iLid: Low-power Sensing of Fatigue and Drowsiness Measures on a Computational Eyeglass

SOHA ROSTAMINIA, University of Massachusetts Amherst
ADDISON MAYBERRY, University of Massachusetts Amherst
DEEPAK GANESAN, University of Massachusetts Amherst
BENJAMIN MARLIN, University of Massachusetts Amherst
JEREMY GUMMESON, University of Massachusetts Amherst

The ability to monitor eye closures and blink patterns has long been known to enable accurate assessment of fatigue and drowsiness in individuals. Many measures of the eye are known to be correlated with fatigue including coarse-grained measures like the rate of blinks as well as fine-grained measures like the duration of blinks and the extent of eye closures. Despite a plethora of research validating these measures, we lack wearable devices that can continually and reliably monitor them in the natural environment. In this work, we present a low-power system, iLid, that can continually sense fine-grained measures such as blink duration and Percentage of Eye Closures (PERCLOS) at high frame rates of 100fps. We present a complete solution including design of the sensing, signal processing, and machine learning pipeline; implementation on a prototype computational eyeglass platform; and extensive evaluation under many conditions including illumination changes, eyeglass shifts, and mobility. Our results are very encouraging, showing that we can detect blinks, blink duration, eyelid location, and fatigue-related metrics such as PERCLOS with less than a few percent error.

CCS Concepts: •**Human-centered computing** →**Mobile devices**; User studies; •**Computing methodologies** →*Interest point and salient region detections*; •**Applied computing** →*Consumer health*;

General Terms: Drowsiness, Fatigue, Blinks, PERCLOS, Eyeglasses, Eyelid

ACM Reference format:

Soha Rostaminia, Addison Mayberry, Deepak Ganesan, Benjamin Marlin, and Jeremy Gummeson. 2017. iLid: Low-power Sensing of Fatigue and Drowsiness Measures on a Computational Eyeglass. *PACM Interact. Mob. Wearable Ubiquitous Technol.* 1, 2, Article 23 (June 2017), 26 pages.
DOI: [10.1145/3090088](https://doi.org/10.1145/3090088)

1 INTRODUCTION

While wearable devices provide insight into a variety of physiological and health conditions, one aspect that has lagged behind is our ability to infer an individual's cognitive state in the natural environment. There is significant need for a device that can continuously monitor fatigue, since this has implications for a wide range of application domains ranging from personal safety to health monitoring. While research in this field has predominantly focused on monitoring drowsy driving to help reduce driving fatalities (there were 846 drowsy-driving-related fatalities in 2014 [7]), it is far from the only scenario where fatigue tracking is beneficial.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Copyright © ACM 2017 2474-9567/2017/6-ART23 \$15.00

DOI: <http://doi.org/10.1145/3090088>

Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies, Vol. 1, No. 2, Article 23, Publication date: June 2017.

Fatigue has been shown to be a predictor of addictive behavior, a measure of quality of life, and a mediator between sleep and health outcomes. Addiction research has shown that sleep deprivation and sustained use of executive function lead to fatigue even in healthy individuals. This impairs self-control which in turn increases substance use, smoking behavior, and alcohol use [68]. Substance use can, in turn, increase daytime fatigue leading to safety hazards or low work performance [23, 62], but ironically, the process of trying to withdraw from addictive substances (even caffeine) can also trigger fatigue and drowsiness [41, 42]. The ability to monitor fatigue in real-time can allow us to predict relapse and trigger timely interventions, which is a key focus of mobile health research (e.g., the NIH Mobile Data-to-Knowledge (MD2K) center [9]). Monitoring fatigue in the natural environment can also complement our ability to monitor circadian rhythms and sleep patterns [10, 32], which can provide a holistic view of sleep, fatigue, and health outcomes. Finally, fatigue measurement in natural settings is also important to understand how to improve quality of life for cancer patients [18, 50], parkinson's patients [35], multiple sclerosis patients [26, 44], as well as increase awareness of mental state among the general public and improve regulation of the levels of stress in our daily lives [15].

The plethora of health and safety conditions that are linked with fatigue have made it the subject of much clinical research on how to estimate the level of fatigue. One modality that has long been known to provide a good measure of fatigue is eye monitoring. Many decades of experimental studies involving eye monitoring have identified that eye closures, blink duration, and blink frequency are the most significant features of interest for predicting the level of fatigue [38, 59, 63, 67].

Despite our understanding of how to measure fatigue, we lack good instruments to measure these eye parameters robustly in natural settings. In constrained environments such as vehicles, the environment can be instrumented with cameras to allow remote monitoring (although robustness to illumination changes and other dynamics remains a challenge). Alternately, commercial eye trackers intended for short-term episodic use can provide such measures and have been available for more than a decade (e.g., SMI [1] and Tobii [6]). But transitioning from technology that works in episodic and controlled settings to natural environments has been challenging. Most wearable eye trackers that provide high-resolution data are bulky and power-hungry, have poor performance in outdoor settings, and are not suitable for continuous daily wear. A promising eye monitoring device that bridges the gap towards daily wear is the JINS MEME [4] eyeglass. This device uses electrooculography (EOG) sensors placed at the nose-bridge. However, this method, as we show in this paper, has shortcomings in measuring subtle eyeball and eyelid movements while being susceptible to noise. Thus, there is a need for a truly wearable device that is low-power, portable, and provides accurate measures of eyelid movement, while being robust to confounders present in everyday scenarios.

Our goal in this work is to develop such a wearable solution. The key questions underlying such a design are robustness and power consumption. From a power perspective, cameras consume significant power if used continuously, particularly at high frame rates of around 100fps that is recommended for extracting fatigue features from the eye [57]. While there have been efforts to reduce this by leveraging sparse sampling techniques [47, 48], these have focused on eye tracking rather than eyelid tracking that we need for monitoring fatigue. From a robustness perspective, the key question is whether a real-time fatigue monitoring system can be robust to face shape, illumination (indoor and outdoor), mobility (stationary vs. walking), eyeglass position (sliding down the nose), and so on. A wearable device needs to be universally applicable under a wide range of conditions.

In this paper, we design a system, iLid, that is able to extract key features of fatigue at low power and high frame rate from a wearable eye tracker. Our contributions are two-fold. First, we develop methods that can dramatically reduce the cost of sensing and processing by sampling a small subset of pixels on an imager (a few columns of pixels) and processing these pixels in real time to extract the salient features for fatigue detection. We develop lightweight classification-based methods to extract blink and eyelid features such as blink duration, blink rate, and eyelid closure patterns. Second, we provide an exhaustive characterization of robustness of the technique under many settings including lighting conditions, eyeglass slippage, and user mobility. We also compare against a

state-of-the-art wearable EOG eyeglass, the JINS MEME [4], to understand the relative strengths and weaknesses of vision-based versus EOG-based fatigue measures.

Our results show that:

- iLid can detect blinks with high precision (above 95%) and high F1 score (about 0.90), blink durations with low error of 2.4%, and PERCLOS with accuracy of 97.5% in both indoor and outdoor conditions. Such low error is within the margin of error for human ground-truth labeling, and exceeds current state-of-the-art performance [13, 31, 56].
- iLid is robust to variability that occurs in mobile settings including outdoor illumination changes, eyeglass shifts, user mobility, and eye state - these confounders are typically not considered in the literature. Typical variability along these axes have negligible impact on the accuracy of estimating different metrics. We also show that iLid is superior to current wearable electrooculography (EOG)-based methods for estimating fatigue measures.
- Our methods can run in real-time on a low-power eyeglass platform with a low-power imager and micro-controller. We show that our system has an end-to-end power consumption of 27mW at 60Hz or 46mw at 100Hz including the cost of running the imager, data acquisition, and computation.

2 RELATED WORK

There is a large body of literature related to general eye tracking, and there has been extensive work in the specific area of cognitive fatigue detection - primarily focused on its application in driver drowsiness detection. Our core contribution is the development and evaluation of a robust on-body fatigue detection system capable of working across a wide spectrum of natural environments. Hence, we focus our literature survey on other relevant efforts and ask two questions: a) how applicable are these methods to mobile scenarios?, and b) how extensively have these methods been evaluated in realistic natural settings?

Due to the size and complexity of the literature on this topic, it is difficult to find simple categories within which all of the related studies fit comfortably. For the purposes of comparison with our work, we classify the related studies into three categories based on the type of sensors used: (1) remote-camera-based methods, which are designed to operate on data from a camera that is not worn on the body but rather mounted in a fixed location, (2) wearable-camera-based methods, which operate on data from a head-mounted camera, and (3) electrooculography-based methods, which measures the movements of the eye via electrodes attached to the skin. We also briefly discuss a fourth category: techniques that are based on eye movement tracking for measuring fatigue via changes in ocular movement patterns, as opposed to tracking the motions of the eyelid.

2.1 Remote-Camera Methods

Many video-based fatigue studies use mounted cameras, generally referred to as “remote” cameras. This automatically limits the applicability to cases in which the user is seated in front of the camera for long periods - generally, working at a desk or operating a vehicle. This means that these methods are not suitable for measuring fatigue ubiquitously, however, we include these studies for completeness and comparison.

There are two key aspects that distinguish this body of work from ours. The first is that since the remote cameras capture not only the eyes but also the entire face and other nearby imagery, the algorithms are considerably more complex. As Table 1 shows, the algorithms often require computationally-intensive feature extraction, segmentation, classification, estimation, and other approaches. At a high level, our pipeline is not very different since we have to follow a similar pipeline from denoising to classification; however, the complexity is significantly less (and accuracies higher) when the imager is directly obtaining an image of the eye at close proximity. Second, most of these devices do not actually measure eyelid location accurately and therefore implement an approximation instead of the true Percentage of Eye Closure (PERCLOS) as defined by the NHTSA, which is the

Work	Algorithm	Autonomous	Evaluation Scenarios	Metrics	Calibration
Garcia <i>et al.</i> [31]	V-J + projection	No	✓ Driving × Illumination	PERCLOS	One-shot
Pauly & Sankar [55]	HOG + SVM	No	× Driving × Illumination	Open/closed	None
Hong & Qin [36]	Haar + projection + motion	Yes 12 fps	✓ Driving × Illumination	Open/closed	None
Xu <i>et al.</i> [70]	LBP + classifier	No	× Driving × Illumination	Open/closed	None
Dasgupta <i>et al.</i> [24]	Haar + LBP + PCA + SVM + transform	No	✓ Driving ✓ Illumination	Open/closed	None
Bergasa <i>et al.</i> [16]	Filter + state machine	No	✓ Driving ✓ Illumination	PERCLOS Blink feat.	None
You <i>et al.</i> [71]	Haar + classifier + template + filter	Yes 30 fps	✓ Driving × Illumination	Open/closed	None
Malla <i>et al.</i> [45]	Haar + Kalman + template	No	× Driving × Illumination	Open/closed	None
Picot <i>et al.</i> [58]	Filter + projection + classifier + energy	No	× Driving × Illumination	Blink feat.	None
Boverie & Giralt [17]	Template + Kalman + motion	No	✓ Driving × Illumination	Blink feat.	Repeated
Grauman <i>et al.</i> [33]	Motion + filter + template	No	× Driving × Illumination	Blink feat.	None
Bacivarov <i>et al.</i> [14]	V-J + AAM	No	× Driving × Illumination	Blink det.	None
Lee <i>et al.</i> [43]	AdaBoost + template + filter + SVM	No	× Driving ✓ Illumination	Blink det.	None
Drutarovsky & Fogelton [25]	V-J + motion + state machine	No	× Driving × Illumination	Blink det.	None
Sukno <i>et al.</i> [64]	Haar + ASM	No	× Driving × Illumination	PERCLOS Blink feat.	None
Chau & Betke [22]	Filter + template + motion	No	× Driving ✓ Illumination	Blink feat.	None
Smith <i>et al.</i> [61]	Filter + motion + template	No	✓ Driving ✓ Illumination	Open/closed	None
Morris <i>et al.</i> [52]	Motion + template	No	× Driving × Illumination	Blink det.	None
Pedrotti <i>et al.</i> [56]	Proprietary + filter	No	× Driving × Illumination	Blink feat.	None
Wu & Trivedi[69]	Particle filter + dynamical modeling	No	✓ Driving × Illumination	Blink feat.	None

Table 1. Comparison of fatigue monitoring studies based on direct-from-video techniques. Parameters insufficient for real-time fatigue detection in a realistic driving setting highlighted in bold.

percentage of frames when the eyes are more than 80% closed excluding the blinks [65]. This is primarily because it is difficult to extract eyelid location from video at a distance in a reliable manner. Thus, these efforts primarily report open and closed states of the eye which is less precise for fatigue detection. Third, many of these methods are not evaluated on real embedded system or under real-world dynamics such as illumination variations. For example, only a couple try to systematically evaluate under different illumination settings.

Work	Algorithm	Autonomous	Evaluation Scenarios	Metrics	Calibration
McIntire <i>et al.</i> [49]	Proprietary + time series	No	× Driving × Illumination	Mean velocity	One-shot
Appel <i>et al.</i> [13]	Filter + random forest	No	✓ Driving × Illumination	Blink	None
Fuhl <i>et al.</i> [28]	Filter + edge detection	No	✓ Driving × Illumination	PERCLOS Blink	None
Jiang <i>et al.</i> [40]	Filter + template	No	× Driving × Illumination	Blink	One-shot

Table 2. Comparison of wearable-camera-based methods. Parameters insufficient for real-time fatigue detection in a realistic driving setting highlighted in bold.

2.2 Wearable-Camera Methods

Using head-mounted video recording systems and computational eyeglasses which do real-time data processing [30] can alleviate the mentioned problems by providing a close-up view of the eye, thereby eschewing the need for face and eye detection. There are many imager-based solutions commercially available such as those offered by SMI [1] and Tobii [6]. There are significantly fewer studies that explore blink detection and fatigue on wearable eye trackers. This is most likely because commercial head-mounted eye tracking systems implement blink feature metrics by default, so the few studies on this topic are either attempting to improve the performance of these systems or implement blink features on systems that do not have them pre-installed. The literature is summarized in table 2, see section 2.1 above for a detailed description of the categories.

On the surface, it seems that these commercial wearable devices are an excellent candidate for our purposes. However, there are a number of significant practical challenges with the use of these devices for continuous fatigue measurement. The first is that commercial wearable gaze trackers are well-known in the literature to fail under varying illumination conditions, especially outdoors [39, 66]. This makes them a poor choice for continuous gaze measurement in any environment other than an indoor setting. The second major issue is that the devices tend to be bulky and uncomfortable, with a battery life rarely exceeding 2 hours and never exceeding 4 [8]. Third, these devices are mostly passive recording systems intended to store data for post-hoc processing since the algorithms used generally require desktop-grade processing power. Such devices are incapable of autonomous fatigue detection and therefore cannot provide live notifications of fatigue onset, which is a necessity for safety-critical applications.

2.3 Electrooculography (EOG) Methods

There are a number of techniques for evaluating fatigue using physiological markers instead of video of eye movements. These include electroencephalogram (EEG) [34] and heart-rate variability (HRV) [54], but most of these techniques require attaching an array of complex sensors to the body and thus are not suitable for use in natural environments. The notable exception is electrooculography (EOG). EOG measures eye movements by leveraging the fact that human eye is a magnetic dipole, so electrodes placed on the skin near the eye can detect a change in electric potential as the eye moves. It is feasible for ubiquitous use as it requires fewer sensing points than other methods, and thus can be wrapped into a wearable form factor (generally eyeglasses). Bulling *et al.* have recently shown EOG metrics to be useful for recognizing different types of activities via eye movement patterns [19, 21]. There have been a number of studies attempting to measure fatigue solely through EOG metrics [29, 37, 46, 53, 60], but achieving high accuracy requires placing additional electrodes on the face to ensure high-quality data. Steps towards miniaturizing the sensors have been made in the literature [20].

On the commercial side, an intriguing EOG-based device that is designed to be wearable is the JINS MEME eyeglass [4]. The EOG electrodes are embedded in the nose bridge of the eyeglass, allowing it to monitor EOG patterns during eye and eyelid movements. While this is an intriguing new device that can enable new applications, studies of the EOG signals from the JINS MEME have also shown that it is limited in the resolution of tracking eye and eyelid movements [12]. For example, the device has difficulty in tracking vertical eye movements. In terms of eyelid closures, the EOG signal for full blinks seem to be sufficiently prominent that it can be extracted. But fine-grained measures and more subtle changes are hard to discern since the signal does not explicitly distinguish between eye movements and eyelid movements. Therefore, a robust measure of the percentage of eye closure is difficult to extract from the signal. §5.3 provides an evaluation with JINS MEME to confirm these hypotheses.

2.4 Fixation / Saccade Methods

Lastly, we note that there is intriguing work on using very fine-grained eye movements - known as fixations and saccades - to detect fatigue. At a high level, fixations are periods when the eye is relatively motionless, whereas saccades are the rapid movements that occur when the eye changes orientation. Certain features of these movements are known to be reliable indicators of the onset of fatigue.

However, there is not yet a body of work on using these metrics for any type of natural environment, even driving. Saccade-based methods require extremely accurate eye position measurement, which is difficult to do under uncontrolled lighting, and extremely high measurement rates (250 fps or more), which wearable eye trackers have not yet reached. Techniques that use fixation patterns are less technically restrictive to implement, but also require that the user be presented with a specific stimulus so that their responses can be measured, which is obviously not conducive to passive measurement in a natural setting. Ahlstrom et al. [11] performed a fatigue-measurement study for drivers using these techniques, but due to the aforementioned limitations they were restricted to testing the drivers' fatigue state in a controlled immediately before and after the driving event. We do not provide related studies in this area since there is not a strong comparison to be made with our work, but we refer the reader to [51] for an excellent literature review of the state-of-the-art on this topic.

2.5 iShadow: A computational eyeglass

The work described in this paper also builds on our own prior work where we described the design and implementation of a novel computational eyeglass, iShadow [47, 48]. The key technical contribution was the use of sparse sampling to detect gaze, eye saccades, and pupil dilation. iShadow [47] uses a neural network-based sampling to detect gaze at 10fps, while consuming 70 mW of power, and follow-on work on CIDER [48] improved on this work to estimate pupil center (for saccades) and pupil dilation while using a combination of sparse neural-network based sampling together with a more optimized cross-sampling approach.

There are several key distinctions between the work in this paper and our prior work. First, this paper focuses on eyelid-related measures that are relevant to fatigue, whereas prior work focused on eye-related measures such as gaze, saccades, and pupil dilation. The difference in the final objective changes our entire pipeline, starting from the sampling method all the way to the signal processing and classification techniques. Only the sensor platform remains the same - we leverage the sub-sampling capabilities offered by iShadow but borrow no other aspect of the pipeline from prior work. Second, this work places considerable emphasis on robustness in real-world conditions and evaluates under a plethora of scenarios including illumination changes, mobility, and eyeglass shifts. This emphasis is significantly more pronounced than our prior efforts. Third, the shift from pupil-based measures to eyelid-based measures means that we had to collect a substantial amount of data exclusively for this effort. As we describe in §4, previous datasets we collected with iShadow were at lower frame rate to capture full frames, but these ended up missing blinks. Therefore, this meant an entirely new data collection effort for this paper.

Fig. 1. The upper eyelid detection pipeline: 1) 4 columns of pixels near the center of the image are sampled, 2) pixel values are horizontally projected, 3) values are median filtered, 4) values are convolved with a box filter for edge detection, 5) specular reflection is detected and removed (top), 6) pupil is detected and removed (bottom), 7) the upper eyelid position is estimated.

3 iLid SYSTEM DESIGN

The main challenge in designing iLid is managing the tension between a) the high-rate sensing and processing needed for drowsiness estimation, b) the power- and resource-constrained nature of an eyeglass-form factor device, and c) the need to ensure robustness under dynamic real-world settings. Since eye closures can be fleeting (the minimum closure duration is about 150ms), prior work has suggested high frame rates of close to 100 fps to ensure that few blink events are missed [57]. The high frame rate results in significant power consumption at the imager as well as substantial computational load on the device to continually process frames and execute the fatigue detection algorithm. For example, a typical low-power imager operating at 100fps consumes 120mW [3], and generates data at 30 MBps. Handling such high volume of data and performing complex vision algorithms requires substantial processing capability, which requires high-end processors which add another 100 - 200mW to the overall cost. Storing raw data for future processing is equally, if not more, expensive. Thus, the overall power consumption of such a device would be a few hundred milliwatts, making it unsuitable as a wearable. While we clearly need to reduce how much data is generated by the imager and how much computation needs to be performed, it is also critical to ensure robustness to real-world settings. Thus, the key challenge that we face is achieving high-rate operation at low power, while ensuring robustness to real-world dynamics.

iLid addresses this challenge using a signal processing and classification pipeline comprising sampling optimizations, domain-aware noise removal and filtering operations, template extraction from the time-series data, and lightweight classification. All stages are optimized for power and computational efficiency, while being carefully tuned for high robustness.

3.1 Upper Eyelid Detection

The first stage in our computational pipeline is upper eyelid detection (shown in Figure 1). Since we need to operate at roughly 100 fps, this stage has to be both effective and efficient. This requires that we optimize both

how we use the imager as well as how we compute on the data from the imager. Our pipeline is therefore simple but carefully tuned to remove specific sources of noise that we observe.

Sub-sampling the image. Reducing the amount of data sampled at the imager is crucial to minimizing power consumption and increasing frame rate. Our eyelid detection pipeline limits the sampling to only a block of 4 columns in the middle of the image, in which the upper eyelid is roughly in its highest position. Such column sampling makes intuitive sense since we are focused on the eyelid; the use of a few columns rather than a single column also helps to reduce noise that may be observed due to intrinsic noise in a low-power imager, specular reflection of the NIR LED on the image, and other such considerations.

Once the columns have been sampled, a basic pipeline for detection involves three steps: a) a simple denoising method like median filtering, b) convolution with a 1D edge detection method to detect edges, and c) identifying the most likely edge corresponding to the upper eyelid position. But in order to make this pipeline robust, we need to address a few noise sources.

Dealing with specular reflections and pupil area. The 1D edge detection filter’s output typically peaks in the region corresponding to the upper eyelid position due to the transition from the eyelid skin and eyelashes to the iris and/or sclera area; however, there are occasions when the specular reflection from the NIR LED and/or the pupil generates a peak higher than the upper eyelid.

In a majority of cases, we can simply use prior knowledge of the size of the specular reflection and pupil to filter them out. We know that the specular reflection from the LED is highly localized and only a few pixels wide; we also know that the pupil is less localized and between 15 and 25 pixels wide. We can leverage this to filter out these artifacts — for example, the specular reflection has a particular form of a short interval containing a lowest peak followed by a highest peak in the filtered signal. Figure 1 shows an instance where the specular reflection is removed on the top row, and pupil is removed on the bottom row.

We also experimented with increasing the number of columns used, and found that it can be useful in a small number of instances where ambient NIR levels are very low and the specular reflections have a strong footprint. Since the iShadow platform has an NIR photodiode to measure external NIR levels, this can potentially be leveraged to dynamically change the number of columns sampled (while sacrificing some frame rate in the process). An example is shown in Figure 2, where aggregating over more columns results in a cleaner signal. In our experiments, we find that this optimization helps only in a small fraction of our data, so we do not dynamically adjust the number of columns in our current computational pipeline. But it provides an avenue for further optimization and personalization.

Dealing with flicker noise. Another noise issue that we face is flicker noise in the imager. We observe flickering in the imager’s video, which itself is caused by small differences between the frame rate of the imager and the frequency of the lights in the environment. The result is high frequency noise spikes in the output of the eyelid detector. In order to deal with this problem, the output of the eyelid detector is passed through an appropriately tuned low-pass filter to remove such noise. Figure 3 shows the result of applying the low-pass filter to the output of the eyelid detector.

3.2 Blink Detection

Once the upper eyelid is detected, the next stage is to determine blinks (shown in Figure 4). This stage operates on a window of time series of eyelid positions extracted from the previous stage. Blink features and blink rate are not only important measures on their own, but they are also integral to calculating the extent of eye closure while the blinks are not occurring. Indeed, the standard PERCLOS metric is defined as the percentage of frames when the eyes are more than 80% closed *excluding the blinks* [65].

Blink Profile Template. The key idea in our method is that, while blinks appear to vary across individuals and is influenced by the state of alertness, there is also significant commonalities in their general structure. Eye

Fig. 2. Increasing the number of sampled columns to ameliorate the specular reflection problem: Figure (a) shows the noisy eyelid detector output in the existence of specular reflection when sampling only 4 columns of pixels, while Figure (b) shows the output of the eyelid detector for the same data segment after increasing the number of sampled columns to 11.

Fig. 3. Filtering the eyelid detector output: The high frequency noise in the raw eyelid detector output (a) is removed by applying a low-pass filter (b).

blinks share the same pattern of a fast eye closure followed by a relatively slower eye opening (shown in Figure 5). On average, eye closure takes about 60ms with a maximum velocity of 350mm/s and the eye opening takes about 120ms with a maximum velocity of 150mm/s [17]. While the general pattern of a blink does not change with the user's cognitive state, the actual duration of a blink could take from 100ms to more than 600ms depending on the person and their alertness condition with 200ms being the normal duration of a blink [17]. As a result, it would be reasonable to approximate the blink profile with the same pattern scaled horizontally with different values in order to model different blink durations. Prior work has also discretized blink durations into three levels — fast, normal, and slow blinks having a duration of 100ms, 200ms, and 600ms respectively [17]. We find out that templates corresponding to these discrete blink durations are sufficient to capture most variations in our data, and also allow us to compactly represent blink patterns.

Template Matching. The fact that eyeblinks can be placed into a few categories based on their duration and pattern makes template matching an ideal method for blink detection. For each of the three templates (fast, normal, and slow depicted in Figure 5), we define 5 keypoints which correspond to a specific duration, slope pattern, and height for the blink shape. The keypoints' relative position in each blink profile is based on experimental blink data and interpolated using cubic spline method to achieve a smooth curve resembling actual

Fig. 4. Blink detection pipeline: 1) A moving window selects samples from the eyelid position sequence, 2) the template matching scores over the window are generated, 3) the scores are used as feature vectors for a logistic regression classifier that detects blinks, 4) detected blinks after eliminating redundant detections.

Fig. 5. The three blink templates representing, from left to right, slow, normal, and fast blinks. The templates are horizontally scaled versions of the same blink profile defined by the five key points A through E which have been interpolated using cubic splines. Points A and E determine the duration of the blink T . Point B sitting at $0.25T$ on the horizontal axis determines the height of the blink H . Points C and D which give a slide pattern to the blink profile in the eyeopening phase of the blink have relative coordinates of $(0.5T, 0.45H)$ and $(0.9T, 0.02H)$, respectively.

experimental blink data. The blink profile is then scaled horizontally to create three templates with durations of 100ms, 200ms, and 600ms for fast, normal, and slow blink categories respectively.

The template matching process is a simple normalized cross-correlation computation. We normalize the eyelid position time-series to remove differences in the height of the blink across individuals. Then we sweep the eyelid data with each of the templates and calculate the dot product of each subsection with the corresponding template to obtain a similarity score $c(u)$ corresponding to the interval $[u, u + \text{template length}]$ of the eyelid position time-series with each template. This is shown below:

$$c(u) = \frac{1}{\text{len}(T)} \sum_{x=u}^{u+\text{len}(T)-1} \frac{(f(x) - \bar{f}(u)) (T(x-u) - \bar{T})}{\sigma_f \sigma_T} \quad (1)$$

where $f(x)$ is the eyelid position time series, $T(x)$ is the template, and σ_f and σ_T are standard deviations of f and T respectively. Also, $\bar{f}(u)$ and $\bar{T}$ represent the DC values of eyelid position data and the template respectively.

Fig. 6. Blink duration fix: The rough position of the blinks is first detected(a) and then a local search is performed to locate the ending points of the blinks(b).

Blink detection as classification. The template matching stage gives us the result of matching with three templates — fast, normal, and slow blinks. Given the scores, we then need to detect whether a blink occurred i.e. we need to map from a continuous output (matching scores with the three templates) to a categorical output (blink vs. not-blink). We can view the problem of detecting blinks as a standard binary classifier learning problem. The input is a short window of scores from the templates as shown in Figure 4 (window length is seven in our implementation), and the output is whether the window corresponds to a blink or not.

While we can apply any binary classification model, a linear logistic regression classifier [27] is particularly attractive due to its low computational complexity. A logistic regression classifier, in our case, involves 21 multiplications and 21 additions, which can be easily computed on a microcontroller.

Given a feature vector $\mathbf{x}$ consisting of D features, the binary logistic regression classifier returns the probability that the feature vector belongs to the positive class. Letting Y represent the label for the instance $\mathbf{x}$, logistic regression computes the class probability as shown below. θ is a length D vector of feature weights. $D = 21$ in our case given the fact that we are using a window of size seven for each of the three template scores. The classifier has a linear decision boundary specified by the weights θ .

$$P(Y = 1) = \frac{1}{1 + \exp(-(\theta^\top \mathbf{x} + b))} \quad (2)$$

The default classification rule when using linear regression is to predict that the data case belongs to the positive class if $P(Y = 1) > 0.5$. Learning the weights of the logistic regression classifier is accomplished by maximizing the log likelihood of the training data using numerical optimization [27]. This is a continuous, convex optimization problem with no constraints. It can be solved using any gradient-based optimizer. Given a data set $\mathcal{D} = \{(y_n, \mathbf{x}_n)\}_{1:N}$, the log likelihood function is defined as shown below. We assume the labels for the two classes are -1 and 1 .

$$\mathcal{L}(\theta, b|\mathcal{D}) = - \sum_{n=1}^N \log(1 + \exp(-y_n(\theta^\top \mathbf{x}_n + b))) \quad (3)$$

The above pipeline comprising template matching followed by classification provides two key advantages. The template matching stage allows us to take advantage of domain knowledge regarding the shape of blinks. This makes the system more robust to spikes in the data that can be mistaken as blinks by alternate methods. The classification stage allows us to map from the continuous measures that we get from the template matching to a categorical output of blinks or not blinks.

Since the classifier can detect multiple overlapping time windows as blinks, the nominated blinks are then pruned based on a minimum time window condition to remove redundant blink detections.

Blink Features Extraction. So far we have only detected the location of eye blinks. But we also want to extract features from the blinks, for example, the exact duration of the blink or the height of the blink. In order to do so, a window of eyelid position time-series around each of the detected blinks is smoothed with a mean filter and then a local search is performed starting at the lowest value in the window (i.e. the point where the eye is closed). The time-series is searched both forwards and backwards to find the first points, which satisfy either of the following two conditions in each search direction: the first condition is for the slope value to cross zero and the second one is for the height to reach a margin of the upper eyelid position baseline. The first points, in each search direction, that meet either of these conditions are chosen as the blink ending points. Figure 6 shows this process. First, the points corresponding to blinks are detected and then the blink duration is extracted by running the local search around the detected points.

3.3 Drowsiness Estimation

Once blinks are detected and their durations are measured, the corresponding frames are removed from the upper eyelid position time series in order to measure drowsiness or fatigue. Many measures of drowsiness are possible based on measures of eyelid position, blink occurrences, and blink duration that we extract in our computational pipeline. We illustrate using a common measure to measure drowsiness, Percentage of Eye Closures, or PERCLOS. For PERCLOS estimation, we need to estimate the number of frames in which the eye is more than 80% closed compared to when the eye is fully open. Therefore, we need to have a baseline of the eyelids' position when the eye is fully open vs. fully closed. In order to do this, we perform a one-time calibration for each user where we ask the user to fully open and close their eyes. This gives us one image from each eye state — fully open and fully closed. Note that it may be possible to avoid calibration by assuming that the largest variation in eyelid position corresponds to the open and closed states. We performed explicit calibration in our current system since we noticed that some subjects tend to not fully close their eyes while they blink. For example, as you can see in Figure 7, the only times when the subject's eyes are fully closed during blink are the two blinks that happened at time indices of about 1600 and 1675. During the calibration the subjects are asked to fully close and then open their eyes. The minimum and maximum upper eyelid positions computed from the eyelid detector algorithm are then used to estimate the position of the lower eyelid as well as the size of the eye in the fully open state, and used to compute the percentage of eye closure. With this information, we can measure PERCLOS as the fraction of the frames in which the upper eyelid position is lower than the calculated threshold (Equation 4) over a one minute window.

Fig. 7. Upper eyelid location time series

Fig. 8. Eyeglass platform containing an eye-facing imager, two NIR LEDs, and a PCB board with the micro-controller, Bluetooth, and other modules on the left, as well as the battery board on the right.

$$Threshold = \text{Upper eyelid baseline} - 0.8 \times (\text{Upper eyelid baseline} - \text{Lower eyelid baseline}) \quad (4)$$

4 IMPLEMENTATION AND DATA COLLECTION

We now briefly describe the hardware we use in our evaluation and hardware-related considerations that impacted our data collection efforts.

Hardware platform. We implement iLid on the iShadow computational eyeglass platform that we have designed and described in prior work (shown in Figure 8) [47, 48]. The platform uses the Stoneman Vision Chip produced by Centeye, Inc. [2] as its imager. The Stoneman camera is a low-power, grayscale, logarithmic-pixel imager that consumes 3mW and has a resolution of 112x112 pixels. The key feature of the imager is that it allows a random-access interface through which specific pixels can be selected and acquired. The random access capability allows us to reduce the cost of acquiring and digitizing pixels to just the ones that are selected. It also allows us to increase frame rate to desired levels — even though the Stoneman can only run at around 30fps when acquiring all pixels, we can operate at much higher frame rates when sub-sampling pixels. The iShadow eyeglass uses the STM32L151 microcontroller, which is a very low power ARM Cortex M3 processor [5] with 48KB of memory. All the algorithms described in this paper are implemented directly on this micro-controller. Note that while the iShadow version shown in Figure 8 has imagers that project a little from the eyeglass, a newer version of the hardware has imagers that are more integrated into the eyeglass frame.

Dataset Collection. Our evaluation includes 5 different datasets that capture different parameter, confounder, and variability settings. All of these datasets were collected under Institutional Review Board approval. We discuss the parameters of the individual datasets as and when we refer to them in the evaluation, but highlight one issue that we faced in our data collection efforts.

A key issue we faced was that we could not capture full 112x112 eye images for ground truth labeling since the frame rate was too low (roughly 10fps when capturing full images and writing to SD card). This also meant that data collected and labeled in our prior work [47, 48] was not useful for us since the frame rate was 10fps, and it missed many blinks.

Instead, we captured 11 columns of a partial eye image (i.e. 11x112 pixels) — only four of these columns are needed for most of the algorithms outlined in this paper (except the optimization to deal with specular reflections as outlined in §3.1), but the 11 columns helped to better visualize the frames and hand-label the eyelid positions and blink instances. We could capture frames at a rate of 60Hz using this method which is sufficient to extract even fast blinks. In a live system, since we only need to sample less than half the number of columns, we can

go more than twice as fast, which would be ideal for PERCLOS measurement based on [57]. However, this modification also meant that we collect all data anew.

5 EVALUATION

We now turn to an evaluation of our system. Our evaluation is exhaustive and includes evaluation of diverse aspects ranging from robustness to comparison against an alternate wearable eyeglass. We structure the evaluation as follows:

- *Aggregate results:* We begin with aggregate results that capture how well we are able to extract eye closure parameters under different conditions for sixteen subjects. These aggregate results show that our techniques perform very well overall.
- *Robustness to variability:* We then present careful experiments under many different conditions including eyeglass shift, illumination changes, and wearer mobility, to show robustness of our techniques to such variations.
- *Comparison against JINS MEME:* We present a comparison against another low-power eyeglass, JINS MEME, to understand the difference in performance between EOG and vision-based methods in detecting eye closure patterns.
- *Implementation benchmarks:* Finally, we present our implementation benchmarks for implementing our algorithm on the iShadow platform that we described in earlier work [47, 48].

5.1 Aggregate Results

Our first set of results provide an aggregate view of the performance of iLid. In this set of experiments, we focus on getting a diverse dataset across indoor and outdoor settings, and different genders and ethnicities.

5.1.1 Datasets and Ground Truth Labeling. We collected several datasets to validate our methods:

- **Indoor Dataset (Fixed illumination):** This dataset was collected from 16 subjects, 10 male and 6 female with average age of 22 from four different ethnicities including White-American, Middle-Eastern, Indian, and Asian. The experiment consists of the subjects watching a short animated movie for 5 minutes while wearing the computational glasses. In order to induce dynamic and challenging eye situations, a dynamic-themed animation is chosen and the subjects are asked to follow an object in the movie. Partial images of the eye (11 columns of pixels) are captured in 60Hz frame rate resulting in 18,000 frames per subject and 288,000 frames in total for all 16 subjects.
- **Outdoor Dataset (Uncontrolled illumination):** The same experiment is conducted on the same 16 subjects but in an outdoor environment with uncontrolled and varying illumination conditions for different subjects, resulting in 288,000 frames of partial eye images.
- **Ground Truth Labeling:** For ground truth, we hand-labeled the upper eyelid location in each frame using the following approach. The distance between the upper and lower eyelid is divided into 10 equally spaced segments that are roughly two or three pixels wide (depending on the open and shut levels obtained from calibration). The number of segments was chosen based on the fact that the thickness of the eyelid itself is around two pixels, so it is difficult to label at a finer granularity. The labels were confirmed by a second individual to ensure correctness.

5.1.2 Blink Detection. We first evaluate the performance of the blink detection algorithm in terms of its ability to recognize blink instances. This is evaluated using precision, recall and F1 score. We use leave-one-out cross validation — the classifier is trained using the data from 15 subjects (both indoor and outdoor datasets combined for training) and tested on the last subject and this set is rotated so that the classifier is tested on all 16 subjects

Dataset	Precision	Recall	F1 score
indoor	0.96	0.85	0.90
outdoor	0.95	0.84	0.89

Table 3. Blink detection accuracy (over 1760 blink instances).

Fig. 9. Breakdown of each analytic component's contribution to overall results. The baseline represents basic eyelid detection (ED) + logistic regression classification, the second bar includes specular reflection and pupil removal (SPF) and the final bar includes templates (TM) (i.e. the entire pipeline).

in both indoor and outdoor datasets. Our dataset from 16 subjects included 643 blink instances in the indoor dataset and 1117 blink instances in the outdoor dataset.

Table 3 shows that we achieve high precision (roughly 0.95) and high recall (roughly 0.85) for both indoor and outdoor datasets. The blink detector shows virtually the same performance for both indoor and outdoor situations, which validates its ability to operate effectively across different illumination conditions that one might encounter in a real-world setting.

We now breakdown the contributions of individual components of the pipeline to the overall result to see how much each of them matters towards the overall result. Figure 9 shows the breakdown – we start with an approach that only uses eyelid detection based on the first peak observed (after median filtering) together with logistic regression to detect eye closure, then we add specular reflection removal and pupil removal components, and finally we add the templates as features to the classifier. We see that each of these methods is important for the overall results, with most of the improvements coming in terms of improved recall. Each of the stages adds about 10% or more to the recall and corresponding improvements in the F1-score. Thus, both design components play a crucial role in the strong aggregate results.

We further breakdown the result across individual subjects in Figure 10, which depicts the F1 score for each subject in both indoor and outdoor separately. This helps us understand whether there is significant variability across individuals. While the results are generally consistently high across subjects, the numbers for Subject 6 seems to show a significant difference in performance between the outdoor and indoor environment. We found the reason to be simply the fact that the subject had very few blinks (less than 10) in the 5-minute experiment period; we do not have enough data for the individual to draw conclusions regarding the accuracy of our methods.

5.1.3 Blink duration estimation. Once blinks are detected, we extract blink duration, which is a key feature for detecting fatigue. We evaluate the accuracy of blink duration estimation by comparing the actual and measured

Fig. 10. Blink detector performance for individual subjects.

Fig. 11. Blink duration measurement error CDF.

duration of the correctly detected blinks. The error is calculated in frame numbers and converted to milliseconds given the constant 60Hz frame rate of the camera.

Figure 11 shows the error CDF of measured duration of correctly detected blinks. Our methods perform extremely well in this regard — more than 95% of the detected blinks had error of less than 17ms (i.e. less than 10% of a typical 200ms blink). This equals the distance of two consecutive frames given the 60Hz frame rate of the imaging device, so we are within one frame of the actual blink duration estimate. Overall, the mean relative error for blink duration was measured as only 2.4%, thus, we can estimate blink duration accurately.

5.1.4 PERCLOS estimation. In this evaluation, we look at how well we determine PERCLOS. Since PERCLOS is essentially a measure of the percentage of eye closure, we divide frames into different stages of eye closure, and see how well we can detect these stages. In other words, we partition the ground truth eyelid position data into frames where the eye was between 50%-60% closed, 60%-70% closed, and so on until 90%-100% closed. Then we evaluate the performance of our eyelid detector in classifying all the frames based on the detected level of eye closure. We then evaluate the precision and recall for classifying the eye closure in each of these partitions.

Fig. 12. Eyelid detector performance in different eye closure states.

Figure 12 depicts eyelid detector's performance in five different eye closure states corresponding to the eye being 50-60%, 60-70%, 70-80%, 80-90%, and 90-100% closed. This gives us an idea of how the error is distributed in different eye states. The F1 score of the detection in each of the 5 classes is shown for both indoor and outdoor datasets across all subjects. As can be seen in the figure, the detector shows excellent performance across all the eye closure states. Of particular interest to PERCLOS estimation is the accuracy for detecting when the eye is more than 80% closed. The accuracy of this detection is 97.5%, which shows that our methods are effective for determining PERCLOS.

5.2 Robustness to variability

Our next set of results look at the robustness of our techniques to different conditions. The dataset in this section is collected from five subjects, 3 males and 2 females with average age of 23 under a variety of conditions.

5.2.1 Robustness to illumination. To look at the sensitivity to outdoor illumination, we look at results at different times of day to see if the results show variation. Note that while the results in §5.1 were under different illumination conditions, this set of results provides a more systematic comparison under different outdoor illuminations. The experiment is conducted outdoors at 3 different times during the day for each subject. The three sessions are held at 10AM, 3PM, and 6PM, when the total intensity of IR light in the near infrared range (NIR) was measured as $10\text{w}/\text{cm}^2$, $5\text{w}/\text{cm}^2$, and $0.1\text{mw}/\text{cm}^2$ respectively using a light meter equipped with the same IR filter as that of the iShadow eyeglasses. In each 5 minute session the subject are asked to watch a video clip while wearing the iShadow eyeglasses.

The results are shown in Table 4 and Figure 13, which depict the blink detection and eyelid detection performance respectively. The results show that blink detection has almost the same performance during different times of the day and hence different illumination conditions. The eyelid detection also shows robustness to illumination changes specifically for cases when eye closure percentage is higher than 60%. The lower performance of the eyelid detector in low light is due to the fact that the contrast between the pupil and iris becomes sharper in such situations. This causes errors in the eyelid detection specifically when the eye is in about 50% closure where the pupil is usually present. Despite this issue, the positive is that measures such as PERCLOS are defined for higher eye closure levels (80%), where performance does not degrade.

5.2.2 Robustness to eyeglass displacement. One issue that we need to address in mobile settings is eyeglass shifts. Small shifts of the eyeglass down the nose-bridge is a common occurrence, particularly when an individual is mobile. To verify this, we look at sensitivity to eyeglass shifts across five individuals. We construct spacers of

Dataset	Precision	Recall	F1 score
10AM	0.99	0.88	0.93
3PM	0.92	0.97	0.94
6PM	0.97	0.89	0.93

Table 4. Blink detection performance in different illumination conditions.

Fig. 13. Eyelid detection performance in different illumination conditions.

different lengths (0.5cm, 1cm, and 1.5cm) and placed these spacers between the eyeglass and the forehead. This allowed us to verify performance under a controlled set of distances. Note that typical eyeglass shifts are much less than 0.5cm - typically only a few millimeters. We use larger distances to understand how the performance degrades as distance increases. The experiment is conducted under controlled lighting conditions. The experiment consists of 4 sessions for each subject as they are asked to watch a video clip for 5 minutes while wearing the iShadow eyeglasses equipped with a specific spacer in each session.

The performance of the blink detection module is shown in Figure 14. As it can be seen, the performance decreases almost linearly as the distance of eyeglasses to their normal position increases. However, the slope of the graph, i.e. the sensitivity is small enough so that even at a spacer size of 1.5cm, which relates to the case of the eyeglasses positioned almost at the tip of the nose the blink detection achieves an F1 score of more than 0.7. The reason for this decrease of performance is mostly due to the fact that the height of the blinks become smaller as the eyeglasses get further from the eye which makes it harder to detect the originally small blinks.

Figure 15 shows the performance of eyelid detector module in detecting eyelid positions corresponding to different levels of eye closure with respect to degree of displacement caused by using different spacers. As expected, the eyelid detection performance reduces with increase in the distance of the eyeglass from the eye. This decrease in performance, however, is not significant for displacement values equal or less than 1cm. Even for displacement value of 1.5cm the performance remains satisfactory for eye closure percentages more than 70%, which is our region of interest given the definition of PERCLOS. The general reason for decrease of performance with larger displacements is the fact that the total number of pixels containing the eye reduces as the eyeglasses are distanced from the eye, which leads to lower signal to noise ratio.

It should be noted that, although the eyelid detection performs robustly with respect to displacement, the PERCLOS measurement would be affected by the displacement in that the total size of the eye in pixels as well as the lower eyelid position changes as the camera goes further from the eye. However, we find that for typical

Fig. 14. Blink detection performance versus different eyeglass displacement values.

Fig. 15. Eyelid detection performance for different eyeglass displacement values.

displacements less than 0.5cm, this change is negligible (less than 2% of the eye diameter) given the current resolution of the iShadow imager, therefore the effect on PERCLOS estimation is low.

5.2.3 Robustness to mobility-induced variability. Since our algorithms are intended for the mobile scenario, one important question is whether performance degrades when the user is mobile. To understand this, we performed a controlled experiment where each participant wore the iShadow platform while walking on a treadmill and investigated our algorithm's blink detection and eyelid detection performance. The experiment consists of two 5 minute sessions for each subject. In one session the subjects are asked to walk on a treadmill with a speed of 2 miles/hour while wearing the iShadow eyeglasses. The other session includes the same subjects watching a 5-minute video clip while sitting on a chair and wearing the iShadow platform.

Table 5 and Figure 16 compare the blink detection and eyelid detection performances respectively in the stationary and mobile situations. The results show no significant superiority of one over the other, which demonstrates that our algorithms perform well in the presence of natural displacements and vibrations caused by mobile situations.

5.2.4 Robustness to drowsiness-induced variability. Our experiments so far were done with subjects who were not drowsy, hence they did not demonstrate natural changes in eyelid location and blinks due to fatigue. To understand sensitivity to this parameter, we performed a small additional study in a driving simulator. We only

Dataset	Precision	Recall	F1 score
stationary	0.96	0.82	0.88
mobile	0.94	0.79	0.86

Table 5. Blink detection performance compared in mobile and stationary scenarios.

Fig. 16. Eyelid detection performance compared in mobile and stationary scenarios.

Dataset	Precision	Recall	F1 score
morning session	0.94	0.78	0.85
night session	0.89	0.75	0.81

Table 6. Blink detector sensitivity to drowsiness-induced variability.

had three of the five subjects complete this study due to logistical reasons. The experiment consists of two 10-minute driving sessions in a driving simulator setup. In order to induce fatigue, the first driving session was held in the morning when the subjects are expected to be alert and conscious. To make sure that this is the case, the subjects were chosen based on the results of a questionnaire which validated that they could be considered as “morning people”. The subjects are then asked not to take a nap or drink caffeinated drinks until the second driving session in the evening of the same day.

Figure 17 compares our algorithm’s measured PERCLOS value to the ground truth for one of the subjects in the night session. The graph shows how our method could closely follow the actual PERCLOS pattern. Moreover, Table 6 shows a comparison between iLid’s blink detection performance for the morning session versus the evening session across three subjects. Our results show that the performance of blink detection which is required for PERCLOS measurement is not affected by the potential changes in eye closure pattern induced by drowsiness.

5.3 Comparison against JINS MEME

We now turn to evaluation against a recently released wearable eyeglass, the JINS MEME, that is also designed to be low-power and provide measures of the eye. This platform uses Electrooculography (EOG) using three electrodes in order to read eye movements in horizontal and vertical directions. As we described in §2.3, EOG has limitations in terms of capturing fine-grained movements of the eyelid, and is more appropriate for saccadic

Fig. 17. PERCLOS results for the driving simulator experiment. The vertical axis represents PERCLOS value and the horizontal axis shows time in seconds. The measured PERCLOS values is represented with the blue solid line and the dashed orange line shows the ground truth.

movements of the eye. While this may be true in principle, our goal in this section is to quantify the gap between EOG-based and vision-based methods for fatigue measurement on a wearable device.

Our experiments is conducted on 5 subjects, 3 males and 2 females, and consists of 3 sessions of 5 minutes in length for each subject. In the first session, the subjects are asked to watch a video clip for 5 minutes while wearing computational eyeglasses. In the second session the subjects recite a pre-defined conversation, and the last session consists of the same subjects walking on a treadmill with a speed of 2 miles/hour. The same experiment is conducted for both the JINS MEME and iShadow. The EOG time-series measured by the JINS MEME and the eyelid location time-series measured by our eyelid locator running on iShadow are both logged to be used for more analysis and comparison offline. In order to extract the ground truth for JINS MEME experiments in terms of the blink instances, close-ups of the subjects are filmed and time-synchronized with the JINS MEME measurements.

One issue that we faced was how to perform a fair comparison between the two approaches. Ideally, we would like our experiment to compare the intrinsic information regarding eyelid closures contained in the EOG stream versus the pixel stream. In other words, we want to make sure that the comparison is not influenced by some uneven feature extraction or preprocessing method for either of the two approaches.

In order to realize this goal and achieve a fair comparison, we consider two time-series streams: a) the EOG stream from the JINS MEME consisting of 4 pairs of channels (vertical, horizontal, left, and right), and b) the Eyelid detection stream output by our eyelid detection module (§3.1). We then train an ensemble classifier to detect blink instances on these streams, taking care to optimize the hyperparameters for each of them separately. For both streams, 300ms windows of the time-series are chosen as feature vectors. The performance is then evaluated on a leave-one-out basis in that the classifiers are trained with the data from 4 subjects and then tested on the 5th one.

The results shown in Figure 18 clearly demonstrate the superiority of iLid as a representative of imager-based computational eyeglasses compared to JINS MEME as an EOG-based computational eyeglass. (Note that the performance of iLid is lower than our results reported in earlier sections because we do not leverage our optimized classification pipeline.) First, the overall performance of the imager-based classifier significantly exceeds the other, particularly in cases where there is co-occurring activity involving muscle movements. For example, when the subjects are talking (Figure 19c), other muscles on the face are being used that create electrical signals and add to the signals coming from eyeball movements. Since the EOG signal is simply an addition over all electrical

Fig. 18. Comparing blink detection performance of a preliminary version of iLid with JINS MEME. The performance of both platforms are shown in three different scenarios. Error bars represent standard error.

activity, it becomes very difficult to separate these components. Second, the EOG-based classifier is also highly sensitive to mobility. While investigating the reason for this vulnerability, we found out that any movement induces significant amount of noise on the EOG readings from JINS MEME, which consequently makes the problem of eyelid detection difficult. Such noise is more significant when the source of movement is closer to the eye, but even gross body movements have a significant effect on the signal.

So far, we have looked at blink detection but how about measures of the eyelid like PERCLOS that is essential for fatigue detection. EOG is fundamentally not suited for these measures since it basically captures the eyeball movements while other movements such as that of the eyelid add as noises to the main signal. This means that there exists no mapping from the EOG signal to the position of the eyelid even though one might be able to detect transient movements of the eyelid based on the EOG signal (as shown in Figure 19). As a result, the EOG signal intrinsically cannot differentiate between the eyelid being 80% closed versus being wide open since both appear as a flat line.

5.4 Implementation on iShadow

We now evaluate the power consumption of our algorithm when it executes on the iShadow platform. There are various components that consume power in our system including the base power consumed by the micro-controller (MCU) and the imager, the power consumed by the NIR (near-infrared) LEDs, the power consumed for pixel acquisition from the camera, and the power consumed for the computation stages (eyelid detection, template matching, classification, and PERCLOS estimation). We measure the power consumption of our system by running the end-to-end system on the eyeglass platform, and measure the power using a DAQ running at a 10KHz sampling rate.

Since we use the ultra-low-power STM32L151 MCU and the low-power Stonyman imager, the base power consumption is generally negligible since the MCU and the imager consume very little in deep sleep mode. We measured the current draw over the baseline power to be 14mA at 3.3V i.e. 46mW while operating at 100Hz. The power consumption drops proportionally with frame rate, so if we reduce the frame rate to 60Hz, this reduces to about 27mW. Most of the power is consumed by the MCU itself since the two NIR LEDs operate in low voltage mode, and the imager consumes only a few milliwatts. We expect that these numbers can be further reduced through various optimizations at the MCU (including addressing some hardware kinks in the iShadow device as well as using more power-efficient micro-controllers), but this is already promising in terms of having a device

Fig. 19. JINS MEME's EOG signal output. The signal consists of 4 channels of vertical (V), horizontal (H), left (L), and right (R). Figure (a) shows the EOG signal corresponding to a blink when the person is completely stationary. Figure (b) shows the EOG signal relating to 4 different states of the eye. Segment 0 shows a period when the eyes are wide open. Segment 1 corresponds to a case when the eyes are slowly closed which is followed by segment 2, in which the eyes are kept closed. In segment 3 the eyes are slowly opened again. As it can be seen the EOG signal is the same for segments 2 and 0 which relate to steady cases when the eyes are kept closed and open respectively. Figure (c) also depicts the cross talk on the EOG signal induced by talking.

capable of continuous operation for several hours on a wearable eyeglass. Using a 570mAh battery, such as the type used on Google Glass, the platform could run up to two days with the current settings.

6 CONCLUSIONS AND FUTURE WORK

In summary, this work is the first to show that we can extract a variety of features that are pertinent for fatigue detection including rate of blinks, blink duration and percentage of eye closure in real-time and low-power on a wearable eyeglass. Our methods are accurate across a range of individuals, and robust to illumination conditions, eyeglass shifts, and user mobility. In addition, our system operates at rates upwards of 60fps while consuming only 27mW of power in the process. Blink and eye closure features are the cornerstone of fatigue and drowsiness detectors that can operate in the real world. We believe that this work paves the way for a regular spectacles form-factor device that has built-in ability to monitor cognitive state in real-time.

There are many directions that we are continuing to explore beyond the scope of this paper. One of these is conducting long term user studies in more natural environments. This would allow us to extract eye-related features at different levels of fatigue during the day, allowing us to observe eye-related patterns at longer timescales than what we have been able to do currently. This can allow us to extract higher level information such as the relation between the measured fatigue and blink rate with the behavior patterns of the subject. We are also exploring clinical ideas that build on this work, including detection of fatigue level in cancer patients. Such a longitudinal study would also reveal if there are instances of eyeglass shifts and displacements that we may not have addressed in this paper. An example could be rotations of the eyeglasses around the vertical axis if people put pressure on it while resting their head in their hand. These situations may be important since they relate to fatigue. We will look at these cases in our future work.

ACKNOWLEDGMENTS

We are very grateful to Tadashi Shimizu and Yuji Uema from JINS Inc., for taking the time to discuss the design of the JINS MEME, as well as its strengths and weaknesses for measuring eyelid-related measures. We are also very grateful to Siby Samuel and Tracy Zafian for helping us in using the driving simulator. We also thank

anonymous reviewers for their suggestions to improve this work. This research was partially supported by NSF grants #1551850, as well as NIH grant #1U54EB020404.

REFERENCES

- [1] 2007. SensoMotoric Instruments (SMI): Eye and gaze tracking systems. (28 Nov. 2007). <http://www.smivision.com/en.html>.
- [2] 2014. Current Centeye Vision Chips. <http://www.centeye.com/>. (22 Feb. 2014).
- [3] 2014. OmniVision: CMOS VGA (640x480) Image Sensor. (27 Feb. 2014). <http://www.ovt.com/products/sensor.php?id=146>.
- [4] 2015. JINS MEME Eyewear. (14 Oct. 2015). <https://jins-meme.com/en/>.
- [5] 2015. STM3232-bitARMCortex MCUs - STMicroelectronics. (24 Mar. 2015). Retrieved September 10, 2016 from <http://www.st.com/stm32>.
- [6] 2015. Tobii Glasses: Mobile Eye Tracker for real world research. (15 Oct. 2015). <http://www.tobii.com>.
- [7] 2016. Research on Drowsy Driving. (23 Jan. 2016). <http://www.nhtsa.gov/Driving+Safety/Drowsy+Driving/scope-of-the-problem>.
- [8] 2016. SMI Eye Tracking Glasses. (2016). <http://www.eyetracking-glasses.com/products/eye-tracking-glasses-2-wireless/technology/>.
- [9] 2017. Center of Excellence for Mobile Sensor Data-to-Knowledge. (2017). <https://md2k.org>.
- [10] Saeed Abdullah, Elizabeth L Murnane, Mark Matthews, Matthew Kay, Julie A Kientz, Geri Gay, and Tanzeem Choudhury. 2016. Cognitive rhythms: unobtrusive and continuous sensing of alertness using a mobile phone. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing*. ACM, 178–189.
- [11] Christer Ahlstrom, Marcus Nyström, Kenneth Holmqvist, Carina Fors, David Sandberg, Anna Anund, Göran Kecklund, and Torbjörn Åkerstedt. 2013. Fit-for-duty test for estimation of drivers' sleepiness level: eye movements improve the sleep/wake predictor. *Transportation research part C: emerging technologies* 26 (2013), 20–32.
- [12] Rizwan Ahmad. 2016. Understanding the Language of the Eye: Detecting and Identifying Eye Events in Real Time via Electrooculography. (2016). <http://www.escholarship.org/uc/item/53n1p2z5>
- [13] Tobias Appel, Thiago Santini, and Enkelejda Kasneci. 2016. Brightness-and motion-based blink detection for head-mounted eye trackers. In *Proceedings of the 2016 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct*. ACM, 1726–1735.
- [14] Ioana Bacivarov, Mircea Ionita, and Peter Corcoran. 2008. Statistical models of appearance for eye tracking and eye-blink detection and measurement. *IEEE transactions on consumer electronics* 54, 3 (2008).
- [15] P. L. Benitez, G. H. Kamimori, T. J. Balkin, A. Greene, and M. L. Johnson. 2009. Modeling fatigue over sleep deprivation, circadian rhythm, and caffeine with a minimal performance inhibitor model. *Meth. Enzymol.* 454 (2009), 405–421.
- [16] Luis Miguel Bergasa, Jesús Nuevo, Miguel A Sotelo, Rafael Barea, and Maria Elena Lopez. 2006. Real-time system for monitoring driver vigilance. *IEEE Transactions on Intelligent Transportation Systems* 7, 1 (2006), 63–77.
- [17] Serge Boverie and Alain Giralte. 2008. Driver vigilance diagnostic based on eyelid movement observation. *IFAC Proceedings Volumes* 41, 2 (2008), 12831–12836.
- [18] J. E. Bower. 2014. Cancer-related fatigue—mechanisms, risk factors, and treatments. *Nat Rev Clin Oncol* 11, 10 (Oct 2014), 597–609.
- [19] Andreas Bulling, Daniel Roggen, and Gerhard Tröster. 2008. It's in your eyes: towards context-awareness and mobile HCI using wearable EOG goggles. In *Proceedings of the 10th international conference on Ubiquitous computing*. ACM, 84–93.
- [20] Andreas Bulling, Daniel Roggen, and Gerhard Tröster. 2009. Wearable EOG goggles: Seamless sensing and context-awareness in everyday environments. *Journal of Ambient Intelligence and Smart Environments* 1, 2 (2009), 157–171.
- [21] Andreas Bulling, Jamie A Ward, Hans Gellersen, and Gerhard Tröster. 2011. Eye movement analysis for activity recognition using electrooculography. *IEEE transactions on pattern analysis and machine intelligence* 33, 4 (2011), 741–753.
- [22] Michael Chau and Margrit Betke. 2005. *Real time eye tracking and blink detection with usb cameras*. Technical Report. Boston University Computer Science Department.
- [23] Elizabeth J. Corwin, Laura Cousino Klein, and Kristin Rickelman. 2002. Predictors of Fatigue in Healthy Young Adults: Moderating Effects of Cigarette Smoking and Gender. *Biological Research For Nursing* 3, 4 (2002), 222–233. DOI: <http://dx.doi.org/10.1177/109980040200300407> arXiv:<http://dx.doi.org/10.1177/109980040200300407> PMID: 12184665.
- [24] Anirban Dasgupta, Anjith George, SL Happy, and Aurobinda Routray. 2013. A Vision-Based System for Monitoring the Loss of Attention in Automotive Drivers. *IEEE Trans. Intelligent Transportation Systems* 14, 4 (2013), 1825–1838.
- [25] Tomas Drutarovsky and Andrej Fogelton. 2014. Eye blink detection using variance of motion vectors. In *European Conference on Computer Vision*. Springer, 436–448.
- [26] John D. Fisk, Amanda Pontefract, Paul G. Ritvo, Catherine J. Archibald, and T.J. Murray. 1994. The Impact of Fatigue on Patients with Multiple Sclerosis. *Canadian Journal of Neurological Sciences / Journal Canadien des Sciences Neurologiques* 21, 1 (001 002 1994), 9–14. DOI: <http://dx.doi.org/10.1017/S0317167100048691>
- [27] J. Friedman, T. Hastie, and R. Tibshirani. 2001. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer Series in Statistics, Vol. 1.
- [28] Wolfgang Fuhl, Thiago Santini, David Geisler, Thomas Kübler, Wolfgang Rosenstiel, and Enkelejda Kasneci. 2016. Eyes wide open? eyelid location and eye aperture estimation for pervasive eye tracking in real-world scenarios. In *Proceedings of the 2016 ACM International*

- Joint Conference on Pervasive and Ubiquitous Computing: Adjunct*. ACM, 1656–1665.
- [29] Kyosuke Fukuda, John A Stern, Timothy B Brown, and Michael B Russo. 2005. Cognition, blinks, eye-movements, and pupillary movements during performance of a running memory task. *Aviation, space, and environmental medicine* 76, 7 (2005), C75–C85.
 - [30] X. Y. Gao, Y. F. Zhang, W. L. Zheng, and B. L. Lu. 2015. Evaluating driving fatigue detection algorithms using eye tracking glasses. In *2015 7th International IEEE/EMBS Conference on Neural Engineering (NER)*. 767–770. DOI: <http://dx.doi.org/10.1109/NER.2015.7146736>
 - [31] I Garcia, Sebastian Bronte, Luis Miguel Bergasa, Javier Almazán, and J Yebes. 2012. Vision-based drowsiness detector for real driving conditions. In *Intelligent Vehicles Symposium (IV), 2012 IEEE*. IEEE, 618–623.
 - [32] N. Goel, M. Basner, H. Rao, and D. F. Dinges. 2013. Circadian rhythms, sleep deprivation, and human performance. *Prog Mol Biol Transl Sci* 119 (2013), 155–190.
 - [33] Kristen Grauman, Margrit Betke, James Gips, and Gary R Bradski. 2001. Communication via eye blinks-detection and duration analysis in real time. In *Computer Vision and Pattern Recognition, 2001. CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on*, Vol. 1. IEEE, I–I.
 - [34] Martijn Haak, Steven Bos, Sacha Panic, and IJM Rothkrantz. 2009. Detecting stress using eye blinks and brain activity from EEG signals. *Proceeding of the 1st driver car interaction and interface (DCII 2008)* (2009), 35–60.
 - [35] E. Havlikova, J. Rosenberger, I. Nagyova, B. Middel, T. Dubayova, Z. Gdovinova, J. P. Van Dijk, and J. W. Groothoff. 2008. Impact of fatigue on quality of life in patients with Parkinsons disease. *European Journal of Neurology* 15, 5 (2008), 475–480. DOI: <http://dx.doi.org/10.1111/j.1468-1331.2008.02103.x>
 - [36] Tianyi Hong and Huabiao Qin. 2007. Drivers drowsiness detection in embedded system. In *Vehicular Electronics and Safety, 2007. ICVES. IEEE International Conference on*. IEEE, 1–5.
 - [37] Chin-Shun Hsieh and Cheng-Chi Tai. 2013. An improved and portable eye-blink duration detection system to warn of driver fatigue. *Instrumentation Science & Technology* 41, 5 (2013), 429–444.
 - [38] Michael Ingre, Torbjörn Åkerstedt, Björn Peters, Anna Anund, and Göran Kecklund. 2006. Subjective sleepiness, simulated driving performance and blink duration: examining individual differences. *Journal of Sleep Research* 15, 1 (2006), 47–53. DOI: <http://dx.doi.org/10.1111/j.1365-2869.2006.00504.x>
 - [39] Amir-Homayoun Javadi, Zahra Hakimi, Morteza Barati, Vincent Walsh, and Lili Tcheang. 2015. SET: a pupil detection method using sinusoidal approximation. *Frontiers in neuroengineering* 8 (2015).
 - [40] Xianta Jiang, Geoffrey Tien, Da Huang, Bin Zheng, and M Stella Atkins. 2013. Capturing and evaluating blinks from video-based eyetrackers. *Behavior research methods* 45, 3 (2013), 656–663.
 - [41] Hughes JR and Hatsukami D. 1986. Signs and symptoms of tobacco withdrawal. *Archives of General Psychiatry* 43, 3 (1986), 289–294. DOI: <http://dx.doi.org/10.1001/archpsyc.1986.01800030107013>
 - [42] Hughes JR, Higgins ST, Bickel WK, and et al. 1991. Caffeine self-administration, withdrawal, and adverse effects among coffee drinkers. *Archives of General Psychiatry* 48, 7 (1991), 611–617. DOI: <http://dx.doi.org/10.1001/archpsyc.1991.01810310029006>
 - [43] Won Oh Lee, Eui Chul Lee, and Kang Ryoung Park. 2010. Blink detection robust to various facial poses. *Journal of neuroscience methods* 193, 2 (2010), 356–372.
 - [44] I. S. Lobentanz, S. Asenbaum, K. Vass, C. Sauter, G. Klsch, H. Kollegger, W. Kristoferitsch, and J. Zeithofer. 2004. Factors influencing quality of life in multiple sclerosis patients: disability, depressive mood, fatigue and sleep quality. *Acta Neurologica Scandinavica* 110, 1 (2004), 6–13. DOI: <http://dx.doi.org/10.1111/j.1600-0404.2004.00257.x>
 - [45] Amol M Malla, Paul R Davidson, Philip J Bones, Richard Green, and Richard D Jones. 2010. Automated video-based measurement of eye closure for detecting behavioral microsleep. In *Engineering in Medicine and Biology Society (EMBC), 2010 Annual International Conference of the IEEE*. IEEE, 6741–6744.
 - [46] R Martins and JM Carvalho. 2015. Eye blinking as an indicator of fatigue and mental load-a systematic review. *Occupational Safety and Hygiene III* (2015), 231–235.
 - [47] Addison Mayberry, Pan Hu, Benjamin Marlin, Christopher Salthouse, and Deepak Ganesan. 2014. iShadow: Design of a Wearable, Real-time Mobile Gaze Tracker. In *Proceedings of the 12th Annual International Conference on Mobile Systems, Applications, and Services (MobiSys '14)*. ACM, New York, NY, USA, 82–94. DOI: <http://dx.doi.org/10.1145/2594368.2594388>
 - [48] Addison Mayberry, Yamin Tun, Pan Hu, Duncan Smith-Freedman, Deepak Ganesan, Benjamin M. Marlin, and Christopher Salthouse. 2015. CIDER: Enabling Robustness-Power Tradeoffs on a Computational Eyeglass. In *Proceedings of the 21st Annual International Conference on Mobile Computing and Networking (MobiCom '15)*. ACM, New York, NY, USA, 400–412. DOI: <http://dx.doi.org/10.1145/2789168.2790096>
 - [49] Lindsey K McIntire, John P McIntire, R Andy McKinley, and Chuck Goodyear. 2014. Detection of vigilance performance with pupillometry. In *Proceedings of the Symposium on Eye Tracking Research and Applications*. ACM, 167–174.
 - [50] T. R. Mendoza, X. S. Wang, C. S. Cleeland, M. Morrissey, B. A. Johnson, J. K. Wendt, and S. L. Huber. 1999. The rapid assessment of fatigue severity in cancer patients: use of the Brief Fatigue Inventory. *Cancer* 85, 5 (Mar 1999), 1186–1196.
 - [51] Inga Meyhöfer, Veena Kumari, Antje Hill, Nadine Petrovsky, and Ulrich Ettinger. 2016. Sleep deprivation as an experimental model system for psychosis: Effects on smooth pursuit, prosaccades, and antisaccades. *Journal of Psychopharmacology* (2016), 0269881116675511.

- [52] Tim Morris, Paul Blenkhorn, and Farhan Zaidi. 2002. Blink detection for real-time eye tracking. *Journal of Network and Computer Applications* 25, 2 (2002), 129–143.
- [53] TL Morris and James C Miller. 1996. Electrooculographic and performance indices of fatigue during simulated flight. *Biological psychology* 42, 3 (1996), 343–360.
- [54] M Patel, SKL Lal, Diarmuid Kavanagh, and Peter Rossiter. 2011. Applying neural network analysis on heart rate variability data to assess driver fatigue. *Expert systems with Applications* 38, 6 (2011), 7235–7242.
- [55] Leo Pauly and Deepa Sankar. 2015. Detection of drowsiness based on hog features and svm classifiers. In *Research in Computational Intelligence and Communication Networks (ICRCICN), 2015 IEEE International Conference on*. IEEE, 181–186.
- [56] Marco Pedrotti, Shengguang Lei, Jeronimo Dzaack, and Matthias Rötting. 2011. A data-driven algorithm for offline pupil signal preprocessing and eyeblink detection in low-speed eye-tracking protocols. *Behavior Research Methods* 43, 2 (2011), 372–383.
- [57] A. Picot, A. Caplier, and S. Charbonnier. 2009. Comparison between EOG and high frame rate camera for drowsiness detection. In *Applications of Computer Vision (WACV), 2009 Workshop on*. 1–6. DOI : <http://dx.doi.org/10.1109/WACV.2009.5403120>
- [58] Antoine Picot, Alice Caplier, and Sylvie Charbonnier. 2009. Comparison between EOG and high frame rate camera for drowsiness detection. In *Applications of Computer Vision (WACV), 2009 Workshop on*. IEEE, 1–6.
- [59] R. Schleicher, N. Galley, S. Briest, and L. Galley. 2008. Blinks and saccades as indicators of fatigue in sleepiness warnings: looking tired? *Ergonomics* 51, 7 (2008), 982–1010. DOI : <http://dx.doi.org/10.1080/00140130701817062> arXiv:<http://dx.doi.org/10.1080/00140130701817062> PMID: 18568959.
- [60] Robert Schleicher, Niels Galley, Susanne Briest, and Lars Galley. 2008. Blinks and saccades as indicators of fatigue in sleepiness warnings: looking tired? *Ergonomics* 51, 7 (2008), 982–1010.
- [61] Paul Smith, Mubarak Shah, and Niels da Vitoria Lobo. 2003. Determining driver visual attention with one camera. *IEEE transactions on intelligent transportation systems* 4, 4 (2003), 205–218.
- [62] M. D. Stein and P. D. Friedmann. 2005. Disturbed sleep and its relationship to alcohol use. *Subst Abuse* 26, 1 (Mar 2005), 1–13.
- [63] J. A. Stern, D. Boyer, David J. Schroeder, United States., and Civil Aeromedical Institute. 1994. *Blink rate as a measure of fatigue [microform] : a review / J.A. Stern, D. Boyer, D.J. Schroeder*. U.S. Dept. of Transportation, Federal Aviation Administration, Office of Aviation Medicine ; Available to the public through the National Technical Information Service Washington, D.C. : Springfield, Va. i, 12 p. ; pages. <http://purl.access.gpo.gov/GPO/LPS84339>
- [64] Federico M Sukno, Sri-Kaushik Pavani, Constantine Butakoff, and Alejandro F Frangi. 2009. Automatic assessment of eye blinking patterns through statistical shape models. In *International Conference on Computer Vision Systems*. Springer, 33–42.
- [65] Louis Tijernia, Stoltzfus Duane Gleckler, Mark nad, Scott Johnston, Michael J. Goodman, and Walter W. Wierwille. 1999. *A Preliminary Assessment of Algorithms for Drowsy and Inattentive Driver Detection on the Road*. Technical Report. Springfield VA.
- [66] Marc Tonsen, Xucong Zhang, Yusuke Sugano, and Andreas Bulling. 2016. Labelled pupils in the wild: a dataset for studying pupil detection in unconstrained environments. In *Proceedings of the Ninth Biennial ACM Symposium on Eye Tracking Research & Applications*. ACM, 139–142.
- [67] W. W. Wierwille, S. S. Wreggit, C. L. Kirn, L. A. Ellsworth, and R. J. Fairbanks. 1994. *Research on Vehicle-Based Driver Status/Performance Monitoring; Development, Validation, and Refinement of Algorithms For Detection of Driver Drowsiness*. Technical Report. Springfield VA.
- [68] Ann Williamson. 2007. Predictors of Psychostimulant Use by Long-Distance Truck Drivers. *American Journal of Epidemiology* 166, 11 (2007), 1320. DOI : <http://dx.doi.org/10.1093/aje/kwm205>
- [69] Junwen Wu and Mohan M Trivedi. 2007. Simultaneous eye tracking and blink detection with interactive particle filters. *EURASIP Journal on Advances in Signal Processing* 2008, 1 (2007), 823695.
- [70] Cui Xu, Ying Zheng, and Zengfu Wang. 2008. Efficient eye states detection in real-time for drowsy driving monitoring system. In *Information and Automation, 2008. ICIA 2008. International Conference on*. IEEE, 170–174.
- [71] Chuang-Wen You, Nicholas D Lane, Fanglin Chen, Rui Wang, Zhenyu Chen, Thomas J Bao, Martha Montes-de Oca, Yuting Cheng, Mu Lin, Lorenzo Torresani, and others. 2013. Carsafe app: Alerting drowsy and distracted drivers using dual cameras on smartphones. In *Proceeding of the 11th annual international conference on Mobile systems, applications, and services*. ACM, 13–26.

Received February 2017; revised March 2017; accepted June 2017