering Wheel Angle		Wavelet 41	85.4	
Steering Wheel Angle	Wavelet 12	85.2		

Table 3. Top 20 facial and vehicle features for minor accident prediction

We found that the facial features most predictive of major and minor accidents differed greatly; most of the top major accident features were movements of points around the mouth and eyes, while over 25% of the top minor accident features were around the nose. This differs significantly from previous works, where one features around the eyes and mouth were used to predict dangerous driving states [3-4]. Less of a difference existed between the top major accident car features and top minor accident car features; in both cases the steering wheel angle and steering inputs made up over 75% of the top 20 features.

We also found the most useful facial feature statistics varied across accident type; wavelets proved the most useful statistics for major accident prediction, whereas simple minimums and maximums were more informative for minor accident prediction.

## CLASSIFICATION AND RESULTS

We experimented with numerous state of the art classifiers to predict driving accidents, including Bayesian Nets, Decision Tables, Decision Trees, Support Vector Machines, Regressions, and LogitBoost simple decision stump classifiers [7]. We evaluated our classifiers according to Cohen's kappa, which corrects for the degree of agreement between a classifier's predictions and reality by considering the proportion of predictions that might occur by chance [8]. This measure has been shown to be more robust than simple measures such as hit rate or overall accuracy [8]. Values range from zero to one with a score of zero implying completely random classification and a one implying perfect classification. Generally, a kappa score of greater than 0.2 is considered statistically significant [9].

### *A. Predicting Minor Accidents*

We first attempted to predict just minor accidents (i.e. centerline crossings, tickets, and road-edge excursions) in order to determine the face's role in minor accident predictions. Our minor-accident datasets consisted of 806 instances: 179 minor accident instances and 627 non-accident instances. We trained five classifiers, a Support Vector Machine classifier with a poly-kernel, a LogitBoost classifier with the weak classifier to be a simple decision stump, a Multilayer Perceptron Neural Net, a Decision Table, and a Logistic Regression. We built all classifiers for these datasets using the publicly available tool Waikato Environment for Knowledge Analysis (WEKA) [10] and validated our models using a ten-fold cross validation. The LogitBoost classifier provided the highest kappa statistic of the five sampled classifiers. Figure 4 presents the results of the LogitBoost classifications across pre-accident intervals ranging from one to four seconds pre-accident and using one to ten seconds of data.

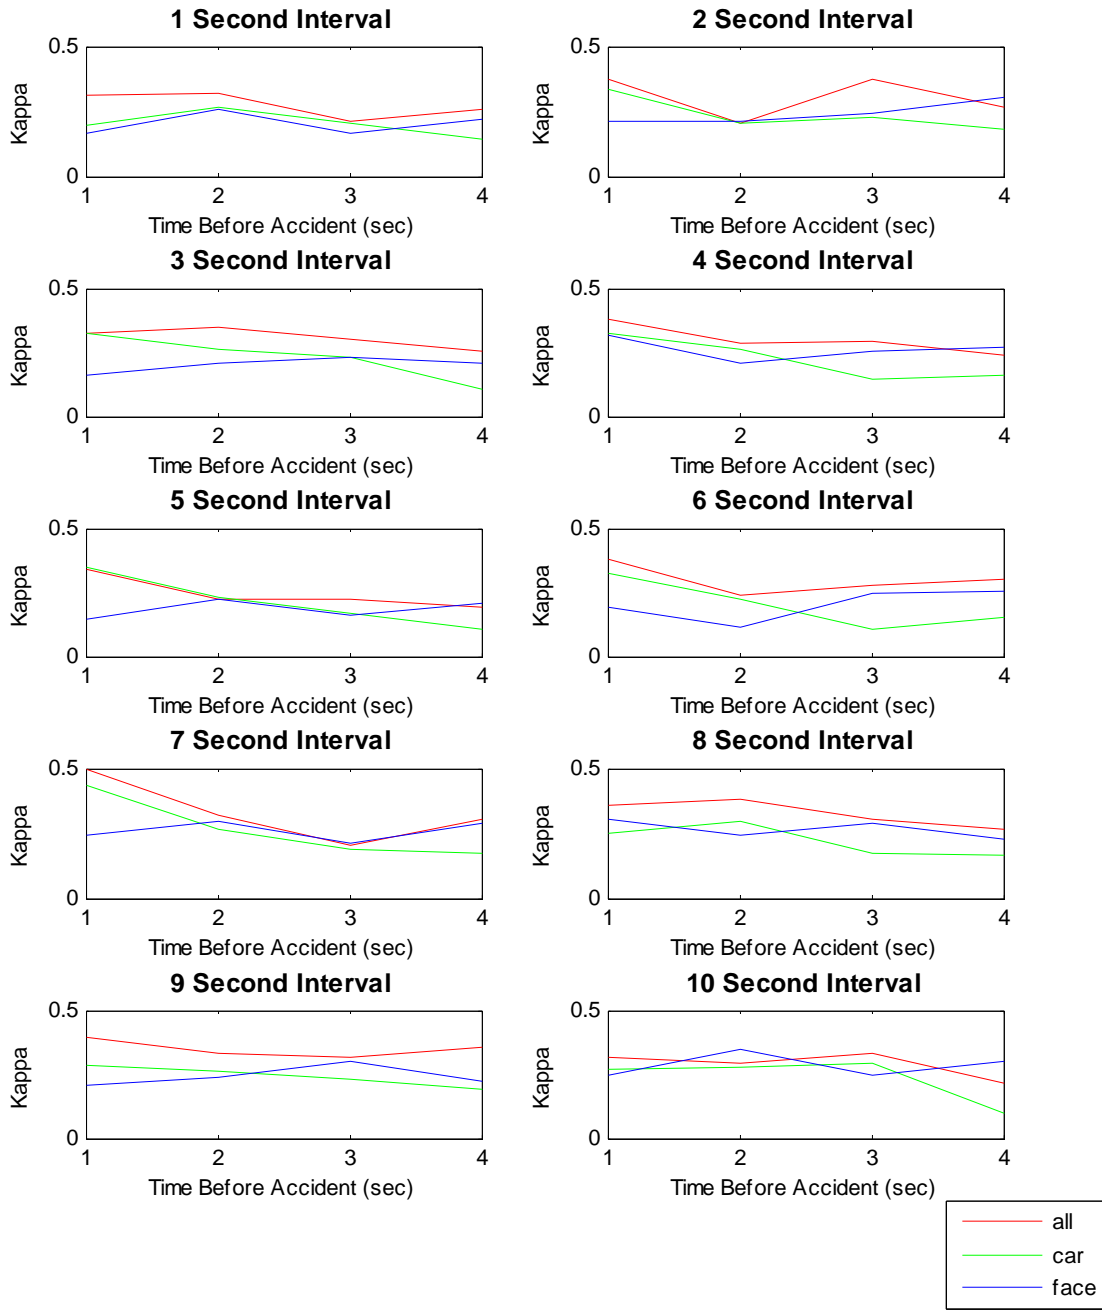


Fig. 4. Performance of classifiers using face, car, and all inputs to predict minor accidents one to four seconds before accidents occur using statistics calculated over one to ten second long intervals of data. Kappa values over 0.2 are considered statistically significant.

An interesting trend appeared when analyzing the various pre-accident intervals used in minor accident prediction; while the car features proved more useful in predicting the accidents close to the accident (one to two seconds before), the face features proved more predictive longer before the accident (three to four seconds before). In fact, in all of the intervals we analyzed, classifiers using the facial features

outperformed the classifiers using the car features at four seconds prior to the accident. Furthermore, at four seconds pre-accident, these classifiers outperformed the classifiers using all of the features in four of the ten intervals. This suggests that facial features, especially those around the nose, could significantly improve the accuracy of driver safety systems further prior to accidents, thus allowing drivers more time to react and prevent the accident.

In order to further analyze the temporal trends in predictive accuracy of our classifiers, we plotted ROC curves depicting true vs. false positives for the classifiers using all features, only car features, and only facial features, at one to four seconds before accidents occurred. We chose to use seven seconds of data for each plot given that our highest accuracies occurred using seven seconds of data. We present these plots in Figure 5.

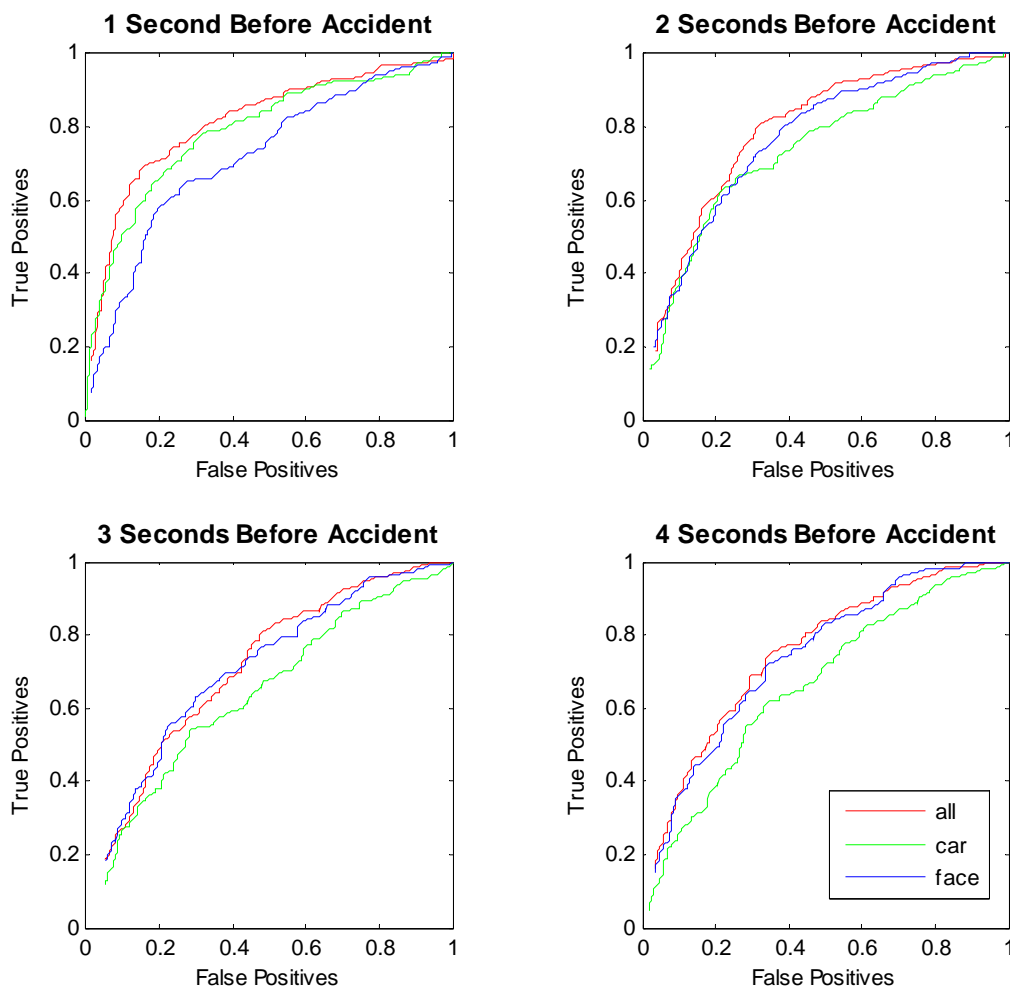


Fig. 5. ROC plots of classifiers using all, car, and face features to predict minor accidents one to four seconds before accidents occur

We note that the ROC curves for face improve to be almost equal to the curves using all the data by four seconds prior to the minor accidents. This implies that by four seconds prior to the accident the signal that provides the bulk of predictive power comes from the face. We also note that in all cases the ROC curves using all the features provides the best overall tradeoff between true and false positives. From this we conclude that the face provides a signal absent within the vehicular features and that and that this signal continues further before the accidents than the signal provided by the car features.

### *B. Predicting Major Accidents*

We next attempted to predict just major accidents (i.e., hitting objects or pedestrians) in order to determine the face's role in major accident predictions. Our major-accident datasets consisted of 758 instances: 131 major accident instances and 627 non-accident instances. We again trained five classifiers, a Support Vector Machine classifier with a poly-kernel, a LogitBoost classifier with the weak classifier to be a simple decision stump, a Multilayer Perceptron Neural Net, a simple Decision Table, and a Logistic Regression. We built all classifiers for these datasets using the publicly available tool Waikato Environment for Knowledge Analysis (WEKA) [10] and validated our models using a ten-fold cross validation. Again, the LogitBoost classifier provided the highest kappa statistic of the five sampled classifiers. Figure 6 presents the results of the LogitBoost classifiers across pre-accident intervals ranging from one to four seconds pre-accident and using one to ten seconds of data.

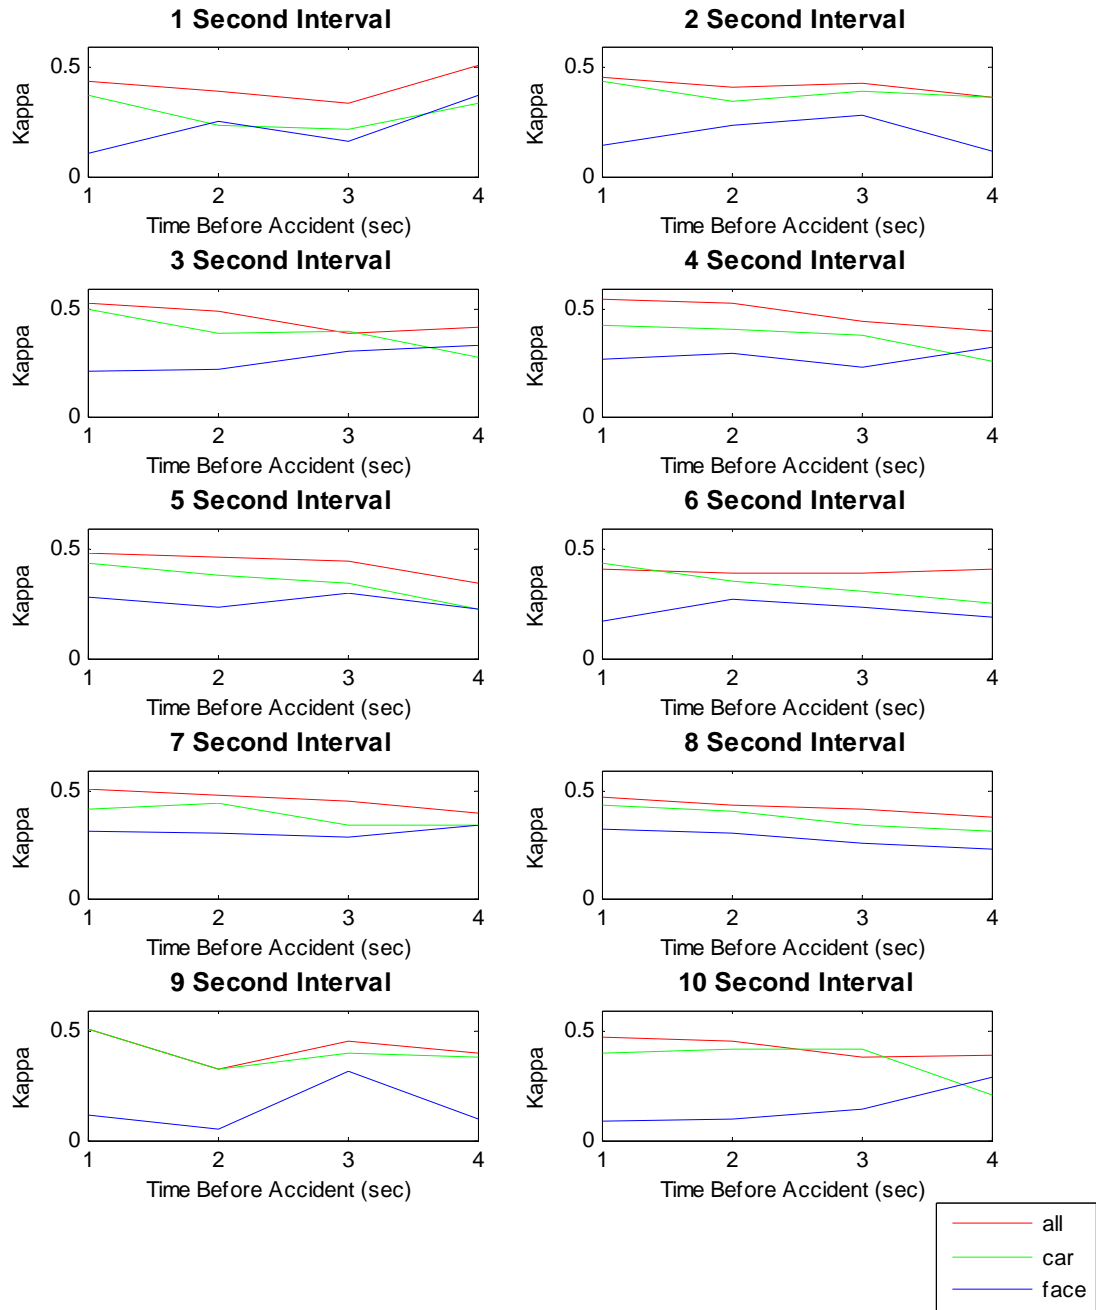


Fig. 6. Performance of classifiers using face, car, and all inputs to predict major accidents one to four seconds before accidents occur using statistics calculated over one to ten second long intervals of data. Kappa values over 0.2 are considered statistically significant.

Unlike in the case of minor accidents, where we observed a marked increase in the predictive power of facial features as the time before the accident increased, we see less of a trend in the predictive power of the face features over the car features at various time intervals for major accidents. However, again, we

note that the classifiers that use facial features in combination with the car features consistently provided a higher classification kappa than either the classifiers that used facial features alone or the classifiers that used car features alone. This suggests that the vehicle features may demonstrate more signal than facial features in the case of major accidents, but that facial features can still be used to improve the overall performance of the classifiers.

### *C. Predicting Minor and Major Accidents*

We next attempted to predict minor and major accidents together. We did this by combining all major, minor, and non-accident instances into one large dataset and creating classifiers on this comprehensive dataset. We trained five classifiers, a Support Vector Machine classifier with a poly-kernel, a LogitBoost classifier with the weak classifier to be a simple decision stump, a Multilayer Perceptron Neural Net, a simple Decision Table, and a Logistic Regression and validated our models using a ten-fold cross validation. Once again the LogitBoost classifier provided the highest kappa statistic of the five sampled classifiers. Thus Figure 7 presents the results of the LogitBoost classifications across pre-accident intervals ranging from one to four seconds pre-accident and using one to ten seconds of data.

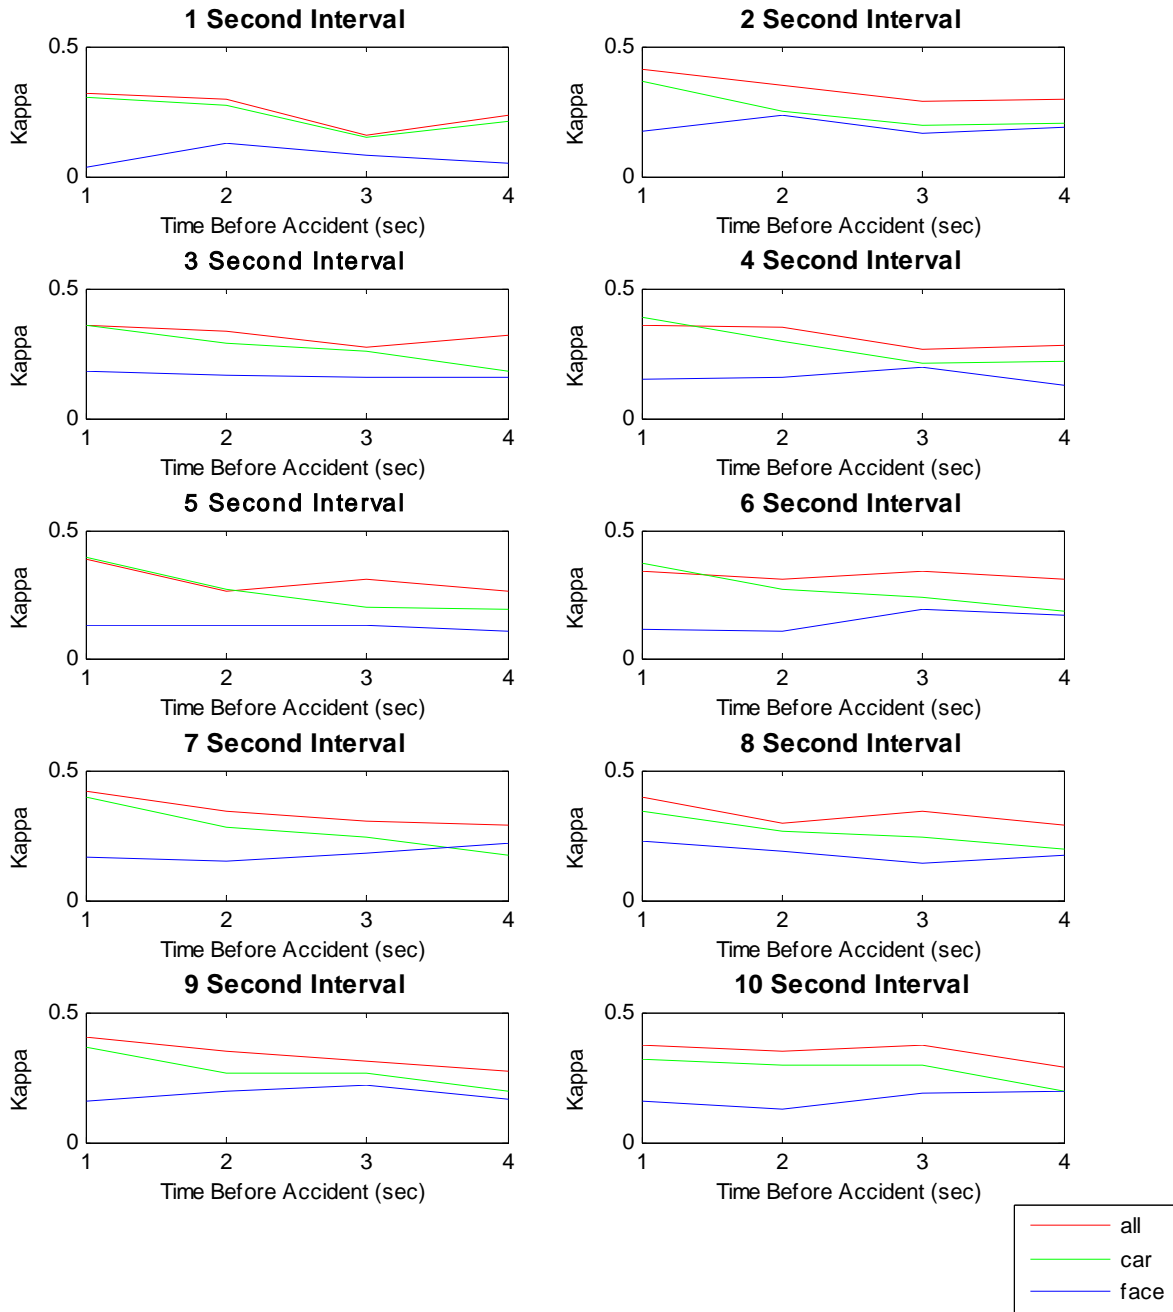


Fig. 7. Performance of classifiers using face, car, and all inputs to predict major and minor accidents one to four seconds before accidents occur using statistics calculated over one to ten second long intervals of data. Kappa values over 0.2 are considered statistically significant.

As in the case of both the minor and major accident predictions, the classifiers using all the features performed the best across sampled intervals. We also observe that the performances of the classifiers using the facial features alone tend to remain steady even to four seconds pre-accident, whereas the performances of the classifiers using only the car features tend to fall off. For a closer view of these trends we plotted



ROC curves depicting true vs. false positives for the classifiers using all features, only car features, and only facial features, at one to four seconds before accidents occurred. We did this for each class (major and minor) in isolation. In this way we could examine whether the facial signals were more predictive of major or minor accidents across the varying pre-accident time intervals. We chose to use seven seconds of data for each plot given that our highest accuracies occurred using seven seconds of data. We present these plots in Figures 8 and 9.

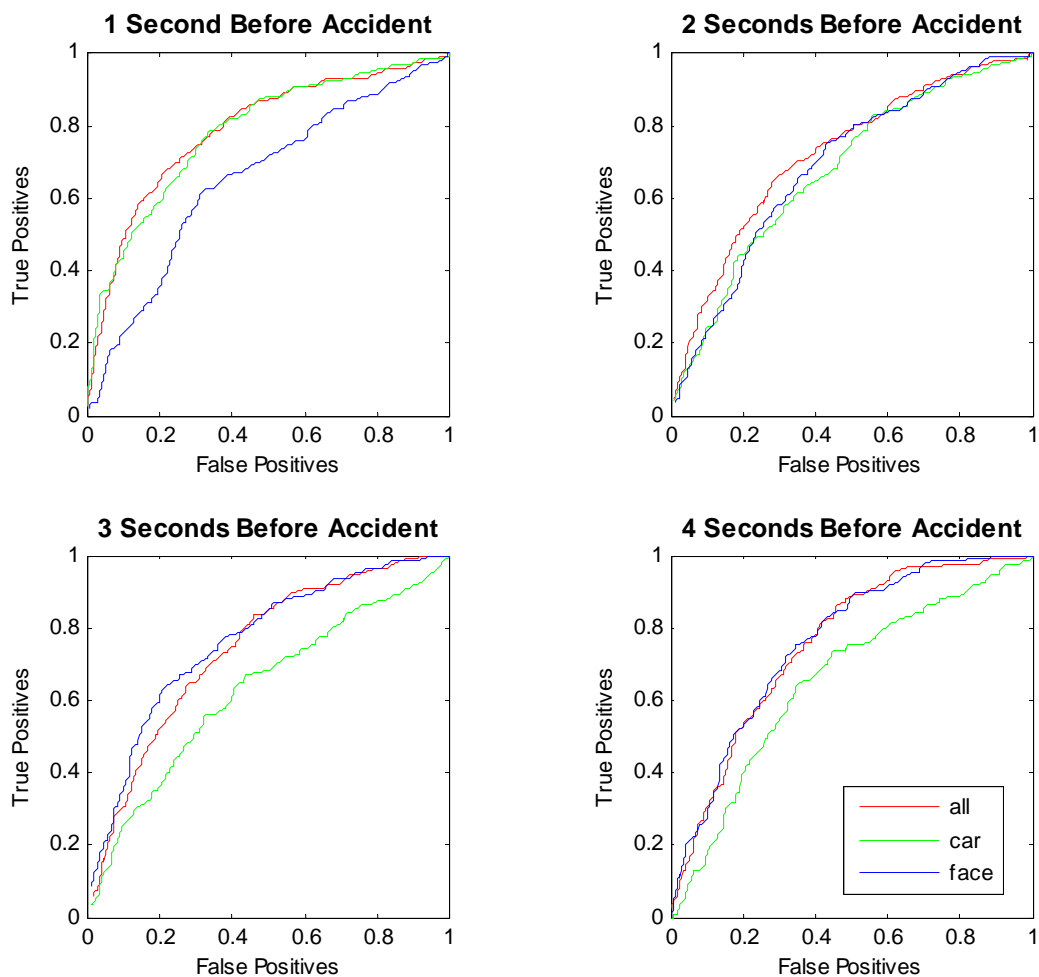


Fig. 8. Minor accident ROC curves using all, car, and face features at varying pre-accident time intervals.

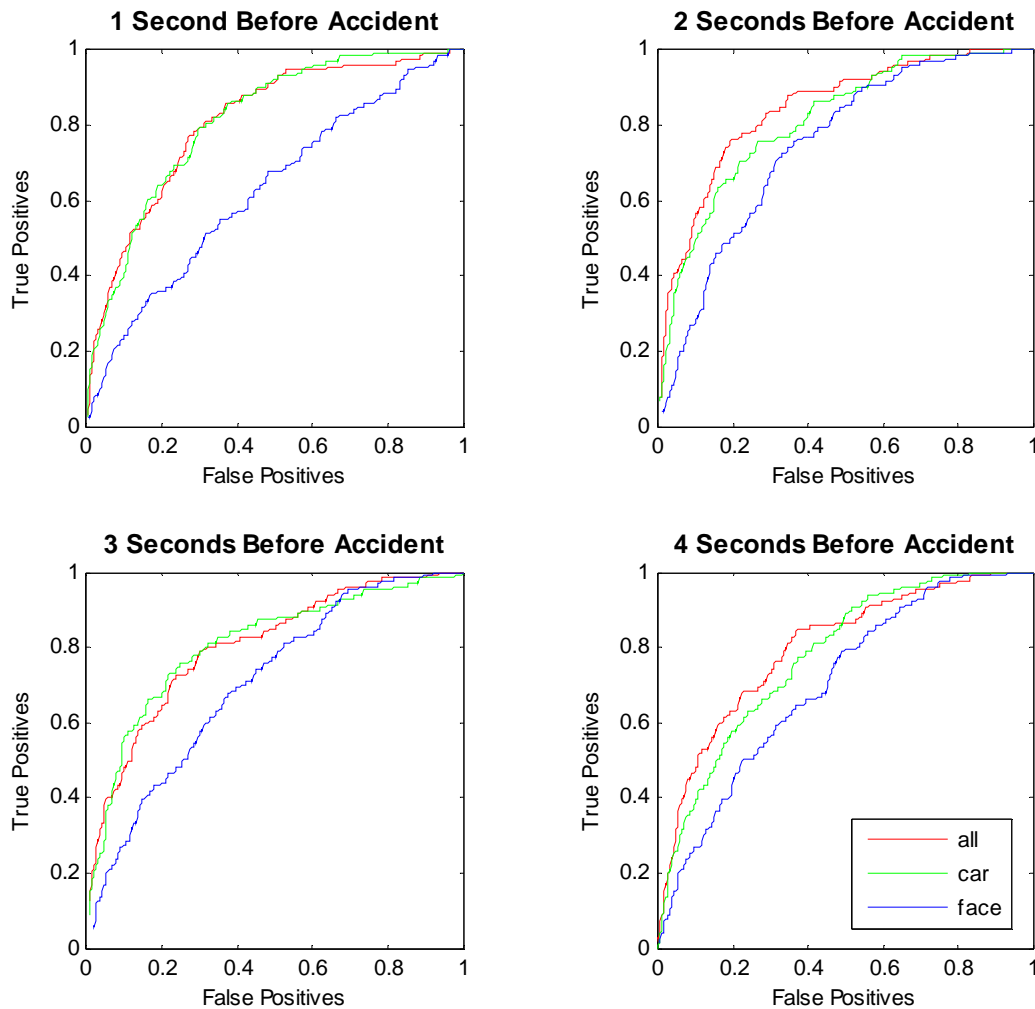


Fig. 9. Major accident ROC curves using all, car, and face features at varying pre-accident time intervals.

When viewing the ROC curves for the minor accident predictions it becomes apparent that the face accounts for most of the predictive accuracy of the classifiers after three seconds prior to the accidents; the classifiers using only the facial features perform essentially the same as the classifiers using all the features in combination. However, the predictive accuracy for major accidents appears to come from primarily the vehicle features. This confirms the results we saw in our binary classifiers, where the facial features proved more helpful in predicting minor accidents than major accidents. Overall, this suggests that important signals for accident prediction exist in drivers' faces up to four seconds prior to accidents, and that these signals are strongest in the case of minor accidents.

## LIMITATIONS

Despite the encouraging results of our study, there are several limitations to note. First, we refrained from running this study on the road because the study design was deemed too hazardous to be run in the physical world until we had first run the study in a simulator. Thus, our study ignores important parameters such as vehicle motion and surround effects, which affect driver's perception and reaction to situations. This limits the generalizability of the study findings, and further work must be done to determine the impact of these parameters.

Second, although this study thoroughly investigated ways of sensing impending accidents on the road, it did not investigate exactly what a pervasive system could do in order to prevent that accident; the system could notify the driver, take action on its own, or do some combination of the two. The actions that the system takes to prevent the accident are open for future work.

Third, we did not actually implement our system in real-time. Thus we could not analyze the false alarm rate over time for our classifiers. However, given that Neven Vision is capable of processing streaming facial video at a rate of 30 frames per second, and that our feature calculations and classifier predictions can be made in under 500 ms using MATLAB on a standard 2.5 GHz Intel Processor, we are confident that our system is capable of being run in real-time.

Fourth, as with any statistical model, these results are limited to the specific features that this study included in the original models and to this particular data set. If we had sensed other aspects of the driver (e.g., heart rate) or any other part of the driver-environment system, we might have generated very different models for predicting impending driver accidents. Thus, future work would benefit from including a wider range of sensor data to improve the accuracy of such driver safety support systems. Similarly, generating more data sets that include other populations of participants and other driving contexts would improve the generalizability of the study.

## CONCLUSION AND FUTURE WORK

Integrating information from the environment (e.g., weather, traffic, hazards) with information from the car (e.g., speed, acceleration) and the driver (e.g., facial features) proves a promising way to leverage today's sensor technologies to actively support driver safety. By identifying the most useful features for major and minor accidents and exploring the best temporal windows in which to use those features, we show that facial features can improve the predictive of classifiers up to four seconds prior to driving accidents. This suggests that the use of facial feature sensors along with vehicle sensors can improve the performance of active driver support systems in the future.

## ACKNOWLEDGMENTS

We would like to thank Dan Merget for software development, Rabindra Ratan for contributing on setting up the driving simulator, Alexia Nielsen, Karmen Morales, Julia Saladino, and Anna Ho for running this experiment, and David Jabon, Grace Ahn, Doug Blumeyer, Steffi Paepcke, Jeff Wear, Kathryn Segovia, and Jesse Fox for helpful comments on earlier drafts of this paper. This work was funded in part by NSF grant SES-0835601.

## REFERENCES

1. U.S. Department of Transportation, "Traffic Safety Facts 2006: A compilation of motor vehicle crash data from the fatality analysis reporting system and the general estimates system," National Highway Traffic Safety Administration Report DOTHS 810 818. Washington, DC, 2006.
2. M. Satyanarayanan, "Pervasive Computing: Vision and Challenges," *IEEE Personal Comm.*, vol. 8, no. 4, 2001, pp. 10-17.
3. Zhiwei Zu, Qiang Ji, and Peilin Lan "Real time non-intrusive driver fatigue monitoring," *IEEE Trans. Veh. Technol.*, vol. 53, 2004, pp. 1052-1068.
4. J. A. Healy and R. W. Picard, "Detecting Stress During Real-World Driving Tasks Using Physiological Sensors," *IEEE Transactions on Intelligent Transportation Systems*, vol. 6, no. 2, 2005, pp. 156-166.
5. R. W. Allen et al., "A Low Cost PC Based Driving Simulator for Prototyping and Hardware-in-the-Loop Applications," *SAE Paper No. 98-0222*, Spec. Pub. 1361. 1998.
6. H. Neven, NevenVision [Computer Software], NevenVision, Inc., Santa Monica, CA, 2003.
7. J. Friedman, T. Hastie, and R. Tibshirani, "Additive logistic regression: A statistical view of boosting," *Annals of Statistics*, vol. 28, 2000, pp. 337-407.
8. A. Ben-David, "What's wrong with hit ratio?," *IEEE Intelligent Systems*, vol. 21, no. 6, 2006, pp. 68-70.
9. J. Landis and G. Koch, "The measurement of observer agreement for categorical data," *Biometrics*, vol. 33, 1977, pp. 159-174.
10. H. I. Witten and E. Fank, *Data Mining: Practical Machine Learning Tools and Techniques*, Morgan Kaufmann, 2005.

## Related Work in Active Driver Safety Systems

Much work has been done in the area of holistic vehicle sensing for active driver safety, including systems for monitoring vehicle environment, vehicle state, and more recently driver state. Systems developed to monitor vehicle environment include pedestrian and obstacle detectors, lane-guidance systems, rear-bumper proximity sensors, blind-spot car detectors, automatic windshield wipers, and surround imaging systems for parking assistance<sup>1-6</sup>. Systems for monitoring vehicle state include systems to track vehicle location via GPS and accelerometers and other sensors to monitor driving speed, steering wheel angle, breaking, and acceleration<sup>7-8</sup>. Systems for monitoring driver state include frameworks that gauge driver fatigue, drowsiness, or stress level<sup>9-14</sup>. In the current paper, we extend previous work by using computer vision algorithms to directly map specific facial features to unsafe driving behavior. We use a comprehensive set of raw facial feature points, including points around the nose that are absent in prior works. Furthermore, we do not infer any specific mental states such as fatigue, but rather implement a more empirical approach that uses machine learning algorithms to find and use the facial features that are the most correlated with accidents. In addition, we identify important trends in predictive accuracy for each feature subset at various temporal windows, showing how the face can best be used to improve the predictive accuracy of classifiers up to four seconds prior to accidents.

## References

1. L. Li et al., "IVS 05: New Developments and Research Trends for Intelligent Vehicles," *IEEE Intelligent Systems*, vol. 20, no. 4, 2005, pp. 10-14.
2. T. Gandhi and M. Trivedi, "Vehicle Surround Capture: Survey of Techniques and a Novel Omni Video-Based Approach for Dynamic Panoramic Surround Maps," *IEEE Trans. Intelligent Transportation Systems*, vol. 7, no. 3, 2006, pp. 293-308.
3. J. McCall and M. Trivedi, "Video-Based Lane Estimation and Tracking for Driver Assistance: Survey, System, and Evaluation," *IEEE Trans. Intelligent Transportation Systems*, vol. 7, no. 1, 2006, pp. 20-37.
4. M. Bertozzi, A. Broggi, and A. Lasagni, "Infared Stereo Vision-Based Pedestrian Detection," *Proc. IEEE Intelligent Vehicles Symposium*, IEEE Press, 2005, pp. 24-29.
5. Y. Suzuki, T. Fujii, and M. Tanimoto, "Parking Assistance Using Multi-Camera Infrastructure," *Proc. IEEE Intelligent Vehicles Symposium*, IEEE Press, 2005, pp. 106-110.
6. H. Kurihata, T. Takahashi, and I. Ide, "Rainy Weather Recognition from In-Vehicle Camera Images for Driver Assistance," *Proc. IEEE Intelligent Vehicles Symposium*, IEEE Press, 2005, pp. 205-210.
7. J. Waldo, "Embedded Computing and Formula One Racing," *IEEE Pervasive Computing*, vol. 4, no. 3, 2005, pp. 18-21.
8. M. Bertozzi et al. in B. Appolloi et al, (Eds.), *Knowledge-Based Intelligent Information and Engineering Systems*, Springer Berlin, 2007, pp. 704-711.
9. M. Trivedi, T. Gandhi, and J. McCall, "Looking-In and Looking-Out of a Vehicle: Computer-Vision-Based Enhanced Vehicle Safety," *IEEE Transactions on Intelligent Transportation Systems*, vol. 8, no. 1, 2007, pp. 108-120.
10. A. Williamson and T. Chamberlain, *Review of on-road driver fatigue monitoring devices*, NSW Injury Risk Management Research Centre, University of New South Wales, 2005.

11. J. A. Healy and R. W. Picard, "Detecting Stress During Real-World Driving Tasks Using Physiological Sensors," *IEEE Transactions on Intelligent Transportation Systems*, vol. 6, no. 2, 2005, pp. 156-166.
12. A. Doshi and M.M. Trivedi, "On the Roles of Eye Gaze and Head Dynamics in Predicting Driver's Intent to Change Lanes," *IEEE Transactions Intelligent Transportation Systems*, vol. 10, no. 3, 2009, pp. 453-462.
13. L. Fletcher, L. Petersson, and A. Zelinsky, "Road Scene Monotony Detection in a Fatigue Management Driver Assistance System," *Proc. IEEE Intelligent Vehicles Symposium*, IEEE Press, 2005, pp. 484-489.
14. Q. Ji, Z. Zhu, P. Lan, "Real Time Non-intrusive Monitoring and Prediction of Driver Fatigue," *IEEE Transactions on Vehicular Technology*, IEEE Press, 2004, pp. 1052-1068.



## Author Information

Maria E. Jabon  
311 Glenview Rd.  
Glenview, IL 60025  
(650) 269-3521  
[mjabon1@gmail.com](mailto:mjabon1@gmail.com)

Maria Jabon earned her B.A. and M.S. degrees from Stanford University in Electrical Engineering in 2009. While at Stanford she was Lead Systems Engineer in the Stanford University Virtual Human Interaction Lab and a Research Assistant in the Stanford University Flow Cytometry Lab. Currently she is employed as a Software Developer at LinkedIn. Jabon's research interests include machine learning and web interfaces for the visualization of high dimensional data sets.



Jeremy N. Bailenson  
Rm.344 McClatchy Hall  
Stanford, CA 94305  
650.723.0701

[bailenso@stanford.edu](mailto:bailenso@stanford.edu)  
[stanford.edu/~bailenson](http://stanford.edu/~bailenson)  
[vhil.stanford.edu](http://vhil.stanford.edu)

Jeremy Bailenson is founding director of Stanford University's [Virtual Human Interaction Lab](#) and an associate professor in the Department of Communication at Stanford. He earned a B.A. cum laude from the University of Michigan in 1994 and a Ph.D. in cognitive psychology from Northwestern University in 1999. After receiving his doctorate, he spent four years at the Research Center for Virtual Environments and Behavior at the University of California, Santa Barbara as a Post-Doctoral Fellow and then an Assistant Research Professor. Bailenson's main area of interest is the phenomenon of digital human representation, especially in the context of immersive virtual reality. He explores the manner in which people are able to represent themselves when the physical constraints of body and veridically-rendered behaviors are removed. Furthermore, he designs and studies collaborative virtual reality systems that allow physically remote individuals to meet in virtual space, and explores the manner in which these systems change the nature of verbal and nonverbal interaction.



440 N Wabash #4508  
Chicago, IL 60611  
(312) 307-5130  
[pontikakis@gmail.com](mailto:pontikakis@gmail.com)

Emmanuel (Manos) Pontikakis earned his B.Sc. degree in Computer Engineering and Informatics from University of Patras in Greece and his M.S. degree in Computer Science from Stanford University. While at Stanford he was a Research Assistant in the Stanford University Virtual Human Interaction Lab. During his graduate studies he has interned as a Software Engineer for Google and he has been a Research Affiliate for Intelligent Data Understanding Group at NASA Ames Research Center. His research interests include privacy-preserving data mining, machine learning, human computer interaction and quantitative finance. Currently he is employed as a Quantitative Analyst at Citadel Investment Group.



[takayama@willowgarage.com](mailto:takayama@willowgarage.com)

Willow Garage, 68 Willow Road, Menlo Park, CA 94043  
(650) 475-9764

Leila Takayama is a research scientist at Willow Garage, studying human-robot interaction. Dr. Takayama completed her PhD in the Communication between Humans and Interactive Media (CHIME) lab at Stanford University in 2008. She also holds a PhD minor in Psychology from Stanford, MA in Communication from Stanford (2008), and BAs in Psychology and Cognitive Science from UC Berkeley (2003). She studies people's responses to, interactions with, and interactions through personal robots, focusing on issues of embodied cognition and perceived agency. She is a member of the Association for Psychological Science, Cognitive Science Society, and ACM Special Interest Group on Computer-Human Interaction.



Clifford Nass

Department of Communication  
Room 300E, McClatchy Hall  
Stanford University  
Stanford, CA. 94305-2050  
nass@stanford.edu  
(650) 723-5499

Clifford Nass earned his Ph.D. and M.A. in Sociology, and his B.A. *cum laude* in Mathematics at Princeton University. Nass is currently the Thomas M. Storke Professor at Stanford University; he has been a professor at Stanford since 1986. His primary appointment is in Communication, but he also has appointments by courtesy in Computer Science; Education; Science, Technology, and Society; Sociology; and Symbolic Systems (cognitive science). Nass's research focuses on (laboratory and field) experimental studies of social-psychological aspects of human-interactive media interaction. Nass founded and directs the Communication between Humans and Interactive Media (CHIME) Lab. Nass is also co-director of the Kozmetsky Global Collaboratory and Auto-X at Stanford University. Research interests include automobile interfaces, mobile interfaces, human-robot interaction, cognitive and social effects of multitasking, adaptive and personalized systems, and voice interfaces.

