Behavior Understanding for Arts and Entertainment

ALBERT ALI SALAH, Boğaziçi University
HAYLEY HUNG, Delft University of Technology
OYA ARAN, Idiap Research Institute
HATICE GUNES, Queen Mary University of London
MATTHEW TURK, University of California, Santa Barbara

This editorial introduction complements the shorter introduction to the first part of the two-part special issue on Behavior Understanding for Arts and Entertainment. It offers a more expansive discussion of the use of behavior analysis for interactive systems that involve creativity, either for the producer or the consumer of such a system. We first summarise the two articles that appear in this second part of the special issue. We then discuss general questions and challenges in this domain that were suggested by the entire set of seven articles of the special issue and by the comments of the reviewers of these articles.

CCS Concepts: • **Applied computing** → **Fine arts**; • **Human-centered computing** → *HCI design and evaluation methods; Interaction techniques; Interaction design;*

Additional Key Words and Phrases: Behavior analysis, interactive arts, human-environment interaction, visual arts, affective computing, social and nonverbal behaviors

ACM Reference Format:
Albert Ali Salah, Hayley Hung, Oya Aran, Hatice Gunes, and Matthew Turk. 2015. Behavior understanding for arts and entertainment. ACM Trans. Interact. Intell. Syst. 5, 3, Article 12 (September 2015), 10 pages.
DOI: http://dx.doi.org/10.1145/2817208

1. INTRODUCTION

Automated methods for human behavior understanding can make a major contribution to interactive intelligent systems. The present issue of TiiS includes the second part of a special issue that tackles the challenges and opportunities associated with human behavior understanding in arts and entertainment, which was the focus theme of the fourth edition of the International Workshop on Human Behavior Understanding (HBU) in 2013 [Salah et al. 2013b].

The special issue deals with the problem of automatic analysis of human behavior during interactions that occur in the context of arts and entertainment. For example, an artist may use such a behavior sensing system in real time as a part of the artistic creation process. Interactive technology can analyze the behavior of the artist and provide a tool, for instance, to enhance or to monitor the creation of the artwork. It may analyze the behavior of the audience, in order to let the interactive artwork respond to changes in the audience behavior in a dynamic way. Such technology can also be used in an off-line fashion in order to better understand the interaction dynamics and to design

This work was supported in part by the Scientific and Technological Research Council of Turkey (TUBITAK) under grant number 114E481 and by the NSF award #IIS-1219261.
Authors' addresses: A. A. Salah, Department of Computer Engineering, Boğaziçi University, 34342 Bebek, Istanbul, Turkey; H. Hung, Faculty of Engineering, Mathematics, and Computer Science, Mekelweg 4, Delft, 2628CD, The Netherlands; O. Aran, Idiap Research Institute, Rue Marconi 19, PO Box 592, CH-1920 Martigny, Switzerland; H. Gunes, School of Electronic Engineering and Computer Science, Mile End Road, E1 4NS, London, United Kingdom; M. Turk, Department of Computer Science, University of California, Santa Barbara, CA 93110, USA.
 2160-6455/2015/09-ART12 $15.00
DOI: http://dx.doi.org/10.1145/2817208

improved solutions. On the entertainment side, it is well known how gesture-sensing technology transformed the gaming domain [Schouten et al. 2011]. Novel behavior sensing modalities create new interaction possibilities, as well as new challenges.

The special issue received a wide range of submissions from both aspects of the problem (i.e., for real-time behavior sensing and for off-line analysis). It includes works that describe interactive systems and provide theoretical and technical discussions.

In our brief introduction to the first part of the special issue [Salah et al. 2015], we summarized the five articles in that part. In this introduction to the second part of the special issue, in addition to summarizing the remaining articles, we reflect on the open issues that were raised by the entire set of articles of the special issue, including the discussions initiated by the expert reviewers. We hope that these reflections will help to inform further explorations of this intriguing application domain for interactive intelligent systems.

2. ARTICLES IN THIS SECOND PART OF THE SPECIAL ISSUE

2.1. Interactive Visuals as Metaphors for Dance Movement Qualities

The analysis of human behavior during artistic performance is informative. Fdili Alaoui, Bevilacqua, and Jacquemin's work aims to design and evaluate interactive reflexive visuals for movement qualities during a dance performance. They developed an interactive installation called Double Skin/Double Mind (DS/DM), which is designed to train dancers on the movement basics used by a professional dance company. This system uses silhouette features and extracts higher-level features like verticality, extension, and periodicity to model movement qualities. A number of visual metaphors act as suitable representations for these qualities. In a series of user evaluations, the article evaluates the dancers' experience of the interactive visuals in the context of a dance workshop and training. This approach is promising for the design of new interactive systems for the performing arts.

2.2. Quantitative Study of Music Listening Behavior in a Smartphone Context

Smartphones provide excellent sensing opportunities for behavior analysis during everyday activities. Since they are carried mostly on the owner's body, they liberate the sensing from its dependence on a fixed sensor location. Additionally, some entertainment activities are directly experienced on the smartphone, and these are particularly amenable to real-time modeling and personalization. In their work on music listening behavior, Yang and Teng investigate the relationship between user factors (i.e., demographics, personality traits, musical background, and musical preference), the context of a user's daily life (i.e., the time, activities, mood, social context, and location) and the music listened to. The article considers a set of eight activities, namely waking up, exercising, traveling in a vehicle, walking, reading, working at the office, eating, and going to sleep. Systematic investigation of these factors enables new applications like automatic tagging of music and brings new possibilities to activity classification.

3. OPEN ISSUES

In this section, we discuss some of the open issues and challenges suggested by the articles of the special issue and by the reviewers' responses to the manuscripts. While the articles do address most of these issues, we find it beneficial to formulate them explicitly here.

An important source of problems is the interdisciplinary nature of the domain, and as in most interdisciplinary work, differences in concepts and terminology between collaborating domains become an issue. As an example, the way in which the concept of *emotion* (as opposed to *affect, mood*, or *feeling*) is used by computer scientists in affective computing applications is often criticised by psychologists and emotion theorists, and computer scientists are accused of not making the finer distinctions required by the more traditional meaning of the term [Gunes and Hung 2015; Gunes and Schuller 2013]. This example is discussed further in Section 3.3 below. The remaining issues are discussed under the headings of arts and new technology, critical content and meaningful interactions, analysis of affect, experimenting in creative domains, and multimodality.

3.1. Arts and New Technology

The article by Vezzani, Lombardi, Pieracci, Santinelli, and Cucchiara in the first part of the special issue described the technology for a sensing floor and touched upon its possible applications [Vezzani et al. 2015]. This work raises the age-old question of how we reconcile new technology with new media art forms [Manovich 2001]. How does behavior-sensing technology contribute to a different perspective on art?

The questions of how art should be defined and how it should be interpreted are central questions of art history [Gombrich 1995]. As was noted by Crary, the creation of an artwork is not independent of the observer (i.e., the audience), and human culture has developed (or better still, co-evolved) ways of creating and appreciating the arts over many centuries [Crary 1992]. The *Umwelt* of an artwork defines (to a certain extent) the conditions of its production, preservation, valuation and meaning. New technology enables more physically interactive arts, which is different than the passive audience's interaction with an exhibited or performed work of art. This development results in a refocusing of the relationship between the artwork and the audience.

Through the use of technology, it becomes possible to use cognitive mechanisms of interaction (like mimicry [Castellano et al. 2012], emotional contagion [Samadani et al. 2012], perspective taking [Trafton et al. 2005], and imitation[1]) more extensively, and the interaction afforded by artworks (or games, for that matter) is not restricted to the kinematics of the physical design. In connection with the sense-making process for artworks in general, attention to making sense of the artistic production—which, in the case of interactive art, includes the ways in which the artwork was intended to interact with members of the public—has been quite important [Foreman-Wernet et al. 2014]. Consequently, interactive technology requires a simultaneous re-evaluation of artistic production and the practice of arts appreciation [Höök et al. 2003].

A related question is that of what technology ultimately leads to a shift in art development trends. A clear example of this type of shift occurred when painting technology changed in Europe and artists could move from egg-tempera-based paint to oil-based paint [Mayer 1991]. This move completely revolutionized the practice of painting. By being able to paint on canvas (instead of on plastered surfaces in palaces and churches), the artists could produce artworks for the bourgeoisie, and the entire art market changed. Many old technologies (which are not even seen as technologies today) were new at one point in time. In the context of interactive intelligent systems, can such revolutions still occur? If so, in which directions?

We do not seek to answer this question here, but let us consider the sensing floor and the possibilities of interaction that it affords. The cheap production of sensors makes this technology accessible and affordable. The range of possible forms of sensing through this modality is broad: It is certainly possible to sense the presence of the audience, but potentially also soft biometrics such as the weight of the subject, his or her activity level and age (up to a point; it is possible to tell a child from an adult most of the time), and the personal history of interaction with the sensed space. It takes the

[1]See for instance Daniel Rozin's interactive mirror series at http://www.smoothware.com/danny/, in which the sculptures react in real time to the viewer's movements.

notion of the smart sensing of affective responses of an audience via their spontaneous physiological [Cupchik et al. 2009; Wang et al. 2014] or movement [Martella et al. 2015] behavior beyond existing approaches.

Other technologies, like the CAVE, create different possibilities [Cruz-Neira et al. 1993]. The CAVE enables surround-screen projection-based virtual reality, and it creates a vast visualization space, into which dynamically updated content can be placed. We tend to see possibilities and affordances through each new technology, but once the combinations of technology are also considered, the expressive potential of the artist is vastly increased. This is not to say that new media are limitless; each technology brings its own limits. The CAVE is still a virtual environment that requires specialist equipment to be experienced; and by design, it is an exclusive medium, which makes co-experiencing difficult. Part of the general challenge is to overcome these limitations by integrating technologies that complement each other or make up for each other's limitations.

In addition to critical content (discussed in the next section) and aesthetic value, the spectacle aspect is also boosted through technology. It becomes possible, for instance, to turn a large public artifact like the London Eye into an interactive and social spectacle [Morgan and Gunes 2013]. Behavior analysis can not only incorporate awareness of the reactions from the audience into the experience of the spectacle itself [Salah et al. 2013a]; it can also provide means for behavior change in artists' creative pursuits [Morgan et al. 2015a]. It is by no means a coincidence that many artists explore new ways of expression and interaction, gaining mastery in one or more technology along the way, and collaborative projects (see for instance Numediart[2]) are initiated to bring artists and engineers together. We discuss some issues related to such exploration next.

3.2. Critical Content and Meaningful Interactions

One function of an artwork is to represent an idea, a communication from the artist to the audience, or a trigger to point the audience's attention to a critical issue. Contemporary art theory attaches great value to such critical content of an artwork, apart from its beauty or aesthetic value [Foster et al. 2004]. The physical experience of the artwork plays a role in the memorability, intelligibility, and strength of such a message. Artists can spend years perfecting an artwork, turning a critical idea into a meaningful communication that will touch the audience. Can automated methods help to speed up that process or to improve the quality and strength of experience for the audience?

The work by Grenader, Gasques Rodrigues, Nos, and Weibel [Grenader et al. 2015] on the VideoMob installation in the first part of the special issue touched on the challenges of in-the-wild analysis, ecological validity, generalization, and the importance of the installation-development-design cycle. But it also raised questions about critical content.

The balance between interaction as a gimmick and meaningful, artistic interaction is subtle and very challenging to create. To what extent is the sense-making part of the process important? How does it change in the context of interactive art? Is it a distracting or an important part of any interactive work? To what extent is the perception of control/causality important? Does the observer's attention focus more on the interaction and less on the critical question of the piece? Can the interaction be designed to force the audience to muse about the critical content?

We take a step back to consider these issues. While the Computer Art movement dates back to the 1960s [Reichardt 1969], until very recently a significant bulk of it consisted of explorations into the possibilities of computers and related peripheral technologies. The critical content was not so much in focus, to the extent that eminent art historians

[2] http://www.numediart.org/.

did not hesitate to label it as "consistently and boringly predictable, theoretically shallow, critically naive, technologically positivist, politically and socially disengaged [if not] outright fascistic and spiritualist", and declare "that whenever someone mentions 'new media,' I always want to reach for my gun" [Preziosi 2006]. While Computer Art never really became part of mainstream artistic production, it has grown to an extent that cannot be ignored any more. Computer Art, by its very nature, is interdisciplinary and hybrid, and it is always on the lookout for new technologies and possibilities to add to its production palette [Akdag Salah 2008]. These characteristics make it non-stationary and difficult to analyze with the traditional tools available to art historians. Additionally, they also signal that the exploration part will not simply disappear, as the technology progresses at a faster pace than artists can assimilate. While it is not possible to ignore the clash between expression forms accepted in the mainstream arts and the hybrid forms created through novel technologies, the mainstream is gradually changing, and digital art venues are growing in importance.

Today, an electronic arts portal called DeviantArt[3] has over 35 million registered members and attracts over 65 million unique visitors per month, and over 160,000 original art works are uploaded every day to its collection.[4] This is but one example where computer art is created, promoted and consumed at a rate much higher than that of traditional art forms.

3.3. Analysis of Affect

Computer technology is used not only in the production of artworks but also in their analysis. An example of computer technology applied to more traditional artistic products is the automated analysis of paintings, with regard to both the content of the paintings [Crowley and Zisserman 2014] and the affective responses that they evoke [Sartori et al. 2015].

Affect analysis *in* paintings and art, like affective analysis *of* art, raises a lot of questions and discussion. Reviewers of the article titled "Affective Analysis of Professional and Amateur Abstract Paintings Using Statistical Analysis and Art Theory" [Sartori et al. 2015], who were experts in emotions and aesthetics, argued that artists and curators in general would find the presented approach to be reductionist more than scientific. This view was due in part to the fact that the computational approach presented was not about emotion; instead it had more to do with positive and negative feelings, which are typically used for reducing the valence dimension into the classes "positive" and "negative". Additionally, the processes that create positive and negative affect were not thought to have been adequately considered, though they could potentially be understood through consultation with design teachers. Finally, no emphasis was given to experimental and exploratory psychology that focuses on how objects and experiences are influenced by and have an influence on affect (e.g., curiosity and arousal). Another argument concerned the assumption that beauty or pleasure are *located* in areas of a painting that a computer can discern. The discussion went on to propose the idea of aesthetics being about spatial and even temporal relationships between art and audience and not about specific locations. The final argument put forward was that art is not created with machines in mind.

Although the authors included clarifications about feelings (pleasant/unpleasant) vs. emotions and aesthetics in their work, it was clear that to fully address the challenging issues raised by the reviewers, the study would have had to be redone from scratch, taking into account the issues raised. As a result, the authors explained how relevant studies have approached the perception and study of emotions from a computational

[3]http://www.deviantart.com/.
[4]Usage figures as of July 2015.

perspective and how their approach, as well as similar approaches, can help curators and other domain experts. In summary, there still seem to be gaps in the terminology of feelings vs. emotions, and different communities seem to use the terms differently (computational vs. psychology or arts literature). Additionally, not all aspects of emotions can be measured using the same sensors; for instance, the arousal dimension is known to be better communicated with nonvisual signals such as voice or with physiological signals [Gunes and Schuller 2013; Zeng et al. 2009]. This channel-dependence may be one of the main reasons for people working in the visual arts to focus on the valence dimension.

Art researchers, and perhaps to some extent artists, seem to be concerned that computational systems developed for analyzing art are aiming to replace the critical decision making of curators and artists. This idea evokes further arguments that a computational system cannot be trained to make critical decisions *correctly* and that it is not possible to reduce art to a series of rules and principles because art is not about generalizations. Similar concerns were raised about neuroaesthetics, which attempted to discover basic features that could determine aesthetic experiences, ultimately reducing it to brain activity [Akdag Salah and Salah 2008].

Overall, we (as computer science researchers) are reminded that we need to be more deeply aware of the relevant theories of art and emotion prior to conducting our research, and that ideally we should attempt to include researchers in arts, art theory and art history into our study design and discussions. For the visual arts in particular, we would like to stress that a computational tool, while not useful for an in-depth analysis of a painting in terms of its historical significance, aesthetic value, or critical content, may still be useful for work with large collections of artworks. Online repositories and platforms like deviantArt, Spotify, Flickr, NetFlix, Pinterest, Etsy, and Artsy contain many individual pieces of art, and computational tools are useful for finding and accessing certain paintings and enabling consumers to search artworks (information retrieval), to receive recommendations based on what they have previously liked (recommender systems), and to visualize trends and patterns (cultural analytics) [Manovich 2009].

Much of what we said about affect in the arts domain can be applied to the entertainment domain as well. In the latter domain, affective analysis pertains to visual aesthetics, dramatic structure, temporal structure, and different sensors incorporated into the experience [El-Nasr et al. 2010]. Movies and video games can be viewed as emotional artifacts [Sykes and Brown 2003]. Interaction technology can be used to eliciting emotions but also to analyze them through the immersive experiences created by the entertainment industry. To give a concrete example, Smeaton and Rothwell monitored and recorded physiological reactions of people as they viewed films [Smeaton and Rothwell 2009] and proposed to automatically highlight important segments for indexing and retrieval. In order to perform this experiment in realistic conditions, they used a sensor-equipped and controlled cinema-like environment. This setup exemplifies the challenges that researchers confront when they try to ensure the ecological validity of experiments, which we discuss next.

3.4. Experimenting in Creative Domains

Investigation of preferences in a domain that is highly idiosyncratic, like music listening behavior in the work of Yang and Teng in this part of the special issue, exemplifies the challenging nature of creative domains. Confrontation with the artistic content in this instance (i.e., via smartphones) is entirely personal and customized. In order to map the music listening preferences to the context of the activity and the background of the subject, the music domain, the activity domain, and the subject need to be characterized with sufficient richness. While Yang and Teng make an admirable effort to investigate the mutual influence of these factors, they note that with a limited number of subjects and controlled experimental conditions it is not possible to derive universal conclusions. More generally, to what extent can we generalize the results of user studies conducted in creative domains?

In particular, several reviewers remarked that a monetary reward changes spontaneous behavior when a system is being tested with users. This factor introduces a bias into the experimental setting that is hard to quantify and measure. Especially if the monetary reward is not fixed but rather tied to the amount or quality of system use, the subjects' behavior can be expected to change. Also, being monitored alters subjects' perception and awareness in what is supposed to be a personal experience of art appreciation. Finally, novelty is an important factor in interactions with interactive intelligent systems. Engagement often decreases as the novelty of the system wears off [Ros and Demiris 2013], but this decrease may not be visible during the time window of the study.

When *experimenting* in creative domains, the size and the public aspect of the installation have a direct impact on how much *experimentation* one can really perform. Working with a large public artifact like the London Eye and turning it into an interactive and social spectacle brought these issues into play [Morgan and Gunes 2013]. In order to let the artistic design be guided by the actual installation and intended experience for the audience, various compromises regarding research considerations had to be made (e.g., questionnaires or interviews could not be conducted to obtain the so-called ground truth or labels). These compromises have direct impact on how the recorded data could be used, how much automated analysis could be done, and what types of research questions could be reliably answered.

Computer science and affective computing researchers are accustomed to asking questions about evaluation, reliability, and accuracy. When it comes to the use of large public artistic installations, such questions cannot always be answered in the usual ways. The assessment of research at this scale, involving systems used by the public outside of lab conditions, should be performed with this point in mind.

3.5. Multimodality

One of the research questions that is being investigated over and over in the area of behavior analysis is how multimodal data fusion improves automatic recognition. Experiments reported in various application domains point to different results and conclusions. On the one hand, multimodal data are reported to help automatic recognition by compensating for missing information and improving recognition accuracy via complementarity. On the other hand, multimodal data do not always provide the best results, or they may ultimately not be needed to achieve a certain level of accuracy.

Mood detection in videos is one such example. Although the authors of "In the Mood for Vlog: Multimodal Inference in Conversational Social Video" report that multimodal features perform better than single channel features, they also point to the fact that all available channels are not always needed to discriminate mood in videos accurately [Sanchez-Cortes et al. 2015]. Similar insights have been obtained by other researchers who investigated collaborative music making using affective and behavioral sensors [Morgan et al. 2015b]; for example, not all signals were relevant and beneficial for automatically analysing and predicting creativity.

In the work of Baur, Mehlmann, Damian, Gebhard, Lingenfelser, Wagner, Lugrin, and André, a multimodal system was proposed to analyze and facilitate the interpretation of social signals automatically in a bidirectional interaction with a conversational agent [Baur et al. 2015]. The authors note that current sensor technology does not allow the detection of subtle behavioral cues in realistic conditions. Often, fully automatic analysis is difficult, and manual annotation or alignment of at least some signals is required for reliable and robust analysis of the results. For event-based annotation, knowing when an event is really an event and not just a spurious movement is an important consideration. There are statistical techniques used to improve such annotations [Raykar et al. 2010]. Nonetheless, multimodality also makes annotation much more difficult, so the use of multiple channels requires justification as well as proper discussion of its limitations.

4. CONCLUDING REMARKS

This two-part special issue aims to encourage research about the challenges and opportunities associated with human behavior understanding in arts and entertainment. The concrete application examples that we discuss are interactive art installations that sense and respond to their viewers in novel ways, systems that analyze user behavior during creative and entertaining activities for contextual modeling and improved services, novel sensory modalities that allow the creation of sensing spaces, and new affordances for creative applications.

We highlight several issues of behavior analysis in arts and entertainment, without prioritizing one issue over the other. Each issue we discuss comes up in one or more articles of the special issue and receives some discussion therein, which we have aimed to extend here. We have also asked a number of questions that we are not yet able to answer fully, in the hope that these questions will serve as useful discussion points for the community.

REFERENCES

Almila Akdag Salah. 2008. *Discontents of Computer Art: A Discourse Analysis on the Intersection of Arts, Sciences and Technology*. Unpublished doctoral dissertation, UCLA.

Almila Akdag Salah and Albert Ali Salah. 2008. Technoscience Art: A Bridge Between Neuroesthetics and Art History? *Review of General Psychology* 12, 2 (2008), 147.

Tobias Baur, Gregor Mehlmann, Ionut Damian, Patrick Gebhard, Florian Lingenfelser, Johannes Wagner, Birgit Lugrin, and Elisabeth André. 2015. Context-Aware Automated Analysis and Annotation of Social Human-Agent Interactions. *ACM Transactions on Interactive Intelligent Systems* 5, 2, Article 11 (June 2015), 33 pages. DOI:http://dx.doi.org/10.1145/2764921

Ginevra Castellano, Matteo Mancini, Christopher Peters, and Peter W. McOwan. 2012. Expressive Copying Behavior for Social Agents: A Perceptual Analysis. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans* 42, 3 (2012), 776–783.

Jonathan Crary. 1992. *Techniques of the Observer: On Vision and Modernity in the Nineteenth Century*. MIT press, Cambridge, MA.

Elliot J. Crowley and Andrew Zisserman. 2014. In Search of Art. In *Computer Vision—ECCV 2014 Workshops*. Springer, 54–70.

Carolina Cruz-Neira, Daniel J. Sandin, and Thomas A. DeFanti. 1993. Surround-Screen Projection-Based Virtual Reality: The Design and Implementation of the CAVE. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques*. ACM, 135–142.

Gerald C. Cupchik, Oshin Vartanian, Adrian Crawley, and David J. Mikulis. 2009. Viewing Artworks: Contributions of Cognitive Control and Perceptual Facilitation to Aesthetic Experience. *Brain and Cognition* 70, 1 (2009), 84–91.

Magy Seif El-Nasr, J. Morie, and Anders Drachen. 2010. A Scientific Look at the Design of Aesthetically and Emotionally Engaging Interactive Entertainment Experiences. In *Affective Computing and Interaction: Psychological, Cognitive and Neuroscientific Perspectives, Information Science Publishing*. IGI Global, 281–307.

Lois Foreman-Wernet, Brenda Dervin, and Clayton Funk. 2014. Standing in Two Worlds Looking at an Art Exhibition: Sense-Making in the Millennial Generation. *The Journal of Arts Management, Law, and Society* 44, 2 (2014), 101–117. DOI:http://dx.doi.org/10.1080/10632921.2014.905813

Hal Foster, Rosalind Krauss, Yve-Alain Bois, Benjamin H. D. Buchloh, and David Joselit. 2004. *Art Since 1900: Modernism, Antimodernism, Postmodernism*. Thames & Hudson, London.

Ernst Hans Gombrich. 1995. *The Story of Art*. Phaidon, London.

Emily Grenader, Danilo Gasques Rodrigues, Fernando Nos, and Nadir Weibel. 2015. The VideoMob Interactive Art Installation: Connecting Strangers Through Inclusive Digital Crowds. *ACM Transactions on Interactive Intelligent Systems* 5, 2, Article 7 (July 2015), 31 pages. DOI:http://dx.doi.org/10.1145/2768208

Hatice Gunes and Hayley Hung. 2015. Emotional and Social Signals: A Neglected Frontier in Multimedia Computing? *IEEE MultiMedia* 22, 2 (2015), 76–85.

Hatice Gunes and Björn Schuller. 2013. Categorical and Dimensional Affect Analysis in Continuous Input: Current Trends and Future Directions. *Image and Vision Computing* 31, 2 (2013), 120–136.

Kristina Höök, Phoebe Sengers, and Gerd Andersson. 2003. Sense and Sensibility: Evaluation and Interactive Art. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'03)*. ACM, New York, NY, USA, 241–248. DOI:http://dx.doi.org/10.1145/642611.642654

Lev Manovich. 2001. *The Language of New Media*. MIT press, Cambridge, MA.

Lev Manovich. 2009. Cultural Analytics: Visualising Cultural Patterns in the Era of "More Media". *Domus* (2009).

Claudio Martella, Ekin Gedik, Laura Cabrera Quiros, Gwenn Englebienne, and Hayley Hung. 2015. How Was It? Exploiting Smartphone Sensing to Measure Implicit Audience Responses to Live Performances. In *Proceedings of ACM Multimedia*. In press.

Ralph Mayer. 1991. *Artist's Handbook of Materials and Techniques: Revised and Updated*. Viking Books, New York.

Evan Morgan and Hatice Gunes. 2013. Human Nonverbal Behaviour Understanding in the Wild for New Media Art. In *Human Behavior Understanding: Proceedings of the Fourth International Workshop*, Albert Ali Salah, Hayley Hung, Oya Aran, and Hatice Gunes (Eds.). Springer, Heidelberg, 27–39.

Evan Morgan, Hatice Gunes, and Nick Bryan-Kinns. 2015a. The LuminUs: Providing Musicians with Feedback on the Gaze and Body Motion of their Co-Performers. In *Proceedings of the International Conference on Human-Computer Interaction*.

Evan Morgan, Hatice Gunes, and Nick Bryan-Kinns. 2015b. Using Affective and Behavioural Sensors to Explore Aspects of Collaborative Music Making. *International Journal of Human-Computer Studies* 82 (2015), 31–47.

Donald Preziosi. 2006. Afterword: Artifice and Interactivity. Art in the Age of Technological Seduction. *media-N* 2, 3 (2006), 73–77.

Vikas C. Raykar, Shipeng Yu, Linda H. Zhao, Gerardo Hermosillo Valadez, Charles Florin, Luca Bogoni, and Linda Moy. 2010. Learning From Crowds. *The Journal of Machine Learning Research* 11 (2010), 1297–1322.

Jasia Reichardt. 1969. *Cybernetic Serendipity: The Computer and the Arts*. Praeger, Westport, CO.

Raquel Ros and Yiannis Demiris. 2013. Creative Dance: An Approach for Social Interaction Between Robots and Children. In *Human Behavior Understanding: Proceedings of the Fourth International Workshop*, Albert Ali Salah, Hayley Hung, Oya Aran, and Hatice Gunes (Eds.). Springer, Heidelberg, 40–51.

Albert Ali Salah, Hayley Hung, Oya Aran, and Hatice Gunes. 2013a. Creative Applications of Human Behavior Understanding. In *Human Behavior Understanding: Proceedings of the Fourth International Workshop*, Albert Ali Salah, Hayley Hung, Oya Aran, and Hatice Gunes (Eds.). Springer, Heidelberg, 1–14.

Albert Ali Salah, Hayley Hung, Oya Aran, and Hatice Gunes. 2013b. *Human Behavior Understanding: Proceedings of the Fourth International Workshop*. Springer, Heidelberg.

Albert Ali Salah, Hayley Hung, Oya Aran, Hatice Gunes, and Matthew Turk. 2015. Brief Introduction to the Special Issue on Behavior Understanding for Arts and Entertainment. *ACM Transactions on Interactive Intelligent Systems* 5, 2, Article 6 (July 2015), 3 pages. DOI:http://dx.doi.org/10.1145/2786762

Ali-Akbar Samadani, Rob Gorbet, and Dana Kulić. 2012. Gender Differences in the Perception of Affective Movements. In *Human Behavior Understanding: Proceedings of the Third International Workshop*, Albert Ali Salah, Javier Ruiz-del Solar, Cetin Mericli, and Pierre-Yves Oudeyer (Eds.). Springer, Heidelberg, 65–76.

Dairazalia Sanchez-Cortes, Shiro Kumano, Kazuhiro Otsuka, and Daniel Gatica-Perez. 2015. In the Mood for Vlog: Multimodal Inference in Conversational Social Video. *ACM Transactions on Interactive Intelligent Systems* 5, 2, Article 9 (June 2015), 24 pages. DOI:http://dx.doi.org/10.1145/2641577

Andreza Sartori, Victoria Yanulevskaya, Almila Akdag Salah, Jasper Uijlings, Elia Bruni, and Nicu Sebe. 2015. Affective Analysis of Professional and Amateur Abstract Paintings Using Statistical Analysis and Art Theory. *ACM Transactions on Interactive Intelligent Systems* 5, 2, Article 8 (June 2015), 27 pages. DOI:http://dx.doi.org/10.1145/2768209

Ben A. M. Schouten, Rob Tieben, Antoine van de Ven, and David W. Schouten. 2011. Human Behavior Analysis in Ambient Gaming and Playful Interaction. In *Computer Analysis of Human Behavior*, Albert Ali Salah and Theo Gevers (Eds.). Springer, Heidelberg, 387–403.

Alan F. Smeaton and Sandra Rothwell. 2009. Biometric Responses to Music-Rich Segments in Films: The CDVPlex. In *Seventh International Workshop on Content-Based Multimedia Indexing*. IEEE, 162–168.

Jonathan Sykes and Simon Brown. 2003. Affective Gaming: Measuring Emotion Through the Gamepad. In *Extended Abstracts of the 2003 Conference on Human Factors in Computing Systems*. ACM, 732–733.

J. Gregory Trafton, Nicholas L. Cassimatis, Magdalena D. Bugajska, Derek P. Brock, Farilee E. Mintz, and Alan C. Schultz. 2005. Enabling Effective Human-Robot Interaction Using Perspective-Taking in Robots. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans* 35, 4 (2005), 460–470.

Roberto Vezzani, Martino Lombardi, Augusto Pieracci, Paolo Santinelli, and Rita Cucchiara. 2015. A General-Purpose Sensing Floor Architecture for Human-Environment Interaction. *ACM Transactions on Interactive Intelligent Systems* 5, 2, Article 10 (June 2015), 26 pages. DOI:http://dx.doi.org/10.1145/2751566

Chen Wang, Erik N. Geelhoed, Phil P. Stenton, and Pablo Cesar. 2014. Sensing a Live Audience. In *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 1909–1912.

Zhihong Zeng, Maja Pantic, Glenn Roisman, Thomas S. Huang, and others. 2009. A Survey of Affect Recognition Methods: Audio, Visual, and Spontaneous Expressions. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 31, 1 (2009), 39–58.

Received August 2015; accepted August 2015