Eye Tracking for Personal Visual Analytics

Kuno Kurzhals and Daniel Weiskopf ■ *University of Stuttgart*

Eye tracking can help record massive amounts of data about the distribution of visual attention in various scenarios. Such data could support nonexpert user self-reflection and self-insight. However, challenges arise when eye tracking is applied to everyday situations and personal visual analytics.

Eye tracking for the analysis of gaze behavior is common in many scientific fields and marketing research. So far, the high cost of eye-tracking hardware, a result of the requirements of high precision and reliability for research measurements, has prevented wide application in personal, nonprofessional scenarios. However, this situation has been changing as affordable consumer hardware has become more widely available. The established eye-tracking vendors have developed consumer versions of portable eye-tracking hardware that can be used on any monitor or TV. The development is not restricted to stationary eye-tracking devices, and "how to build your own eye-tracking glasses" instructions can even be found in various publications.[1] Therefore, easy-to-use mobile eye tracking integrated in wearable glasses is already available.

In combination with the industry's interest in intelligent glasses, we expect that wearable mobile eye tracking will be available for everyone in the near future. The main purpose of this development is the use of eye tracking as a device for human-computer interaction—for example, to adapt the user interface. However, we see a great opportunity in using such hardware for personal analytics as well. How can users of intelligent glasses recapitulate their viewing behavior, understand their interactions with others and the environment, or just have fun with their personal data?

The possible application scenarios for personal eye tracking cover diverse fields. With the additional information about the user's visual attention, important events in the video database can be extracted to allow users to re-experience these events. Possible scenarios might include applications that support self-reflection and self-insight[2] via video analysis with gaze information. This could involve analyzing interaction logs for personal relations with others, vigilance optimization during driving situations, or cognitive activity recognition that can be applied for quantified-self scenarios.[3] For example, users could set a goal to read at least 10,000 words a day and then monitor their reading behavior and time spent on reading texts. Also, recommender systems could generate catalogs of interest based on the objects that attracted the user's attention. Viewing behavior could also be analyzed to present similar suggestions for future media consumption. The time spent on a personal visual analytics application strongly depends on the scenario. For example, users who benefit from the analysis for health or social reasons will be more motivated to spend time with the application than users who browse recorded data just for fun.

With the changes in technology and new applications, new opportunities and challenges for data analysis will arise. Mobile eye tracking produces massive amounts of complex data because it both produces spatiotemporal information of eye gazes and provides video recordings of the person's environment. Without such video information, we are missing the semantic context of the gaze data; we would not be able to relate visual attention to objects in the environment or to other people

 Published by the IEEE Computer Society

with whom the user interacts. In other words, we are facing a coupled analysis problem: analysis of spatiotemporal gaze data and video analysis. Each of these analysis problems comes with challenges of its own, even in a professional setup. (See the related discussions of professional visual analytics for eye tracking[4] and video[5] for more details.) The combined analysis problem raises even more interesting questions for personal visual analytics. In particular, with the vast amount of personal video data with gaze information being available, standard personal information visualization methods—such as replaying the videos—will fail. Therefore, new visual analytics methods must be developed to filter and summarize information that is important to the user.

In this article, we discuss how eye tracking fits into the design space of existing personal visual analytics applications as well as the special requirements and research perspectives of personal eye tracking. Because this personal information can be applied for the user's self-reflection, it also fits into the concept of personal informatics.[6] In this context, the focus of our discussion is on the analysis of personal data rather than on data acquisition. As one example of the visualization of personal eye-tracking data, we present a new approach, the areas of interest (AOI) cloud, to display information about the distribution of attention across multiple videos. With our technique, AOIs (which might be objects or people) can be displayed in an annotated overview using a representation similar to a tag cloud. Additional rings on the AOIs allow for easy navigation through several videos to examine time spans that received the user's attention.

Current Use of Mobile Eye Tracking

Mobile eye tracking is often used for user studies that do not restrict the participants as much as a laboratory experiment under controlled conditions might. Figure 1 shows a typical example of mobile eye tracking. In this scenario, to investigate the viewing behavior of people in a supermarket, participants perform a shopping task while wearing eye-tracking glasses. The glasses record eye movements and a video of the participant's field of view. To analyze the recorded data, statistical methods (in particular, statistical inference for hypothesis testing) and/or visualization are used. However, statistical methods cannot be applied as easily as in laboratory studies because of the less controlled environment and stimuli. For mobile eye-tracking scenarios "in the wild," changing conditions exacerbate the statistical comparison of multiple participants. Therefore, qualitative visual evaluation of the data is often required.

Figure 1. Mobile eye tracking in a supermarket scenario. The person in the front is wearing glasses with integrated mobile eye tracking. In this consumer study, the participants' eye movements are recorded when they perform shopping tasks. The picture was taken as part of a research project on mobile, unconstrained eye tracking.

For qualitative and quantitative analysis of mobile eye-tracking data, most techniques rely on the definition of AOIs to relate the stimuli contents between participants. Unfortunately, these analysis methods require extensive manual processing and labeling; there are only a few automatic computer vision techniques that detect and recognize trained objects from a database to generate AOIs.[7] Today's analysis methods for mobile eye tracking are restricted to professional users and require extensive work to setup the experiments and postprocess the data recorded. Therefore, these methods cannot be applied directly for personal visual analytics.

Personal Eye Tracking

Here, we investigate how personal eye tracking can be categorized in the general context of personal visual analytics and what special requirements and challenges have to be considered for applications of personal eye tracking.

In the Context of Personal Visual Analytics

To apply eye tracking in a personal context, we will first investigate the design dimensions of personal visual analytics and how an application for personal eye tracking fits in. To this end, here we examine the classification introduced by Dandan Huang and his colleagues,[8] which consists of four categories: *data*, *context*, *interaction*, and *insight*.

The scope of the recorded *data* is a combination of data about oneself and data about other people.

Data about oneself is recorded by gaze information and by the video camera of the eye-tracking device that captures data about the environment, including other people. This data is personal and has to be handled with care. Under the assumption that eye-tracking devices will become more and more comfortable in the future and comparable to wearing regular glasses, the effort to record data will be reduced to sensor recording only. Current eye-tracking devices still require elaborate calibration procedures that increase the effort to record data. Regarding the controllability of the data acquisition, the user has partial control over whether to record the surrounding.

The *influence* context of mobile eye-tracking analysis is mainly personal, functioning to inform the user wearing the eye-tracking device. However, other people will often be involved in the recorded data, so the user could communicate extracted events through social media to involved persons—for example, to recapitulate parts of a conversation. The design context of an application depends on the scenario. In the example we describe in the next section, the application to examine the recorded data is designed by the researcher. However, users can freely organize the components of the visualization, such as to arrange groups of people or extract and summarize important personal events in an easily accessible visual representation. For scenarios with automatic data analysis (such as recommender systems), predefined representations of the results should be sufficient in many cases.

Personal eye tracking will cover much longer time spans than traditional eye-tracking experiments, requiring more time-compressed visual representations.

The degree of attentional demand for *interaction* also depends on the scenario. In cases when the analysis is performed automatically and the user just has to choose between different results (for example, recommended media), the attentional demand will be low. For the analysis of personal encounters, the user must focus attention on the visualization to investigate interesting events in the data, so a high attentional demand is required. The high explorability of the data in the application allows users to investigate multiple video streams simultaneously for interesting events that received much attention during the recording of the data.

Apart from technical issues, fully automatic analysis of the data can only be applied in a subset of scenarios and for preprocessing. An analysis of subjective events cannot be automated, and it requires the user to make conclusions based on the data. Also, the degree to which extracted *insight* from the application can influence future actions varies. In the best case, examination of the recorded data leads to an identification of self-defined misbehavior that can be avoided in the future. For example, a close friend may have received less attention than the user would consider appropriate. Now aware of this situation, the user can then spend more time with this person to strengthen their friendship.

Special Requirements

For the personal analysis of mobile eye-tracking data, we have to consider certain aspects that differentiate personal from professional visual analytics. From our perspective, the following characteristics and requirements of personal eye tracking are most relevant.

In professional eye tracking, the accuracy of the analysis is critical because research results, product design, security-relevant decisions, or other factors rely on the quality of the analysis. For example, both recall and precision of pattern recognition in the eye-tracking data are highly relevant. Fortunately, personal eye tracking is less critical in terms of analysis accuracy. Therefore, some leeway exists when designing personal visual analytics.

Personal eye tracking will cover much longer time spans than traditional eye-tracking experiments, requiring more time-compressed visual representations. Similarly, different reasoning artifacts are relevant.[9] For example, patterns in the transitions between fixations are of lesser interest than events or objects extracted from the eye-tracking data (such as people with whom the person interacted). Specific aspects of tasks for personal eye tracking will be complemented by general observations for casual visualization.[10]

Because personal eye tracking focuses on identifying relevant events or objects, it benefits from linking those to semantic information and embedding them into the context of “outside” information. For example, people identified as being important could be associated with information from their Web profiles.

Like any personal visual analytics application, the design of the visual interface has to be easy to use for nonexpert users. The design should be intuitive and not require a steep learning curve. The automatic processing for the analysis should be ro-

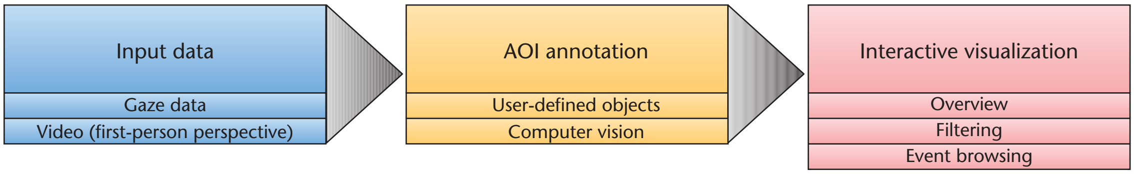

Figure 2. Data processing pipeline. Eye-tracking and video data need a semantic annotation of AOIs. The annotated data can then be displayed in an interactive visualization for video event browsing based on the distribution of attention on AOIs.

bust so that there is little or no need for the user to interfere and fine-tune data mining or computer vision techniques. Similar to many of the apps in mobile personal use on smartphones, visual analytics software for personal eye tracking will most likely be application-specific. In contrast, professional tools tend to be generic so that they can work with any study setup.

Personal visual analytics has to incorporate mechanisms to protect privacy because potentially sensitive information is recorded from the environment. Therefore, the analysis needs to be designed to work with the principle of data minimization (for example, to work with video recordings in which faces of persons or license plates of cars are modified to make them unrecognizable). Also, high data security is required to protect the user's personal gaze data.

These aspects will be critical in the design of appropriate visual interfaces and the development of automatic analysis techniques to be integrated within visual analytics. We expect that personal eye tracking will come with many challenging research questions related to design, interaction techniques, visualization, computer vision, pattern recognition, and semantic modeling. Although there is substantial research in these areas, we believe that the personal perspective will require us to devise new variants of existing techniques or develop completely new ones. To illustrate the personal visual analysis of eye-tracking data, we implemented a prototype for a commonly representative scenario: the analysis of a user's personal encounters.

Personal Encounter Analysis Case Study

The analysis of interactions between people plays an important role in psychology and cognitive science.[11] For a private user, the analysis of personal encounters can also be interesting, be it a self-reflection of social behavior or just for re-experiencing situations that received much attention.

In our example scenario, the user was wearing eye-tracking glasses during coffee breaks, a recurring event over one week. During the coffee breaks, groups of between three and six people, including the person wearing the eye-tracking glasses, gathered to discuss miscellaneous themes. The recordings during these breaks lasted between three and nine minutes with a varying set of participants. All participants agreed to be recorded on video if their faces were anonymized. Considering the privacy issues discussed earlier, this was an important prerequisite for all participants. We also agreed not to include the recorded audio in any form of publication of the data. One coffee break participant (P1) did not agree to be recorded in any form, so P1 sat next to the user wearing the eye-tracking glasses so as not to be visible to the camera, and therefore, P1's face was not annotated as an AOI. This situation exemplifies the issues that occur when other people are recorded on video and that have to be considered for personal eye-tracking applications.

Automatic preprocessing of this data requires an algorithm to detect faces in the videos, store them in a database, and recognize the faces when they reappear. In this scenario, the faces are the AOIs. Compared with other tasks in computer vision, this can be performed without much user interaction because there is no semantic gap that requires human interpretation of situations. The user might identify a person once, while the rest of the data is processed automatically. With the information about which faces can be seen in the videos and where they appear, an attention measure can be calculated by the AOIs of faces and the eye-tracking data. Although computer vision approaches can nowadays be applied for automatic segmentation and classification of such events,[12] we decided to showcase our example with manually annotated data because current automatic approaches often face difficulties with changing environment conditions, as in our case.

Data Processing

The visualization of personal eye-tracking data requires a preprocessing step that is necessary to map gaze data to semantic AOI information. Figure 2 shows the processing pipeline the data has to pass before it can be displayed in an interactive visualization for data exploration.

Assuming the user wears a set of eye-tracking glasses during an arbitrary occupation, two types of recorded data are of special interest. An integrated

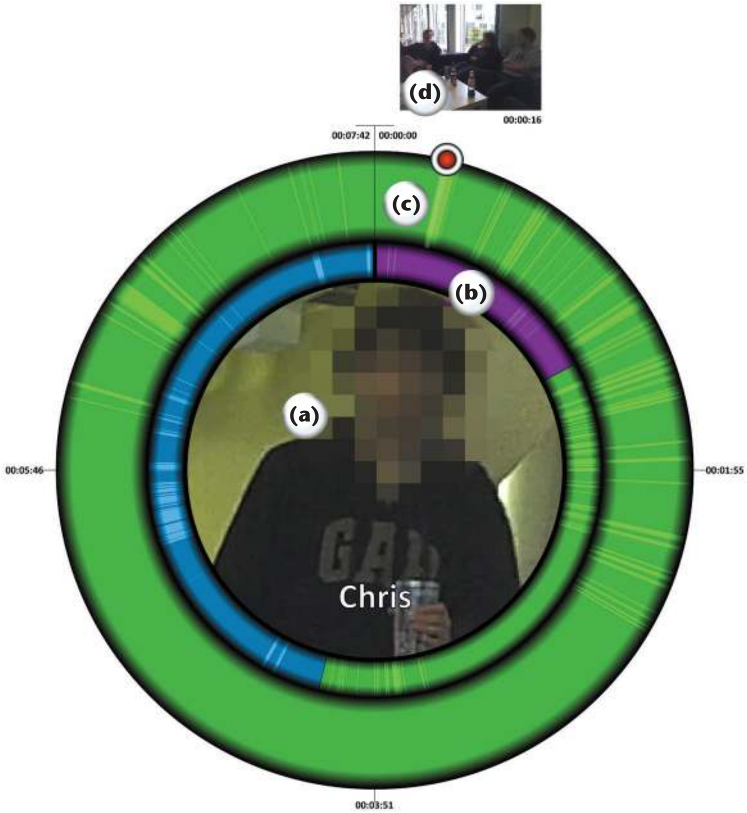

Figure 3. Visualization of one AOI. (a) This representative image of a person includes a name label, and the radius indicates the attention spent on the person. (b) The inner ring has segments for all the videos the person appeared in. (c) The outer ring shows the currently selected video. (d) Reference images can be created with markers on the outer ring.

camera records a video from a first-person perspective comprising most of the user's field of view. The video data serves as a foundation for semantic interpretations of the user's viewing behavior. The eye-tracking hardware can map eye-gaze positions to the coordinate system of the video. Following the eye-mind hypothesis,[13] we can assume that the fixated regions in the video were those to which the user's attention was directed. In combination with the video images, semantic interpretations can be derived.

To collect aggregated information about how much attention was directed to a particular object, it has to be annotated for a semantic mapping of gaze data to this object. By defining an object as an AOI, attention metrics can be aggregated even over several recordings that contain the same object. Depending on the user's interest, the AOIs can consist of a set of tools that are used during a work task or of the people the user interacts with, as in our example. Because this annotation of AOIs is task-specific, an automatic computer vision approach will be not sufficient in most cases. We suggest a semiautomatic approach where the user can define interesting objects once and the detection and tracking of these objects will then be performed automatically. Although computer vision approaches still need improvements to work in everyday situations, the semantic gap[14] that requires user input can be closed by such an interactive approach. For our example, we annotated the data manually to show how the interactive visualization works with a ground truth annotation. After the annotation, the processed data consists of AOI information about when and where an object appears in the videos and how much attention was directed to this object.

AOI Cloud Visualization

To visualize the distribution of attention on AOIs, common visualization principles such as an overview and interactive filtering of the data have to be available. For personal eye-tracking data, the overview of all AOIs and how much attention was spent on them play an important role. The interactive visualization has to meet the requirements that we discussed earlier for personal eye tracking and enable the user to browse the recorded video data for events and time spans where attention was spent on a specific object.

In our visualization approach, the annotated people (or AOIs) are represented as circles consisting of a representative image and an inner and outer ring (see Figure 3). Radial visualization approaches are applied in cases where hierarchical structures, relationships among disparate entities, or as in our case, time series data have to be displayed in a dense representation.[15] We decided to use a radial approach because of its accessibility for novices,[16] possibilities for fast interactions, and its compact representation of the temporal dimension on the rings that can be interpreted by using a clock metaphor.

The radius of the circle can be determined by an appropriate attention metric. In this example, we applied the total amount of gaze points on the person from all videos. Notice that other metrics such as transition counts between AOIs or mean fixation durations could also be applied, depending on the analysis question. Hence, our visualization approach is independent from the applied metric.

Because some people appear only in one video and others in three, the difference between the attention of the AOIs with the lowest and highest values can be large. This leads to extreme differences in the size of the circles, resulting in the problem that at least one of the AOIs is either too small or too big to be readable. Hence, we used a logarithmic scaling of the metric to adjust the visualization for a better representation of all AOIs. The representative image of a person is determined by the first appearance in the data. Alternative approaches to determine the representative image based on a special event in the data or on a profile image from social networks could also be applied.

The inner ring consists of segments that each represent a video containing the AOI. Hence, the inner circle contains all videos where the AOI appeared, and the size of a segment is determined by the relative length of the corresponding video. Segments in the inner ring are connected to the outer ring by identical colors. To visualize when attention was spent on the AOI, we use an approach similar to AOI timeline visualizations known from various approaches in this field.[17] Time spans without attention on the AOI are displayed darker, whereas time spans with attention are displayed with full brightness. This way, important events can be identified efficiently by directly selecting the time spans with attention on the AOI. Notice that approaches with AOI timelines usually consider only one video. In our approach, multiple video stimuli are combined in one visualization to investigate the data more efficiently. To distinguish between the different videos, we use an HSV (hue-saturation-value) color scheme where neighboring segments receive colors with a distant hue.

By selecting a segment of the inner ring, a second ring appears outside representing the selected segment zoomed over the whole ring. Time scales for the start and end of the video as well as for the quarters help the user to navigate clockwise through the video. Initially, one marker is available on the rim of the outer ring. It can be moved around the ring to navigate through the video. A thumbnail image next to the marker shows the currently selected frame as a reference to the video content. By clicking on the thumbnail, the corresponding video appears in a separate player window and can be played back directly at the selected position. The user can also create additional markers to select multiple events of potential interest to compare them or just summarize the gist of important interactions with the person in this video. With this approach, the user can generate a set of interesting events that can be assessed simply by clicking on the corresponding thumbnails.

The complete dataset can finally be visualized with items for each AOI that can be arranged in a layout similar to a tag cloud.[18] Important AOIs are placed in the center of the cloud, and less important AOIs appear in the outer regions. This setup makes our visualization accessible because tag clouds are familiar to most users and already established in everyday life. Selected items appear in the foreground, while the other items can be faded out. From that point on, the user is free to rearrange all the items to build groups or rank people based on subjective criteria. For example, a user could rank people based on friendship relations and investigate if their received attention relates to this ranking.

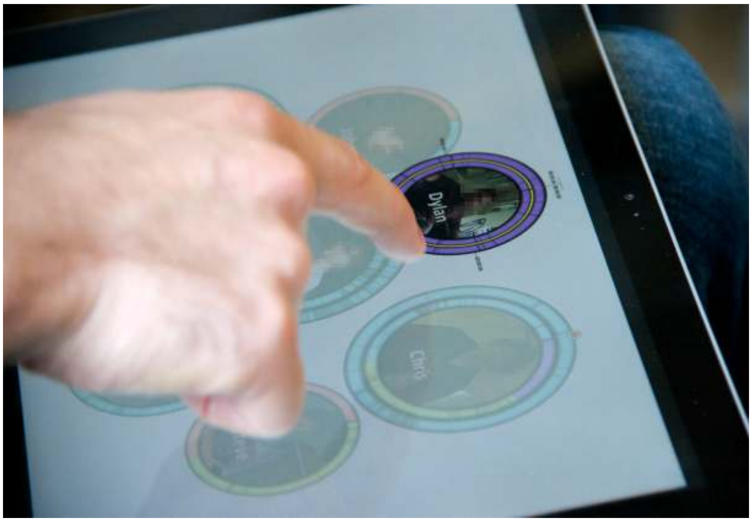

Figure 4. The touch-friendly design of the AOI cloud allows for analysis of the data on mobile devices such as tablets.

The time spans when a person received attention are easily accessible by the inner and outer rings. By adding markers to the outer ring, the user can define interesting events in the data and directly play back the corresponding video. With this approach, we simplify the exploration of multiple video sources in an easy-to-understand interactive visualization.

Due to the touch-friendly design of our visualization, users can also examine their data on the go on mobile devices (see Figure 4). This enables an easier integration of the application into the everyday life of the user, which is important for the long-term use of the application.

Use Case

Figure 5 shows a summarization of four videos from the coffee break dataset. Two videos (green and purple) are from the same session because the constellation of people changed after the first record ended. Altogether, eight individuals participated in the breaks and received different amounts of attention from the user wearing the eye-tracking glasses.

The user organized the participants in three groups based on the amount of attention they received:

- *Group 1*: Dylan and Russel appeared just once in different videos. Both received less attention than the others, especially Russel, who was sitting next to the user and only received attention when he was talking because the user had to turn to look at him. Dylan was watched when he was not talking because he was sitting in front of the user. Both people could have received a similar

Figure 5. AOI cloud for eight people over four videos. The items are freely adjustable and can be arranged by the user. In this example, three groups were created: Group 1 (Dylan and Russel) received the least attention; group 2 (Anya, John, and Steve) received medium attention; and group 3 (Jack, Oliver, and Chris) received the most attention.

amount of attention as those in group 2 if they had appeared in another video and if Russel had been seated in a better position.

- *Group 2:* Anya, John, and Steve appeared in two videos and were watched occasionally by the user. Steve could also be shifted to group 1 because he received little attention during his attendance in the coffee break.
- *Group 3:* Jack, Oliver, and Chris received most of the attention, although the distribution of attention depended on the constellation of people. For example, Oliver received a lot of attention in video 3 (green), when Chris, Steve, and he were present. During this coffee break, Chris left the room for half of the time (see markers at 00:02:35 and 00:05:59), at which points the main attention was on Oliver. In video 2 (blue), Oliver received less attention. In this video, as well as in video 1 (red), Jack was the attention catcher. Because Jack talked most of the time in both videos, the user gave him a good deal of attention. Hence, he received most of the attention although he was only present in two videos.

In this coffee break example, we can see that the amount of attention people received strongly depends on their position in the room, their active participation in discussions, and the other people attending at the time. People who talked less and required the user to turn to see them received less attention, especially when an attention-catching person was present. Thus, if the user would like to give more attention to some of the people from groups 1 or 2, talking with these people outside the coffee breaks when an attention-catching person such as Jack is not present might be an option.

Discussion

The AOI cloud provides an accessible approach to investigating the personal distribution of attention over several videos. The visualization approach is not restricted to people and could be applied to an arbitrary set of objects, assuming that it is possible to annotate the objects.

Although the most important AOIs will always be in the center of the initial cloud, a large number of AOIs and videos might reduce the readability of the visualization. Therefore, the scalability of our approach can be improved with additional filtering of the AOIs and video segments. By thresholding the attention values, AOIs that received less attention than the given threshold could be removed from the visualization. The same approach could be applied to the video segments of an AOI.

The presented visualization approach focuses on the analysis of individual relations between the user and other individuals. For future extensions, an analysis of group interactions would be beneficial for a reflection on personal social activity. By adding new options for examining attention changes between different people and how these changes correlate with people's activities, we could cover a comprehensive set of personal analysis interests.

Mobile eye tracking comprises most scenarios that can be achieved with head-mounted cameras or head tracking. Its main advantage lies in the additional gaze information. That is, when multiple objects are in the center area of a recorded image, we can derive detailed information about particular objects. A typical example could be a person looking at a picture collection. In this case, it would not be possible to identify the specific picture of interest without determining the user's gaze position. In addition, because the focus of this research is on personal scenarios, designing interfaces to combine mobile eye-tracking data with existing applications for personal visual analytics would be desirable.

To extend the possibilities of personal eye tracking in the near future, the challenges linked to the requirements we discussed here must be addressed. First, to increase accuracy, self-calibrating approaches need to be developed. Current techniques still rely on calibration procedures that are not feasible for a personal application. Also, managing the influence of uncontrolled lighting conditions in the environment introduces problems that require further research.

Second, defining areas or objects of interest by solely relying on computer vision might be hard to achieve in the near future. Arbitrary user-defined queries (for example, searching all cars in the videos of the database that received the user's attention) are required to process the recorded data to its full extent. Semiautomatic approaches and crowdsourcing could bridge the semantic gap in automatic approaches. Hence, visual analytics could help support such semiautomatic analysis.

Lastly, regarding cognitive processes, the interpretation of the gaze data itself has to be considered. Current approaches using cognitive modeling and machine learning to predict and classify gaze behavior (for example, detecting arousal or vigilance) need further development to provide more information than just distributions of attention. In our example, this information could be applied to weight the AOI circles. Additional information from measured pupil dilation can be included because current eye-tracking devices already record this data and preliminary work to correlate pupil changes with emotional states already exists. Supplementary sensors (such as heart rate sensors) can also provide such information and are already combined with mobile eye tracking.

Acknowledgments

This work was funded by the German Research Foundation (DFG) as part of the Scalable Visual Analytics Priority Program (SPP 1335). We thank Albrecht Schmidt of the University of Stuttgart for the discussion during the development of this work, and our reviewers for their constructive feedback.

References

1. R. Mantiuk et al., "Do-It-Yourself Eye Tracker: Low-Cost Pupil-Based Eye Tracker for Computer Graphics Applications," *Advances in Multimedia Modeling,* K. Schoeffmann et al., eds., LNCS 7131, Springer, 2012, pp. 115–125.
2. J.G. Hixon and W.B. Swann, "When Does Introspection Bear Fruit? Self-Reflection, Self-Insight, and Interpersonal Choices," *J. Personality and Social Psychology*, vol. 64, no. 1, 1993, pp. 35–43.
3. K. Kunze et al., "Activity Recognition for the Mind: Toward a Cognitive 'Quantified Self,'" *Computer*, vol. 46, no. 10, 2013, pp. 105–108.
4. G. Andrienko et al., "Visual Analytics Methodology for Eye Movement Studies," *IEEE Trans. Visualization and Computer Graphics*, vol. 18, no. 12, 2012, pp. 2889–2898.
5. B. Höferlin et al., "Scalable Video Visual Analytics," *Information Visualization*, vol. 14, no. 1, 2015, pp. 10–26.

Semiautomatic approaches and crowdsourcing could bridge the semantic gap in automatic approaches.

6. I. Li, A. Dey, and J. Forlizzi, "A Stage-Based Model of Personal Informatics Systems," *Proc. SIGCHI Conf. Human Factors in Computing Systems*, 2010, pp. 557–566.
7. T. Toyama et al., "Gaze Guided Object Recognition Using a Head-Mounted Eye Tracker," *Proc. Symp. Eye Tracking Research and Applications* (ETRA), 2012, pp. 91–98.
8. D. Huang et al., "Personal Visualization and Personal Visual Analytics," *IEEE Trans. Visualization and Computer Graphics*, vol. 21, no. 3, 2014, pp. 420–433.
9. J.J. Thomas and K.A. Cook, eds., Illuminating the Path: The Research and Development Agenda for Visual Analytics, IEEE CS Press, 2005.
10. D. Sprague and M. Tory, "Exploring How and Why People Use Visualizations in Casual Contexts: Modeling User Goals and Regulated Motivations," *Information Visualization*, vol. 11, no. 2, 2012, pp. 106–123.
11. A. Navab et al., "Eye-Tracking as a Measure of Responsiveness to Joint Attention in Infants at Risk for Autism," *Infancy*, vol. 17, no. 4, 2012, pp. 416–431.
12. R.S. Jasinschi et al., "Integrated Multimedia Processing for Topic Segmentation and Classification," *Proc. 2001 Int'l Conf. Image Processing*, 2001, pp. 366–369.
13. M.A. Just and P.A. Carpenter, "A Theory of Reading: From Eye Fixations to Comprehension," *Psychological Rev.*, vol. 87, no. 4, 1980, pp. 329–354.
14. A. Smeulders et al., "Content-Based Image Retrieval at the End of the Early Years," *IEEE Trans. Pattern*

Analysis and Machine Intelligence, vol. 22, no. 12, 2000, pp. 1349–1380.

15. G. Draper, Y. Livnat, and R. Riesenfeld, "A Survey of Radial Methods for Information Visualization," *IEEE Trans. Visualization and Computer Graphics*, vol. 15, no. 5, 2009, pp. 759–776.
16. G. Draper and R. Riesenfeld, "Who Votes for What? A Visual Query Language for Opinion Data," *IEEE Trans. Visualization and Computer Graphics*, vol. 14, no. 6, 2008, pp. 1197–1204.
17. T. Blascheck et al., "State-of-the-Art of Visualization for Eye Tracking Data," *Proc. EuroVis State of the Art Reports*, 2014; http://dx.doi.org/10.2312/eurovisstar.20141173.
18. F.B. Viéas and M. Wattenberg, "Timelines Tag Clouds and the Case for Vernacular Visualization," *Interactions*, vol. 15, no. 4, 2008, pp. 49–52.

Kuno Kurzhals *is a research assistant working in the Visualization Research Center at the University of Stuttgart. His research interests include eye tracking, visualization, visual analytics, and computer vision, especially developing new visualization methods to analyze eye-movement data from dynamic stimuli. Kurzhals has a Diplom in computer science from the University of Stuttgart. Contact him at kuno.kurzhals@vis.uni-stuttgart.de.*

Daniel Weiskopf *is a professor of computer science in the Visualization Research Center and the Visualization and Interactive Systems Institute at the University of Stuttgart. His research interests include visualization, visual analytics, GPU methods, computer graphics, and special and general relativity. Weiskopf has a doctor rerum naturalium (PhD) in physics from Eberhard-Karls-Universität Tübingen. Contact him at weiskopf@visus.uni-stuttgart.de.*