Seminar Distributed Systems: Assistive Wearable Technology

Stephan Koster
Bachelor Student
ETH Zürich skoster@student.ethz.ch

ABSTRACT
In this seminar report, we explore the field of assistive wearable technology. To that end, we examine three directions of research, each represented by a research paper.

General terms: Performance, Design, Reliability, Human Factors

Keywords: Wearable computing, activity recognition, eye tracking, industry

INTRODUCTION
Assistive wearable technology is a subset of wearable technology. As such, devices in the area are usually designed to be attached to the user's body. This has a number of consequences. For instance, the devices need to be physically small so their bulk does not inconvenience the user. Also, energy use needs to be low because big, heavy batteries again inconvenience the user. The sensor data are collected and processed by the wearable devices without any deliberate action by the user.
Of course, these devices need to be justified by some kind of application. A common approach to making use of sensor data is to give the application an indication of what the user is doing and in what environment. The application can then provide some kind of context sensitive help to the user.

Wearable EOG Goggles
As an example of the sensor technology that is emerging in the field, we take a look at the work of Tessendorf et al. In the paper "Wearable EOG goggles: eye-based interaction in everyday environments" [2], a novel device for eye tracking is discussed. It functions on the well-known principle of Electrooculography (measurement of the electric field around the eyes) and brings this kind of sensor into a package suitable for wearable technology applications.

Eye Movement
In order to understand the workings of an eye tracker, we must first quickly introduce some background knowledge about eye movement. The human eye moves in specific patterns of fixations where both eyes stay fixed on a reference point, and saccades where the eyes make a fast movement from one reference point to the next. The exact pattern is largely involuntary. Different visual environments and activities such as reading lead to different patterns of fixations and saccades.
Eye movement can be electronically tracked in three ways: *a*) The oldest and still most accurate method involves applying a special contact lens that contains a mirror or other feature that can be precisely tracked. *b*) A newer development uses cameras and visual computing algorithms to reconstruct the gaze direction. Often, camera setups will illuminate the eye with infra red, so the camera can track both the dark iris and the bright spot on the back of the eye where the infrared light gets focused by the lens. Using these two points, the axis of the eye can be determined. Obviously, cameras to track the eyes must have a free line of sight to the eyes, which poses some problems if the user is to carry them around on his body. *c*) The third method of tracking eye movement is electrooculography (EOG) which the Tessendorf paper focuses on. EOG is much less precise than the other two methods, so some potential applications like for example using the eyes as a mouse pointer are not possible. Still, the method provides enough information about eye movement patterns to determine what the user is doing.

EOG
Electrooculography or EOG is the recording of the electrical field around the eyes by means of electrodes. The resulting waveforms are called an electrooculogram. Because the eyes create an electrical field along the seeing axis when they are exposed to light, suitably placed electrodes allow the reconstruction of the approximate gaze direction. The strength of the electrical field varies with the intensity of light entering the eyes. A device that tries to deduce the gaze direction therefore needs to record ambient light level and adjust the signal processing. Blinks can be detected in the electrooculogram as well.
Previous EOG devices were meant for medicinal purposes, to find irregularities in eye movement patterns and the adaption of the electrical field strength to differing light levels. These medical devices were not portable and could only be mounted and operated with the help of a professional.

Figure 1: Wearable EOG Goggles

Making EOG Wearable
The device introduced in the paper integrates a set of dry electrodes on a pair of standard lab safety goggles as can be

seen in Figure 1. This reduces setup time and comfort enormously compared to wet electrodes which have to be covered in gel and stuck to the skin with tape. Also, the entire data processing happens in a credit-card sized embedded device that can be carried on a belt. The goggles also contain a light sensor for calibration purposes and an an IMU (inertial measurement unit) to track head tilt. In the here presented paper, the head tilt sensor is not used.

Eye Gesture Experiment

A case study was carried out to assess the possibility of using the device to track deliberate gestures performed with the eyes. For this setup, the sensor data are translated into a string of characters, with each character representing a fixation in a certain orientation or a movement in a certain direction. The user tries to perform the gestures by looking at spots on the corner of a computer screen in a specified sequence, as seen in Figure 2. To recognize a gesture, a simple string matching algorithm tries to match the predefined gesture strings to the continuous stream from the device. The in the experiment was sitting in front of a computer screen displaying a mark on each corner. The user was then prompted to look at the marks in a specific order which the system was supposed to recognize as an eye gesture.
Test subjects reported that entering the gestures was easy to learn but tiring. The authors remark that this is common with all kinds of novel input devices. A potential application for eye gestures would be for paralyzed patients who could use the system to communicate.

Figure 2: Eye Gesture Experiment

Context Awareness

In another paper, "What's in te Eyes for Context Awareness?" [1] the same device is used to determine the activity a user was performing. Here, the sensor data stream is processed into a feature set which is then presented to an SVM classifier which tries to figure out what activity the user is involved in. In a series of case studies, the group tries to find out how well certain activities can be recognized by this system. In one experiment, the ability of the system to recognize the user reading text on a sheet of paper in a variety of environments was tested. Figure 3 shows the achieved recall and precision in this experiment. As the authors expected, the system works best when the user is sitting and worst when walking.
Other case studies showed that certain office tasks can be distinguished from each other. Finally, the paper explores the possibility to detect weather the user is remembering a picture he is looking at or not.

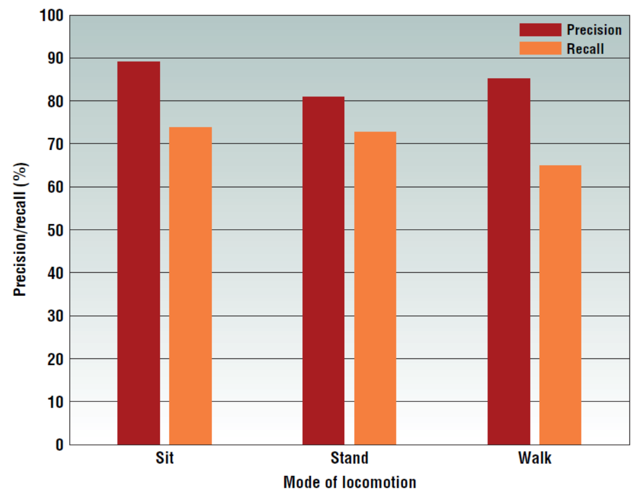

Figure 3: Recognize Reading Activity: Reliability

Comments and Criticism

In these papers, the authors do not present a real application that actually gives any benefit to a user. We hope that applications of this technology emerge as its capabilities are explored.
The way the eye gesture experiment was performed, many important questions remain open, mainly how to deal with the issue of distinguishing a deliberate eye gesture from normal eye movement (the so called Midas touch problem) if eye gestures were used in an everyday context. This problem must somehow be addressed should the system be used for controlling any kind of application in a less controlled setting. Another downside is that users reported performing the eye gestures was tiring. This may be acceptable for a system that helps paralyzed patients communicate who will hopefully adapt to the strain with practice. For an assistive wearable computing application that is supposed to remain in the background and take cognitive load off of the user, an interface that physically tires the eyes is not acceptable.
The potential of a wearable EOG system to detect the user's activities seems to be much more promising, especially the tantalizing prospect of gaining insight into cognitive processes of the user. Still, consumers will only be interested in such a device if there are applications whose benefit outweighs the drawbacks of wearing these electrodes on the face.

Improving Hearing Aids

In the paper "Recognition of Hearing Needs from Body and Eye Movements to Improve Hearing Instruments" [4], Tessendorf et al describe a way of improving modern hearing aids with the help of additional wearable sensors.

Modern Hearing Aids
Modern hearing aids feature a number of uni- and omni directional microphones as well as configurable digital signal processing capabilities to tailor the hearing experience to the needs of the user. To simplify the settings for the user, the hearing aid usually provides a small number of predefined hearing profiles. The industry standard seems to be a set of four settings:

a) "Speech" is designed to make human speech sound as natural as possible. It is also the standard setting in quiet situations.

b) "Speech in noise" sacrifices natural sound for better understandability of a conversation partner in a noisy environment. This profile emphasizes directional microphones (pointing forwards) and applies some audio filters.

c) "Noise" tries to reduce the distraction from a noisy environment.

d) "Music" is optimized to faithfully reproduce sound sources with a high dynamic range, like music.

Modern high-end hearing aids can autonomously switch between these hearing profiles based on analysis of the sounds recorded by the microphones. Tessendorf et al note that this purely sound based approach fails in situations where the acoustic environment is similar, yet the need of the user is different.

Figure 4: Sensors to Recognize Hearing Need

Data Capture
In order to resolve ambiguities and generally improve recognition of hearing needs, the user is fitted with a variety of sensors mounted on a jacket, see Figure 4. There are a total of nine IMUs mounted to a harness to track movements and orientation of body and limbs. Another IMU is mounted to the back of a hat. Around one eye of the user, a simple EOG is set up with four electrodes. Finally, there is a microphone attached to the user's throat. All sensors except the microphone were compared against each other with regard to their worth in distinguishing hearing needs.

Data Processing
Since this project focused more on evaluating different sensors than on creating a realistic system, all data processing happens off line. The raw sensor data, including audio from the hearing aid, is continually stored to a laptop connected by a bundle of cables. The classification algorithms can then work on the stored data stream.

Experiment
An experiment was performed to quantify the benefits of using multiple kinds of sensors to distinguish hearing environments which the classical solution (based on sound only) has trouble keeping apart. The specific challenge to be tested is determining weather "speech in noise" or "noise" should be used.
To create a realistic and reproducible environment, the scenarios are all played out in a quiet office but the audio channels get overlayed with constant office noise. This makes sure that different runs do not have disparaging results just because the background noise level was different. The two scenarios played out are *a*) Subject tries to work on a task while a coworker at the same table has a conversation with a "disturber". The system is supposed to select the noise profile. See Figure 5 *b*) Subject talks with the coworker or the disturber. The system should select the speech in noise profile. See Figure 6 The way the experiment was set up, the system only has to distinguish between two hearing needs. Figure 7 shows the success rates when different sets of sensors were used to distinguish the two situations. When multiple sets of sensors are used, there is actually a separate SVM classifier working on each sensor and the final result comes from a majority vote between the SVMs. When the vote is tied, the system keeps the judgment of the last round of classification. This voting system explains why the result sometimes gets worse when more sensors are included. If all data were fed to a single SVM, the result should in theory never deteriorate with the addition of more data.
The way the system was set up for this experiment, using all IMUs on body and head yielded the best results, followed by eye motion only and the single IMU on the head. This is interesting, because adding a single IMU to the user's head is by far the most convenient way of gathering data presented here, yet still yields a significant improvement on the the state of the art.

Figure 5: Scenario 1: Trying to work

Outlook and Criticism
This paper shows the potential for additional sensors to improve the automatic selection of hearing needs in modern hearing aids. Results show that even a single IMU attached to the head can significantly improve the performance of exist-

Figure 6: Scenario 2: Conversation

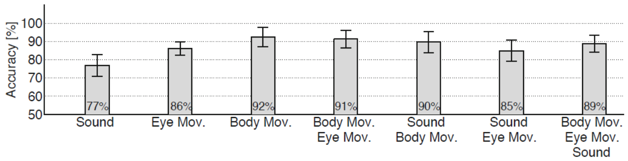

Figure 7: Accuracies of Different Sensors

ing systems in specific situations. With further miniaturization of sensor technology, it is easily conceivable that such sensors might be included directly in the hearing aids.
We would like to see a prototype of the system that performs the data processing and classification on line with the hearing aid actually active, so a user study with hearing impaired participants could be performed. This could clear doubts about weather the the reported improvements in distinguishing some hearing needs translate into a truly improved hearing experience.

Activity Tracking in Car Manufacturing

In the paper"Wearable Activity Tracking in Car Manufacturing' [3]' by Stiefmeier et al, new applications for wearable sensors in car manufacturing are discussed. Specifically, they introduce two tasks that could be improved, the so called "Learning Island" where trainees get prepared for the assembly line and the final step of the assembly line, the quality control. These tasks were identified as potential target for improvement in cooperation with a real Skoda car manufacturing plant. Some settings from the factory were recreated in a laboratory so the experimental systems could be tested without interruption of the facility.

Learning Island

The learning island is a part of the examined factory where new workers are trained and tested until they are ready to start work on the actual assembly line. The most important part of this learning island is a specially prepared car on which different parts can be installed and removed over and over again. Trainees get introduced to new assembly steps in theory lessons before they can practice on the training island under supervision. Once the trainees are judged to have mastered all required assembly steps, they can start work on the assembly line.

Sensors on Learning Island

In order to model a task, a finite state automaton (FSM) was used. Figure 8 shows the FSM for installing the head lamps, a task representative for many other manufacturing steps. Edges in the graph correspond to assembly actions detected by the sensors, nodes to the changing configuration of the car. Because this model is has zero fault tolerance, it will report the task as failed as soon as the trainee makes a single

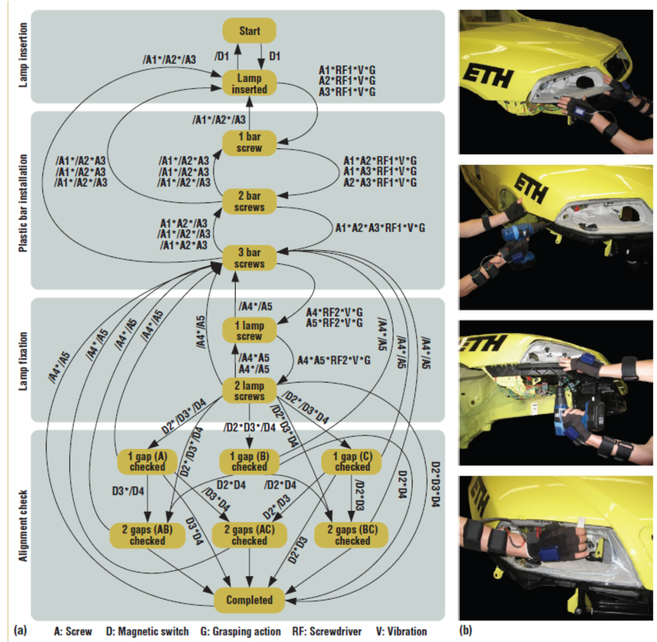

Figure 8: FSM Task Model: Installing Head Lamps

deviation from the model or the sensors fail to pick up a step. It is therefore important that each assembly step is picked up with near 100% reliability to avoid false negatives. The sensors used reflect that requirement.
Figure 9 shows the wearable sensors that were used. A single IMU on the back of the hand detects when a power tool hits its torque limiter, causing the hand to shake. The bracelet on the forearm contains force sensitive resistors, so it detects when it is deformed. This effect is used to recognize a firm grip on a tool. In the glove itself, between thumb and index finger, there is a small RFID reader. All tools needed for the task have had an RFID tag added, so the reader can recognize which tool is currently being held in the hand.
These sensors alone are not sufficient to determine for example which screw hole was used and some (particularly the bracelet) are not reliable enough to drive the FSM task model. These shortcomings are addressed by a variety of sensors installed on the training car. Magnetic switches in the car frame detect the presence of parts that need to be added with excellent reliability. Near screw holes, force sensitive resistors are glued to the metal. When a screw is tightened, the metal around the screw hole is deformed slightly which the sensitive force sensitive resistors detect. The magnetic switches work out of the box, while the force sensitive resistors at the screw holes need calibration. The paper notes that installing and calibrating all these sensors on the training car takes at least half a day of work by an expert and is therefore expensive.

Applications on Learning Island

The goal of the system on the learning island is to improve the training process. An obvious improvement is that trainees can now practice some assembly steps without supervision and still get a feedback on their performance. This frees the time of the (expensive) instructors.

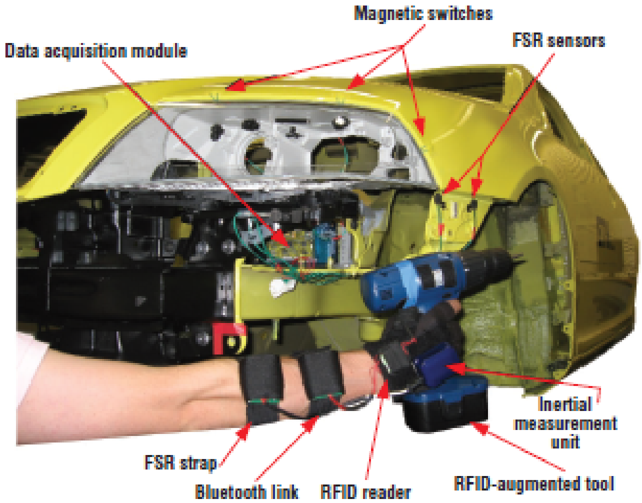

Figure 9: Sensors for Learning Island

Stiefmeier et al also note that an evolution of their system could eliminate theory sessions and instead lead the trainees through the assembly process step by step the first time and immediately report errors and give guidance. This would lead to a similar learning experience as current flight simulators.

Quality Check

Another area where wearable sensors could be used is the final step of car assembly, the quality check. In the quality check station, workers go through a checklist and verify the end product confirms to specifications. This involves checking the function of doors, hood and trunk as well as measuring the proper alignment of assembled parts with special checking tools. The research team identified 46 distinct checking tasks in the process.
The proposed sensor and data processing system is supposed to recognize and distinguish these tasks without deliberate input by the worker.

Sensors in Quality Check

Figure 10: Sensor Vest for Quality Check

Because the quality check happens on production vehicles, the approach to instrumentation used on the learning island is not feasible in this scenario. Instead, workers wear a jacket with integrated sensors as displayed in figure 10. This jacket includes a total of seven IMUs placed on body and arms. The sleeves of the jacket are lined with a special fabric containing multiple force sensitive resistors. These sleeves can measure the bending of the elbows. The jacket also contains two tags in the shoulder area which can be located in relation to four base stations in the work area by a commercial system, a so called "ultra wide band system".
With all these sensors combined, the system is powerful enough to create a rough model of the worker wearing it including his absolute position in the working area as can be seen in figures 11 and 12. The authors note that the precision of the ultra wide band location system is decreased markedly on the real assembly line because the four base stations have to be placed farther apart so as not to obstruct work. In a test, the system correctly identified 74% of distinct checking activities, but with the restriction that the system only checked for 6 out of 46 distinct activities.

Figure 11: Checking Door Function

Figure 12: Checking Filler Cap

User Acceptance Study

To find out if the sensor jacket is a device that could be used in a commercial application, a user study was performed in which workers wore the jacket on the real assembly line. Workers reported that the sensors were not stopping them from doing their tasks, yet were still clearly noticeable at all times and needed some getting used to.

Applications in Quality Check

The paper mentions two main ideas on how an activity recognition system could improve the work flow at the quality control station. First, the system should raise a warning when any checking steps were missed. Second, the current pen and paper checklists could be replaced with some kind of portable electronic system into which faults can be entered. The activity recognition system comes into play by presenting the correct page of the checklist to the worker. A future system could maybe even recognize when the worker has found a fault and offer the worker the option of confirming with a single button push or gesture. This would permit the system to correct false categorization and continually improve the classifiers.

Summary

Assistive wearable technology is an area of active research. Research groups continue to introduce new sensor modalities and potential applications. Data collected by the sensors is commonly used to deduce the user's activity or environment. This information is then used to enable context sensitive applications.

We have looked at a novel device for sensing eye movement by Bulling et al [1] [2]. They focus on the sensor more than on the potential applications.

We have then discussed the paper by Tessendorf et al [4] who found an interesting application for context information in selecting settings for a hearing aid. A number of sensors are used to improve the existing selection algorithms which chooses a hearing profile solely on acoustic input.

Finally, we have discussed the paper by Stiefmeier et al. [3] In that paper, applications for wearable computing in an industrial setting, namely car manufacturing, are discussed. They show how a wearable system could improve training and quality assurance.

REFERENCES

1. A. Bulling, D. Roggen, and G. Troster. What's in the eyes for context-awareness? *Pervasive Computing, IEEE*, 10(2):48–57, 2011.

2. Andreas Bulling, Daniel Roggen, and Gerhard Tröster. Wearable eog goggles: eye-based interaction in everyday environments. In *CHI '09 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '09, pages 3259–3264, New York, NY, USA, 2009. ACM.

3. Thomas Stiefmeier, Daniel Roggen, Georg Ogris, Paul Lukowicz, and Gerhard Tröster. Wearable activity tracking in car manufacturing. *IEEE Pervasive Computing*, 7(2):42–50, April 2008.

4. Bernd Tessendorf, Andreas Bulling, Daniel Roggen, Thomas Stiefmeier, Manuela Feilner, Peter Derleth, and Gerhard Tröster. Recognition of hearing needs from body and eye movements to improve hearing instruments. In *Proceedings of the 9th international conference on Pervasive computing*, Pervasive'11, pages 314–331, Berlin, Heidelberg, 2011. Springer-Verlag.