

Affective Computing

By **BJÖRN W. SCHULLER**^{ID}, *Fellow IEEE*
Guest Editor

MATTI PIETIKÄINEN^{ID}, *Life Fellow IEEE*
Guest Editor

Affective computing's history goes back over 25 years, with its roots, such as in facial expression recognition or emotional speech analysis and synthesis, going back even earlier in time. As a field, it has been one of the earliest adopters of deep learning

and is in its nature massively multimodal, spanning from audio over text, images, video, physiological, to haptical and further data. At the same time, it spans from analysis to synthesis and is massively based on signal processing and machine learning. Affective computing is also interdisciplinary as are a few other fields, from computer science and engineering to psychology, philosophy, medicine, and many other disciplines such as linguistics, musicology, or arts. In fact, it concerns all of us, as it is all about emotions, and

emotions are all about life, arts, and being. Accordingly, this special issue in *PROCEEDINGS OF THE IEEE* will likely be of interest to any reader and beyond. In particular, the field lacks an all-encompassing high-profile snapshot that provides a perfect entry-point provided by leaders in the field. This is what this special issue aims to provide: A holistic tutorial and an overview of affective computing, from analysis through synthesis to applications, aimed at new readers interested in this field as well as those that are experts in the field.

Note also that affective computing has never been as “big” as it is now. The journal of the field, *IEEE TRANSACTIONS ON AFFECTIVE COMPUTING*, co-founded by the lead guest editor of this issue (Björn Schuller), has risen in Impact Factor from “one-plus” to “ten-plus” in its ten plus years of existence. Companies are popping up more than ever, and the established ones have become serious players, such as Affectiva, audEERING, and RealEyes, to name but three. Major enterprises such as Amazon, Apple, Baidu, Facebook, Google, Huawei, Microsoft, and Samsung—again to name but a few—are hiring and investing heavily in the field at the very moment. We are currently witnessing a massive interest to go large in the next months or so with applications known to the broad public, when the next generations of Alexa, Siri, or video games,

The articles in this special issue cover four major subfields in affective computing, namely affect analysis, affect synthesis, applications, and ethics.

and simply anything computing with a sophisticated user interface will sense and react to our emotions. Be excited about this future! We are happy to have brought together the leading experts in the field who contributed to this first-of-its-kind special issue (in terms of completeness and timeliness), and we are certain this issue will be an utmost valuable resource in the field.

The articles of this issue cover different aspects of affective computing, including speech and language, facial expressions, images, physiological signals, brain-computer interfaces, touch, affect conversion, social functions of machine expressions, affective game computing and learning, and ethical considerations.

In summary, the articles cover five major subfields in affective computing, namely, affect analysis, affect synthesis, applications in affective computing, and ethics, presented in this order. The detailed review highlights are as follows.

I. AFFECT ANALYSIS

The first article in the issue [A1] starts the affect analysis part of this special issue. It provides an overview of the state-of-the-art in speech emotion recognition. The authors thereby focus on how to make the technology available to everyone. This is reflected in the

context of responsible and human-centered applications. The authors stress the importance thereof as they believe the lack of such a focus is so far limiting the wide adoption of speech emotion recognition in the real world.

The second article [A2] is a contribution that focuses on the content side of language: dealing with sentiment analysis. The authors highlight the huge requirement of data as a main limiting factor to automatic sentiment analysis. Therefore, in this tutorial, they introduce a label-efficient view on sentiment, opinion, and emotion recognition from text—in particular in view of deep learning as the paradigm of the hour.

The third contribution [A3] provides facial expression analysis based on facial action units. The authors outline three important key assets of a system able to recognize these: the ability to capture local action unit-centered features, modeling of co-occurrence of certain action units, and exploiting a change in the appearance of these from frame to frame. They discuss challenges in each of these aspects and introduce an end-to-end approach to the problem.

Continuing with facial expression, the fourth article [A4] investigates the issue on the “micro”-level, focusing on “involuntary, fleeting, and subtle facial expression” as appearing, for example, in case of emotion regulation and hiding. The authors provide a detailed survey on the topic, including a psychological and computer science perspective and the state-of-play in this exciting and emotion-telling subfield.

While the lion’s share of visual affect analysis focuses on the literature on the human and in particular its face, the survey in [A5] broadens the scope of general visual media. From a method point of view, the authors not only focus on deep learning but also include the psychological view alongside the artistic view in addition to the engineering one. This sets this article apart from the other papers, as it also deals with the affect in general media not necessarily related to a human.

With affect not only impacting the behavior, but also physiological responses, the next tutorial [A6] introduces this side of affective computing. The particular advantage according to the authors is the reliability, as conscious manipulation of these signals is close to impossible over longer time stretches. The article deals with practical aspects of assessment in everyday life and the challenges coming with this setting, including noisy data and self-reported learning targets.

Sensing a particular type of physiological signal, brain–computer interfaces (BCIs) measure activity in the central nervous system allowing it to analyze but also impact the emotional state. In this light, the tutorial in [A7] guides the reader into the basics of BCIs leading into “closed-loop” affective BCIs that include the chain from analysis to brain stimulation. The authors further discuss applications, data, algorithms, and privacy aspects.

The final survey in this section [A8] deals with touch in human interaction as a marker of emotion. The authors look into the epistemology, main findings, and ever-present challenges in this subfield’s literature. They include both touch in human–computer interaction, such as with robots or virtual humans, and computer-mediated human–human communication. Databases and algorithms are further presented before concluding on ethical aspects.

II. AFFECT SYNTHESIS

The second part of this special issue is dedicated to the synthesis of affect. The first contribution [A9] covers the modality of speech. In particular, the synthesis and conversion of speech targeted to an emotion are described. Deep learning is the focus from a method’s point of view, and an overview of these approaches is provided. The authors further explain the current approaches in a nutshell and do not leave out the latest trends.

In addition to speech, visual and motion behavior generation is crucial in virtual humans and social robots. The authors in [A10] first stress the

relevance of the ability to incorporate this aspect into human–machine interaction. They further link emotional expression and social function from a psychological perspective. Then, approaches toward the automatic generation of according expressions in facial expressions, language, and touch alongside the voice are presented. Again, ethical aspects are named.

III. APPLICATIONS IN AFFECTIVE COMPUTING

Two examples illustrate application scenarios in the third part of this special issue. The first application [A11] deals with engagement, which the authors describe as hard to measure, but also to sustain. The authors provide a multicomponential definition of engagement, review its conceptualization, and introduce its detection within context. Thereby, they feature “what works” and where problems persist. In conclusion, the authors discuss pro-proactive and reactive ways of boosting engagement in learning applications.

Discussing entertainment as an application use case, the authors of [A12] survey affective computing in (electronic) gaming. Alongside the state-of-play, principles, approaches, and tools, the authors discuss the affective loop: game affect elicitation, sensing, detection, and adaptation. They further define taxonomies in this context and highlight data and its collection in gaming. As in other chapters, limitations and perspectives are provided for the reader.

IV. ETHICS

This special issue is concluded with a crucial summary of ethical considerations on affective computing. While these have been given repeatedly throughout this special issue at the end of the articles, in [A13], one finds a deeper and bundled insight. Dangers include oversimplification of affective states, lack of grounding in reality, potential addiction to affective systems and implicit dependence on those controlling the systems, and potential manipulation among others. ■

APPENDIX: RELATED ARTICLES

- [A1] C.-C. Lee, T. Chaspari, E. M. Provost, and S. S. Narayanan, "An engineering view on emotions and speech: From analysis and predictive models to responsible human-centered applications," *Proc. IEEE*, vol. 111, no. 10, pp. 1142–1158, Oct. 2023, doi: [10.1109/JPROC.2023.3276209](https://doi.org/10.1109/JPROC.2023.3276209).
- [A2] S. Zhao, X. Hong, J. Yang, Y. Zhao, and G. Ding, "Toward label-efficient emotion and sentiment analysis," *Proc. IEEE*, vol. 111, no. 10, pp. 1159–1197, Oct. 2023, doi: [10.1109/JPROC.2023.3309299](https://doi.org/10.1109/JPROC.2023.3309299).
- [A3] J. Yang, Y. Hristov, J. Shen, Y. Lin, and M. Pantic, "Toward robust facial action units' detection," *Proc. IEEE*, vol. 111, no. 10, pp. 1198–1214, Oct. 2023, doi: [10.1109/JPROC.2023.3257542](https://doi.org/10.1109/JPROC.2023.3257542).
- [A4] G. Zhao, X. Li, Y. Li, and M. Pietikäinen, "Facial micro-expressions: An overview," *Proc. IEEE*, vol. 111, no. 10, pp. 1215–1235, Oct. 2023, doi: [10.1109/JPROC.2023.3275192](https://doi.org/10.1109/JPROC.2023.3275192).
- [A5] J. Z. Wang et al., "Unlocking the emotional world of visual media: An overview of the science, research, and impact of understanding emotion," *Proc. IEEE*, vol. 111, no. 10, pp. 1236–1286, Oct. 2023, doi: [10.1109/JPROC.2023.3273517](https://doi.org/10.1109/JPROC.2023.3273517).
- [A6] Y. S. Can, B. Mahesh, and E. André, "Approaches, applications, and challenges in physiological emotion recognition—A tutorial overview," *Proc. IEEE*, vol. 111, no. 10, pp. 1287–1313, Oct. 2023, doi: [10.1109/JPROC.2023.3286445](https://doi.org/10.1109/JPROC.2023.3286445).
- [A7] D. Wu, B.-L. Lu, B. Hu, and Z. Zeng, "Affective brain–computer interfaces (aBCIs): A tutorial," *Proc. IEEE*, vol. 111, no. 10, pp. 1314–1332, Oct. 2023, doi: [10.1109/JPROC.2023.3277471](https://doi.org/10.1109/JPROC.2023.3277471).
- [A8] T. Olugbade, L. He, P. Maiolino, D. Heylen, and N. Bianchi-Berthouze, "Touch technology in affective human–, robot–, and virtual–human interactions: A survey," *Proc. IEEE*, vol. 111, no. 10, pp. 1333–1354, Oct. 2023, doi: [10.1109/JPROC.2023.3272780](https://doi.org/10.1109/JPROC.2023.3272780).
- [A9] A. Triantafyllopoulos et al., "An overview of affective speech synthesis and conversion in the deep learning era," *Proc. IEEE*, vol. 111, no. 10, pp. 1355–1381, Oct. 2023, doi: [10.1109/JPROC.2023.3250266](https://doi.org/10.1109/JPROC.2023.3250266).
- [A10] C. M. de Melo, J. Gratch, S. Marsella, and C. Pelachaud, "Social functions of machine emotional expressions," *Proc. IEEE*, vol. 111, no. 10, pp. 1382–1397, Oct. 2023, doi: [10.1109/JPROC.2023.3261137](https://doi.org/10.1109/JPROC.2023.3261137).
- [A11] B. M. Booth, N. Bosch, and S. K. D'Mello, "Engagement detection and its applications in learning: A tutorial and selective review," *Proc. IEEE*, vol. 111, no. 10, pp. 1398–1422, Oct. 2023, doi: [10.1109/JPROC.2023.3309560](https://doi.org/10.1109/JPROC.2023.3309560).
- [A12] G. Yannakakis and D. Melhart, "Affective game computing: A survey," *Proc. IEEE*, vol. 111, no. 10, pp. 1423–1444, Oct. 2023.
- [A13] L. Devillers and R. Cowie, "Ethical considerations on affective computing: An overview," *Proc. IEEE*, vol. 111, no. 10, pp. 1445–1458, Oct. 2023.