

ENGAGE: Early Insights in Measuring Multi-Device Engagements

Rachel Jones, Sarah Clinch, Jason Alexander, Nigel Davies, Mateusz Mikusz

School of Computing & Communications

Lancaster University, Lancaster, UK

r.jones1 | s.clinch | j.alexander | n.davies | m.mikusz @ lancaster.ac.uk

ABSTRACT

Users are increasingly accessing content through a complex device eco-system involving both public and private screens. Traditional research into display eco-systems has focused on developing new multi-screen applications and on techniques for understanding how interactions and activities such as shopping flow across screens and devices. There has been relatively little research into the more fundamental question of how users actually engage with multiple screens and in particular how levels of engagement can be systematically monitored. In this paper we describe our early experiences with ENGAGE – a toolkit designed to help researchers explore user engagement across multiple devices.

Author Keywords

multi-device engagement; audience tracking; multi-device applications.

ACM Classification Keywords

H.5.m. Information Interfaces & Presentation (e.g. HCI): Misc.

INTRODUCTION

We are rapidly moving towards a world in which displays are ubiquitous. Users have a wide range of devices on which they can view content—from personal devices such as Google Glasses and smartphones, through conventional laptops and desktops to pervasive digital signage. Crucially users are choosing to use many of these devices in parallel—surfing the web while watching TV or glancing up from their laptop at a nearby digital sign. As a result, users and hence applications, must function in complex multi-screen eco-systems.

To date research on multi-screen eco-systems has focused on two distinct themes. The first is the development of applications executing across multiple screens, e.g. to display private information on a personal device while showing non-sensitive information on a larger, public display [3]. The second relates to analytics that track interactions over multiple devices

in order to help improve web analytics in the face of multi-device interaction [10]. However, a more fundamental question is *how do users actually engage with individual screens in a multi-screen eco-system?* Such engagement is, of course, likely to be related to a range of contextual factors, e.g. the content shown, and the tasks being undertaken by the user.

Digital signage manufacturers have long recognised the need to produce analytics that attempt to capture engagement and a wide range of video analytics tools exist (e.g. [12, 13, 15]). However, understanding the level of engagement users have when confronted with multiple devices is becoming increasingly important. How, for example, does a user's phone ringing impact on their engagement with public displays—do they look up and hence see more content or do they focus on their smartphone? Do users find displays more or less engaging if the content they can see flows across multiple displays or is it better to have distinct content on each display?

Gaining insight into user engagement requires tools for capturing engagement levels across multiple displays being used concurrently. The key issue for researchers is: how should these tools be designed and built? For example, are specific new hardware and software systems required or can existing components be repurposed? What sampling rates are required for accurate understanding of user engagement? What constitutes engagement with a device such as a mobile phone (holding, viewing, calling etc.)? To answer these questions requires a systematic program of research that studies many facets of engagement and attempts to build robust solutions that work in a wide variety of contexts. Such a program of research is likely to involve the creation of many probes that produce data points within the overall design space. In this paper we report on early experiments to try and measure engagement levels across multiple devices and provide one such data point. We make three specific contributions:

1. the ENGAGE system: an investigation using existing commonplace software and hardware components to measure user engagement with three distinct types of screen,
2. a description of a generalised architectural approach for measuring user engagement, and,
3. a description of tests designed to measure the effectiveness of these engagement tools in a laboratory setting.

There are numerous audience behaviours that may constitute engagement – for example, subconscious glances; actively switching visual attention; or interacting through touch, ges-

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

PerDis '15, June 10 - 12, 2015, Saarbruecken, Germany.

Copyright © 2015 ACM. ISBN 978-1-4503-3608-6/15/06?\$15.00

DOI: <http://dx.doi.org/10.1145/2757710.2757720>

tures and other mechanisms. For the purpose of this paper, we focus on visual interaction (i.e. looking at the display) as a sign of “engagement” and develop methods to analyse this visual engagement with and between displays. We use image processing techniques to recognise when an individual is facing a particular screen with the assumption that when they are viewing the screen, they are engaged with its content. We use the period of interaction as a key measure of engagement, with content which is interacted with for long periods of time having a higher level of engagement. Despite our focus on visual engagement, we believe that the framework presented in this paper can easily support other engagement measures through the development of additional “engagement sensors”.

This paper describes the design and implementation of our prototype, ENGAGE, and our evaluation results. We do not argue that the tools we describe provide the solution for measuring user engagement—rather we believe that the results will be useful to researchers wishing to explore multi-screen engagement. In particular, we hope to fuel debate on the type of testing methodology that would be appropriate for future generations of such tools while providing a common architectural framework within which these tools can be developed.

RELATED WORK

Measuring engagement and tracking audience behaviour is an important area of study when evaluating public displays [1]. Intel’s AIM suite [13] provides visual analytics that enable display owners to track their audience in real-time, gathering comprehensive demographic information about passers-by. Like ENGAGE, AIM uses video streams to anonymously detect faces and then uses classifiers to derive demographic information. A field trial, combining data from AIM with point-of-sale for measuring the effectiveness of targeted advertising, has demonstrated the accuracy of the system [14]. Other commercial systems also provide audience tracking (e.g. IBM’s video analytics [12]), and can use audience information to influence behaviour in order to meet certain business goals (e.g. Scala [15, 17]). Within the research domain, similar approaches support displaying personalised adverts and applications on the screen [8]. Furthermore, by fusing sensor and location data from multiple devices and data sources (smartphones and on-screen video analytics) with personal data shared on social networks, researchers have created context-aware displays and applications that allow the content to “follow” the user across multiple screens [2, 9].

In addition to simply describing and tracking passersby, visual analytics have also been used to measure the actual levels of engagement for those in front of displays. For example, Hernandez et al., used a video camera to capture 47 participants in a “naturalistic” living room and manually classified user’s engagement with a television screen into four groups of attention: high, medium, low and none – their classifiers used face alignment technology to recognise face distance and angle, head roll, head size and head position and from these they derived their levels of engagement [11]. Other approaches have looked at non-facial behaviours, for example by using depth video cameras to capture gestures of passers-by interacting with an interactive public display [18]. Another

powerful approach for measuring engagement is through eye-tracking. Wedel and Peters used gaze tracking to investigate if the number of eye fixations reflected how well subjects engaged with full-page advertisements within a magazine [20]. More recently, Dalton et al. [7] used eye-tracking to detect engagement with screens in a shopping mall.

The systems described above have each attempted to capture engagement with a single display or physical object. In contrast, this work is concerned with the measurement of engagement with multiple devices used either in parallel or in close succession. Prior work in this space does exist but is typically focussed on laboratory settings in which custom hardware for tracking is acceptable. For example, Cauchard et al. [6] measured switches in visual context between a phone screen and projected display, but required the user to wear eye-tracking glasses and equipped the setting with IR markers. By contrast our work aims to develop an approach that may have future applications in real deployments, and as such is based on cheap, simple, readily-available tools and hardware.

A final area of prior work relates to our use of OpenCV [16] for face recognition. This library provides an implementation of Viola and Jones’ real-time object detection algorithm [19] which has been shown to achieve a 95% accuracy rate. In this work we focus not on measuring the accuracy of OpenCV, but instead explore its use to support the development of tools for tracking multi-device engagements in the wild.

DESIGN AND IMPLEMENTATION

Overall Architecture

To begin our exploration into multi-device engagement data, we have built a system called ENGAGE. As shown in Figure 1, our system is divided into two core components: *Engagement Sensors* and an *Engagement Analytics Service*.

Engagement Sensors represent software and hardware components deployed either on a device or in the surrounding environment that are able to capture and report levels of engagement. Examples of such sensors might include eye tracking devices, key-stroke detectors or video analytics feeds. Our prototype system explores the use of a subset of sensors to monitor content and viewers that are targeted at three distinct classes of display: personal computers (e.g. laptops, desktops), public displays, and personal mobile devices (e.g. smartphones). We propose that additional sensors should be developed independently of those described below to support additional devices (i.e. smart watches) using appropriate technology whilst feeding to the same Engagement Analytics Service. Development of additional sensors could also allow support for tracking a wider range of engagement behaviours (e.g. change of trajectory when passing a public display, keyboard interactions with a laptop).

The Engagement Analytics Service is a set of back-end components to which Engagement Sensors can report their measurements. The Engagement Analytics Service provides data storage and analytics for sensor measurement data.

Our realisation of these two core components is described in the following sections.

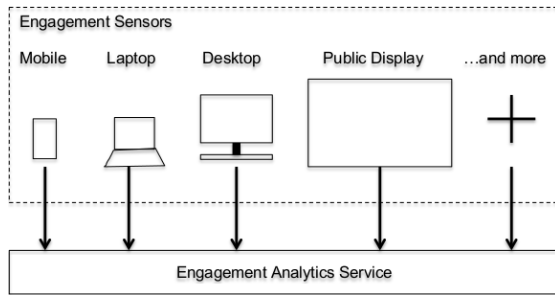


Figure 1: The ENGAGE Architecture.

Engagement Sensors: Detecting Engagement

Computers and Public Displays

Our first set of sensors are designed to measure engagement on desktop computers, public displays and laptops and run as a Java program, gathering the data that represents a ‘device interaction’ analytics event.

The developed toolkit uses a webcam and the OpenCV library [16] to determine the number of content viewers and the duration of their engagement. For some devices user input may cause a change in content (e.g. when operating a phone or laptop) and so for these classes of devices interaction is also counted when the content is seen to change.

OpenCV object detection is used to identify observers – we assume that when a human is facing and looking at the display, they are engaged with the display content. Our toolkit captures images from either a built-in or external webcam (as selected through the user interface) and looks for observers using existing frontal face and eye classifiers: “20x20 profile face detector” [4] and “22x5 Eye pair detector computed with 7000 positive samples” [5]. The differing nature of displays and their human interactions are best tackled by varying the classifiers used. For this reason, our interface provides the user with a drop-down list of device types to choose from upon set-up—each device type maps to a classifier optimised for a typical setup (e.g. distance from screen, body angle). For example, public display interaction occurs at a greater distance to that of a computer. For this reason when ENGAGE runs on a public display it attempts to identify frontal faces, whereas on a laptop or desktop the toolkit confirms the classification of a viewer by cropping the detected face region and checking for a pair of eyes. Classifiers used by the system can be easily changed and extended, and we anticipate that future work would focus on the implementation of a training system for space-specific engagement classifiers.

Mobile Devices

In addition to detecting engagement with conventional displays we also wished to explore engagement with mobile devices. To reduce the use of limited resources, OpenCV face detection is not implemented.

To the user, the ENGAGE for Android system is a simple app, with a single ‘dashboard’. There are two options presented: a toggle switch to enable and disable the monitor and a button to change the owner’s ID.

Using Android Alarm Manager, the application makes regular checks to find the focused application and whether or not the screen is lit. If the activity has changed, and the device is lit, then we assume the user is interacting with the mobile device and send a device interaction hit to the Engagement Analytics node. We have also implemented checks for audio interaction such as music and phone calls, these are counted as interaction events.

Engagement Analytics: Making Sense of Engagement

In order to capture the results of our Engagement Sensors we make use of an existing analytics backend that runs on Google App Engine. We have developed a specific set of procedures to process device interaction and engagement.

All nodes communicate with the Engagement Analytics Service through HTTP posts with JSON content. Sensors nodes can either send their interaction events in real-time, one interaction object at a time, or may choose to batch upload interactions – batching interactions allows a device to save resources during busy periods and instead send an array of events when resource use is lower (e.g. when a user is no longer interacting). Each device engagement is represented by a set of data values including: a client ID that links the data to an individual user or public content provider, and the details of the content which was interacted with. Each communication with Engagement Analytics contains an Engagement Analytics Service ID to determine what is done with the received data. All real-time engagement events are automatically tagged with the appropriate interaction time by the Engagement Analytics Service; for batch uploaded data an analytics timestamp object is included with each interaction event in order to override this default interaction time.

The Engagement Analytics Service runs a series of validation checks on data before saving to the Google Datastore. The resulting data can then be analysed using a range of tools.

LAB-BASED EVALUATION

Methodology

In order to assess the effectiveness of the ENGAGE systems in their ability to follow a user’s engagement with multiple devices, we conducted a small-scale lab-based experiment. In this experiment each participant was given a task to complete that was designed such that engagement with multiple devices was required, with a predicted pattern of movement between devices. The study was intended to measure the ability of the ENGAGE systems to register a user’s engagement and follow their flow from device to device as they completed the task. The multi-device environment was designed to imitate that of a public space with personal and public devices, the devices and their content were designed and positioned in a way that encouraged the subjects’ interactions to shift between the three devices accompanied by an obvious physical movement. Forcing physical movement helped improve the laptop and public display sensors’ classification of engagement and was particularly useful when coding the ground truth from the video recording.

Our experimental setup is shown in Figure 2. Each participant was seated on a swivel chair with a laptop on a desk 1.25 metres to their left and a 50" display (showing the display output of a second laptop) 2.89 metres in front of them. Each participant held an Android mobile phone. The laptop driving the public display ran Mac OS 10.10.1 (Yosemite), and the other laptop ran Mac OS 10.9.5 (Mavericks)—both ran the Mac implementation of the ENGAGE software. The public display was instrumented with a web camera with a resolution of 1280 x 720 and the laptop used its on board camera to capture video at the same resolution. The smartphone ran Android 4.3 and the Android implementation of the ENGAGE software. Participants were asked to remain seated and not to move the chair, laptop or large display but were free to turn their chair or to move the mobile phone as they wished.




Figure 2: Experimental Setup.

The experimental task required a participant to first read a written experimental brief on the large display, and then watch a series of eighteen videos each with a duration of 6–125 seconds (mean duration 62.17 seconds; median 52.50 seconds; total combined duration 18 minutes 39 seconds). The videos were randomly ordered into a playlist, and this same ordered playlist was then used for each participant. Videos were played alternately on the laptop and large display beginning with the laptop (i.e. one video on the laptop, followed by one on the large display; overall nine videos on each of the laptop and large display). Using a mixture of video durations allowed us to ensure that transition points between devices were not able to be predicted by our participants. Transitions between videos were achieved using a cross-fade to black and the screen on the ‘unused’ device remained black for the duration of the video being shown on the other device. Between the brief and the eighteen videos, we had a total of nine content transitions from display to laptop, plus nine from laptop to display.

While watching the videos, participants were also required to complete a series of 36 ordered multiple choice quiz questions delivered using the Android smartphone. The smartphone showed a single question at a time; completion of a question resulted in the immediate display of the next question. Each quiz question related the video content (thus encouraging engagement with both the quiz and the videos): two questions per video, delivered in the same order as video

playback. Quiz questions required no additional knowledge beyond the video content and drew on both audio and visual aspects of the videos. Both videos and quiz questions were deliberately selected to require no prior skill or particular interest. Participants were advised that they could complete the quiz questions at any time during video playback.

Throughout the experiment, the ENGAGE systems ran on each device and sent data through the Engagement Analytics Service running on Google App Engine so that data could be viewed and analysed. In order to capture ground truth we additionally videoed the study and had the experimenter manually code the engagement captured in the video.

We recruited 5 participants through our own personal social networks. Participation was voluntary and no compensation was provided for participating. Our participants were roughly gender balanced (3 male, 2 female), under 25 years of age, with a high degree of computer literacy. The overall duration of the study was typically around 25 minutes including consent, briefing and debriefing.

Results

Throughout the experiment, the ENGAGE system reported that a participant was or was not engaging with each of the devices. Our groundtruth (based on the manual video coding) also allowed us to identify exactly which device (or no device) had the participants attention at each point in time. We compared these two datasets for each participant in order to identify the correlation between the reports from ENGAGE and the groundtruth. The results for our five participants are shown in Figure 3.

Looking at the mobile and laptop measurements, we sampled the data into 1 second chunks and compared the groundtruth engagement count for each device with the ENGAGE report for that same device. From this we calculated overall accuracy (ratio of identical reports to sample size) and the false positive and negative rates (as a proportion of sample size). Across our participants we see that the ENGAGE software on the laptop accurately counts engagement in 67.35–79.70% (median 73.38%) of cases. When an accurate engagement count is not provided, we find that the laptop ENGAGE software is more likely to fail to classify a viewer (i.e. a false negative) than to incorrectly report someone who is not engaging. The software over-reports engagement with the device in 3.77–12.04% (median 10.20%) of cases, and under-reports in 8.88–28.88% (median 16.41%) of cases. Accuracy of the ENGAGE software on the public display is slightly poorer. The ENGAGE software on this device accurately counts engagement in 51.33–68.69% (median 63.42%) of cases. As on the laptop, the public display software more frequently reports false negatives than false positives; the public display ENGAGE software over-reports engagement with the device in 4.89–23.18% (median 7.76%) of cases, and under-reports in 23.55–32.83% (median 25.49%) of cases.

The ENGAGE software for the laptop and display continuously report changes in engagement counts. By contrast, the mobile software only provides spot readings that indicate that an engagement was present at any point in time. For

Session I: ENGAGEMENT AND ATTENTION




Figure 3: Plots showing the engagement periods for each device over the course of the experiment (one plot per participant). Darker lines show the engagement periods as measured by the ENGAGE software and the lighter lines ground truths for the same devices.

this reason, we calculate the accuracy of the mobile software slightly differently. Looking again at our 1 second samples, we compare the groundtruth and ENGAGE reporting—we calculate overall accuracy of the ENGAGE reports by dividing the number of identical reports by the total number of reports made by the ENGAGE software. Across our participants we see that the ENGAGE software on the mobile provides an accurate count of engagement in 41.67–94.44% (median 80.00%) of its reports. This means that the software incorrectly reports an engagement when there is none in 5.56%–58.33% (median 20.00%) of cases. However if we look at the number of ENGAGE reports as portion of sample size, we note that the ENGAGE mobile software is unable to provide an engagement count for around 96.59–97.01% (median 96.96%) of the sampled time.

These statistics provide an initial indication of the quality and variability of ENGAGE's classification over three device types. We note however, the impact of a small sample on the validity of these results—in this case, exploring these statistics is less about validating our deployment, and more about consideration of the process for evaluating such systems.

DISCUSSION

One of the key questions we wished to explore with the ENGAGE toolkit was how we should structure engagement capture systems. Our basic approach of simple, low-cost hardware and software sensors together with a more complex back-end provided a suitable framework for the ENGAGE prototype. This is perhaps unsurprising – this basic approach has also been used in systems such as Google's web analytics. However, unlike web analytics in which most developers using existing tracking code for engagement researchers we would expect them to develop their own engagement sensors. Having a simple framework with the majority of the work being conducted by the back-end considerably simplifies this process and is, we believe, likely to lead to a reduced barrier to the proliferation of engagement sensors.

The experimental design for evaluating our prototype toolkit proved non-trivial. Ground truth is clearly required for any successful engagement evaluation and yet this requires capture and subsequent analysis of large quantities of video material. This can lead to a significant overhead in reviewing the material as well as raising significant ethical issues of trials that take place outside closely controlled lab settings.

We note that it is not possible to simply leave users to interact with multiple devices and then assess the accuracy of the engagement measures, as such unstructured interaction is likely to lead to unreproducible results and will not test all engagement combinations. For example, we were keen to explore whether our toolkit was able to capture transitions between devices and hence we needed an experimental setup that caused subjects to engage with devices in predictable sequences. We believe that the methodology described could be useful for a range of engagement trials. Using a commonly agreed methodology for engagement toolkits would also enable researchers to report on a common set of metrics. For example, we report overall accuracy, number of false positives, and the number of unknown states and/or false negatives as

applicable. Unfortunately, for true reproducibility we would also need to capture a wide range of contextual information such as ambient lighting conditions and typical viewing distances. The community will need to develop a common way of describing these experimental set-ups to enable effective comparison between tools.

With privacy being a common concern when implementing face-tracking systems, it is important that the systems are designed such that the sensors can be used in public settings without violating the privacy of passers-by to gain better chance of being accepted in many scenarios such as schools and shopping centres. Once in place, the simple public engagement sensors should collect data in a anonymous manner and prevent storage of the grabbed web-cam feed. As for personal devices, individuals can expect to use a unique and unpredictable ID to retrieve their engagement analytics.

Currently the system is not capable of identifying an individual within a public space of multiple users. If we designed and extended the system to possess a face re-recognition feature, the ENGAGE system could link an individual's ID to the viewer and the interaction would be found within the individual's engagement reports. This would allow more depth for public display owner's reports as they would then be able to know the number of unique viewers and gain more insight into their content distribution. The main issue evident when considering re-recognition is the privacy of users and the size of database needed to store the classifier for each individual. Other options for individual identification include radio-frequency identification (RFID) or matching up GPS locations between devices, although both of these options are likely to be less accurate.

Finally, we note that our experiments have been largely positive—it appears that simple engagement sensors implemented using common freely available software are able to provide a first approximation of engagement levels in controlled settings.

CONCLUSIONS

Understanding engagement with multiple displays becomes increasingly important as users access more complex display eco-systems. Our work represents an early experiment into techniques for measuring multi-device engagement. Our experiences with using easy-to-develop sensors within a generalised architecture for engagement capture have been largely positive and we have shown that we are able to capture engagement with reasonable accuracy in the laboratory. We believe that the architectural framework and evaluation methodology will be of value to a wide range of researchers interested in exploring engagement with public displays.

ACKNOWLEDGMENTS

This research is partially funded through the Future and Emerging Technologies (FET) programme within the 7th Framework Programme for Research of the European Commission, under FET grant number: 612933 (RECALL), and was made possible with the support of a Google Faculty Research Award and a Google Cloud Credits Award.

REFERENCES

1. Alt, F., Schneegaß, S., Schmidt, A., Müller, J., and Memarovic, N. How to evaluate public displays. In *Proceedings of the 2012 International Symposium on Pervasive Displays*, PerDis '12, ACM (New York, NY, USA, 2012), 17:1–17:6.
2. Beach, A., Gartrell, M., Xing, X., Han, R., Lv, Q., Mishra, S., and Seada, K. Fusing mobile, sensor, and social data to fully enable context-aware computing. In *Proceedings of the Eleventh Workshop on Mobile Computing Systems & Applications*, HotMobile '10, ACM (New York, NY, USA, 2010), 60–65.
3. Berger, S., Kjeldsen, R., Narayanaswami, C., Pinhanez, C., Podlaseck, M., and Raghunath, M. Using symbiotic displays to view sensitive information in public. In *Proceedings of the Third IEEE International Conference on Pervasive Computing and Communications*, PerCom 2005 (March 2005), 139–148.
4. Bradley, D. 20x20 profile face detector, 2000.
5. Castrillón Santana, M., Déniz Suárez, O., Hernández Tejera, M., and Guerra Artal, C. Encara2: Real-time detection of multiple faces at different resolutions in video streams. *Journal of Visual Communication and Image Representation* (April 2007), 130–140.
6. Cauchard, J. R., Löchtfeld, M., Irani, P., Schoening, J., Krüger, A., Fraser, M., and Subramanian, S. Visual separation in mobile multi-display environments. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*, ACM (2011), 451–460.
7. Dalton, N., Collins, E., and Marshall, P. Display blindness? looking again at the visibility of situated displays using eye tracking. In *Proceedings of CHI '15*, ACM (2015).
8. Farinella, G., Farioli, G., Battiato, S., Leonardi, S., and Gallo, G. Face re-identification for digital signage applications. In *Video Analytics for Audience Measurement*, C. Distante, S. Battiato, and A. Cavallaro, Eds., Lecture Notes in Computer Science. Springer International Publishing, 2014, 40–52.
9. Gillian, N., Pfenninger, S., Russel, S., and Paradiso, J. Gestures everywhere: A multimodal sensor fusion and analysis framework for pervasive displays. In *Proceedings of the 3rd International Symposium on Pervasive Displays* (2014).
10. Google Inc. About the user id feature—connect multiple devices, sessions, and engagement data to the same visitors. <https://support.google.com/analytics/answer/3123662>.
11. Hernandez, J., Liu, Z., Hulten, G., DeBarr, D., Krum, K., and Zhang, Z. Measuring the engagement level of tv viewers. *2013 10th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG) 0* (2013), 1–7.
12. IBM Intelligent Video Analytics. Tech. rep., IBM Corporation, June 2013.
13. Intel Corporation. Intel AIM Suite. <https://aimsuite.intel.com/>.
14. A Report on a Field Trial of Anonymous Video Analytics (AVA) in Digital Signage. Tech. rep., Intel Corporation, 2011.
15. Menger, S. In-store messages optimized by analytics – case studies and results – pri journal of retail analytics. <http://scala.com/in-store-messages-optimized-by-analytics-case-studies-and-results-pri-journal-of-retail-analytics/>, May 2014.
16. <http://opencv.org>.
17. Patzer, M. Combining predictive analytics and digital signage to boost retail performance. <http://www.retailtouchpoints.com/executive-viewpoints/1569>, 2012.
18. Tomitsch, M., Ackad, C., Dawson, O., Hespanhol, L., and Kay, J. Who cares about the content? an analysis of playful behaviour at a public display. In *Proceedings of the 3rd International Symposium on Pervasive Displays* (2014).
19. Viola, P., and Jones, M. Rapid object detection using a boosted cascade of simple features. In *Computer Vision and Pattern Recognition, 2001. CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on*, vol. 1 (2001), I–511–I–518 vol.1.
20. Wedel, M., and Pieters, R. Eye fixations on advertisements and memory for brands: A model and findings. *Marketing Science* 19, 4 (Oct. 2000), 297–312.