EngageMeter: A System for Implicit Audience Engagement Sensing Using Electroencephalography

Mariam Hassib[1,2], Stefan Schneegass[2], Philipp Eiglsperger[1], Niels Henze[2], Albrecht Schmidt[2], Florian Alt[1]

[1]LMU Munich – Ubiquitous Interactive Systems Group – *{firstname.lastname}*@ifi.lmu.de
[2]University of Stuttgart – VIS – *{firstname.lastname}*@vis.uni-stuttgart.de

ABSTRACT
Obtaining information about audience engagement in presentations is a valuable asset for presenters in many domains. Prior literature mostly utilized explicit methods of collecting feedback which induce distractions, add workload on audience, and do not provide objective information to presenters. We present EngageMeter – a system that allows fine-grained information on audience engagement to be obtained implicitly from multiple brain-computer interfaces (BCI) and to be fed back to presenters for real time and post-hoc access. Through evaluation during an HCI conference ($N_{audience}$=11, $N_{presenters}$=3) we found that EngageMeter provides value to presenters (a) in real-time, since it allows reacting to current engagement scores by changing tone or adding pauses, and (b) post-hoc, since presenters can adjust their slides and embed extra elements. We discuss how EngageMeter can be used in collocated and distributed audience sensing as well as how it can aid presenters in long term use.

Author Keywords
Physiological Sensing; Audience Feedback; EEG; BCI.

ACM Classification Keywords
H.5.2 Information Interfaces and Presentation

INTRODUCTION
Presenting in front of an audience is an integral part of everyday work in academia, education, and industry. Presenters communicate their latest results and explain new topics and ideas to colleagues and interested parties using slide-based presentations. Gaining feedback from the audience is important to ensure information is delivered and to keep the audience engaged and attentive. However, reliable and fine-grained feedback from the audience is hard to collect. Mostly presenters use non-verbal cues (e.g., eye contact, posture) to perceive audience engagement, or collect feedback after a performance or presentation using interviews and questionnaires. Recently, researchers proposed using video streams to automatically predict audience feedback through facial expressions [7, 11].

CHI 2017, May 6-11, 2017, Denver, CO, USA.

ACM ISBN 978-1-4503-4655-9/17/05 ...$15.00.
http://dx.doi.org/10.1145/3025453.3025669

Figure 1. Three participants in a presentation in a scientific conference during the real-world study of EngageMeter

Due to wearable devices, such as wrist-worn heart rate sensors or head-mounted consumer brain-computer interfaces (BCIs), rich information about the physical, emotional, and cognitive state of users can be collected. In contrast to current wearable mainstream applications, we extend the use of wearable psycho-physiological sensors beyond personal tracking to the aggregation of information from multiple users with the goal of obtaining audience feedback. Only few systems utilized information from physiological sensors (e.g., SCL) in real-world audience sensing [14, 19]. However they provide only a post-hoc view of raw data to presenters. We propose EngageMeter a system exploiting rich information gathered by BCIs to provide real-time and post-hoc feedback to presenters. We leverage the fact that electroencephalography (EEG) signals from the brain are able to detect shifts in engagement, alertness, and workload [4, 5, 9].

In this note, we first discuss our implicit audience sensing concept and implementation, called EngageMeter. We then present a real-world evaluation of EngageMeter in an HCI conference. Findings stem from interviews conducted with presenters and audience members. The evaluation shows how presenters react in real time to fluctuations in audience engagement by introducing pauses or changing tone. In post-hoc, presenters gained insights about audience engagement on a per-slide basis which allows for adapting and changing their presented material. We conclude with a discussion on using psycho-physiological sensing for implicit audience feedback.

ENGAGEMETER
EngageMeter consists of three components as follows:

Engagement Sensing Component. Fundamental EEG research [13] provided a formula to calculate cognitive engagement using $\alpha(7-11Hz)$, $\beta(11-20Hz)$, and $\theta(4-7Hz)$ frequency bands, where E, representing the engagement index, is

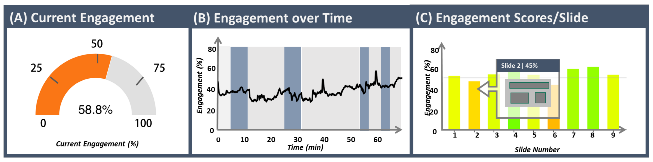

Figure 2. Feedback Views: (A) Current engagement gauge showing normalized audience engagement in percent in real-time. (B) Moving graph showing the audience engagement over time, vertical sections indicate slide changes, this view is shown in real-time and post-hoc. (C) Slide scores view showing the average engagement score per slide in post-hoc.

calculated as: $E = \frac{\beta}{\alpha+\theta}$ (1). The index reflects visual processing and sustained attention [4] and is able to identify changes in attention related to external stimuli [4, 12]. We built on top of prior research that utilize consumer EEG headsets which proved their success in detecting cognitive engagement in the learning domain [2, 8, 10, 15, 16] as well as in other domains [1, 20, 18]. We use the Neurosky Mindwave headset[1] (see Figure 1), a light-weight, dry-electrode EEG device. It collects EEG data at 512 Hz from the frontal cortex (FP1) according to the 10–20 positioning system. This brain region is related to learning and cognitive states such as engagement [4, 6]. To collect and process EEG signals, we developed an Android application that connects to the Mindwave via Bluetooth. We apply a Fast Fourier Transform to the raw signal to extract the relevant frequency bands (β, α, θ) averaged over 1 second. We calculate the 1-second engagement index E. To filter the signal from muscle (e.g., blinking) artifacts, we used a five second sliding window approach as proposed by Szafir and Mutlu [16]. We then smooth the engagement index using an Exponentially Weighted Moving Average to pick up general engagement trends and further remove movement artifacts [16]. This outputs a smoothed engagement index per device per 5 seconds E_{smooth} sent to the administration component.

Administration Component. Composed of a web server and a database, the server provides presenters with a front end to create new sessions. When the presenter starts the presentation, she/he chooses a session name, description, and type. We offer two different session types: (1) calibration and (2) recording sessions. Calibration is done since EEG signals are highly person-dependent. Hence a generic solution would not be possible for calculating an average engagement score depicting the entire audience without accounting for person dependencies. It is an extra session conducted before the start of the presentation to determine maximum and minimum values for engagement through low (e.g., relaxation) and high (e.g., solving visual puzzles) engagement-inducing tasks [17]. Based on the minimum E_{min} and maximum E_{max} engagement scores collected from each calibration session, we calculate a normalized engagement score between 0 and 100 as $E_{norm} = \frac{E_{smooth}-E_{min}}{E_{max}-E_{min}} * 100$ (2) similar to Vi et Al.'s work[17]. Recording sessions use the normalized engagement scores from each audience member to calculate an average engagement between 0 and 100 for all audience members at the same time presented to presenters via the *Feedback Client*. The database stores normalized engagement scores, session data, and timestamps of the slides of the presented material.

[1] http://neurosky.com/biosensors/eeg-sensor/

Feedback Client. The normalized engagement score is sent from the administration component together with the slide timestamps to the front-end web client where it is visualized. Any registered user (administrator or audience) can gain access to the feedback by signing in and providing a session ID. Thus, EngageMeter allows the feedback to be seen by multiple presenters (only averaged scores of all audience members) or the audience themselves (who can access personal and average audience scores). The design of the feedback client was informed through two pre-studies with six presenters. A real-time and a post-hoc view are provided.

The *real-time view* shows current audience-averaged normalized engagement score represented as a gauge (cf. Figure 2,A) and a moving line graph with the average engagement over time, where vertical sections indicate slides (cf. Figure 2, B). The engagement gauge provides presenters with a quick view that can be comprehended in short glances. It shows the averaged engagement of the audience as percentage and reflects the color in the gauge between red (0%) and green (100%). The moving graph gives a holistic view of the presented material so far, showing slide lengths and variations of engagement during each slide which presenters can use if they are pausing.

The *post-hoc view* shows a line graph with normalized audience-averaged engagement over time (cf. Figure 2, B), and individual slide scores in a bar chart (cf. Figure 2, C). Each slide score is calculated as the average of normalized engagement over the slide duration. Presenters can upload their slide deck and presentation audio through the interface and can see visual slide previews and replay the audio.

REAL-WORD EVALUATION OF ENGAGEMETER

To evaluate the concept of EngageMeter, we deployed it in a real world setting during a large HCI conference. In particular, we focus on three keynotes given by experienced presenters over the course of three days.

Participants and Presenters

We recruited 11 participants from the audience to take part in our study (8 males, 3 females) aged between 24 and 28 years ($M = 25.2$, $SD = 2.14$). All participants were graduate students from computer science, HCI, or psychology and were attendees of the conference. All participants attended keynotes 1 and 3, whereas two participants missed keynote 2. They received 20 Euros for participation. The three keynote speakers (2 female) were experienced presenters. Two were academic researchers, one was an industry professional.

Procedure

The day before the conferences started, we invited the participants to the venue and briefed them about the study. They signed informed consent forms and filled in a demographic questionnaire. We introduced them to the overall system and the BCI in particular. We conducted a calibration, consisting of two sessions – a relaxation session, common to BCI studies [8, 9, 17], and a visual puzzle solving session which was proven to increase engagement scores to almost double that of the relaxation task during prestudies. Each session lasted for five minutes. We used both calibration sessions to determine a minimum and maximum engagement index per participant

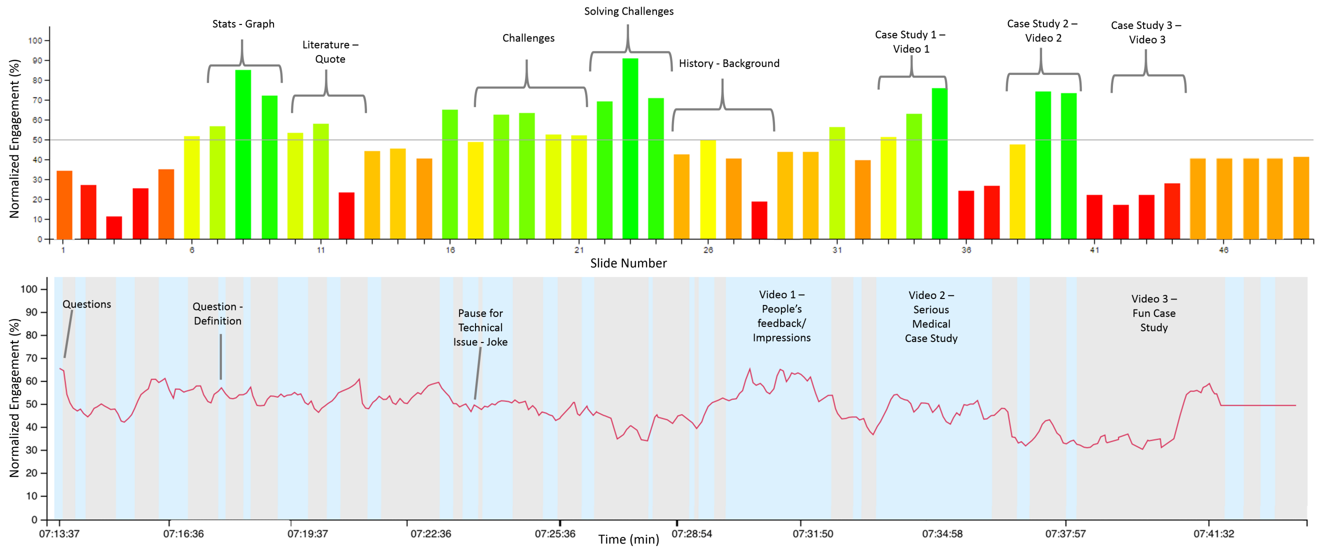

Figure 3. Keynote 2 results: Top graph shows slide scores, slide durations are not depicted, x-axis represents slide numbers. Bottom graph shows the overall engagement levels of the audience members which presenters saw in real-time and at the end of the talk, vertical sections show slide changes. Labels show presentation sections (e.g. videos) or presenter actions (e.g. questions). The two graphs comprise the post-hoc view of EngageMeter.

and develop the normalized engagement range used during the keynotes. All three keynotes took place in the following three consecutive days between 8 and 10 am in a large lecture hall with more than 300 attendees. Participants were free to choose where they sat and were instructed to start the system at the beginning of the talks. EngageMeter recorded the engagement index of each participant during each keynote. We briefed the three keynote speakers about the study a-priori and introduced the system. We placed a laptop on the side of the podium where each speaker gave her/his talk so that they can easily perceive it from their standing position. The three keynotes presented topics related to HCI and Information Technology and had different durations, ranging from 35-45 minutes. We conducted semi-structured interviews with each keynote speaker after their talks, as well as with the participants (i.e., audience) after the conference ended. We gathered participants' subjective feedback after each keynote. They rated on a 7-point Likert item after each keynote their engagement during the presentation (1=not engaged at all, 7=very engaged). Furthermore, we interviewed participants to gather qualitative feedback on aspects they liked or did not like in the presentation.

RESULTS

Subjective and Measured Engagement

We analyzed participants' subjective Likert scale ratings and the measured normalized engagement for each keynote. The first talk with a subjective engagement score of $Med = 5$, had a median measured engagement of 37%. Keynote 2 scored the highest subjective engagement with $Med = 6$ and median measured engagement of 60%. Finally, the third keynote scored $Med = 4$ and a 40% measured engagement.

We present the measured engagement line graph and slide scores of keynote 2 as it had the highest measured engagement score and received the highest rating by participants (Figure 3). As can be seen, each part of the talk was perceived differently by the audience. The top graph depicts the slide scores and interesting themes are shown. The bottom graph shows the measured normalized engagement. Slide changes and durations are depicted by the vertical sections and interesting points in the presenter's talk are indicated as well.

The presenter asked several questions during her talk – some were rhetorical questions after pausing and two with a call for action in the beginning of her talk (Slides 1 to 3). This has been positively acknowledged by all participants in their comments after the keynote(cf. Figure 3, Slides 1,2,3,9,13). The presenter paused between different parts of the talk and when a technical issue arose (cf. Figure 3, Slides 21-22) she joked and talked whilst solving the issue which can be seen to sustain audience engagement. She presented three videos in the second half of the talk which increased the audience engagement after a phase of history and background information.

Presenters' Real-Time View Feedback

The three presenters differed in their opinions about the utility of the real-time view and how they actually used it in their talks. Presenter 2 said *"I loved it!"*, when we asked her about her feedback on the real-time view. She stated *"every now and then I would look at it and if it was low I would slow down or clarify my words"*. Presenter 1 said that she was entirely immersed in her talk and did not use the real-time view. She said *"I was so in the zone"*. Presenter 3 stated *"I hardly looked at it at all. The large number of audience and the situation made me not want to check it out."*

We asked presenters if they found the gauge or moving graph more useful in real-time and at which points they used each. Presenter 1 said that the real-time feedback could overwhelm presenters especially if the feedback is negative, however, she would use it in trying different aspects while presenting repeatedly with students. Presenter 2 said both the gauge and moving graph are optimal and not overwhelming. She found that two views are the optimal number in real-time feedback. She stated *"because there is this concept of delay, you are communicating* *your talk, your words (are) a little bit ahead, then the audience reacts, and the line graph takes into account this delay. The gauge was useful to see it moving back and forth, if I saw that it was going down red it is an immediate call to action for the speaker to do something"*. Presenter 3 preferred the moving graph and said that it is easier to interpret because with the gauge he needs to think of the previous engagement levels and what has happened to cause the increase/decrease.

Presenter 2 stated that using EngageMeter in real-time is very useful if she is talking to a large audience, especially if there is a language difference where there is a higher chance of losing the audience. Presenter 3 stated that the real-time view can be shared with the audience as well. He said that *"If the audience also sees the graph I am seeing then we could interact and comment about this."*. He suggested using an ambient display that everyone can simultaneously see. Presenters 1 and 3 mentioned that the real-time feedback is more useful in presentations given repeatedly, like courses with the same audience. Presenter 3 said that he would try something different every time and see how that affects the audience, for example, by asking questions or changing the way of presenting a slide.

Presenters' Post-Hoc View Feedback

All three presenters found the post-hoc view useful and informative. They would use it in giving repetitive talks such as in course lectures. Presenter 3 mentioned he would put in anchors into his slides. When he reaches these anchors he would use different presentation elements each time giving the talk (e.g., ask the audience or small (group) exercises) and compare the engagement depending on the used elements post-hoc. Presenter 1 said it will be interesting to compare how the attention spans vary in different contexts, for example, when asking students in a course to not use any external devices at all or asking them to use their devices (e.g., laptop/phone) as they would normally use it during a lecture.

All three presenters mentioned that comparing the post-hoc view over multiple talks will provide useful insights and suggestions. Presenter 2 said that *"It will all start to blend in together after you have given many talks, you can then start to form recommendations about the things that worked well"*. Presenter 1 said that she is interested in correlating how she felt after the given talk to the post-hoc measured engagement of the audience. She suggested writing down how she felt, recording the context such as the audio, and description of the room over a large number of talks. Presenter 3 found slide scores more useful than the line graph because it provides sufficient detail about the slide itself. He stated *"generally I would not be interested in anything beyond the slide scores"*. On the other hand, presenters 1 and 2 both agreed that the post-hoc view can be extended with more context by including not only slide previews and audio, but video of the talks. Presenter 2 also mentioned that she would like to see a summary report showing statistics of the highest and lowest scoring slides.

DISCUSSION AND IMPLICATIONS

We designed EngageMeter to provide presenters with a fine-grained overview of the collective audience engagement. One challenge is interpreting the measured engagement. Our presenters mainly considered moments of low engagement as a call-to-action. However, high engagement levels might not be desirable in every case. A continuously high level of engagement can over-challenge the audience and may affect learning outcomes [21]. As a result, engagement sensing systems can also be designed to notify presenters about excessively high engagement levels. Finding the level of engagement that both fits the audience as well as the presentation poses further challenges which need to be tackled in future work.

Through EngageMeter's post-hoc slide and audio view, presenters stated that they can track the effect of their changes over time. This is not possible using other explicit sensing methods without putting significant effort on the audience. Capturing additional external context using other environment-based sensors, for example, eye-trackers to know where the audience are looking, would further enrich the provided feedback.

Our approach worked in a real-world context. However, in a long-term study, such as over the course of a semester, personal day-to-day mood, timing, and other external aspects can have an effect on the engagement value. In this case, future work could look into developing a dynamic calibration protocol that adapts to individual changes per session. Furthermore, improving the algorithm for removing motion artifacts can further increase the signal quality.

Few implicit audience sensing systems provide real-time feedback from multiple audience members at the same time (e.g., [3]). Most systems provide post-hoc feedback with one-to-one cardinality where in many cases the audience (sender) was also the receiver of the information (i.e., in case of online learning) [8, 15]. Explicit sensing systems provide more support for the presenter. With EngageMeter, we cover additional sender-receiver cardinalities. The system can be used by one or multiple audience members and presenters at the same time, and provide support to presenters in real-time and post-hoc. These opportunities for support became apparent through our interviews with presenters. For example, presenters stated that real-time feedback is useful in contexts such as meetings, rehearsal talks, MOOCs, and live online presentations. They saw a major advantage for talks in front of a foreign audience or where they use a second language. In this case, implicit audience sensing can provide an objective feedback on the perceived complexity and audience engagement during the talk, regardless of cultural differences. In such cases identifying the engagement from other sources, such as mimics, may be difficult. Presenters also expressed their interest in using EngageMeter over a long period of time in talks they give repeatedly, e.g., courses and seminars. Thus, they could enhance content based on feedback from previous presentations.

CONCLUSION

In this work, we reported on EngageMeter, an implicit audience sensing system utilizing EEG signals to provide real-time and post-hoc feedback about audience engagement. EngageMeter is a scalable system and can be used in different contexts including meetings (i.e., work environment), classic as well as flipped classrooms, public speeches, etc. EngageMeter was evaluated in a real-world study during a conference. The evaluation of the system revealed opportunities for using EEG for audience engagement in real world scenarios.

ACKNOWLEDGMENTS

This work was partly conducted within the Amplify project which received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 683008).

REFERENCES

1. Yomna Abdelrahman, Mariam Hassib, Maria Guinea Marquez, Markus Funk, and Albrecht Schmidt. 2015. Implicit Engagement Detection for Interactive Museums Using Brain-Computer Interfaces. In *Proceedings of the 17th International Conference on Human-Computer Interaction with Mobile Devices and Services Adjunct (MobileHCI '15)*. ACM, New York, NY, USA, 838–845. DOI: http://dx.doi.org/10.1145/2786567.2793709

2. Marvin Andujar and Juan E. Gilbert. 2013. Let's Learn!: Enhancing User's Engagement Levels Through Passive Brain-computer Interfaces. In *CHI '13 Extended Abstracts on Human Factors in Computing Systems (CHI EA '13)*. ACM, New York, NY, USA, 703–708. DOI: http://dx.doi.org/10.1145/2468356.2468480

3. Louise Barkhuus and Tobias Jørgensen. 2008. Engaging the Crowd: Studies of Audience-performer Interaction. In *CHI '08 Extended Abstracts on Human Factors in Computing Systems (CHI EA '08)*. ACM, New York, NY, USA, 2925–2930. DOI: http://dx.doi.org/10.1145/1358628.1358785

4. Chris Berka, Daniel J Levendowski, Michelle N Lumicao, Alan Yau, Gene Davis, Vladimir T Zivkovic, Richard E Olmstead, Patrice D Tremoulet, and Patrick L Craven. 2007. EEG correlates of task engagement and mental workload in vigilance, learning, and memory tasks. *Aviation, space, and environmental medicine* 78, Supplement 1 (2007), B231–B244.

5. Laurent George and Anatole Lécuyer. 2010. An overview of research on "passive" brain-computer interfaces for implicit human-computer interaction. In *International Conference on Applied Bionics and Biomechanics ICABB 2010 - Workshop W1 "Brain-Computer Interfacing and Virtual Reality"*. Venise, Italy. https://hal.inria.fr/inria-00537211

6. Alan Gevins, Michael E Smith, Harrison Leong, Linda McEvoy, Susan Whitfield, Robert Du, and Georgia Rush. 1998. Monitoring working memory load during computer-based tasks with EEG pattern recognition methods. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 40, 1 (1998), 79–91.

7. J. Hernandez, Zicheng Liu, G. Hulten, D. DeBarr, K. Krum, and Z. Zhang. 2013. Measuring the engagement level of TV viewers. In *2013 10th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*. 1–7. DOI: http://dx.doi.org/10.1109/FG.2013.6553742

8. Jin Huang, Chun Yu, Yuntao Wang, Yuhang Zhao, Siqi Liu, Chou Mo, Jie Liu, Lie Zhang, and Yuanchun Shi. 2014. FOCUS: Enhancing Children's Engagement in Reading by Using Contextual BCI Training Sessions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 1905–1908. DOI: http://dx.doi.org/10.1145/2556288.2557339

9. Johnny Chung Lee and Desney S. Tan. 2006. Using a Low-cost Electroencephalograph for Task Classification in HCI Research. In *Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology (UIST '06)*. ACM, New York, NY, USA, 81–90. DOI: http://dx.doi.org/10.1145/1166253.1166268

10. Marco Marchesi and Bruno Riccò. 2013. BRAVO: A Brain Virtual Operator for Education Exploiting Brain-computer Interfaces. In *CHI '13 Extended Abstracts on Human Factors in Computing Systems (CHI EA '13)*. ACM, New York, NY, USA, 3091–3094. DOI: http://dx.doi.org/10.1145/2468356.2479618

11. R. Navarathna, P. Lucey, P. Carr, E. Carter, S. Sridharan, and I. Matthews. 2014. Predicting movie ratings from audience behaviors. In *IEEE Winter Conference on Applications of Computer Vision*. 1058–1065. DOI: http://dx.doi.org/10.1109/WACV.2014.6835987

12. Anton Nijholt, Boris Reuderink, and Danny Oude Bos. In *Proceedings of the Third International Conference on Intelligent Technologies for Interactive Entertainment*, Vol. 9.

13. Alan T Pope, Edward H Bogart, and Debbie S Bartolome. 1995. Biocybernetic system evaluates indices of operator engagement in automated task. *Biological psychology* 40, 1 (1995), 187–195.

14. Fernando Silveira, Brian Eriksson, Anmol Sheth, and Adam Sheppard. 2013. Predicting Audience Responses to Movie Content from Electro-dermal Activity Signals. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp '13)*. ACM, New York, NY, USA, 707–716. DOI: http://dx.doi.org/10.1145/2493432.2493508

15. Daniel Szafir and Bilge Mutlu. 2012. Pay Attention!: Designing Adaptive Agents That Monitor and Improve User Engagement. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 11–20. DOI: http://dx.doi.org/10.1145/2207676.2207679

16. Daniel Szafir and Bilge Mutlu. 2013. ARTFul: Adaptive Review Technology for Flipped Learning. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 1001–1010. DOI: http://dx.doi.org/10.1145/2470654.2466128

17. Chi Thanh Vi, Jason Alexander, Pourang Irani, Behrad Babaee, and Sriram Subramanian. 2014. Quantifying EEG Measured Task Engagement for use in Gaming Applications. (nov 2014).

18. Chi Thanh Vi, Kazuki Takashima, Hitomi Yokoyama, Gengdai Liu, Yuichi Itoh, Sriram Subramanian, and Yoshifumi Kitamura. D-FLIP: Dynamic and Flexible Interactive PhotoShow. In *Proceedings of the 10th International Conference on Advances in Computer Entertainment. (ACE'13)*.

19. Chen Wang, Erik N. Geelhoed, Phil P. Stenton, and Pablo Cesar. 2014. Sensing a Live Audience. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 1909–1912. DOI: http://dx.doi.org/10.1145/2556288.2557154

20. Shuo Yan, GangYi Ding, Hongsong Li, Ningxiao Sun, Yufeng Wu, Zheng Guan, Longfei Zhang, and Tianyu Huang. 2016. Enhancing Audience Engagement in Performing Arts Through an Adaptive Virtual Environment with a Brain-Computer Interface. In *Proceedings of the 21st International Conference on Intelligent User Interfaces (IUI '16)*. ACM, New York, NY, USA, 306–316. DOI: http://dx.doi.org/10.1145/2856767.2856768

21. Beste F. Yuksel, Kurt B. Oleson, Lane Harrison, Evan M. Peck, Daniel Afergan, Remco Chang, and Robert JK Jacob. 2016. Learn Piano with BACh: An Adaptive Learning Interface That Adjusts Task Difficulty Based on Brain State. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. ACM, New York, NY, USA, 5372–5384. DOI: http://dx.doi.org/10.1145/2858036.2858388