UNIVERSITÉ DE GENÈVE

Archive ouverte UNIGE
https://archive-ouverte.unige.ch

Article scientifique | Editorial | 2017

Accepted version | Open Access

This is an author manuscript post-peer-reviewing (accepted version) of the original publication. The layout of the published version may differ .

Toward Commercial Applications of Affective Computing

Novak, Domen; Chanel, Guillaume; Guillotel, Philippe; Koenig, Alexander

How to cite

NOVAK, Domen et al. Toward Commercial Applications of Affective Computing. In: IEEE transactions on affective computing, 2017, vol. 8, n° 2, p. 145–147. doi: 10.1109/TAFFC.2017.2676318

This publication URL: https://archive-ouverte.unige.ch//unige:94787
Publication DOI: 10.1109/TAFFC.2017.2676318

© This document is protected by copyright. Please refer to copyright holder(s) for terms of use.

Applied affective computing

Affective computing has come a long way from its humble beginnings. Starting with a seminal paper by Rosalind Picard [1], it has evolved into an interdisciplinary field with its own IEEE Transactions and numerous special issues in other journals. One of the main goals of affective computing is to create machines which can adapt to users' emotions in order to produce more natural and efficient interaction [2]. Emotion recognition is thus a central component of the field, and is based on a variety of measurements (facial expressions, speech, gait patterns, physiology, eye tracking…) that are analyzed using advanced pattern recognition techniques [3], [4]. Furthermore, researchers and entrepreneurs have identified countless possible applications of affect-aware technology, from health and driver monitoring [5] to exercise [6] and computer game adaptation [7].

However, although great scientific advances have been made and many applications have been proposed, only few robust implementations have been presented or validated, and commercial adoption of affect-aware technology has been marginal. One rare example of a successful application is the smile detector used in digital cameras to automatically take pictures when the subject is smiling. This weak adoption of the technology can be attributed to several unsolved challenges in the domain of affective computing.

While the accuracy and robustness of affect recognition has made great advances, it is still not clear how to respond to recognized affective states (e.g. how to provide feedback in a vehicle if the driver is stressed or tired). State-of-the-art machines that respond to recognized states generally do so using very simple if-then rules, and little work has been done on how a machine can effectively modify a user's affective state or how it can show empathy toward the user given many possible actions, user-specific preferences, and potentially unreliable affect recognition. Furthermore, it is often unclear whether the expected user state has been achieved and whether it could have been reached more efficiently. Finally, very few affective computing methods are suitable for the noisy, unpredictable situations that occur outside the laboratory. Without addressing these issues, it is nearly impossible to efficiently close the affective 'loop' and show a measurable, commercially attractive benefit of affective computing.

Contribution

For this special section, our goal was to present studies that go beyond affect recognition in laboratory environments. Thirteen submissions were received and were subject to two rounds of reviews by three experts in the field. Of these thirteen, five passed the first round and were sent back to the authors for revisions based on the reviewers' suggestions. In the second round of review, four articles were finally accepted for publication. Thus, the articles of this special section were selected using a thorough peer-review process and represent the best ideas and methods that will move the field of affective computing toward commercial applications.

The development of methods for affective computing often starts with the collection of human affective expressions. McDuff and el Kaliouby present a large dataset of over two million facial responses collected from 500,000 over a four-year period in response to everyday media content (product ads, political ads, speeches…) using webcams. They present robust ways of automatically labeling the data as well as reliably interpreting it via the facial action coding system. As an example of the method's efficacy, they show how facial responses differ between successful and unsuccessful commercial ads. Furthermore, they discuss potential applications in marketing: ad copy-testing, political polling, move trailer testing, and several other emerging applications. Particular attention is paid to privacy issues in affective computing.

Yüce et al. focus on another very promising application of affective computing: monitoring attention and cognitive load levels in drivers. In addition to a cognitive load database, they present a completely automatic system that can be integrated in a vehicle in order to detect the level of cognitive load via nonintrusive face monitoring. Action units are identified from the drivers' faces and a new type of features is proposed by computing the synchronization among action units. The idea is that synchronized face movements can be relevant indicators of emotion or distraction. Both individualized and subject-independent approaches are tested, showing the efficiency of the proposed approach.

Finally, recommendations are given for how the system could warn drivers when distraction is detected.

Similarly, Zhang et al. also use affective computing to monitor cognitive load in drivers, but for a different purpose: teaching driving skills to people with autism spectrum disorder. Twenty adolescents with autism spectrum disorder are recruited and perform driving assignments of different difficulties over six 1-hour sessions while multiple physiological measurements are taken. To increase environmental validity, participants are allowed to move frequently, resulting in a noisy, realistic dataset. Three information fusion techniques are then investigated to determine how affect recognition can most effectively adapt to each participant's characteristics in order to classify cognitive load. Feature-level fusion is shown to outperform other classification schemes. Again, recommendations are given for how the system could be used in practice.

It is commonly acknowledged that emotions are expressed through many channels including facial expressions, the voice and physiology. It is thus fundamental to develop frameworks that can facilitate the integration and synchronization of this multimodal information as well as place the information in an appropriate context. To address this challenge, Stratou and Morency present a framework that could be used to combine affective and contextual cues in order to infer psychological states in a wide variety of applications. The framework has a modular structure, allowing different measurement modalities, different feature extraction methods and different classifiers to be used. The proposed framework is applied to a specific use case which aims at identifying distress during interviews with people suffering from post-traumatic stress disorder. A virtual human then acts on the identified distress information by changing its behavior.

Discussion and perspectives

The four presented papers address very promising applications of affective computing: driver monitoring, media analysis, and biomedical technology. Going beyond well-established paradigms, they provide methods of individualizing the technology to each user and providing feedback in response to the inferred affective information. As the papers include recommendations for use in practice, we believe that these technologies are now ready to be used in long-term studies that will evaluate their practical benefits – their ability to help people and to deliver a commercially attractive advantage that justifies development costs. We hope that other developers will follow the example set by the papers in this special section, providing more robust technology and giving recommendations for its practical use. In this way, affective computing will continue to mature and become an increasingly important part of our everyday lives.

With this special section, we aim to encourage research in new directions that will allow the development of more applied technologies in affective computing and bridge the infamous "valley of death" that many ideas experience on the road toward commercialization. A first direction is user experience. Affective computing was developed to improve user experience by including emotions in the human-computer interaction loop. However, most studies so far have focused on the recognition of emotions using performance measures such as accuracy without considering user experience. However, depending on the adaptation mechanism, even a system with high affect recognition accuracy could impact user experience negatively (or vice versa), and significant increases in recognition accuracy may not result in significant changes to user experience [8]. It is thus critical to better investigate how machine behavior can be best adapted and what affect recognition accuracy is needed for positive user experience.

Another important research direction is the inclusion of user characteristics and context in the affect recognition and computer behavior adaptation. User characteristics such as age or personality have long been acknowledged as important for affective computing [9], but relatively few affect-aware systems take them into account. One promising example are robotic coaches that adapt their verbal encouragement to the personality of the user [10]. Similarly, context remains underexplored in affective computing: the same affective state may, for example, evoke completely different physiological responses depending on its context, but most affect-aware systems are built for only a single situation and are unable to handle significant contextual changes. If affect-aware technology is to achieve widespread adoption, it must instead be able to function for a wide variety of situations and user types.

Furthermore, user acceptance depends not only on whether the system can recognize and respond to affective states, but also on other factors related to quality of experience: for example, on whether a system is invasive and whether user privacy is guaranteed. State-of-the-art sensors range from cameras to helmets, but none of them are ready for commercial deployment – either because users may not want to be watched in their environment, because wearable devices are not unobtrusively integrated into everyday life, or because users do not know how the measured data are used. It is thus critical to conduct more research on how affective systems can be designed in order to maximize user acceptance. Such research will likely require a 'user benefit' metric that would weigh the disadvantages and disadvantages of affective systems, allowing designers to better compare different designs. To measure this user benefit, it may even be necessary to study the bidirectional affective relationship that may develop between the user and the affect-aware device [11]. For example, particularly elderly people may build an emotional bond with a robot that can measure and reflect their emotions. Similarly, people who interact with computer-controlled characters in virtual reality may develop a bond with these characters if they perceive them as capable of recognizing and expressing emotions.

Finally, from an industrial perspective, the final obstacle on the road to widespread adoption will be a focus on real-world experiments. Successful commercialization will only be possible once sensors and decision-making systems are able to function under rough, dynamic conditions such as rapidly changing light levels and a broad variety of acoustic disturbances in the background. Furthermore, it is only then that we can expect affective computing to improve lives (e.g. workload reduction) or even save them (e.g. driver monitoring).

To conclude this editorial, we would like to thank Prof. Björn Schuller, editor-in-chief IEEE Transactions on Affective Computing, and Ms. Samantha Jacobs, IEEE Administrator, for their support and encouragement during the development of this special section. We would also like to thank the anonymous reviewers who provided constructive feedback on the submitted papers.

Domen Novak, PhD
University of Wyoming

Guillaume Chanel, PhD
Swiss Center for Affective Sciences, University of Geneva

Philippe Guillotel, PhD
Technicolor Research & Innovation

Alexander Koenig, PhD
Reactive Robotics GmbH, formerly with BMW Research & Technology GmbH

References

[1] R. W. Picard, E. Vyzas, and J. Healey, "Toward machine emotional intelligence: Analysis of affective physiological state," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 10, pp. 1175–1191, 2001.

[2] R. Calvo, S. D'Mello, J. Gratch, A. Kappas, R. Calvo, S. D'Mello, J. Gratch, and A. Kappas, "Introduction to Affective Computing," in *The Oxford Handbook of Affective Computing*, Oxford University Press, 2015.

[3] Z. Zeng, M. Pantic, G. I. Roisman, and T. S. Huang, "A survey of affect recognition methods: audio, visual, and spontaneous expressions," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 38, no. 1, pp. 39–58, Jan. 2009.

[4] C. Mühl, B. Allison, A. Nijholt, and G. Chanel, "A survey of affective brain computer interfaces: principles, state-of-the-art, and challenges," *Brain-Computer Interfaces*, vol. 1, no. 2, pp. 66–84, May 2014.

[5] J. A. Healey and R. W. Picard, "Detecting stress during real-world driving tasks using physiological sensors," *IEEE Trans. Intell. Transp. Syst.*, vol. 6, no. 2, pp. 156–166, Jun. 2005.

[6] D. Novak, M. Mihelj, J. Ziherl, A. Olenšek, and M. Munih, "Psychophysiological measurements

in a biocooperative feedback loop for upper extremity rehabilitation," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 19, no. 4, pp. 400–410, 2011.
[7] C. Liu, P. Agrawal, N. Sarkar, and S. Chen, "Dynamic difficulty adjustment in computer games through real-time anxiety-based affective feedback," *Int. J. Hum. Comput. Interact.*, vol. 25, no. 6, pp. 506–529, Aug. 2009.
[8] S. M. McCrea, G. Geršak, and D. Novak, "Absolute and relative user perception of classification accuracy in an affective videogame," *Interact. Comput.*, p. in press, 2017.
[9] C. Conati, "Probabilistic assessment of user's emotions in educational games," *Appl. Artif. Intell.*, vol. 16, pp. 555–575, 2002.
[10] J. Fasola and M. J. Matarić, "Using socially assistive human-robot interaction to motivate physical exercise for older adults," *Proc. IEEE*, vol. 100, no. 8, pp. 2512–2526, 2012.
[11] S. H. Fairclough, "Fundamentals of physiological computing," *Interact. Comput.*, vol. 21, no. 1–2, pp. 133–145, Jan. 2009.