

WISEglass: Smart eyeglasses recognising context

Florian Wahl, Martin Freund, Oliver Amft
ACTLab, Chair of Sensor Technology, University of Passau, <http://actlab.uni-passau.de>
{wahl, freund, amft}@fim.uni-passau.de

ABSTRACT

We investigated how regular eyeglasses could be extended with multi-modal sensing and processing functions to support context-awareness applications. Our aim was to leverage eyeglasses as a platform for acquiring and processing context information according to the wearer's needs. The WISEglass architecture consists of inertial motion, environmental light, and pulse sensors, processing and wireless data transmission functionality, besides a rechargeable battery. We implemented prototypes of WISEglass and evaluated them in three application scenarios: daily activity recognition, screen-use detection, and heart rate estimation. We conducted a daily activity study with nine participants, each wearing WISEglass and recording for one day. When evaluating daily activity recognition, we obtained 77% average accuracy for continuous recognition using Gaussian Mixture Models and classifier reject to ignore null class data. Using the light sensor for detecting screen-use, yielded 80% accuracy. Against a chest-worn ECG reference, our heart rate estimation showed a difference below 10 beats for stationary activities across the full recording day. We concluded that smart eyeglasses provide information from a single measurement spot that is relevant in various context recognition applications.

Keywords

context inference, mobile sensing, smart glasses, activity recognition, eyewear

Categories and Subject Descriptors

I.5.m. [Pattern Recognition]: Miscellaneous

1. INTRODUCTION

Context awareness has opened a vast spectrum of applications that benefit from momentary information on user activity, environment, physiology, and similar. Wearable devices are often key to provide context information, as sensors




Figure 1: WISEglass prototype. Baseboard, battery, accelerometer, gyroscope, and magnetometer were mounted on the outside of a eyeglasses temple. The pulse sensor was mounted on the inside of a temple. The light sensor was integrated onto the bridge.

could be placed comfortably at body locations such as wrist, leg, chest, or the ear. While hearing aids and Bluetooth headsets got broadly accepted as daily accessories, head-worn wearables were rarely proposed for context awareness due to potential obtrusiveness of head-attached devices. Some niche applications include ear-worn computers to manage information [17], or the ear-worn e-AR device to monitor activity and physiology [14]. There is nevertheless substantial context information available around the head [?]. Eyeglasses are regularly worn accessories that have a unique opportunity to carry sensors and process data at the head, thus fill a gap for many context-awareness applications. Moreover, eyeglasses are being worn by millions of people, for improved sight, but also as sunglasses, sports glasses, etc.

Substantial work has been dedicated to developing smart eyewear already. However, functionality was often centred around displaying information in front of the wearers eyes, e.g. for augmented or virtual reality applications [3, 20]. GoogleGlass and others established microinteractions as key feature of smart eyewear and focused on interaction and displaying information using glass-attachable electronics. We believe that instead of displays and direct interaction, smart eyeglasses can be built for context-aware applications, where the focus is on sensing and processing, rather than interaction. Smart eyeglasses would hence focus on acquiring and processing context information relevant to the wearer, but minimise the risk of stigmatising or raising privacy concerns.

Besides the integration of sensing and processing functions, the benefit of smart eyeglasses for different context

recognition applications needs confirmation. Previous research on head-attached devices and smart eyeglasses focused on detecting selected activities, such as walking and reading behaviour using inertial sensors and the Electrooculogram (EOG) [14, 13]. A set of daily living activities was not investigated (see related work for details). Moreover, eyeglasses could be used to acquire environmental and physiological information continuously, if a suitable smart eyeglasses implementation is found.

In this paper, we propose smart eyeglasses as a platform for a variety of context-aware applications. Our aim was to confirm that smart eyeglasses could be built and used instead of various sensors placed at different body positions. We implemented WISEglass with various sensors, including inertial motion, environmental light, and optical heart rate and investigated three application scenarios: daily activity recognition, screen-use detection, and heart rate estimation. This paper provides the following contributions:

1. We present our WISEglass eyeglasses architecture, integrating multi-modal sensing functions corresponding to frequent context types, including user activity, environmental state, and user physiology. We detail the embedded design in regular eyeglasses and data processing.
2. We evaluate WISEglass prototypes in a daily life study with nine individuals that did not wear eyeglasses regularly, including a programme of 20 daily living activities. Based on the study data, we confirm that the sensor placement and wearing across a day is suitable to implement the context recognition tasks for each application scenario.

Towards smart, regular eyeglasses there are many challenges to resolve. Our aim in this work is to establish that context recognition - and thus assistance to wearers - benefits from using smart eyeglasses due to their position at the head. Earlier work on smart eyewear primarily considered interaction and conveying visual information to the wearer. Devices often clipped onto eyeglasses only. Our approach is different, as we focus on single-point, multi-modal sensing and processing, integrated into a typical accessory, and serving different applications. We expect that in a subsequent step, continuous miniaturisation will allow us to even further integrate functionality into unobtrusive eyeglasses.

2. RELATED WORK

Head mounted sensors were considered for context awareness before. Aziz et al. [2] first used e-AR, an ear-worn sensor, for monitoring patients after abdominal surgery. In a study with 20 users they investigated motion patterns in data obtained from two dual-axis accelerometers. They used a pulse oximeter clipped to the users earlobe to monitor heart rate and oxygen saturation. No performance analysis was provided. Atallah et al. [1] compared seven different on-body motion sensor locations while grouping the activities by physical intensity level. At the head, they used e-AR and found that the ear location delivers good results for 4 out of 5 activity levels. In contrast to e-AR, our sensors are worn as eyeglasses. WISEglass features a 3-axis accelerometer, 3-axis gyroscope, 3-axis magnetometer, a RGB light sensor, and a pulse sensor, used to detect a wide range of activities of daily living.

Ishimaru et al. [12] used built-in proximity sensor and accelerometer of Google Glass to classify 5 different activities in a laboratory setting. For user dependent models they achieved an average accuracy of 67% using blink frequency alone, and 82% in combination with the accelerometer. Ishimaru et al. [13] used EOG and acceleration signals to classify reading, typing, eating, and talking activities, recorded from two users. They achieved an average of 70% accuracy using user independent models. In our work, we investigate a full set of 20 activities of daily living grouped into 9 clusters in a field study rather than a laboratory setting. Moreover, our approach yielded comparable performance using user-independent models. Hernandez et. al [10] used the accelerometer, gyroscope, and camera to obtain heart and respiration rate from Google Glass, while participants were asked to remain still. They also attempted to detect emotion of the wearer in combination with an arm-worn sensor [11]. In contrast to Google Glass, which was focussed on micro-interactions and displaying information, we aim at integrating sensing and processing capabilities into regular eyeglasses. We avoid the use of a camera in order to preserve privacy. Furthermore, we gathered measurements in the field instead of a controlled environment, e.g. heart rate is estimated during different activities.

Previous works often used multiple on-body sensing locations. Bao and Intille [4] used bi-axial accelerometers in five different body locations. They detected a set of 20 different daily living activities and achieved accuracies of up to 84% in 20 participants. Measuring in multiple sensing locations complicates the setup process. In contrast, we propose a sensor system integrated into an accessory, thus minimising the required setup to putting on a pair of eyeglasses.

Exposure to light entrains our circadian clock, but requires light measurements close to the eye. Figueiro et al.[8] compared three different wearable sensor systems for light exposure measurement. Besides a wrist-worn device and a button worn at the collar, they found that a head-worn solution performs best as the sensor is close to the eyes. For WISEglass, we integrated a light sensor into the bridge of the glasses, thus optimising placement for light exposure measurement. As blue light is dominant in the entrainment of the circadian clock, it is of interest to know the time and duration of screen use that emits elevated blue light levels.

In the past, a variety of sensors have been placed around the head for different applications. WISEglass integrates existing sensors into *one* wearable accessory that allows wearers to benefit from smart eyeglasses in a wide range of applications, which previously required separate sensing systems.

3. WISEGLASS ARCHITECTURE

WISEglass is an approach to retrofit regular eyeglasses with a multi-modal sensor system for different context-aware applications. WISEglass could be used as a everyday accessory, just as regular eyeglasses are used today, or as a special-application device. For our design investigations, we specifically chose standard of-the-shelf eyeglasses and avoided designing the device from scratch, as to ensure that the typical eyeglasses form factor was maintained.

Processor, flash memory, communications interfaces, power controller, as well as inertial sensors were integrated onto a baseboard, the central unit of WISEglass. In our first prototypes, battery and baseboard were bound together using shrink tubing, however we consider that the units could be

embedded in opposite ends of the eyeglasses frame in future versions.

With the light sensor, we investigated the embedding of components into the glasses’ bridge. The light sensor magnet wiring was routed through miniature holes in the bridge into a milled channel along the top rim over the end pieces ending at the baseboard. We picked the bridge location as it is closest to the eyes and allows us to obtain most accurate light exposure measurements. Environmental light entrains our circadian clock and measuring light exposure is relevant for circadian phase guidance [18]. Another use for light sensor data could be in indoor/outdoor detection.

The pulse sensor was linked to a cable tie to enable users to customise the wearing position. The end of the cable tie was fixed to the temple tip using shrink tubing. A small block of foam was placed between the pulse sensor and the temple of the frame to ensure enough fixation to obtain a relevant signal while maintaining wearing comfort. In further versions, the implementation could be embedded directly into the temple, thus avoiding manual adjustments. Our prototype of WISEglass is depicted in Figure 1. Figure 2 depicts the WISEglass hardware architecture.

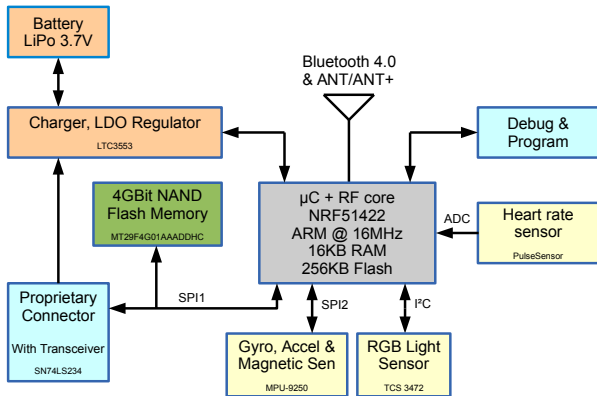


Figure 2: WISEglass hardware architecture. Separate SPI busses are used for communication of the main controller unit (MCU) with the flash memory and the inertial measurement unit to optimise data throughput.

The main controller unit (MCU) was a nRF51422 from Nordic Semiconductor that provides a 32 bit ARM Cortex-M0 core running at 16 MHz with a Bluetooth Low Energy and ANT+ wireless module as a System on Chip. We added a 512 MByte NAND flash to complement on-chip memory when storing sensor data. WISEglass is powered by a 3.7 V, 330 mAh Lithium Polymer battery.

For motion and orientation estimates, a MPU-9250 from InvenSense was used. The MPU-9250 provides accelerometer, gyroscope, magnetometer, all 3-axis inertial sensors, and a digital motion processor in a single package. The inertial sensors were AD-converted into 16 bits. Sensitivity can be configured to optimise accuracy depending on the application. For the accelerometer the ranges ± 2 g, 4 g, 8 g, and 16 g can be selected. The gyroscope offers ± 250 dps, 500 dps, 1000 dps, and 2000 dps. The magnetometer range is fixed to $\pm 4912 \mu\text{T}$.

For light intensity measurements, a TCS3472 from ams was used. The TCS3472 senses light intensities for the red,

green, blue, and clear light spectra separately. The sensor has an integrated IR filter and was connected to the MCU through the I²C interface.

To obtain heart rate we added a PulseSensor [9] which is based on the reflexive photoplethysmography (PPG) principle. We found the temple to be a good location for measuring pulse as it eliminated the need for an ear clip.

4. DATA COLLECTION

4.1 Study methodology

To evaluate our approach, we conducted a study with nine participants (3 female, 6 male, between 20 and 27 years). Participants either did not require prescription glasses or wore contact lenses. Each participant was given a pair of smart glasses. For heart rate reference measurements, participants were fitted with a CamNtech Activwave Cardio ECG device. Annotations were performed by an observer using the ACTLog application for Android [16]. Figure 3 shows the sensor setup for a sample of the activities considered.

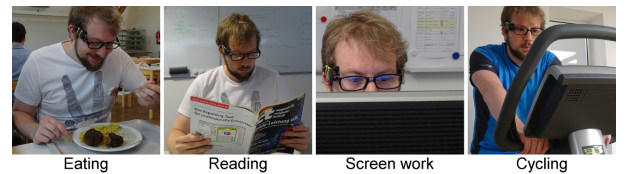


Figure 3: WISEglass worn during a subset of the activities considered in our full-day study protocol. In the protocol a variety of daily activities were covered and subsequently grouped into nine activity clusters for the recognition analysis. During our study we collected 66.08 hours of data from nine participants.

We recorded up to two participants per day. Participants received three complimentary meals during the recording day and a 25 Euro Amazon voucher as compensation.

During recordings we configured the accelerometer to a range of ± 4 g and the gyroscope to a range of ± 500 dps. Accelerometer, gyroscope, and magnetometer were sampled at a rate of 50 Hz. Light measurements were integrated with a time constant of 154 msec yielding a sampling rate of 6.5 Hz and were upsampled to 50 Hz using a latest value strategy. The pulse sensor voltage was sampled at a rate of 50 Hz. The ECG reference was sampled at 200 Hz. From the study, we acquired a total 66.08 hours of data.

4.2 Study protocol

Participants were first introduced to the study protocol and signed an informed consent form. Participants then performed a scripted study protocol as shown in Figure 4. During the exit interview, participants evaluated their experience using the system usability scale (SUS) [6]. The protocol was designed to cover typical daily activities. Table 1 shows how activities were combined to activity clusters for further evaluation.

5. EVALUATION METHODOLOGY

5.1 Data preprocessing

Data streams from the WISEglass, ECG device, and smartphone were aligned according to their time stamps, merged,

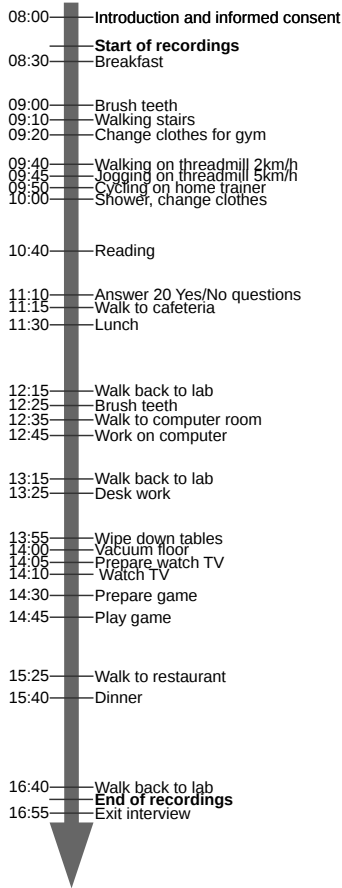


Figure 4: Protocol of the daily activity study. Recordings started before breakfast and ended before the exit interview. Sensors were temporarily removed while participants were in the gym locker room.

resampled to 50 Hz, and cropped to a common time axis. Subsequently, annotations were post-processed by visual inspection using the ACTLab MARKER toolbox for MATLAB.

5.2 Feature extraction

Features were extracted using pandas [15] and numpy [19] libraries of Python. For recognition of daily activities and the detection of screen-use, a sliding window over data samples with a window size of $n = 1500sa$ (30 s) and a step size of $s = 50sa$ (1 s) was used. Feature vectors were standardised and normalised before further processing. The sliding windows were assigned to the class that was most represented in the window. Motion features were computed over each of the three axis of the accelerometer (a_x , a_y , and a_z) and the gyroscope (g_x , g_y , and g_z). We computed the L^2 norm for the acceleration axis a_{norm} for each sample. Features listed in Table 2 were subsequently calculated for a_{norm} , a_x , a_y , a_z , g_x , g_y , and g_z per window.

In order to detect screen-use based on light sensor data, we derived features on the RGB and clear light intensity values named r , g , b , and c respectively. We calculated the median $\text{med}()$ and number of median crossings $\text{mcr}()$ for each color channel per window. To express spectral com-

Table 1: Activity clusters and total duration.

No	Cluster	Activity & total duration [min:sec]
1	Eat	Breakfast [84:43]
		Lunch [156:49]
		Dinner [181:53]
2	Walk	Lab to Bathroom [17:37]
		Lab to gym [36:27]
		On treadmill 2km/h [47:06]
		Gym to lab [35:56]
		Lab to cafeteria [20:11]
		Queuing for lunch [6:28]
		Picking up lunch [13:20]
		Cafeteria to lab [25:28]
		Lab to restaurant [58:36]
Restaurant to lab [55:26]		
3	Brush	Teeth [39:17]
4	Stairs	Walking [19:49]
5	Jog	On treadmill 5km/h [45:27]
6	Cycle	On gym trainer [91:14]
7	Read	A book [276:31]
		Desk work [273:34]
8	Screen	Computer work [253:50]
		Watching movie [187:56]
9	Cleaning	Vacuuming [45:22]
		Wiping tables [45:02]

Table 2: Motion features calculated for a vector of data samples w with the window size n with elements $1 \dots N$, where w_i is the i -th sample in the window. w' and w'' are the first and second derivative of w determined between two neighbouring samples. $\text{zcr}()$ denotes the number of zero crossings.

1. \bar{w}	2. $\sum_{i=1}^n w_i$, if $w_i > 0$	3. $\sum_{i=1}^n w_i^2$
4. $\sigma^2(w)$	5. $\sum_{i=1}^n w_i$, if $w_i < 0$	6. $\sum_{i=1}^n w_i $
7. $\max(w)$	8. $\sum_{i=1}^n w'_i$, if $w'_i > 0$	9. $\sum_{i=1}^n w_i $
10. $\min(w)$	11. $\sum_{i=1}^n w'_i$, if $w'_i < 0$	12. $\text{zcr}(w)$
13. w_1	14. $\sum_{i=1}^n w''_i$, if $w''_i > 0$	15. $\sum_{i=1}^n w'_i $
16. w_N	17. $\sum_{i=1}^n w''_i$, if $w''_i < 0$	18. $\frac{\bar{w}}{\sigma(w)}$
19. $w_N - w_1$	20. $\max(w) - \min(w)$	21. $\sqrt{\sum_{i=1}^n w_i^2}$
22. $\sum_{i=1}^n w'_i$	23. $\sum_{i=1}^n w'_i $	24. $\sum_{i=1}^n w''_i $
25. $\sum_{i=1}^n w''_i$		

position of light, we computed the median of the ratio for each combination of two color channels c_1 , c_2 defined as $\text{crm}(c_1, c_2) = \text{med}(\frac{c_1}{c_2})$.

Heart rate estimation was performed by counting the number of positive peaks in the raw voltage signal of the pulse sensor. For peak detection, we used an existing peak detector [7]. We experimented with different minimum peak distances, ranging from 0 to 25 sa in steps of 5 sa, but found 20 sa to be the best for our needs. As reference heart rate, the estimate provided by the ECG sensor was used. We computed the mean of the heart rate estimate from the ECG sensor per window. The calculation of the heart rate error was performed on windows with no overlap using window size $n = 1500$ samples.

5.3 Feature selection

We applied principle component analysis (PCA) to reduce the number of features. We configured the PCA algorithm to select the minimal number of features, such that at least 99.9% of the variance in the training data was explained

resulting in the selection of 78 out of 175 features for recognising daily activities. For screen-use detection, 9 out of 14 features were selected.

5.4 Cross-validation

To ensure generality and stability of our statistical model, we used Leave-One-Out (LOO) cross-validation on a per-participant basis. In every cross-validation fold the classifier was trained on all available feature instances from eight participants. To evaluate model performance the data of the remaining participant was used.

5.5 Classification with reject option

Classification of the motion-based activity and light-based screen detection was performed using a Gaussian Mixture Models (GMM) classifier as described in [5]. After training one GMM for each class, we computed the probabilities for each test sample. The classifier then selected the class for which the highest probability occurred. To also predict unlabelled data we implemented classifier reject to avoid making assumptions on the data in the null class. First we obtained the likelihoods for each test sample from the GMM model. If the maximum likelihood for a test sample was below the threshold, the classifier rejected it as a null class instance and otherwise selected the class of the GMM yielding the maximum likelihood. A sweep search was performed to find the optimal threshold. We used all maximum likelihoods, which occurred during training in each LOO cross-validation fold and picked the threshold yielding the best F1 score. Results per fold were averaged to show the performance for an optimal F1 score.

5.6 Heart rate estimation performance

To validate heart rate estimation from the pulse sensor, we computed the root-mean-square error (RMSe) between the ECG reference and the pulse sensor. Both signals were filtered using a bandpass filter before peak detection was applied. RMSe computation was limited to data samples, where the ECG heart rate was between 40 bpm and 150 bpm according to the reference estimate. We calculated RMSe for three data subsets to investigate the effect of motion artefacts: (1) All available data. (2) Data from stationary clusters eat, read, and screen. (3) Data from classes stairs, walking, and jogging on treadmill where the participant is assumed to be moving.

6. RESULTS

Here we show the results for recognising user activities using inertial sensor data, detecting screen-use using environmental light, and estimating heart rate using the WISEglass pulse sensor.

6.1 Daily activity recognition

We evaluated the recognition of the activity clusters (Tab. 1) using the GMM classifier and LOO cross-validation. GMM performance was analysed by varying the number of Gaussian mixture components between $c = \{1, 3, 5, 7\}$ per class, as well as analysing the covariance matrix configuration (full or diagonal). All configurations using diagonal covariance matrix yielded similar accuracies, except for the configuration with one mixture only. The full covariance matrices generally perform marginally better (3%) than diagonal, for the same number of mixture components. For our further

analysis, we resided on using a diagonal covariance matrices due to the increased modelling and computation complexity that a full covariance matrix would introduce. Using diagonal covariance matrices, highest accuracy was 77% using 3 GMM mixture components. The lowest performance was 70% using 1 GMM mixture component. Based on the results, we used 3 GMM mixture components per class for all subsequent analyses.

Daily activity recognition varied in accuracy between 70% and 84% depending on the participant considered for testing. The confusion matrix shown in Figure 5 depicts the class assignment averaged over all participants. For the confusion analysis LOO cross-validation results were averaged across folds. The confusion results indicate per-class performances between 80% and 90%, except for “cycle”. The performance of the cycle cluster was only 50% due to confusions with “eat”, “walk”, “brush”, and “clean” clusters. As “cycle” took place at slow pace on a gym trainer (see Fig. 3), only minor head motion occurred, thus resulting in motion patterns similar to other classes. The activity clusters “read” and “screen” show the largest confusions. Based on motion data it seems plausible that “read” and “screen” clusters were confused as they showed similar head motions.

At a first glance, the classifier reject yields a null class accuracy of only 20% and several confusions with other activity clusters. We found that the performance is related to our recording protocol and the diversity of motion in the activity clusters. Our activity clusters were comprehensive on the activities contained in it, thus there was simply not much unlabelled data left with other motion patterns than what was represented in the activity clusters. In the unlabelled episodes, participants were, e.g., walking, sitting at a desk, or remained in another state that often matched a modelled activity cluster.

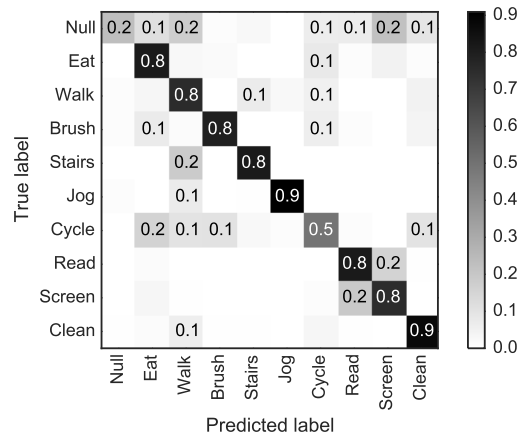


Figure 5: Daily activity recognition: confusion matrix across cross-validation folds. With the exception “cycle” cluster, all activity cluster perform at an accuracy of at least 80%. Null class performance was low due to the nature of the study protocol, where many modelled activities occurred also as unlabelled data, thus resulting in confusions.

6.2 Screen-use detection

We detected screen-use from the light sensor data. Screen-use and its timing are important factors in circadian rhythm

entrainment, as light emitted from screens and received through eyes at night may hamper falling asleep later.

Figure 6 shows mean accuracies for the light-based screen detection per participant using LOO cross-validation and a GMM classifier with diagonal covariance matrices using 3 mixtures per class. An accuracy between 64% and 93% was obtained here, with a mean of 80%. While most folds achieve performance around the mean score, evaluation for participant 4 showed as an outlier.




Figure 6: Screen-use detection: Normalised accuracy per participant (cross-validation fold). Mean accuracy was 80%. Results for participant 4 seems to be an outlier.

6.3 Heart rate estimation

Heart activity was measured in WISEglass using an optical pulse sensor worn at the temple. Similar to the daily activity recognition, there are many context applications that benefit from physiological estimates. Here we focus on heart rate as one important indicator. Figure 7 shows the RMSe for all recordings as well as activity subsets “motion” and “stationary” per participant and averaged. For all activities combined, heart rate estimation RMSe varied between 10 bpm and 19 bpm with a mean of 13 bpm. Motion artefacts contributed to an increased RMSe. For motion activities the mean RMSe increases to 14 bpm while for stationary activities it is reduced by 40% compared to motion activities, to 9 bpm. Motion artefacts were prominent in the heart data from both WISEglass pulse sensor and the ECG reference. During motion it could not be determined which of the systems (WISEglass or ECG reference) provides more accurate data. Thus during motion, the error reported here must be considered as a difference between the two systems only.

6.4 Usability evaluation

The SUS evaluation of all nine participants yielded a mean score of 66.4 of 100. Participants found the system easy to use with 4.1 of 5.0 points on average. The overall score may be negatively biased due to the experiment process (e.g. placing ECG electrodes) and hence consider the easy to use question the most relevant result of the SUS evaluation.

7. DISCUSSION

The results of this work clearly confirm idea behind WISEglass: smart eyeglasses can provide multi-purpose sensing functions for different context recognition applications. Instead of clip-on designs, as Google Glass and others, we emphasise that eyeglasses could integrate sensing and processing functions into typical eyeglasses frames. The per-

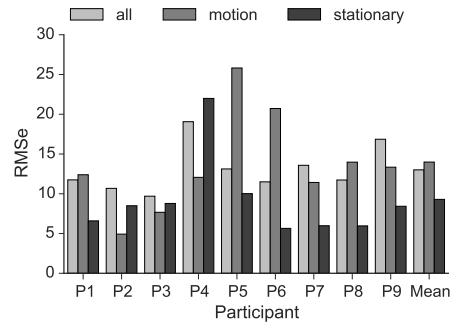


Figure 7: Heart rate estimation: RMSe analysis per participant, where the reference heart rate was between 40 and 150 bpm. RMSe was analysed for three activity clusters: all, motion (“stairs”, “walk”, and “jog”), and stationary (“eat”, “read”, and “screen”).

formance for activity recognition, screen-use detection, and heart rate estimation in this work show that the head-worn measurement positions and the used pattern processing methods are adequate. Compared to a phone, the smart eyeglasses are always at the same body position and in addition can measure HR continuously. Compared to a watch-like devices, smart eyeglasses can measure the light shining into the users eyes. Another advantage is that the activities a watch can detect must somehow include the arm (e.g. watching TV might be difficult to detect), so eyeglasses can cover a broader spectrum of activities. Nevertheless, for daily activity recognition improvements of the null class modelling seem necessary. It is clear that our recordings, although made over full days and including diverse activities do not reflect all conditions under which smart eyeglasses could be used. As a consequence, we had little data that was not already modelled by the activity clusters and thus unknown activities were under-represented in the dataset. This effect shows that activity clusters are robust enough to generalise onto unknown data. In a next step, WISEglass could be used in free living conditions. Then however, less accurate annotations must be expected.

The GMM classification approach demonstrated that user-independent models with good average performance (approx. 80%) can be obtained. Additional GMM parameter selection could be performed on a per class basis to further improve the model fit. Specifically, for the screen-use detection improvements could be realised by combining light and motion data. While we intended to span a wide range of possible applications of smart eyeglasses, improvements to the heart rate estimation method should be investigated, too. Here, a more controlled setting may be helpful to provide detailed RMSe for different movements. The ECG reference data showed motion artefacts due to insufficient electrode attachment. The present work presents a pilot study to confirm benefit of smart eyeglasses. In further investigations, additional participants and additional scenarios and applications for smart eyeglasses shall be investigated to further expand the use cases for smart eyeglasses.

We are aware that our current WISEglass prototype is yet not fully resembling the vision of regular eyeglasses. While we made important steps, e.g. by integrating the light sensor and wiring into the eyeglasses frame, further miniaturisation and integration work is necessary. Moreover, to improve un-

obtrusiveness, the baseboard and battery could be moved to the temple hinges. The pulse sensor could be made smaller and integrated in the temple such that the foam block disappears. Given the reliable data from the pulse sensor, we are confident that such improvements are feasible.

8. CONCLUSION

WISEglass is the first multi-modal sensor system retrofitted to off-the-shelf eyeglasses. The platform offers motion, light, and pulse sensors for use in a wide range of context-aware applications. All sensor modalities were successfully evaluated in three application scenarios, demonstrating that the smart eyeglasses serve as single-position context measurement device. For daily activity recognition, an average accuracy of 77% for distinguishing 9 activity clusters using LOO cross-validation was achieved. The recognition rate is comparable to previous work [12, 13], while in our work more activities and user-independent models were considered. However, most importantly, WISEglass integrates all sensors into one commonly used accessory, thus does not require to wear head-mounted displays, cameras, etc. Screen-use detection achieved a mean accuracy of 80%. Heart rate estimation was evaluated against a chest-worn ECG reference with an average RMSE of 13 bpm for all available data and 9 bpm RMSE during stationary activities confirming that the temple-worn sensor provides an adequate wearing position.

Acknowledgements

We thank Timo Spinde and Christoph Weidemeyer for collecting the dataset, Jakob Kasbauer for contributing to the light data analysis, and all study participants. This research was supported by the Dutch Technology Foundation STW under grant number 12184.

9. REFERENCES

- [1] L. Atallah, B. Lo, R. King, and G.-Z. Yang. Sensor placement for activity detection using wearable accelerometers. In *Body Sensor Networks (BSN), 2010 International Conference on*, pages 24–29. IEEE, 2010.
- [2] O. Aziz, B. Lo, R. King, A. Darzi, and G.-Z. Yang. Pervasive body sensor network: an approach to monitoring the post-operative surgical patient. In *Wearable and Implantable Body Sensor Networks, 2006. BSN 2006. International Workshop on*, pages 4–pp. IEEE, 2006.
- [3] R. T. Azuma et al. A survey of augmented reality. *Presence*, 6(4):355–385, 1997.
- [4] L. Bao and S. S. Intille. Activity recognition from user-annotated acceleration data. In *Pervasive computing*, pages 1–17. Springer, 2004.
- [5] C. M. Bishop. *Pattern recognition and machine learning*. springer, 2006.
- [6] J. Brooke. Sus-a quick and dirty usability scale. *Usability evaluation in industry*, 189(194):4–7, 1996.
- [7] M. Duarte. Notes on scientific computing for biomechanics and motor control. <https://github.com/demotu/BMC>, 2015.
- [8] M. Figueiro, R. Hamner, A. Bierman, and M. Rea. Comparisons of three practical field devices used to measure personal light exposures and activity levels. *Lighting Research and Technology*, 45(4):421–434, 2013.
- [9] Y. Gitman and J. Murphy. Pulse sensor-a wearable device to give your projects a live heartbeat. *Make-Technology on Your Time*, 29:52, 2012.
- [10] J. Hernandez, Y. Li, J. M. Rehg, and R. W. Picard. Bioglass: Physiological parameter estimation using a head-mounted wearable device. In *Wireless Mobile Communication and Healthcare, EAI 4th International Conference on*, pages 55–58. IEEE, 2014.
- [11] J. Hernandez and R. W. Picard. Senseglass: Using google glass to sense daily emotions. In *Proceedings of the adjunct publication of the 27th annual ACM symposium on User interface software and technology*, pages 77–78. ACM, 2014.
- [12] S. Ishimaru, K. Kunze, K. Kise, J. Weppner, A. Dengel, P. Lukowicz, and A. Bulling. In the blink of an eye: combining head motion and eye blink frequency for activity recognition with google glass. In *Proceedings of the 5th Augmented Human International Conference*, page 15. ACM, 2014.
- [13] S. Ishimaru, Y. Uema, K. Kunze, K. Kise, K. Tanaka, and M. Inami. Smarter eyewear: using commercial eog glasses for activity recognition. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication*, pages 239–242. ACM, 2014.
- [14] D. Jarchi, B. Lo, E. Jeong, D. Nathwani, and G.-Z. Yang. Validation of the e-ar sensor for gait event detection using the parotec foot insole with application to post-operative recovery monitoring. In *BSN 2014: Proceedings of 11th International Conference on Wearable and Implantable Body Sensor Networks*, pages 127–131. IEEE, 2014.
- [15] W. McKinney. pandas: a foundational python library for data analysis and statistics. *Python for High Performance and Scientific Computing*, pages 1–9, 2011.
- [16] G. Spina, F. Roberts, J. Weppner, P. Lukowicz, and O. Amft. Crntc+: A smartphone-based sensor processing framework for prototyping personal healthcare applications. In *Pervasive Computing Technologies for Healthcare, 7th International Conference on*, pages 252–255. IEEE, 2013.
- [17] E. Tamaki, T. Miyaki, and J. Rekimoto. Brainy hand: An ear-worn hand gesture interaction device. In *CHI '09 Extended Abstracts on Human Factors in Computing Systems, CHI EA '09*, pages 4255–4260, New York, NY, USA, 2009. ACM.
- [18] F. Wahl, T. Kantermann, and O. Amft. How much light do you get? estimating daily light exposure using smartphones. In *Proceedings of the 2014 ACM International Symposium on Wearable Computers*, pages 43–46. ACM, 2014.
- [19] S. v. d. Walt, S. C. Colbert, and G. Varoquaux. The numpy array: A structure for efficient numerical computation. *Computing in Science & Engineering*, 13(2):22–30, 2011.
- [20] F. Zhou, H. B.-L. Duh, and M. Billinghurst. Trends in augmented reality tracking, interaction and display: A review of ten years of ismar. In *Proceedings of the 7th IEEE/ACM International Symposium on Mixed and Augmented Reality*, pages 193–202. IEEE Computer Society, 2008.