Multimodal Sentiment Analysis in the Wild: Ethical considerations on Data Collection, Annotation, and Exploitation

Björn Schuller[1,2], Jean-Gabriel Ganascia[3],Laurence Devillers[4,5]
[1]University of Passau, Chair of Complex & Intelligent Systems, Passau, Germany
[2]Imperial College London, Department of Computing, London, United Kingdom
[3]University Pierre & Marie Curie, LIP6, Paris, France
[4]CNRS, LIMSI, Orsay, France
[5]University Paris-Sorbonne IV, Paris, France

E-mail: [1]schuller@IEEE.org, [3]jean-gabriel.ganascia@lip6.fr, [4]devil@limsi.fr

Abstract

Some ethical issues that arise for data collection and annotation of audio-visual and general multimodal sentiment, affect, and emotion data "in the wild" are of types that have been well explored, and there are good reasons to believe that they can be handled in routine ways. They mainly involve two areas, namely research with human participants, and protection of personal data. Some other ethical issues coming with such data such as its exploitation in real-life recognition engines and evaluation in long-term usages are, however, less explored. Here, we aim to discuss both – the more "routine" aspects as well as the white spots in the literature of the field. The discussion will be guided by needs and observations as well as plans made during and for the European SEWA project to provide a showcase example.

Keywords: affective computing, sentiment analysis, ethical, legal and social implications (ELSI), data protection

1. Introduction

In the recent time we witness ever-more collection "in the wild" of individual and personal multimodal and increasing amounts of sensorial affect and sentiment data, crowd-sourced annotation by large groups of individuals with often unknown reliability and high subjectivity, and "deep" and partially less supervised learning with limited transparency of what is being learnt, and how applications depending on such data may behave. This renders the ethical, legal, and social implications (ELSI) more crucial than ever before in the field of language and multimodal resources. Accordingly, it makes the related aspects(e.g., privacy, traceability, explainability, validity, etc.) and according responsibility that comes with the collection, annotation, storing, and in particular also exploitation of (human) data of personal affect, behaviour, emotion, opinion, and sentiment a key concern. This comes in particular, as automatic systems increasingly exploit data of (and interact with) humans of all ranges (e.g., children, adults, vulnerable populations) including non-verbal and verbal data occurring in a variety of real-life contexts (e.g., at home, the hospital, on the phone, in the car, in the classroom, or within public transportation) and act as assistive and partially instructive technologies, companions, and/or commercial or even decision making systems.

In contrast to this increased relevance, the body of literature (cf., e.g., [1-5]) dealing with ELSI aspects is hardly in any balance with the number of technical publications found on the topic. Here, we aim to discuss these aspects, guided by a showcase example to provide a basis of discussion: This example will be the Automatic Sentiment Analysis in the Wild (SEWA) European project[1] that set off early in 2015. The project has the goal to advance models and algorithms for machine analysis of facial, vocal, and verbal behaviour, to realise naturalistic human-centric human-computer interaction and computer-mediated face-to-face interaction. It aims at a set of audio and visual spatiotemporal methods for automatic analysis of human spontaneous (as opposed to posed and exaggerated) patterns of behavioural cues including analysis of sentiment and liking. Technologies that can robustly and accurately analyse human facial, vocal, and verbal behaviour (and interactions) in the wild, i.e., in people's everyday life's surroundings, as observed by webcams in digital devices, would have profound impact on both, basic sciences, and the industrial sector. They could open up tremendous potential to measure behaviour indicators that heretofore resisted measurement because they were too subtle or fleeting to be measured by the human eye and ear, would effectively lead to development of the next generation of efficient, seamless opinion mining. Accordingly, one could expect profound impact on business as automatic market research analysis would become possible, and further beyond, recruitment could become more objective and green as travels would be reduced drastically at the same time, however, raising considerable ELSI implications such as whether computer-assisted recruitment is sufficiently reliable. The technology would also enable user-centric human-computer interaction by affective multimodal interfaces, and one could think of interactive multi-party games, and online services such as social TV. A large number of further applications would be enabled such as next generation healthcare technologies by remote monitoring of conditions like pain, anxiety and depression, and alike, to mention but a few examples.
This makes it obvious, what huge responsibility lies in the accuracy of such according recognition engines, their thoughtful implementation, and reasonable communication with regards to their reliability and privacy and individual rights awareness. Furthermore, learning models of human affect, behaviour, and sentiment suitable for machine analysis depends on having suitable data recordings of human behaviour to

[1] http://www.sewaproject.eu/

learn from. Hence, an important aspect of the SEWA project lies in collecting suitable datasets of sufficient labelled examples for building robust tools. Its plan includes the release of a large volume of audio-visual data of human behaviour recorded in the wild together with expert annotations in the form of a publicly available database. The intention behind is to push forward the research in multicultural and multilingual automatic human affect behavioural analysis and user-centric HCI and FF-HCI.

Here, we aim to discuss the good practices and ethical standards and issues at different stages of such collection, annotation, and exploitation of sentiment as collected "in the wild". There are two main ethical issues of concern that we will be dealing with:

- The first concerns the fact that human subjects are involved in the data collection process.
- The second concerns the use of emergent sentiment analysis technologies and their possible applications.

2. Database Collection

There are several guidelines on considerations to be made when collecting data that involves humans. Further, boards and mechanisms to overlook the process are usually in place. To give an example, in the United States of America and similarly across Europe, database collection is strongly governed by a university's institutional or ethical review board. Such a board has to approve ahead of the collection, monitor throughout the collection, and review afterwards what has been collected and potentially distributed in the context of (human) behavioural research. Similar boards are increasingly required and overlooked in connection with (inter-)national public research funding. To stick with our illustrative example, we will outline the process as encountered in the SEWA project data collection. These are given for the sake of completeness, albeit most of the outlined points are common knowledge. According to the ethical standards of human experimentation and to the requirement of the Imperial College Research Ethical Committee (ICREC), the questions relative to data collection can be subdivided into three parts:

Informed consent: A form of consent needs to best be validated by experts before the beginning of the data collection phase. This form, which will be signed by subjects involved in the experimentation, includes the data protocol description, the aim of this experiment, the description of technologies used to capture audio and visual signals and the storage and use of the data. In the SEWA project, this form needed to be translated into six languages by native speakers as the SEWA team collects multilingual and multicultural data. As the raw data will be made available to the scientific community for research purposes, the participants are made aware about the openness of the data when they sign the consent providing different levels of agreement such as usage only within the scientific community or giving consent usage of image and video material in dissemination for the scientific community or broader public. Ideally, this consent form should also explain the benefit to the public arising from the collection and the individuals taking part.

Verification of the harmless nature of the data collection. In the example of the SEWA project, there are no invasive sensors because only acoustic and video signals are used. In fact, such sensors would raise additional concerns, as it is less transparent to participants what kind of information could be contained, as they cannot access it themselves in natural ways. In addition, the video material that will be presented during the experiment is not supposed to be traumatic.

Data storage: Here, proper ethical and legal handling has to be ensured. For example, it needs to be validated whether the data are "sensitive" such as including banking information, the tax information, the health data, etc. In SEWA, the samples collected are twofold, reactions to video ads and face-to-face dyadic conversations. While the data is not sensitive in the sense as described above, anonymous storage needs to be discussed critically. Note that, anonymisation means two different things: First, to break the link between personal data and the persons whom these data are drawn from, and second, to make it impossible to retrieve the persons from their personal data. While it appears quite easy to anonymise data in the first sense, because it is sufficient to remove or to make inaccessible explicit references to the persons, i.e., their name and address, it is by far more difficult to prevent re-identification. Even when explicit references to names and addresses are removed, it appears possible by cross-references on multiple databases to infer names and addresses. Besides, the specific nature of data in the SEWA project, i.e., face images and speech signal, allows a natural re-identification by people who know the persons or using face and speaker recognition. As a consequence, if raw data are stored in the database, it is always possible to recognise people that participate to the project.

Rather than discussing the technical implications of this oversight in further detail, let us now switch to the interesting question what one cannot– or more specifically, examples of what we could not – do when collecting according data, given the named restrictions and further ethical considerations. As the desire of the SEWA project and in fact of most data collection centred around affective and behavioural computing, sentiment analysis, and opinion mining is usually to collect naturalistic data "in the wild" to be as close to the real-world use-case as possible, implications arise in particular from the multimodal nature of data: If one collects audio-visual data in the wild as is the case in SEWA, one potentially collects footage of other individuals not knowingly involved in the recording, or private information such as number plates of cars parked, the inside of private living space, etc. Thus, without ensuring massive resighting, reviewing, and processing of the collected data (such as by hiding others' faces or number plates by black bars or alike), one cannot share such recordings. However, the workload involved may be prohibitive such that, the recordings may at the end not be made in spaces outside of the property of those recorded or empty public spaces. Also, introducing such "hiding" or "blurring" elements may influence the training of machine learning algorithms.

Even more obviously, collecting data without the participants being aware of the recording may be desirable to increase the degree of spontaneity and naturalism, yet, comes at even higher ethical and legal restrictions. Very little data has been collected in the named fields in such a manner up to this point, such as in [6]. There, the local (Austrian) law allowed for private (audio-visual) recordings that may only be used for oneself unless the persons involved gave their consent (including consent given only after recording) for the material to be used for scientific (or other) purposes. In addition, as the recordings in that work took place in a private Supermarket, agreement of the shop-owner was needed in advance. As a consequence, surveillance notes needed to be put into place potentially reducing spontaneity of the behaviour. This example shows that "in the wild" collection comes at considerable efforts, but also limitations in terms of local environment and persons involved and their awareness of being recorded. Further, as the desire is often to cover for a gender, age, and cultural balance in such collection, it may be of critical relevance to decide on the material used for stimulation or induction of sentiment or affect. As a consequence, the effect may be reduced, as certain material or ways of eliciting reactions may not be appropriate to all participants of a database collection. As an example, showing extreme violence to participants may have a strong affect eliciting effect, but may not be appropriate in many cases. Similarly, some religious or political material may be sensitive to some cultural or ethnical groups in the context of sentiment analysis and opinion mining as outlined above.

A related interesting question touches upon how privacy impacts on the ability to understand the collected data. To give an example, in the SEWA project, pairs of subjects have been recorded that briefly discussed commercial spots they first watched by themselves. The precondition during enrolment for the study was that, such a pair has to know each other in advance in order to avoid (usually over-friendly and targeted towards each other rather than the subject of interest of the recordings) "getting to know" behaviour in the short time of the recording. However, owing to privacy restrictions, the full relationship status may not be known or revealed but clearly of interest when interpreting the data as to which part of behaviour shown is related to the content of the commercial or to the person being spoken to.

Further, it seems not trivial to make participants in recordings understand privacy protections leading to the question of according consequences of data collection. A first (comparably minor) "risk" is losing potential participants as they may misinterpret protection such that they refrain from participation despite the data and privacy protection mechanisms being at very high levels. In the example named above taken from [6] of subjects being recorded unknowingly at first, there is a fair chance of losing participants in a study due to their surprise of having been recorded unknowingly at first that might have agreed if they had been told in advance. Then, however, the behaviour would have likely been less spontaneous. However, a more serious risk is of the nature that subjects do not understand all implications if the privacy protection is rather weak.

To conclude this section, we provide a sketch of the collection in the SEWA project. The following balancing of participants was targeted: across age from 18 years onwards by five groups as follows: 18-29, 30-39, 40-49, 50-59, and 60-80 years. Further, as participants were grouped in pairs each knowing each other as described, the couples were balanced in terms of best even distribution per age group by female-female, male-male, and mixed gender. As a target, each age and gender constellations had to appear at least once and ideally, each should appear twice totalling up to at least 30 pairs or 60 individuals per language/culture of collection. Given the difficulty to evenly recruit according pairs across all age groups and classes, some groups such as younger individuals are present slightly more strongly in the final database.

The recording was split into two parts: In the first part, each participant had to individually watch four commercial spots with 60 seconds, each. These were chosen such as to induce different affective states including amusement, empathy, positive sentiment or boredom. A challenge at this point was to select these such as to induce target states and behaviour and at the same time not be offensive or disturbing in any way as described above. Within the second part, the participants communicated via a video-chat software to another participant known to them – on average for 4-5 minutes – regarding the content of the spots just seen each by themselves. The intention behind is to collect further reactions and opinions with respect to the content of the commercial and the product, service or charity appeal shown, which are the highlights of the spots, whether these are appropriate, how they could be improved and alike. Further, this allows for analysis of inter-human behaviour in dyadic conversations.

After obtaining ethical approval for the SEWA experiment internally and from an external ethical advisory board formed by the second and third author of this contribution, the experiment protocol was implemented and again overlooked. Next, a website and service for collection was implemented by partners of the project via which at this point 199 successful data recording sessions took place including 398 participants from the six different language and cultural backgrounds (British, Chinese, German, Greek, Hungarian, and Serbian). This required informed consent forms and the web interface to be translated into each of the six languages involved as named above. These informed them on the funding source, the intended recording and annotation in principle, the foreseen benefit to society coming from the project, and their rights to withdraw recordings at any time besides standard explanations on privacy and protection concerns. It clearly stated that, participation is voluntary, non-participation will not result in any kind of disadvantage, and that termination is possible at any moment during the recording. It also provided a contact address for independent help and information on ELSI implications at the university or responsible body. The participants had to register first on a secure web page, fill in a form of demographic questions and confirm their email address via an email sent to them. With the conductor of the experiment they then had to agree upon an appointment where both partners were

available for around 15 minutes via email. They were instructed to do the recordings at home or any other venue of choice, and that noises are no problem. However, they were not allowed to be in the same room as the partner. This freedom of choice of venue can lead to the above sketched issues of potential inclusion of other individuals or other's property which has to be counter-checked during annotation.

Each participant further used her own PC or notebook with their own webcam and microphone (intentionally) leading to a high "in the wild" variability of recording devices. They were further using their own internet connection. From an ethical point of view, this may be seen as limiting factor given potential exclusion of parts of the population. For the SEWA project, this may be less of a concern given that (most of) the use-cases address data analysis with implication of mostly such individuals that possess and use according infrastructure. This may, however, clearly be different in other studies. The recording was fulfilled via the webpage which was largely self-explanatory. Participants had to log in in time at the agreed upon time slot. Each participant had to fill in a consent form also in print version and sign and send as a scan via email. There, they had the choice to agree to usage for scientific purposes and additionally whether recordings may be used in a public context. A financial reward was given to the participants via bank transfer. Bank data could optionally be communicated via email or phone. Obviously, ethical implications also come with graduation of participants. Here, the amount was chosen small enough to be rather of symbolic nature then risking involving participants that "sell" their data.

The recordings made contain 44 hours of audio-visual footage including a wide range of spontaneous expressions of emotions and sentiment. It seems noteworthy that, due to the technical framework and requirements (higher bandwidth needed, recordings were considered only valid if successful during the first attempt, as otherwise the reactions would not be spontaneous any more) a higher rate of failure exists. To exemplify, for the recordings taken in Germany, 57 sessions were started, 43 pairs attempted, but only 37 pairs successfully recorded in the end. Obviously, this is difficult from an ethical point of view, as some participants could not be included due to technical issues, which may be disappointing to them. Further, the higher number of attempts than pairs shows repeated difficulties at the beginning prior to the recording of interest. This is time consuming for the participants and potentially influences their affect and mood. As a consequence, all efforts were made to avoid such circumstances. There was no pronounced gender effect for within-gender and cross-gender pair differences (i.e., all three constellations of gender grouping occurred equally often).

The collection further included extraction of acoustic, linguistic, and visual features. Acoustic features were extracted in two different sizes of feature space by the open-source openSMILE [7] ComParE and GeMAPSv01a standardised feature sets from all SEWA recordings. Similarly, 49 facial landmarks were automatically tracked. This provides an interesting alternative option of distribution of data: Rather than distributing the full audio-visual recordings, sharing just feature representations for reproduction and comparison of and with scientific findings comes at higher protection of privacy. However, care is needed, as features sampled at short intervals and in high numbers and complementarity may allow for resynthesis of the original (audio-visual) source data to large extents. Further compression such as by (sub-band) vector quantisation may reduce this risk [8].

3. Data Annotation and Release

Data annotation bears its own ELSI pitfalls in particular in the context of crowdsourcing given that, the data will be shown to potentially unknown raters "outside the lab". This may include them watching the material in public spaces in the presence of others, as crowdsourcing increasingly becomes mobile (cf., e.g., [9]). In [10], the authors name the primary concern of 12 researchers questioned in an according study to be privacy-related ranked second after accuracy-related and further concerns such as related to the reliability and the costs involved when it comes to crowdsourced video coding. The authors further suggest blur filters as suited means to better hide the identity of the individuals to be coded. Unfortunately, as one may expect, this does at the same time downgrade also the coders "ability to accurately and reliably code behaviours" [10]. Luckily, however, the decrease was "not as steeply as the identity test"[10]. Accordingly, such methods need to be improved, and similar methods need to be established and evaluated carefully for audio or even textual and further information "blurring" in this context.

For SEWA data annotation, annotation within the lab including the crowdsourcing platform iHEARu-PLAY [11] was successfully used up to this point. The latter provides a gamified approach without monetary compensation. Specific ethical considerations are summarised in detail in [12]. The scheme for the annotation of the data includes continuous assessment along the three primitives or dimensions arousal, valence, and sentiment/liking. A major issue in this respect is to correctly instruct annotators such as to ensure good understanding of the differences between these primitives to warrant high quality annotations. Further, verbal transcriptions including non-verbal vocalisations were made manually and counter-checked in five languages up to this point (excluding Greek). Here, it was necessary to identify native speakers of these languages, each, for the transcriptions to ensure accurate transcription. The results are overall further refined through semi-automatic correction.

A core SEWA dataset (currently 540 representative segments – 90 from each culture group – chosen in balance by high/low arousal, high/low valence, and liking/disliking) is currently further annotated fully in terms of facial landmarks, facial action units (FAUs), mimicry, sentiment, rapport, and template behaviours. Again, this will partially require expert coders – in particular Facial Action Coding System (FACS) certified coders for the FAUs. This shows that only part of the annotation can be distributed to the (partially laymen) crowd. To ensure high privacy standards given the "in the wild" nature of the collection, the SEWA database shall be

used at first within the project consortium to identify potential remaining issues internally during application and use of the data prior to a public release. This release will be via a web-portal allowing for enhanced search functionality to invite also non-technical scientific usage where fast retrieval of specific behaviour is crucial. In fact, we believe a broad release for research only (e.g., by password protected access via secure sites) or potentially even the greater public to be of crucial importance: First of all, giving access to other researchers will avoid double efforts in collection and thus require less participants and annotators. At the same time, this can accelerate progress that may be highly needed such as in the health sector. Further, it increases reproducibility of findings – an often violated key principle of good research. This can be done, e.g., in competitive challenges to increase interest in the data such as in [13] or the MediaEval series. Finally, the data collection often is subsidised by grant money from public sources – thus, the public should best benefit from the efforts and resources should be spent in an utmost efficient manner, only.

4. Exploitation

Data collected and annotated in the context of affective and behavioural computing and sentiment analysis is usually used to train models for applications including analysis and synthesis of emotion, sentiment, and behaviour. In the SEWA project, application of recognition of human sentiment and behaviour includes in particular recommendation systems and face-to-face interaction through a chat roulette social game. In these applications, the data storage is a challenging topic. The system architecture solution proposes local data storage to protect privacy. Similarly important are, however, ethical implications of the actual application. In the project, two focus groups were built to ensure responsible and sensitive discussion: the first includes the Ethical Advisory Board (EAB) of the project – as outlined above, instantiated by the second and third author of this contribution, and members of and industrialValorisation Advisory Board (VAB); the second comprises users and professionals. The following key points are considered of interest by these boards and in discussions:

Recognition, recognisability, and uncertainty: It needs to be ensured that what is being recognised by an automatic system is recognisable at all. One easily falls for the trap of taking it for granted that computerised measurement and classification are objective, as they stem from a technical system. They thus would lead to formalised representation of human emotion and disposition. However, many human phenomena including the above are too complex and ambiguous to allow for (complete) objectification. This comes among others, as higher level individual aspects and context need to be taken into account, but often are not. Proper communication of the recognisable thus is of crucial importance, such as by provision of confidence measures and implementation of benchmark tests such as the Interspeech Computational Paralinguistic Challenges 2009-16 or the Audio/Visual Emotion Challenges 2011-16 (cf. e.g., [13]). Further, the uncertainty has to be protected, i.e., it has to be ensured that certain private spheres are not entered and users of technology are aware of a remaining uncertainty.

Reductionism: Models designed for computational assessment of human emotion, sentiment, and behaviour are often simplified. This bears the danger of unforeseeable implications as the actual problem's complexity is reduced to a potentially insufficient representation.

Effect of erroneous decisions: The harmless character of erroneous decisions has to be ensured in best possible ways. In a recommender system such as envisioned by the SEWA project as one exemplary use-case, the implications may be less severe such as receiving sub-optimal recommendations on the content of potential interest, e.g., music or movies. However, in the second use-case of a social chat-roulette game, implications are more severe: If a system makes wrong assumptions on users (dis-)liking each other, the social implications may be (more) drastic such as (erroneously) made to belief someone dislikes the other. Clearly, however, there are potentially even more critical use-cases such as the above named "green" job interviews where a system may become responsible of someone erroneously not being employed.

As industrial partners and health and security providers increasingly collaborate with scientists rooted in computer science, and electrical engineering in the fields of affective computing, sentiment analysis, opinion mining, behavioural and social signal processing, it is increasingly important to understand what can or cannot be modelled and sensed in an accurate and reliable fashion. It will be important to also further strengthen the collaborative and communicative aspect in this respect.

5. Conclusion

Many ethical issues (evaluation of the sentiment analysis technologies in the wild, possible applications, etc.) need to be addressed when dealing with affective and behavioural corpus collection, annotation, and exploitation. Here, we named key-aspects, and exemplified them in the context of an ongoing European project dealing with "in the wild" collection across six cultures / languages. The idea was to demonstrate by a case study how broader ethical principles can be translated into a concrete policy. However, additional experience can be expected throughout the further runtime of the SEWA project contributing to its detailed policy to be shared. Future implications may be even more challenging, once technical systems become increasingly "conscious" also in emotional ways [14-18].

Acknowledgements

The research leading to these results has received funding from the European Union's Horizon 2020 Programme research and innovation programme under grant agreements Nos. 645094 (SEWA) and 644632 (MixedEmotions), and from the German national BMBF IKT2020-Grant under grant agreement No. 16SV7213 (EmotAsS). The content presented has benefited from discussion with Harald Traue of University of Ulm and further participants of the ELSI networking workshop InterEmotio organised by the German BMBF in Stuttgart/Germany on 26 February 2016. We also thank the members of the SEWA consortium and its industrial Valorisation

Advisory Board. The authors further acknowledge the three (anonymous) reviewers' highly productive suggestions.

Bibliographical References

[1] Reynolds, Carson, and Rosalind Picard. "Affective sensors, privacy, and ethical contracts." Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI 2004 Extended Abstracts on Human Factors in Computing Systems, pages 1103-1106, Vienna, Austria, ACM, 2004.

[2] Reynolds, Carson Jonathan. "Adversarial uses of affective computing and ethical implications." Dissertation. Massachusetts Institute of Technology, 2005.

[3] Goldie, Peter, Sabine Döring, and Roddy Cowie. "The ethical distinctiveness of Emotion-Oriented Technology: Four long-term issues." Emotion-Oriented Systems. Springer Berlin Heidelberg, 2011. 725-733.

[4] Duffy, Brian R. "Fundamental issues in affective intelligent social machines." Open Artificial Intelligence Journal 2.1 (2008): 21-34.

[5] King, Elizabeth W. "The Ethics of Mining for Metadata Outside of Formal Discovery" Penn State Law Review 113 (2008): 801.

[6] Eyben, Florian, Weninger, Felix, Paletta, Lucas, and Schuller, Björn. "The acoustics of eye contact – Detecting visual attention from conversational audio cues" Proceedings 6th Workshop on Eye Gaze in Intelligent Human Machine Interaction: Gaze in Multimodal Interaction, GAZEIN 2013, held in conjunction with the 15th International Conference on Multimodal Interaction, ICMI 2013, pages 7–12, Sydney, Australia, ACM, December 2013.

[7] Eyben, Florian, Weninger, Felix, Groß, Florian, and Schuller, Björn. "Recent Developments in openSMILE, the Munich Open-Source Multimedia Feature Extractor" Proceedings of the 21st ACM International Conference on Multimedia, MM 2013, pages 835–838, Barcelona, Spain, ACM, 2013.

[8] Zhang, Zixing, Coutinho, Eduardo, Deng, Jun, and Schuller, Björn. "Distributing Recognition in Computational Paralinguistics." IEEE Transactions on Affective Computing, 5.4:406–417, October–December 2014.

[9] Nebeling, Michael, To, Alexandra, Guo, Anhong, de Freitas, Adrian A., Teevan, Jame, Dow, Steven P., and Bigham, Jeffrey P. "WearWrite: Crowd-Assisted Writing from Smartwatches" Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI 2016, 10 pages, San Jose, CA, ACM, 2016.

[10] Lasecki, Walter S., Gordon, Mitchell, Leung, Winnie, Lim, Ellen, Bigham, Jeffrey P., Dow, Steven P. "Exploring Privacy and Accuracy Trade-Offs in Crowdsourced Behavioral Video Coding" Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI 2015, 6 pages, Seoul, Republic of Korea, ACM, 2015.

[11] Hantke,Simone, Eyben, Florian, Appel, Tobias, and Schuller, Björn. "iHEARu-PLAY: Introducing a game for crowdsourced data collection for affective computing." Proceedings 1st International Workshop on Automatic Sentiment Analysis in the Wild, WASA 2015, held in conjunction with the 6th biannual Conference on Affective Computing and Intelligent Interaction, ACII 2015, pages 891–897, Xi'an, P. R. China, AAAC, IEEE, September2015.

[12]Hantke, Simone, Batliner, Anton, and Schuller, Björn."Ethics for Crowdsourced Corpus Collection, Data Annotation and its Application in the Web-based Game iHEARu-PLAY" Proceedings of the1st International Workshop on ETHics In Corpus Collection, Annotation and Application, ETHI-CA² 2016, satellite of the10th Language Resources and Evaluation Conference, LREC2016, 6 pages, Portoroz, Slovenia, ELRA, May 2016.

[13] Ringeval, Fabien, Schuller, Björn, Valstar, Michel, Jaiswal, Shashank, Marchi, Erik, Lalanne, Denis, Cowie, Roddy, and Pantic, Maja."AV+EC 2015 - The First Affect Recognition Challenge Bridging Across Audio, Video, and Physiological Data." Proceedings of the 5th International Workshop on Audio/Visual Emotion Challenge, AVEC15, collocated with the 23rd ACM International Conference on Multimedia, MM 2015, pages 3–8, Brisbane, Australia, ACM, October 2015.

[14] Sullins, John P. "Robots, love, and sex: The ethics of building a love machine." IEEE Transactions on Affective Computing, IEEE 3.4 (2012): 398-409.

[15] Torrance, Steve. "Ethics and consciousness in artificial agents." AI& Society 22.4 (2008): 495-521.

[16] Whitby, Blay. "Sometimes it's hard to be a robot: A call for action on the ethics of abusing artificial agents." Interacting with Computers 20.3 (2008): 326-333.

[17] Pontier, Matthijs A., and Johan F. Hoorn. "Toward machines that behave ethically better than humans do." Proceedings of the 34th international annual conference of the cognitive science society, Cognitive Science 12 (2012).

[18] Grinbaum,Alexeij, Chatila, Raja, Devillers, Laurence, Ganascia, Jean-Gabriel, Tessier, Catherine, Dauchet, Max. "Ethics in Robotics Research" Proceedings International Conference on Robot Ethics, ICRE 2015, Lisbon Portugal, October 2015.