

Emotional AI, soft biometrics and the surveillance of emotional life: An unusual consensus on privacy

McStay, Andrew

Big Data and Society

DOI:

[10.1177/2053951720904386](https://doi.org/10.1177/2053951720904386)

Published: 01/01/2020

Publisher's PDF, also known as Version of record

[Cyswllt i'r cyhoeddiad / Link to publication](#)

Dyfyniad o'r fersiwn a gyhoeddwyd / Citation for published version (APA):

McStay, A. (2020). Emotional AI, soft biometrics and the surveillance of emotional life: An unusual consensus on privacy. *Big Data and Society*, 7(1).
<https://doi.org/10.1177/2053951720904386>

Hawliau Cyffredinol / General rights


Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Emotional AI, soft biometrics and the surveillance of emotional life: An unusual consensus on privacy

Big Data & Society
January–June 2020: 1–12
© The Author(s) 2020
DOI: 10.1177/2053951720904386
journals.sagepub.com/home/bds


Andrew McStay 

Abstract

By the early 2020s, emotional artificial intelligence (emotional AI) will become increasingly present in everyday objects and practices such as assistants, cars, games, mobile phones, wearables, toys, marketing, insurance, policing, education and border controls. There is also keen interest in using these technologies to regulate and optimize the emotional experiences of spaces, such as workplaces, hospitals, prisons, classrooms, travel infrastructures, restaurants, retail and chain stores. Developers frequently claim that their applications do not identify people. Taking the claim at face value, this paper asks, what are the privacy implications of emotional AI practices that do *not* identify individuals? To investigate privacy perspectives on soft non-identifying emotional AI, the paper draws upon the following: over 100 interviews with the emotion detection industry, legal community, policy-makers, regulators and NGOs interested in privacy; a workshop with stakeholders to design ethical codes for using data about emotions; a UK survey of 2068 citizens on feelings about emotion capture technologies. It finds a weak consensus among social stakeholders on the need for privacy, this driven by different interests and motivations. Given this weak consensus, it concludes that there exists a limited window of opportunity to societally agree principles of practice regarding privacy and the use of data about emotions.

Keywords

Affective computing, biometrics, consensus, data protection, emotional AI, group privacy

This article is a part of special theme on Big Data and Surveillance. To see a full list of all articles in this special theme, please click here: <https://journals.sagepub.com/page/bds/collections/hypecommerciallogics>

Emotional AI refers to technologies that use affective computing and artificial intelligence techniques to sense, learn about and interact with human emotional life. This paper assesses the privacy implications of these technologies and organizational applications employed to make inferences about emotions, feelings, moods, perspective, attention and intention. While at an embryonic stage, they are becoming increasingly present in everyday objects and practices such as assistants, cars, games, mobile phones, wearables, toys, marketing, insurance, policing, education and border control. They are also being used to regulate and optimize the emotionality of spaces, such as workplaces, hospitals, prisons, classrooms, travel infrastructures, restaurants, retail and chain stores. To explore these developments, this paper asks, what are the privacy implications of emotional AI practices that do *not*

identify individuals? The question arises from commercial and other interests in soft biometrics that entail categorization of bodily traits where a person may not be identified in the process. Emotion is an example of soft biometric trait profiling. To investigate the privacy implications of using non-personally identifying data about people to try to infer emotions, this paper situates its discussion in the context of burgeoning interest in emotional AI and (separately) group-focused privacy. It draws on three sources of empirical

Bangor University, Gwynedd, UK

Corresponding author:

Andrew McStay, School of Music and Media, Bangor University, Gwynedd, LL57 2DG, UK.
Email: mcstay@bangor.ac.uk



insight: firstly, over 100 interviews with data regulators, privacy-oriented NGOs, policymakers, legal actors and diverse industrial and service sectors interested in data about emotions. The second source is a UK national survey conducted on citizens' attitudes towards emotional AI. The third source is a multi-stakeholder workshop conducted to co-create ethical codes of conduct for employing emotional AI. In conclusion, the paper observes a weak consensus on the need for privacy, albeit driven by different motivations. It argues that this temporary consensus should be seized to implement regulation on these emergent technologies.

On emotion sensing

The practice of using computer sensing to interact with emotional life has origins in the 1990s, with the field of affective computing (Picard, 1997, 2007). What McStay (2016, 2018) terms 'emotional AI' and 'empathic media' are made possible through weak, narrow and task-based AI efforts to see, read, listen, feel, classify and learn about emotional life. This involves data about words, images, facial expressions, gaze direction, gestures, voices and the body, that in turn encompasses heart rate, body temperature, respiration and electrical properties of skin. Given that this paper is interested in physiological inferences, input features might be facial expressions, voice samples or biofeedback data. The output is named emotional states that are then used for a given purpose. Applicable machine learning techniques vary, but frequently involve convolutional neural nets (useful for images and system efficiency gains), region proposal networks (useful for multiple object recognition) and recurrent neural networks (that draw on recent past data to determine how they respond to new input data).

Output emotional states are used to enhance interaction with devices and media content; create new forms of toys and entertainment; make experiences more immersive; enhance artistic expression; surveil and enable learning; facilitate self-understanding of moods and well-being; optimise and regulate behaviour in closed spaces (e.g. prisons and travel infrastructures); judge risk (e.g. by providing car insurance companies data about reactivity); surveil and measure emotionality of bounded spaces (such as retail outlets or cities); surveil customers and worker performance; provide emotional reactivity feedback to marketers and facilitate creation/targeting of advertising.

The 'basic emotions' methodology (Ekman and Friesen, 1971, McDuff and el Kaliouby, 2017) that sits behind much emotional AI has been widely critiqued (Andrejevic, 2013; Leys, 2011; Russell, 1994). Indeed, the 2018 AI Now Report debunks it as pseudoscience, linking facial coding with phrenology (Whittaker et al.,

2018). This critique is extreme, although practitioners and vendors of emotional AI recognise that single labels rarely capture complex emotional and affective behaviour (Gunes and Pantic, 2010). The key problem with face-based approaches is that they are based on reverse inference where an expression is taken to signify the experience of an emotion (Barrett et al., 2019). This is problematic because 'similar configurations of facial movements variably express instances of more than one emotion category' (Barrett et al., 2019), which indicates that more detail on the context of the situation is required to understand the emotion. This requires more data and potentially more invasive practices (McStay and Urquhart, 2019).

Application of emotion tracking has progressed from in-house research facilities (such as those used in neuro-marketing to detect responses to adverts) to online and physical contexts. Including emojis (Davies, 2016; Stark and Crawford, 2015), wearables (Lupton, 2016; Neff and Nafus, 2016; Picard, 1997), human-robot interaction (Bryson, 2018), education (Williamson, 2017), retail (Turow, 2017), employee behaviour (Davies, 2015; Grandey et al., 2013) and border control (Sánchez-Monedero and Dencik, 2019), the broader business strategy for emotional AI companies is ubiquitous usage of automated emotion detection in all personal, commercial and public contexts. To suggest a broad rule, if there is any form of value in understanding emotion in a given context, emotional AI has scope to be employed. Companies interested in emotional AI include established companies such as NEC, IBM, Apple, Google, Microsoft and Facebook, but also a long list of smaller companies (such as Eyeris, Sensing-Feeling and Affectiva) seeking to define new markets (for an extended list, see Emotional AI, 2018).

On privacy: Towards a common good

Despite liberal roots in respect for individuality, selfhood, autonomy and control, it is clear that privacy includes these principles but is not synonymous with them. Critically for this paper, privacy is not only an individual right, but a group right because it is a collective good. Further, given what is at stake – bodies, emotions and experience – dignity (for individuals and groups) remains especially important in ubiquitous computing contexts (Edwards, 2016).

A dignity-based understanding diagnoses the problems with passive profiling and using Big Data techniques about emotions for unconscious influence: it is about recognising that phenomenological experience is important, innately worthy and should not be appropriated. Applied to the body and face expressions, this is not moral idling as the blog for the European Data Protection Supervisor (Europe's data protection

authority) states, ‘Turning the human face into another object for measurement and categorisation by automated processes controlled by powerful companies and governments touches the right to human dignity’ (Wiewiórowski, 2019). Dignity also serves to block conceptions of privacy as an indirect expression of other rights, such as property. As Floridi (2016) puts it, the ‘my’ in ‘my data’ is not the same as the ‘my’ as in ‘my car’, because personal, sensitive and intimate information plays a constitutive role of who we are.

On privacy as a common good, Floridi (2014) points out that the right to privacy is not just an individual right, it is a *group right*: it is held by a group as a group, rather than by its members. Momentarily leaving to one side legal discussion of re-identification and mosaic effects, the logic of Big Data profiling is rarely based on targeting individuals. To quote Floridi:

There are very few Moby-Dicks. Most of us are sardines. The individual sardine may believe that the encircling net is trying to catch it. It is not. It is trying to catch the whole shoal. It is therefore the shoal that needs to be protected, if the sardine is to be saved. (2014: 3)

Floridi’s (2014) sardines also connect well with affinity profiling in online behavioural advertising where profiling ‘does not directly infer sensitive data (“special category data”) but rather measures an “affinity” with a group defined by such data’ (Wachter, 2020/forthcoming: 5). Wachter’s focus is on assumed interests as a proxy for personal traits, but the principle of *indirect inference* is interesting because data collected about emotional conditions may be collected by group aggregate, but will impact them personally, albeit not through strictly speaking personal data. The interest in online advertising has parallel with urban smart advertising that makes use of emotional AI (McStay, 2016; McStay and Urquhart, 2019). For example, groups of people that move daily through urban spaces (such as a commuter-line train stations) will by default become identifiable groups clustered by psycho-physiological emotional reactivity. They will occur through intimate objectification (granular assessment of reactivity) and special treatment distinct from commuters in other parts of a city, based on collection of data that they have no control over.

Regulatory context: Emotional AI and soft biometrics

The value of group-oriented body-focused privacy critique becomes clear when cast against European Union (EU) regulation on data protection. The EU is a useful

benchmark as it has the most stringent data protection and privacy standards in the world: thus, if the EU is not adequately prepared to legally address emotion sensing, arguably, nowhere else will be. In general, as stated in Article 4(1) of the General Data Protection Directive (GDPR), EU directives and regulations exist to protect personal data. This is when information can be used to identify or single out a person from others, so they can be treated differently. If the information in question is not personal data, the regulations do not apply. Yet, the EDPS, through Opinion 4/2015 on data, technology and dignity, moves towards group privacy sentiment stating that Big Data ‘should be considered personal even where anonymisation techniques have been applied’, although adding that ‘it is becoming ever easier to infer a person’s identity by combining allegedly ‘anonymous’ data with other datasets including publicly available information’ (2015: 6). While the overall point focuses on personal identity, it is multi-staged, i.e. in the first instance such data should be considered personal upfront (even if not identifying). Following this thinking, if personal in the first instance, then it follows that it will be also sensitive (as per Article 9 of GDPR, requiring explicit opt-in) because it has scope to involve identifying biometrics. However, this is notable opinion rather than law. Oddly, the GDPR makes no reference whatsoever to emotions. Similarly, a proposal for the revised ePrivacy directive rarely mentions emotions (European Commission, 2017). Only recitals 2 and 20 mention emotions although, importantly, recital 2 defines them as highly sensitive. Yet, while introduced, emotions do not appear in the articles of the proposed ePrivacy directive. To an extent this is understandable because ‘emotion’ is an imprecise word (including social media sentiment analysis and Big Data inferencing, such as mood tracking of Spotify usage, as well as ‘soft’ and ‘hard’ biometrics). However, given the increasing role of emotion in data analytics and facilitating human–machine interaction, the absence is still surprising.

Although GDPR makes no reference to emotions, it does address identifying (or hard) biometric data. Article 4(14) defines it thus:

‘biometric data’ means personal data resulting from specific technical processing relating to the physical, physiological or behavioural characteristics of a natural person, which allows or confirm the unique identification of that natural person, such as facial images or dactyloscopic data. (European Commission, 2016: 34)

The laws on personally identifiable biometric data are stringent. Article 9(1) asserts conditions required to process identifying biometric data. The most important for emotional AI is the need for explicit consent (see Article 9(2)a). However, if the data in use cannot

identify an individual or single them out in some way, it follows that the regulation does not apply. This is especially pertinent for out-of-home emotional analytics that do not rely on personal devices to track people.¹ Also relevant is Article 29 Data Protection Working Party's *Opinion 3/2012 on developments in biometric technologies*. This states that biometric data 'by their very nature, are directly linked to an individual' (Article 29 Data Protection Working Party, 2012: 2), also pointing out that this data may be defined in terms of 'biological properties, behavioural aspects, physiological characteristics, living traits or repeatable actions where those features and/or actions are both unique to that individual and measurable' (Article 29 Data Protection Working Party, 2012: 3–4).

Importantly, the Article 29 Opinion states that:

Mention should also be made to the use of the so-called soft biometrics defined by the use of very common traits not suitable to clearly distinguish or identify an individual but that allow enhancing the performance of other identification systems. (Article 29 Data Protection Working Party, 2012: 16)

The Opinion reflects the biopolitical understanding employed in this paper that focuses on commercial rather than governmental dimensions. For example, Article 29 (Article 29 Data Protection Working Party, 2012: 17) says that soft biometrics opens the door 'to uses far different from large scale security applications' and that 'gaming and retail will benefit from an enhanced man-machine interaction allowing more than identification, or categorisation of an individual'. The other two relevant passages in the Opinion are:

Moreover some systems can secretly collect information related to emotional states or body characteristics and reveal health information resulting in a non-proportional data processing as well as in the processing of sensitive data in the meaning of article 8 of the Directive 95/46/EC. (Article 29 Data Protection Working Party, 2012: 17)

More recently it is not only identity that can be determined from a face but physiological and psychological characteristics such as ethnic origin, emotion and well-being. The ability to extract this volume of data from an image and the fact that a photograph can be taken from some distance without the knowledge of the data subject demonstrates the level of data protection issues which can arise from such technologies. (Article 29 Data Protection Working Party, 2012: 21)

Implicit within the soft biometric approach is that by definition it cannot 'single out' a person or device. If it

is able to do so, it tips into the category of hard biometrics. This raises a biopolitical conundrum because, as already argued, such data about emotions is intimate and sensitive, yet not legally personal. This echoes observations about big group-based data profiling that may also take place without identifying individuals and/or using any personal data (Wachter, 2019). This raises urgent governance questions, and privacy conceived as collective violation is left wanting. While principally interested in scope for group discrimination, Mittelstadt's remedy for how group rights should be exercised also applies to use of emotional AI that does not identify in that 'an educated guess can be made about what the group would value if given the opportunity to assemble and act collectively' (2017: 485).

In addition to the need for group privacy protections, there is another issue for concern about emotion tracking. That is, despite being about the body, modern regulation of soft biometrics speaks of data rather than privacy. Whereas the original Data Protection Directive (that GDPR replaces) frequently mentions privacy, in GDPR this is radically downplayed in preference of 'data protection'. Lynskey (2015) suggests comparing Article 8 of the European Convention on Human Rights (ECHR) to the protection offered by data protection law. She observes that data protection law has little to say about an intrusive strip-search, but this would be accounted for by Article 8 of ECHR that demands respect for private and family life. This presents a different picture from that found in data protection legislation that focuses on identification as the primary route to harm. What is clear is that an account of emotional AI based solely on data tracking and identification is a limited one. This paper reasons that as emotion tracking emerges, a broader dignity-based understanding to questions of commercial power and data privacy is required. Given that a quintessential attribute of dignity and humanity is physical, mental and experiential self-determination, how do influential stakeholders conceive of privacy in intimate contexts?

Methods

So far, this paper has identified a lacuna in critical literature on emotional AI and privacy, and in European law regarding soft biometrics and non-identifying usage of data about emotions. To explore what stakeholders think about ethics and emotional AI, this paper draws on insights derived from interviews with relevant stakeholders, a workshop with stakeholders to design ethical codes for using data about emotions and a UK survey to gauge citizen feelings about emotion capture technologies.

Interviews

The first source of insight comprises interviews with stakeholders directly interested in emotional AI. The objective here was to obtain understanding about the composition and scope of the emotional AI industry, what it does, why it does it, where it is heading and what stakeholders see as the main ethical issues. Across 2015–2017, 108 open-ended 1-hour interviews were conducted to elucidate views from industry, policymakers concerned with data and national security, municipal authorities and privacy-oriented NGOs.² Of these 108 interviews, 33 entailed discussion of non-identifying emotional AI, privacy and ethics. Many interviewees prefer not to be named in person or by company, but the categories are as follows:

- *Industry* (including a household name social media company, another global technology company, voice analytics, smart city vendors, market researchers using emotional AI, adtech firms, advertising agencies, a well-known insurance company and an online security firm). Interview number = 21.
- *Data protection NGOs* (including representatives of Privacy International, Electronic Frontier Foundation and Open Rights Group). Interview number = 6.
- *Self-regulation and law* (including an advertising self-regulator, a member of DG Connect, a global legal privacy association and three law firms that handle media and technology cases). Interview number = 6.

This paper focuses on this sample, although it is indirectly informed by contextual norms, values and attitudes encountered in the wider interviewing process. The sampling strategy was based on two factors: knowledge saturation and diversity (Bertaux, 1981). Variety was important because the study sought a broad understanding of interest in emotional AI at this early stage of its application. These included geographically varied organisations from the US, UK, France, Belgium, Estonia, Israel, Russia, United Arab Emirates and South Korea. The size of companies ranged from global (e.g. Alphabet (Verily), Philips, IBM and Facebook) to start-ups (mostly from London and San Francisco).

Interviewees comprised chief executive officers and individuals in strategic positions from the following sectors: advertising and marketing; policing and national security; education; insurance; angel investment; in-car experience and navigation; human resources and workplace management; sports; sex toys and psychosexual therapy; mental health; ethical hacking; art; and media, interactive film and games companies. Each sector was selected on the basis of their current

work in emotion detection, or likelihood of interest in these applications. In addition to industrialists and public sector actors, people working in privacy-friendly NGOs (Electronic Frontier Foundation, Open Rights Group and Privacy International) were interviewed to obtain a critical, policy-oriented perspective. Legal dimensions of emotion capture were explored in interviews with media and technology law firms, and European policymakers in the field of data protection. A multi-tiered consent form allowed interviewees to select a level of disclosure they were comfortable with.

Questions revolved around *opportunities* (trends and growth of emotion-sensitive technology), *rules* (data protection issues) and *harm minimisation* (corporate behaviour, privacy and thoughts about societal acceptance). Qualitative thematic coding was employed to hand code transcripts and identify commonalities of interest across interviews (Miles et al., 2014). Balanced with phenomenological sensitivity to the context in which statements were made, coding entailed noting the recurrence of key words and themes and highlighting statements regarding ethics and emotional AI. Analysis of the 33 transcripts that discussed non-identifying emotional AI followed an adaptive approach (Layder, 1998) to combine pre-existing theory (on non-identifying emotional AI, privacy and critical accounts of biometrics) with coded interview transcripts. These were hand-coded using Cresswell's (1994) approach that entails topic generation, abbreviation and creation of categories. Production of initial codes and thereafter themes were theoretically led by the interest in non-identifying emotional AI. However, in analysis, wider interests soon became apparent, clustered under the following themes: change; method; justification; privacy and regulation; and consensus.

Workshop

The second source of empirical understandings is a creative workshop organised by this paper's author at Digital Catapult, London (16/09/2016).³ It was conducted on the basis that privacy insights about emotional AI could be disclosed through peer-based discussion and activities. Following Veales's (2005) approach to creative workshops, 21 participants interested in data about emotions were asked to design ethical guidelines. This took the form of a list of do's and don'ts. Titled *Emotion Capture and Trust Workshop*, participants included representatives from the technology industry (n=10), the UK's Information Commissioner's Office (n=1), the UK's advertising self-regulator (Committee of Advertising Practice) (n=1), security companies (n=3), NGOs (n=1), a psychologist (n=1), legal ethicists (n=2) and critical

surveillance academics ($n=2$). To facilitate unconstrained discussion, there was no audio-visual recording. Instead, stakeholders designed a list of ‘dos and don’ts’ for working with data about emotions and outcomes in the form of A1-sized tear-off sheets which were kept and photographed. The next and last stage entailed the moderator (this paper’s author) summarizing each group’s view and collapsing these into a shorter set of ethical statements. These were enthusiastically debated and eventually agreed by participants.

Survey

This third tranche of data comprises a demographically representative UK nationwide survey ($n=2068$) conducted in November 2015. The UK was chosen because emotional AI is developing there apace.⁴ Closed-ended questions were used to gauge lay attitudes to potential uses of emotion detection employed in technologies and contexts that citizens are familiar with. These embraced its employment in social media for market research, reactive billboards in advertising, online games, interactive movies and voice-based search (see Table 1 for the specific questions asked). Utilising a multiple-choice format, response options were scaled to reflect industry practices and the research interest in anonymisation. Options were as follows: overall rejection of emotion capture practices (‘not OK’); acceptance of anonymised emotion capture practices (‘OK anonymised’); acceptance of identifying emotion capture (‘OK identifiable’) and a final category for those who do not have a view or do not understand the question (see Table 2 for response options in full).

The survey was executed online via ICM Unlimited, a commercial survey organisation. Online surveys have methodological caveats including difficulties of

presenting complex topics and minimal control over respondents’ condition (attentive or distracted). However, this approach generated a respectable weighted sample of geographical regions, age groups, social classes and gender, while avoiding social desirability bias. This is especially pertinent in privacy-related research that is recognised as having scope for bias (Zureik and Stalker, 2010). On complexity, participants were not provided further information than that contained in the question. This was deemed acceptable because the research was interested in lay responses to plainly stated propositions about emergent technologies.

Key findings from interviews and workshop: A weak consensus on privacy

There were notable differences of opinion on specific aspects and motives among stakeholders with a professional interest in emotional AI. Some of the themes established were to be expected, such as speed of technological *change*, defence and criticism of emotional AI *methods*, but others were less obvious, and it is on these that this paper focuses: namely, *privacy and regulation*; and *consensus*. Indeed, the surprise overall finding is one of *weak consensus* on need for privacy. This refers to a temporary alignment of privacy interests between diverse stakeholders. This alignment is not stable and is unlikely to last because while ethical and privacy interests currently align, the motives of stakeholders vary greatly.

All interviewees with commercial interests in profiling said that it is inevitable that machines will be employed to try to gauge feelings, emotions and intentions. They were confident that emotional AI would increase in scope and prevalence in 5–10 years. Gabi Zijdeveld from Affectiva, a sector-leading

Table 1. Closed-ended questions in UK public survey of emotion detection.

Question
1. Companies use social media data (from Facebook, Instagram, Tumblr, Flickr, Pinterest and others) to understand how we feel about brands. To do this they analyse images we post of ourselves, and others, for brand logos and facial expressions. They do this to understand what types of people are using their brands and what emotions they display. Which of the following best represents your feelings about this?
2. Advertising agencies have developed outdoor ads equipped with cameras that scan onlookers’ faces to work out our emotions towards the ad. If our reactions are not positive the ad changes itself to be more appealing. Which of the following best represents your feelings about this?
3. Electronic games companies are beginning to make use of cameras and wearable technologies to track players’ eyes, heart rates and respiration as they play games. They do this to heighten players’ involvement and entertainment levels. Which of the following best represents your feelings about this?
4. Entertainment companies are developing movies (particularly thrillers) that read viewers’ faces for emotional reactions when viewed on a smartphone, tablet or other device with a camera. They do this to personalize stories for viewers as the content unfolds and provide different storylines. Which of the following best represents your feelings about this?
5. Voice analytics companies collect information about moods and emotions from voice data (not just ‘what’ people say, but ‘how’ they say it). If applied to smartphones and voice-based search (rather than typing), this offers opportunity for devices that can react appropriately to emotions. Which of the following best represents your feelings about this?

Table 2. Response options in UK survey on public views of emotion detection.

Response options

1. I am not OK with my data being collected in this way.
2. I am OK with data collection about my emotions in this way as long as the information is anonymised and cannot be associated with me, my email address, phone number or any other possible means of personally identifying me.
3. I am OK with data collection about my emotional state in this way and OK for this data to be linked with personal information held about me.
4. Don't know.

company that uses facial coding to understand emotion, represents the overall view of interviewees from emotion-based companies. She says:

We believe that this tech will be ubiquitous in the future, maybe five years; we'll see a lot of the tech that we interact with on a daily basis using emotion. We see a lot of tech with human-tech interaction playing out in a digital context. With smart AI systems we need to understand how human emotions factor into this. We believe this is largely missing today and this is a negative thing. (Interview 2016)

This view was echoed in many other interviews with those from the emotion capture industry. For example, in early 2016, Yoram Levanon of Beyond Verbal, a company that specialises in extracting information about emotion from voice, speculated that in five to ten years emotional AI would be embedded in call centres and market research. After ten years, he says that the relationships between people and machines will become closer by means of emotional AI, such as through homecare by robots that understand people and respond appropriately. Others pointed to trends in appearance. For example, Kim Du from Emotiv, that makes EEG headwear, said that in five to ten years devices will become smaller and less visible. Looking further, she suggested that by twenty years sensors will be embedded not just in us, but in everything people interact with. Her vision is one where sensors at work and home feel-into, optimize and intuitively adjust settings, so people do not have to manually initiate things. Although applications differ, each interviewee working directly with psycho-physiological data foresees *ubiquitous feeling-into* by devices and environments. No interviewee deviated from this view.

Each interviewee recognized privacy as a concern yet had different perspectives on why it is an issue. Coming from a health background (although Emotiv's headwear is widely used in market and user experience research), Kim Du states:

It's your brain data. Right now there's no regulations, starting with security about how wearable companies

are collecting the data they're collecting about their users. It's an ethical issue given no policy guidance. (Interview 2016)

Surprisingly (given Silicon Valley's antipathy to regulation), Emotiv were not the only US firm to complain of a lack of clarity. No company sought extra legislation, but *absence of explicit rules* was a recurring theme. On privacy, other interviewees, such as Mike Ambinder from the biofeedback gaming company Valve, pointed to 'opt-in' processes of consent as a primary privacy tool, which supports his value-based argument that Valve intend to 'do right by Valve's users' and only use biometric technologies in a transparent manner to improve gaming experiences (Interview 2016).

A handful of commercially oriented interviewees see opportunity in privacy-by-design techniques. One smartphone app developer interested in emotions and video-calls stated a preference for 'edge computing' because 'Anything that is processed by the cloud entails resale of emotions and an after-market of emotions: this presents dangers' (Anonymised Interview 2016). However, importantly, the interest in privacy ethics is not just ethical, but also entails commercial *self-interest*. Commercially oriented interviewees recognize that alongside other parties, client perceptions need to be managed.

An interviewee from a global technology firm exemplifies the overall view of emotion capture businesses stating, 'the reason why emotion detection has not scaled as quickly as expected has less to do with the technology itself, but reticence of clients and their desire to avoid a PR backlash' (Anonymised Interview 2016). Gawain Morrison of Sensum⁵ echoes this, saying emotions are 'that final personal frontier' and 'Super-tech companies are paranoid about being seen to do the wrong thing, such as facial coding'. In reference to adding emotional AI to existing media, Morrison adds: 'No-one will hurt shareholders or their bottom line with an unnecessary bolt-on. It's a powerful tool and it will require a new company to take it into the setting before it's accepted. The old guys are untrusting' (Interview 2016).

The workshop held at Digital Catapult was primarily organised to create self-generated guidelines for technologists and businesses working with emotional AI. Attendees concluded with a range of overlapping suggestions. While focused issues were discussed regarding psychological suppositions (e.g. reliability of ‘basic emotions’), appropriate time length of data storage, racial and cultural difference regarding emoting, potential for citizen manipulation, the right to be forgotten, scope for pre-crime analytics, repurposing of collected data and user experience, the workshop delegates were asked to agree some basic rules. They said that ‘Do’s’ should include ‘put the person first’; ‘put them in control’; that ‘external and internal guidelines are necessary’ (i.e. regulations) and ‘autonomy and choice’. Unanimously, they also agreed ‘use of data about emotions should be proportionate to the goal’, and that ‘users’ benefit trumps commercial gain’. On ‘Don’ts’, they agreed that those in the business of emotion detection should ‘not be covert’ (regardless of whether identifying or not) and they should not use emotion as the ‘be and end all of profiling’. These are familiar liberal approaches to privacy, involving autonomy, control, capacity for management and transparency. Yet, given that participants were told that an overview of the findings would be published and disseminated by Digital Catapult, an important organization in the UK and European digital industry ecology, the uniformity and strength of recommendations is surprising.

Despite this overall emergent finding of a weak consensus on the need for privacy in emotional AI, there were *outlying attitudes*. These came from the retail and advertising sectors. Few interviewees here were queasy about emotion detection in public spaces if the data is not legally personal data. Realeyes, for example, saw no problem in teaming with the retailer Mothercare⁶ to attempt to track the emotional states of customers and biometric reactions to promotions, store layout and (clearly involving personal data) the behaviour of store assistants (Interview 2016). As also noted in McStay (2016), it was advertising companies (including representatives from Havas, Ogilvy & Mather, DataXu and M&C Saatchi) who argued most clearly that emotionally relevant advertising serves citizens as well as businesses (Interviews 2015–2016). Yvonne O’Brien of Havas, for example, sees ‘opportunities in biometrics and thereafter insight into emotional life’. Later in interview (but still in context of emergent technology), she asserted skepticism about privacy concerns because people continue to use services (such as social media) that are said to be privacy invasive. Rather, for O’Brien, citizens will willingly provide data about emotions to have meaningful interaction with brands (Interview 2016).

Perhaps surprisingly, the weak consensus on privacy (excluding advertising and retail) also includes NGOs working in data protection. While they raised concerns, they did not dismiss using emotions to interact with technologies if liberal privacy principles are respected. On fears, Gus Hosein of Privacy International (views not representative of organisation) observed that emotional AI makes us prone to subtle manipulation. Hosein cites state surveillance of populations and commercial interest in understanding purchase intention. This is realised in China as it uses emotional AI in public spaces such as airports, railway stations and other parts of smart city infrastructure (Wong and Liu 2019). Similarly, Dubai (or ‘Smart Dubai’), for example, analyses sentiment and retail spending trends and employs in-house psycho-physiological measures (using Emotiv headwear) to understand happiness in the region (Anonymised Interview 2016). On emotional AI and soft biometrics, Hosein says this ‘freaks me out’ and ‘we don’t have the legal frameworks’ (Interview 2015). On being asked about emotional AI that does not make use of personal data, he reasons ‘it is still taking something from me... it is still interacting with me... it is interfacing without my say so’.

Yet, Hosein also highlighted enabling aspects of new technologies. Key factors for him are ‘control’, and a ‘say on outcomes’. These principles align with industry views described earlier, indicating that the views of industry and pro-privacy groups are *not oppositional*. Hosein’s were not outlying comments: interviews with Jim Killock and Javier Ruiz from Open Rights Group in 2016 continued this theme. While sceptical of industry claims to anonymisation, on use of emotion detection when matched with identifiable data, Killock spoke of ‘legitimate services’, saying that there is an argument to be made for emotional AI and first-party relationships with services and vendors (also noted by Valve’s Ambinder).

A more antagonistic response was expected. Instead, Killock insisted that liberal principles of control, awareness, meaningful consent and better regulation by centralized institutions are required. An interview with Jeremy Gillula of the US Electronic Frontier Foundation (views not representative of organisation) found Gillula arguing that a person should be aware of what is being collected, adding: ‘If you’re aware that machines are tracking your emotions to make interactions better, that’s OK’. As an example, he says, ‘If my smartphone understands that I’m angry and can provide a response that calms me down, that’s OK’ (Interview 2016). Echoing the smartphone app developer, he said that what matters is that information is secure and that the device owner has control over whether data is shared to the cloud. His criticisms were less about commercial uses of emotional AI,

but governmental and policing awareness of moods (such as anger), accountability to devices that pertain to read emotions, algorithmic biases in interpreting of emotions, and how information about emotions might be used in courts, such as in a divorce subpoena.

On regulation, Gillula diverged from UK NGOs. Reflecting EFF's libertarian roots, he is wary of increased government regulation because technology evolves quicker than law. Yet, he adds that companies should be open to independent auditing of security and that they should have a legal duty to tell users how their data is processed. Omer Tene (Vice President of Research and Education of International Association of Privacy Professionals (IAPP)) made a similar argument. Stating that technological progress is inevitable, he suggests that at best privacy advocates can try to adjust social norms and policies to limit or moderate it (Interview 2016). Speaking of his own views rather than IAPP's, he says that it is difficult to see government curtailing development of emotional AI. Recognising potential privacy issues, he believes emotional AI will grow and be widely used. Tene says that new types of data-intensive technology always trigger privacy hot-spots, adding, 'One juncture you're bound to bump into is children'. This point was made in reference to emergent tracking of emotion in schools. Tene says that, 'I think there'll be angst from parents on the impacts on opportunities of their kids down the road'.

To summarise, the interviews and workshop found that the emotional AI sector is likely to grow rapidly within the next five or so years. A hard-core element (advertising) sees little wrong with collecting emotional data in bulk because it reasons that greater consumer understanding will lead to better service from marketers. However, the majority consensus evident across commercial, regulatory and NGO sectors is that more ethical industrial practices (such as opt-in consent) will win public trust; and that there is a creepy factor to be overcome. What, however, do citizens actually think of these emerging practices? Are they party to the emerging weak consensus?

UK survey findings

Reported feelings from UK citizens showed little variance across different media and technological forms (sentiment analysis, out-of-home advertising, gaming, interactive movies, voice and mobile phones). Similarly, gender, social class and region did not produce noticeable differences. As detailed in Table 3, the overall mean averages for *all* forms of emotional AI are as follows: half of UK citizens (50.6%) are 'not OK' with it in any form; almost a third (30.7%) are 'OK' with it if the application does not personally identify them; less than a tenth (8.3%) are 'OK' with having

Table 3. Overall UK citizen feelings in 2015 about emotion detection.

Statement	Number of people (n=2068)	Percentage
1 (Not OK)	1028	50.6%
2 (OK/no personal identification)	687	30.7%
3 (OK/personal identification)	163	8.3%
4 (Don't know)	190	10.4%

data about their emotions connected with personally identifiable information and a tenth (10.4%) do not know. At this level, there is no reason to suggest that UK citizens, taken as a totality, share the weak consensus among other stakeholders regarding privacy-friendly approaches to emotional AI.

However, as reported in Table 4, age was *the* sole factor in attitudinal differences to emotional AI. For example, the mean average of 18–24s 'not OK' with emotional AI is almost a third (31.4%), whereas the overall figure of all age groups is just over half (50.6%). This upward trend of being 'not OK' with emotional AI ends with the over 65s whose overall mean average is 66.3% 'not OK'. In effect, this means that the oldest people sampled are more than twice as likely to be 'not OK' with emotional AI than the youngest people. Conversely, if 'non-identifying OKs' and 'identifying OKs' for 18–24s are combined, this means 56.8% of 18–24s are 'OK' with some form of emotional AI. Looking to future generations, there is no obvious reason why this figure should decrease given increase in use of networked devices and services (Ofcom, 2019). However, this does not mean that younger people are 'OK' with having data about their emotions linked with personally identifying data. While older people are certainly less keen on being identified when having data about emotions processed (only 1.7% are 'OK' with this), 18–24s show little indication of being comfortable with emotion capture practices that identify them either (only 13.7% are 'OK' with this).

This paper speculates that younger people are more open to novel forms of engagement with technology but remain wary of identifying processes. On this basis, there is a case to be made that younger people might be party to a weak consensus on the basis of control-based accounts of privacy. Nevertheless, it should be kept in mind that the 18–24s are a small percentage of an overall suspicious citizenry. On attitudinal differences between age groups, the survey methodology shows its weaknesses. However, speculatively, the generation most open to emotion detection was born between 1991 and 1997, when the web emerged as a mass medium. Further, according to

Table 4. Using age to segment UK citizen feelings about emotion detection.

Statement	18–24 (n=248)	25–34 (n=331)	35–44 (n=393)	45–54 (n=351)	55–64 (n=310)	65+ (n=434)
1) Not OK	31.4%	38.7%	48.6%	48.1%	62.0%	66.3%
2) OK/no personal identification	43.1%	35.5%	29.0%	32.3%	25.2%	24.0%
3) OK/ personal identification	13.7%	13.9%	11.7%	8.0%	3.4%	1.7%
4) Don't know	11.8%	12.0%	10.7%	11.6%	9.5%	8.0%

UK findings by media regulator Ofcom (2019), this generation displays the highest levels of internet usage per week; is highly likely to use lots of websites/apps; is most likely to access the internet via smartphones and shows high overall levels of interactive media use. It is also most likely to be very confident about staying safe online, least likely to have read terms and conditions thoroughly, yet most likely to have changed social media settings of specific sites to be more private. Thus, although younger people are open to new experiences, this should not be mistaken for not caring about privacy.

Discussion: A critical window for the weak consensus

This paper has accounted for the growth of emotional AI that assesses bodies for indication of emotional states. Its first contribution is that it has pushed forward the privacy debate by diagnosing an over-emphasis on identification in data privacy regulation and omission of non-identifying soft biometric data about emotional life.

Its second is the finding that over half of UK citizens are ‘not OK’ with the principle of emotion detection (identifiable or otherwise) and that this has no current remedy in UK or European law. Especially, when privacy is seen in terms of bodily integrity (Nussbaum, 1999), this takes us beyond questions of personal data and identification, to dignity and the right for people to have self-determination over their own bodies. Privacy as dignity does not hinge on identification and, as a result, may be extended to groups, especially given interest in objectifying emotions without opt-in consent for commercial gain. On what should be done to promote dignity in relation to emotional AI and soft biometrics, if regulatory action is deemed necessary, there is scope to do this within the framework of GDPR as Article 9(4) grants capacity for Member States to maintain or introduce further conditions or limitations with regard to the processing of genetic data, biometric data or health data. As a result, argued through the prism of emotional AI and [soft] biometric data, this paper echoes Wachter’s call for ‘new protections based on holistic notions of “data

about people” and group conceptions of privacy’ (2020: 55).

Third and finally, there is the unusual consensus on privacy among industry and data protection NGOs, providing opportunity for regulatory change. Although older people (over-55s) were ‘not OK’ with emotional AI, and therefore not party to the weak consensus, younger people appear party to the weak consensus on the basis of a control-based account of privacy (that is, they are happy to interact with technologies but want a meaningful say over the process), also suggesting that antagonism will reduce in future decades. Interviewees working with voice, facial coding and body measurement data about emotions all said they ‘feel strongly’ about *opt-in* processes of consent. One might see this as an obvious tactic to brush-off a critical interviewer, but the scope for a high creepy factor means that interviewees also have to manage reputational concerns of existing and potential clients, as well as regulators. This means control-oriented opt-in is not just the dignity-enhancing thing to do, but also the self-interested thing to do. With the exception of the advertising sector, the convergence and consensus between industry interviewees and NGOs regarding identifying and non-identifying emotion detection was surprising. While data protection NGOs expressed concern about data privacy and social manipulation, they did not dismiss the premise of emotional AI if citizens have meaningful control over processes. However, again, this weak consensus only exists because the creepy factor associated with emotional AI currently acts as a brake on industry. The technology industry is keen to foreground privacy credentials while it is expedient, but the fragile alignment of interests is unlikely to last. Given that industrial stakeholders see that these technologies will become more popular in the early 2020s, this paper concludes that *now* is the time to take advantage of this unusual consensus by regulating use of both identifying and group data about emotions and related psycho-physiological behaviour.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was supported by the UK's Arts and Humanities Research Council [grant number AH/M006654/1].

ORCID iD

Andrew McStay  <https://orcid.org/0000-0001-8928-3825>

Notes

1. For legal examination of device-level empathic media, see Clifford (2017) who offers a thorough analysis of emotion capture in relation to the GDPR and ePrivacy Directive. Uniquely, he also draws attention to the Unfair Commercial Practices Directive, arguing for a precautionary approach to technologies that may unduly influence rational economic behaviour and decision-making.
2. I do not draw directly on each of the interviews in this paper, but they are utilized more fully in Emotional AI: The Rise of Empathic Media (McStay, 2018). However, they have indirectly shaped suppositions, assertions and argumentation in this paper.
3. Digital Catapult advises, nurtures and shares learning about start-up digital projects.
4. Excluding Northern Ireland.
5. Sensum uses multiple biometric means to understand audiences' emotional responses to various media.
6. The store specialises in products for expectant mothers and children.

References

- Andrejevic M (2013) *Infoglut: How Too Much Information Is Changing the Way We Think and Know*. New York: Routledge.
- Article 29 Data Protection Working Party (2012) Opinion 3/2012 on developments in biometric technologies. Available at: http://ec.europa.eu/justice/data-protection/article-29/documentation/opinion-recommendation/files/2012/wp193_en.pdf (accessed 30 October 2019).
- Barrett LF, Adolphs R, Marsella S, et al. (2019) Emotional expressions reconsidered: Challenges to inferring emotion from human facial movements. *Psychological Science in the Public Interest* 20(1): 1–68.
- Bertaux D (1981) From the life-history approach to the transformation of sociological practice. In: Bertaux D (ed.) *Biography and Society: The Life History Approach in the Social Sciences*. London: Sage, pp.29–45.
- Bryson JJ (2018) Patency is not a virtue: The design of intelligent systems and systems of ethics. *Ethics and Information Technology* 20: 15–26.
- Clifford D (2017) Citizen-Consumers in a Personalised Galaxy: Emotion Influenced Decision-Making, a True Path to the Dark Side? CiTiP Working Paper Series, 31/2017, SSRN [online]. Available at: <https://ssrn.com/abstract=3037425> (accessed 23 January 2020).
- Cresswell JW (1994) *Research Design: Qualitative and Quantitative Approaches*. Thousand Oaks: Sage Publications.
- Davies W (2015) *The Happiness Industry: How the Government & Big Business Sold Us Wellbeing*. Verso: London.
- Davies W (2016) How are we now? Real-time mood-monitoring as valuation. *Journal of Cultural Economy* 10(1): 34–48.
- Edwards L (2016) Privacy, security and data protection in smart cities: A critical EU law perspective. *European Data Protection Law Review*. Available at: papers.ssrn.com/sol3/papers.cfm?abstract_id=2711290 (accessed 30 October 2019).
- Ekman P and Friesen WV (1971) Constants across cultures in the face and emotion. *Journal of Personality and Social Psychology* 17(2): 124–129.
- Emotional AI (2018) Useful links. Available at: emotionalai.org/useful-links (accessed 30 October 2019).
- European Commission (2016) REGULATION (EU) 2016/679 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). Available at: ec.europa.eu/justice/data-protection/reform/files/regulation_oj_en.pdf (accessed 30 October 2019).
- European Commission (2017) Proposal for a regulation on privacy and electronic communications. Available at: ec.europa.eu/digital-single-market/en/news/proposal-regulation-privacy-and-electronic-communications (accessed 30 October 2019).
- European Data Protection Supervisor (2015) Opinion 4/2015 Towards a new digital ethics Data, dignity and technology. Available at: https://edps.europa.eu/sites/edp/files/publication/15-09-11_data_ethics_en.pdf (accessed 23 January 2020).
- Floridi L (2014) Open data, data protection, and group privacy. *Philosophy & Technology* 27: 1–3.
- Floridi L (2016) On human dignity as a foundation for the right to privacy. *Philosophy & Technology* 29: 307–312.
- Grandey AA, Diefendorff JM and Rupp DE (2013) Bring emotional labor into focus. In: Grandey AA, Diefendorff JM and Rupp DE (eds) *Emotional Labor in the 21st Century: Diverse Perspectives on Emotion Regulation at Work*. New York: Routledge, pp.3–27.
- Gunes H and Pantic M (2010) Automatic, dimensional and continuous emotion recognition. *International Journal of Synthetic Emotions* 1(1): 68–99.
- Layder D (1998) *Sociological Practice: Linking Theory and Social Research*. London: Sage.
- Leys R (2011) The turn to affect: A critique. *Critical Inquiry* 37(3): 434–472.
- Lupton D (2016) *The Quantified Self*. Cambridge: Polity.
- Lynskey O (2015) *The Foundations of EU Data Protection Law*. Oxford: Oxford University Press.
- McDuff D and el Kaliouby R (2017) Applications of automated facial coding in media measurement. *IEEE Transactions on Affective Computing* 8(2): 148–160.

- McStay A (2016) Empathic media and advertising: Industry, policy, legal and citizen perspectives (the case for intimacy). *Big Data & Society* 3(2): 1–11.
- McStay A (2018) *Emotional AI: The Rise of Empathic Media*. London: Sage.
- McStay A and Urquhart L (2019) ‘This time with feeling?’ assessing EU data governance implications of out of home appraisal based emotional AI. *First Monday*. Available at: firstmonday.org/ojs/index.php/fm/article/view/9457/8146 (accessed 23 January 2020).
- Miles MB, Huberman AM and Saldana J (2014) *Qualitative Data Analysis*. London: Sage.
- Mittelstadt B (2017) From individual to group privacy in big data analytics. *Philosophy & Technology* 30: 475–494.
- Neff G and Nafus D (2016) *Self-Tracking*. Massachusetts: MIT Press.
- Nussbaum MC (1999) *Sex and Social Justice*. Oxford: Oxford University Press.
- Ofcom (2019) Children and parents: Media use and attitudes report. Available at: ofcom.org.uk/__data/assets/pdf_file/0024/134907/children-and-parents-media-use-and-attitudes-2018.pdf (accessed 30 October 2019).
- Picard RW (1997) *Affective Computing*. Cambridge: MIT.
- Picard RW (2007) Toward machines with emotional intelligence. Available at: affect.media.mit.edu/pdfs/07.picard-EI-chapter.pdf (accessed 30 October 2019).
- Russell JA (1994) Is there universal recognition of emotion from facial expression? A review of the cross-cultural studies. *Psychological Bulletin* 115(1): 102–141.
- Sánchez-Monedero J and Dencik L (2019) The politics of deceptive borders: ‘Biomarkers of deceit’ and the case of iBorderCtrl. *arXiv*. Available at: arxiv.org/abs/1911.09156 (accessed 3 January 2020).
- Stark L and Crawford K (2015) The conservatism of emoji: Work, affect, and communication. *Social Media + Society* 1(2): 1–11.
- Turow J (2017) *The Aisles Have Eyes: How Retailers Track Your Shopping, Strip Your Privacy, and Define Your Power*. New Haven: Yale.
- Veale A (2005) Creative methodologies in participatory research with children. In: Greene S and Hogan D (eds) *Researching Children’s Experience: Approaches and Methods*. London: Sage, pp.253–272.
- Wachter S (2019) Data protection in the age of big data. *Nature Electronics* 2(1): 6–7.
- Wachter S (2020/forthcoming). Affinity profiling and discrimination by association in online behavioural advertising. *Berkeley Technology Law Journal* 35(2). Available at: papers.ssrn.com/sol3/papers.cfm?abstract_id=3388639 (accessed 30 October 2019).
- Whittaker M, Crawford K, Dobbe R, et al. (2018) AI Now Report 2018. Available at: ainowinstitute.org/AI_Now_2018_Report.pdf (accessed 30 October 2019).
- Wiewiórowski W (2019) Facial recognition: A solution in search of a problem? European Data Protection Supervisor. Available at: edps.europa.eu/node/5551 (accessed 30 October 2019).
- Williamson B (2017) Moulding student emotions through computational psychology: Affective learning technologies and algorithmic governance. *Educational Media International* 54: 267–288.
- Wong S-L and Liu Q (2019) Emotion recognition is China’s new surveillance craze. *Financial Times*. Available at: ft.com/content/68155560-fbd1-11e9-a354-36acbbb0d9b6 (accessed 3 November 2019).
- Zureik E and Stalker LLH (2010) The cross-cultural study of privacy. In: Zureik E, Stalker LLH, Smith E, et al. (eds) *Surveillance, Privacy and the Globalization of Personal Information*. Montreal: McGill-Queen’s University Press, pp.8–30.