Article | Platform Biometrics

Jeremy W. Crampton

Newcastle University, UK
jeremy.crampton@ncl.ac.uk

Abstract

This article identifies and analyses the emergence of platform biometrics. Biometrics are measurements of behavioral and physical characteristics, such as facial expressions, gait, galvanic skin response, and palm or iris patterns. Platform biometrics not only promise to connect geographically distant actors but also to curate new forms of value. In this piece I describe Microsoft Face, one of the major facial biometric systems currently on the market; this software promises to analyze which of seven “universal” emotions a subject is experiencing. I then offer a critique of the assumptions behind the system. First, theories of emotion are divided on whether emotions can be reliably and measurably expressed by the face. Second, emotions may not be universal, nor are there likely only seven basic emotions. Third, I draw on the work of Rouvroy and Berns (2013) to identify emotion-recognition technologies as a classic example of algorithmic governance. To outcome algorithmic governance is to enable the subject to creation and govern surveillance. Platform biometrics will therefore provide a key component of surveillance capitalism’s appropriation of human experience (neuro-liberalism).

iBorderCtrl

In 2016, the European Union awarded €4.5m to a consortium of researchers and industry partners to develop strengthened border-control technologies. The winning bid, which calls itself iBorderCtrl, proposes to use an AI-driven avatar that will conduct an interview with the traveler while assessing facial biometrics of micro-expressions to determine whether the person is lying or telling the truth. If the traveler is deemed to be lying, then the avatar will become more stern and the tone of its voice will start to question the traveler more sharply. The avatar will refer those who are believed to be lying to a human for further interviews, while those believed to have told the truth will be allowed to pass through the border. In 2018, the system was trialed at borders in Hungary, Latvia, and Greece (Kelly 2018).

Developed at Manchester Metropolitan University, the Automated Deception Detection System (ADDS) in iBorderCtrl is based on an earlier system called Silent Talker. As is typical of facial emotion recognition technologies, Silent Talker draws on the controversial work of Paul Ekman, a psychologist who claimed to have isolated the seven “universal emotions” (happiness, sadness, anger, surprise, fear, disgust, and contempt) in facial expressions (Ekman, Friesen, and Hagar 2002). Ekman believed that “micro-expressions” or what the iBorderCtrl project calls “biomarkers of deceit” could reveal a person’s true feelings no matter how closely they tried to control their face. As a proponent of facial recognition once put it: you can uninstall apps on your phone, but you can’t turn off your face (Warnock 2015).

Platform Biometrics

The iBorderCtrl project is an instance of the rise of “platform biometrics.” Biometrics are measurements of behavioral and physical characteristics, including facial expressions, gait, galvanic skin response, palm or iris patterns, and many others. As I shall show, biometrics aim not only to personally identify individuals but also to detect and draw actionable inferences about personality, intent, emotional state, social conformity, sexual orientation, and many other formerly private attributes, positing that they manifest bodily. The term “platform” has diverse meanings, including a computer architecture but also a politics (Gillespie 2013). As such, platforms do more than link online entities together; they curate this connectivity to create new forms of value (Richardson forthcoming). Biometrics aim to extract “behavioral surplus,” or the human experience not necessary to keep apps and software running, which can then be traded via the platform (Zuboff 2019). Biometrics are a vital component of such platform surveillance.

By understanding biometrics in this sense, we can understand the huge leaps that have taken surveillance beyond simply seeing and recognizing, to categorizing and inferring a subject’s innermost nature. Whether this is done fairly or not—whether the machine learning and the algorithms are based on sound science and data—is to a large extent black-boxed by these platforms’ claimed egalitarian structure (Pasquale 2015).

In this paper I take as a case study Microsoft’s Azure “Face Cognitive Service.” MS Face offers facial emotion recognition that categorizes the array of facial expressions according to the seven universal emotions. However, MS Face is only one of many platform biometrics solutions which, together, are instituting new regimes of algorithmic governance. Other major systems include Amazon’s Rekognition, as well as Palantir, WeSee, Affectiva, iMotions, Face++, Facebook, Siri, and Amazon Alexa. Both Amazon and Palantir have sold behavior-prediction technology to US Immigration and Customs Enforcement (ICE) and police departments in the United States (Hao 2018). In one widely reported test, the American Civil Liberties Union (ACLU) found that Rekognition incorrectly matched 28 faces of members of Congress with people in a crime database (Snow 2018). Pointing to racial bias in crime data, the ACLU has called on Amazon to stop selling Rekognition to communities vulnerable to surveillance (Dvoskin 2018).

In its review of many of these systems, the AI Now Institute warned of the dangers these systems present to mass surveillance due to its faulty emotion and affect recognition, its lack of scientific evidence, and its irresponsible ethics. In their words, “Affect recognition attempts to read inner emotions by a close analysis of the face and is connected to spurious claims about people’s mood, mental health, level of engagement, and guilt or innocence” (AI Now 2018: 8). This followed a May 2016 report from the US Government Accountability Office (GAO) about the FBI’s use of face-recognition technology. The GAO said that the FBI has not released a Privacy Impact Report for five years prior to 2018 despite being required to do so by law. The GAO also revealed, for the first time in public, that the FBI could access over 412 million images, mostly acquired through special arrangements with state driver licensing offices or police departments. A recent study of the reliability of facial recognition revealed that the software interprets emotions differently, based on the person’s race (Rhue 2018). The study ran 400 pictures of basketball players through two emotion recognition software services, Face++ and MS Face. It found that no matter how much black players were smiling, they were assigned more negative emotions.

Taken together, these findings are concerning because research has determined that as few as four geolocation points are enough to uniquely identify 95 percent of individuals in a crowd (de Montjoye et al. 2013). When coupled with affect recognition, the prospect of surveillance at a scale that identifies individuals and their innermost personality becomes a reality.

How It Works

The algorithms behind MS Face and its competitors rest on the Facial Action Coding System (FACS) developed by the US psychologist Paul Ekman in 1978 and revised by him, Wallace V. Friesen, and Joseph

C. Hager (Ekman, Friesen, and Hager 2002). FACS breaks down the face into a lengthy list of categories of movement or position known as Action Units (AU). These AUs, for example “lip suck” or “lid droop” are then scaled according to criteria of intensity from A (trace) to E (maximum). Figure 1 (from Microsoft competitor iMotions) shows two Action Units that comprise the facial expression for joy (AU 6 “cheek raiser” plus AU 12 “lip corner puller”). For a detailed account of FACS, see Gates (2011).

Figure 1. AU 6 (L) and AU 12 (R). Source: iMotions.com

FACS is based on the work of the Swedish anatomist Carl-Herman Hjortsjö (1914-1978), who is best known today for his book *Man's Face and Mimic [Facial] Language* (Hjortsjö 1970). Hjortsjö developed a systematic taxonomy of the muscular movements that create facial expressions. Although his book acknowledges the prior work of others in drawing together “both psyche and soma” (Hjortsjö 1970: 8), including Ernst Kretschmer (1888-1964), Johan Kaspar Lavater (1741-1801), and the anatomist Guillaume Benjamin Duchenne (1806-1875), Hjortsjö found their work to be far too generalized and not based on a specific system. As Gates has argued, this earlier work cannot easily be divorced from outmoded and pseudo-scientific theories such as phrenology; and in fact is part of the history of the pathologizing of deviation (Gates 2011).

Hjortsjö developed a set of 24 facial emotional expressions to isolate what he called the basic “letters” of facial expressions: an uplifted eyebrow or downturned mouth. These could then be combined to form the “words” of emotion, such as fear, skepticism, or anger. He was interested in the relations between somatic/bodily expression, emotion, and mental constitution (see Figure 2).

Hjortsjö was aware that his work had apparent connections to physiognomy, the study of the relation between facial expression and the qualities of a person, which is now recognized as a pseudo-science. In rejecting physiognomy and phrenology, Hjortsjö said that the mistake they made was to infer not a general disposition but specific qualities.

Today, however, Hjortsjö’s warning that the system may not work at the individual level is largely ignored; indeed, the initial goal of facial emotion recognition systems is to provide an analysis of a specific individual. From here, digital media platforms shape and make legible subjects of digital control “across multiple types of technologically mediated scales” (Stark 2018: 207). In 2018, for example, Snapchat was granted a US patent for a notional system for “representing emotions of attendees...to determine an aggregate mood of the attendees of the event” using selfies taken at festivals, concerts, or political events (Chang 2018). Although their patent referred to mobile technologies, such as smartphones, the concept could adequately cover sufficiently high-resolution CCTVs in other public spaces (a march in Washington DC, a Parisian boulevard, or perhaps Tiananmen Square) and be coded to detect an uptick in anger or fear.

Ekman’s work often underpins the literature on computer vision, for example claiming that facial expressions provide a “direct link to the emotional state experienced by the sender” (Huang et al. 2016). In other words the “face is equipped to lie the most and leak the most” (Ekman and Friesen 1969); the face is where one’s innermost nature plays out, whether intentional or not. This theory has come to be known as the Basic Emotion model (Fischer 2013).

Figure 2. Hjortsjö's examples of facial expressions. Source: Hjortsjö (1970).

MS Face

The MS Face service is part of Microsoft's suite of cloud-based services known as Azure and is their response to Amazon Web Services (AWS). Microsoft claims that over 95% of Fortune 500 companies use Azure. MS Face represents a significant entry into the combination of cognitive or psychological assessments and artificial intelligence or algorithmic psychometrics (Stark 2018). Perhaps the most notable example is the case of Cambridge Analytica, which revived psychographic data profiling. Following the election of Donald Trump, the company claimed to have been a factor in his election in that it categorized people not just on the basis of their demographics or where they lived (geodemographics) but also by their scores on psychological personality scales obtained from Facebook and third-party data brokers. As Tufekci (2014) noted, such granular target marketing is becoming a central concern of such computational politics.

Given a target image (or frame from a video), MS Face will draw on its training data set to match the facial expression to one of the universal emotions (plus "neutral"). In Figure 3, I ran MS Face on an image of Paul Ekman; the results indicate 91% confidence happy, and 8.7% confidence neutral. These results are meaningful, Microsoft argues, because "emotions are understood to be cross-culturally and universally communicated with particular facial expressions" (Microsoft 2018); they also conform to Ekman's Basic Emotion model. I return to this claim below.

Figure 3. Microsoft Azure's Face classifier for image of Paul Ekman. Source: Microsoft (2018).

MS Face, like all other facial recognition technologies, is trained on special datasets that teach the algorithms the desired outputs. Known as machine learning (ML), the training has the goal of identifying and extracting patterns from the data. Most ML algorithms work by using either supervised or unsupervised learning. In supervised learning, the goal is to get the algorithm to map values from an example instance to the value of another attribute, known as the target attribute, of that example. So, in the case of emotion recognition, the goal is to find the best match between the facial expression and a value (one of the seven universal emotions). Given that there are so many possible ways the face can express itself, the algorithm cannot try all possible matches. Instead, it employs preferences, known as the algorithm's learning bias (Kelleher and Tierney 2018). Identifying which of these learning biases will work best for a data set is accomplished experimentally by trying different algorithms. The use of open source training data sets makes this challenge considerably more tractable for data scientists. Images in these datasets have already been categorized and indexed. One of the earlier such datasets, known as the Carnegie-Mellon Pose, Illumination and Expression (PIE) dataset, was created in 2000 (Sim, Baker, and Bsat 2003). It contains images of 68 individuals, in different poses, for a total of about 41,000 images. Datasets such as these have become the "Psych 101" of vision machine learning.

In unsupervised learning, the algorithms do not have a target attribute; therefore, they save time in not necessitating the creation of an indexed and categorized training dataset. Instead, the algorithms are tasked with finding regularities in the data, which is known as cluster analysis. Clusters are groupings of similar inputs (perhaps facial expressions) that are found to occur together. The goal of cluster analysis is to reduce the total variation in a dataset to a more manageable number (to the seven emotions, for example). To use the FACS model as an example, the cluster algorithm may find that AU-12 and AU-6 occur together very often.

In both supervised and unsupervised learning, one of the most powerful applications of the result is to make a prediction when working on new instances that are not in the training data set. In the case of facial emotional recognition, then, the trained computer model will predict which emotion the subject is experiencing. The emotion prediction model can also be required to provide a reliability estimate of how confident it is in the prediction (typically by using weighted linear regression).

The Microsoft Face API (application programming interface) is a cloud-based solution that analyses images (still or video) and "returns their emotions"; in this case, the six Ekman emotions plus neutral. The API cannot currently be modified for supplemental emotions. Users can subscribe to their services for a fee (typically \$1 per 1,000 API calls). Face API is part of Azure Cognitive Services (formerly Project Oxford), by which it means artificial intelligence and machine learning—in essence algorithmic models. These services often have applications for identifying the user's current and future condition; they also provide access to "smart" technologies developers may be too costly or time consuming to do themselves. For

example, Face API offers to detect faces in the image and return their position within that image, pick out facial “landmarks,” such as face shape, nose and eyes, find similar-looking faces from a dataset, group similar faces together, identify if two images are of the same person, and so on. Their speech API offers to identify the intent of the user (for example, to make a phone call). As shown in Figure 3, Face API returns an identification (“happiness” or “anger”) and a degree of confidence between 0 and 1. Taken as a whole, Microsoft Azure Cognitive Services, including its Face API, is an example of a platform as these services connect users and software services, curate value (the output of the algorithms), and enable that value to be placed on a marketplace or otherwise made actionable. In their literature, Microsoft slides easily from “face recognition” to “emotion recognition” within the same API. But there is a huge step between these two things.

Why Emotion Recognition Is Such a Huge Step

The step from facial recognition to emotion recognition and from making inferences about the nature of a subject’s feelings or nature might seem an insignificant one. But we should be very careful about accepting this move. There are several reasons for this.

First, there is the question of whether emotion is meaningfully and measurably expressed on the face. In fact, there are competing theories of emotion and how these might manifest in the body. Does an emotion first occur consciously in response to a stimulus (stimulus-conscious emotion), or does it occur subsequently following an autonomic nervous system response? Although the first option may appear intuitive (e.g., we see a snake, feel fear, and then have bodily reaction; e.g., shaking, or flight) it is not supported by evidence. If the second option (first advanced in the nineteenth century by William James and Carl Lange) is correct, then our bodies react first (e.g., by shaking) and then we, ourselves, interpret each different type of reaction as a different emotion (Cannon 1927).

The James-Lange theory raises two impediments for automated affect recognition. On one hand, it implies that we learn to associate emotions from different bodily reactions and, therefore, presumably could learn them differently, in different cultural contexts. (Indeed, this is exactly what a competing theory, the Lazarus-cognitive mediational theory, proposes.) On the other hand, it would still be critical to reliably physiologically disambiguate between different emotional responses, such as fear and anger. However, making this distinction has proven very difficult in practice: “there is no one-to-one correspondence between a behavior and an emotion category” (Barrett 2012: 414).

Although the body may not exhibit clear and distinct boundaries between emotions, perhaps it is nevertheless true that emotions are clearly separable? This is a necessary assumption for facial emotion recognition because if it were not true, then the goal of identifying which emotion was being felt would fall apart. And this is so even if we grant that subjects may experience more than one emotion at a time, as with MS Face’s assessment of Ekman in Figure 3. However, there is evidence that facial expressions can themselves affect the experience of emotions. For example, when instructed to wrinkle their noses, subjects reported odors to be more unpleasant. In a recent paper exploring this facial-feedback theory, the author noted the effects of injections of botox (which has a temporarily paralyzing effect on the facial muscles) significantly altered subjects’ ratings on a standard mood scoresheet (Lewis 2012; see also Davis et al. 2010). In other words, far from reflecting emotions in the face, the face is at least in part responsible for *generating* them. Modern theories of emotion have emphasized that emotions require the perceivers of emotional arousal to have a socially shared set of conceptual knowledge (Barrett 2012). This is because emotions are both socially constructed and biologically manifest.

The second question is whether emotion is universal. A claim advanced by Microsoft’s Face and other emotion recognition literature is that not only are emotions discriminable but they are also universal. But what does this mean? Does it mean that there are the seven, among many others, that are experienced by all humans (what we might call weak universality)? Or does it mean that these emotions are always expressed the same way, on the face, by all human cultures (and perhaps even non-human primates; strong universality,

also known as discrete emotion theory)? Ekman's argument is for the latter (1992). In one study, for example, Ekman and colleagues presented 35-mm slides of 18 Caucasian faces to subjects in ten different cultures and asked them to assess these faces in terms of the seven universal emotions (Ekman et al. 1987). The results indicated very high (mean >90 percent) intersubjective agreement for happiness, surprise, and sadness, with high (60 percent–91 percent) agreement for fear, anger, and disgust. Interestingly, and not discussed by the authors, their data show that the those of the Japanese culture had the lowest agreement.

This latter finding comports with other research that shows differences between “closed” expressions of emotion (such as the hidden smile, a common practice among women in Japan and some other Asian countries, where the hand is placed to obscure the mouth) and “open” expression of emotions. In other cultures and situations, it is also taboo to “drop” a smile (consider a job interview, for example). Thus, even in open-emotion cultures, such as in North America or Europe, people may be socially pressured to smile (as many women can no doubt attest). Many other cultural differences have been observed that challenge the strong universality model, such as the valence placed on anger—is it a “good” emotion (righteous anger, often associated with men) or a “bad” emotion (loss of control and rational thought)? However, it is also unclear whether the weak universality model holds true. Later in his career, Ekman himself hypothesized that there were further basic emotions beyond the basic emotions and that the face did *not* contain the totality of nonverbal communication (voice tone and gestures are also important channels of communication). More recent psychological research indicates there may be as many as 27 distinct emotions, not all of which will manifest on the face (Cowen and Keltner 2017). Anthropologists and psychologists also dispute that there is no cross-cultural variation in the interpretation of facial expressions (Barrett 2017). For example, smiling (teeth-bearing) in our closest primate relatives often constitutes an act of aggression.

The third point is a more conceptual one, but perhaps the most important of all. Surveillance is not simply observation, but a process of assigning identity: you are not seen but rather are *seen as*...terrorist, threat, black, etc. The particular kind of surveillance we're concerned with here, then, is not just that surveillance that recognizes an individual (sometimes known as personally identifying information or PII). Facial recognition technologies capture an image of a face and match it to a database of known individuals. Both the UK and the US maintain such a database. In the UK, in January 2018, Baroness Williams (Minister for Countering Extremism) revealed there were 12.5 million searchable facial images in the Police National Database, many of whom were not involved in criminal proceedings (HM Government 2018). As mentioned earlier, in the US, the FBI can access over 412 million images, most of which can be acquired through state driver licensing offices or police departments. Further, the Bureau did not release privacy impact reports over a five-year period, despite being required to do so by law.

Drawing on the work of Rouvroy and Berns (2013) we can identify platform biometrics as a form of “algorithmic governance” that can bring together automated regimes of decision-making with new forms of data surveillance and modeling. These authors identify three successive stages:

1. Big Data are collected to form “data doubles” or models. For example, facial movements become statistical data. This is a process which both separates the individual from their data (data extraction) but is also a representation of the individual (modeling).
2. Knowledge is then produced via automated procedures from these data. For example, machine learning is developed from training data sets of facial expressions to characterize them as one or more of the seven universal emotions. The goal here is the “absolute objectivity” to score against metric indices that can be widely applied. Categories of emotion (e.g., anger, joy) are used to impute how an individual is feeling by how well that individual scores against the metric.
3. Action is taken on behaviors: the goal here is to anticipate and, if necessary, modify individuals' behaviors. For example, in a public space facial expressions that are marked as “angry” are tagged for further surveillance; facial expressions tagged as “happy” are surveilled less.

It is this third step that poses the most concerning developments. Facial emotion recognition technologies, and platform biometrics, more generally, perform a specific governmental function. A platform curates the creation of value and via surveillance extracts behavioral surplus. Emotion recognition biometrics, such as MS Face, do so by adopting the Basic Emotion model to assert actionable (valuable) emotion scores. There is already work that claims to identify ontological aspects of personality or identity using biometrics, such as whether subjects can be identified as gay, or straight (Wang and Kosinski 2018), or criminal (Wu and Zhang 2017). In China, police have already tested “facial recognition glasses” that help identify criminals (by identifying them against a database) and have begun the implementation of facial recognition in schools, where students are scanned every thirty seconds to gauge which of the seven emotions they are displaying in order to tell if they are paying attention (Connor 2018; Lee 2018). China’s well-known social credit system, although not yet fully implemented, will also reputedly draw on biometric surveillance (Ma 2018). One might reasonably estimate that what occurs in China will re-occur in western democracies within five to ten years.

Biometrics is also forming an important part of the smart city movement, aimed at making cities healthier and more efficient. Among other proposals, it could be used to monitor employee well-being, or detect shopper frustrations at long check-out lines, or anger in crowds via sentiment analysis of social media (see Gates 2011 for a range of examples).

Conclusion

Although digital surveillance is not new, I would argue that platform biometrics and emotional recognition technologies, such as Microsoft’s Face API, bring Rouvroy and Berns’ third stage, that of the creation of the surveillantly governed subject, into full fruition. It is at this stage that a shift occurs from surveillance as epistemology (what is seen is known) to surveillance as ontology (what is “seen as” comes into being). Given the likelihood of the continued field deployment of biometrics, careful disambiguation of their social impacts needs to be parsed out.

References

- AI Now. 2018. *AI Now Report 2018*. New York: AI Now Institute.
- Barrett, Lisa Feldman. 2012. Emotions are real. *Emotion* 12 (3):413-429. doi: 10.1037/a0027555.
- Barrett, Lisa Feldman. 2017. *How Emotions are Made: The Secret Life of the Brain*. Boston: Houghton Mifflin Harcourt.
- Cannon, Walter B. 1927. The James-Lange Theory of Emotions: A Critical Examination and an Alternative Theory. *The American Journal of Psychology* 39 (1/4):106-124.
- Chang, Sheldon. 2018. *Determining a Mood for a Group*, edited by US Patent Office. USA: Snap Inc.
- Connor, Neil. 2018. Chinese School Uses Facial Recognition to Monitor Student Attention in Class. *Daily Telegraph*. <https://www.telegraph.co.uk/news/2018/05/17/chinese-school-uses-facial-recognition-monitor-student-attention/>
- Cowen, Alan S., and Dacher Keltner. 2017. Self-Report Captures 27 Distinct Categories of Emotion Bridged by Continuous Gradients. *Proceedings of the National Academy of Sciences* <https://doi.org/10.1073/pnas.1702247114>.
- Davis, Joshua Ian, Ann Senghas, Fredric Brandt, and Kevin N. Ochsner. 2010. The Effects of BOTOX Injections on Emotional Experience. *Emotion* 10 (3):433-440.
- de Montjoye, Yves-Alexandre, César A. Hidalgo, Michel Verleysen, and Vincent D. Blondel. 2013. Unique in the Crowd: The Privacy Bounds of Human Mobility. *Scientific Reports* 3: 1376.
- Dwoskin, Elizabeth. 2018. Amazon Is Selling Facial Recognition to Law Enforcement—for a Fistful of Dollars. *The Washington Post*, May 22. https://www.washingtonpost.com/news/the-switch/wp/2018/05/22/amazon-is-selling-facial-recognition-to-law-enforcement-for-a-fistful-of-dollars/?utm_term=.86b03a385975 [accessed January 15, 2019].
- Ekman, Paul. 1992. An Argument for Basic Emotions. *Cognition and Emotion* 6 (3-4):169-200.
- Ekman, Paul, and Wallace V. Friesen. 1969. Nonverbal Leakage and Clues to Deception. *Psychiatry* 32 (1): 88-106.
- Ekman, Paul, Wallace V. Friesen, and Joseph C. Hager. 2002. *Facial Action Coding System*. Salt Lake City, UT: Research Nexus.
- Ekman, Paul, Wallace V. Friesen, Maureen O’Sullivan, Anthony Chan, Irene Diacoyanni-Tarlatzis, Karl Heider, Rainer Krause, William Ayhan LeCompte, Tom Pitcairn, Pio E. Ricci-Bitti, Klaus Scherer, Masatoshi Tomita, and Athanasios Tzavaras. 1987. Universals and Cultural Differences in the Judgments of Facial Expressions of Emotion. *Journal of Personality and Social Psychology* 53 (4):712-717.
- Fischer, Shannon. 2013. About Face: Emotions and Facial Expressions May Not Be Related. *Boston Magazine*, June 25.
- Gates, Kelly A. 2011. *Our Biometric Future. Facial Recognition Technology and the Culture of Surveillance*. New York: NYU Press.

- Gillespie, Tarleton. 2013. The Politics of “Platforms.” In *A Companion to New Media Dynamics*, edited by John Hartley, Jean Burgess, and Axel Bruns, 407-416. Malden, MA: Wiley-Blackwell.
- Hao, Karen. 2018. Amazon is the Invisible Backbone Behind ICE’s Immigration Crackdown. *MIT Technology Review*, October 22.
- Hjortsjö, Carl-Herman. 1970. *Man’s Face and Mimic Language*, translated by W.F. Salisbury. Lund, Sweden: Studentlitteratur.
- HM Government. 2018. *Biometrics Strategy and Forensic Services*. London: HM Government.
- Huang, Xiaohua, Jukka Kortelainen, Guoying Zhao, Xiaobai Li, Antti Moilanen, Tapio Seppänen, and Matti Pietikäinen. 2016. Multi-Modal Emotion Analysis from Facial Expressions and Electroencephalogram. *Computer Vision and Image Understanding* 147:114-124. doi: <http://dx.doi.org/10.1016/j.cviu.2015.09.015>
- Kelleher, John D., and Brendan Tierney. 2018. *Data Science, MIT Press Essential Knowledge*. Cambridge, MA: MIT Press.
- Kelly, Ross. 2018. “Lie Detecting” AI Border Guards to be Deployed at EU Airports. *Digit*. <https://digit.fyi/ai-border-guards-eu-ibordercntrl/> [accessed November 15, 2018].
- Lee, Don. 2018. At this Chinese School, Big Brother Was Watching Students—and Charting Every Smile or Frown. *LA Times*. <http://www.latimes.com/world/la-fg-china-face-surveillance-2018-story.html>
- Lewis, Michael B. 2012. Exploring the Positive and Negative Implications of Facial Feedback. *Emotion* 12 (4): 852-859.
- Ma, Alexandra. 2018. China Has Started Ranking Citizens with a Creepy “Social Credit” System—Here’s What You Can Do Wrong, and the Embarrassing, Demeaning Ways They Can Punish You. *Business Insider*, 29 October. <https://www.businessinsider.com/china-social-credit-system-punishments-and-rewards-explained-2018-4> [accessed December 1, 2018].
- Microsoft. 2018. Azure Face. <https://azure.microsoft.com/en-us/services/cognitive-services/face/>. [accessed November 15, 2018].
- Pasquale, Frank. 2015. *The Black Box Society: The Secret Algorithms that Control Money and Information*. Cambridge, MA: Harvard University Press.
- Rhue, Lauren. 2018. Racial Influence on Automated Perceptions of Emotions. *SSRN*, November 9. <https://dx.doi.org/10.2139/ssrn.3281765> [accessed November 9, 2018].
- Richardson, Lizzie. Forthcoming. Digital and Platform Economies. In *International Encyclopedia of Human Geography*, edited by Audrey Kobayashi. Amsterdam and New York: Elsevier.
- Rouvroy, Antoinette, and Thomas Berns. 2013. Algorithmic Governmentality and Prospects of Emancipation: Disparateness as a Precondition for Individuation through Relationships? *Réseaux* 1 (177).
- Sim, T., S. Baker, and M. Bsat. 2003. The CMU Pose, Illumination, and Expression Database. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 25 (12):1615-1618.
- Snow, Jacob. 2018. Amazon’s Face Recognition Falsely Matched 28 Members of Congress with Mugshots. July 26. <https://www.aclu.org/blog/privacy-technology/surveillance-technologies/amazons-face-recognition-falsely-matched-28> [accessed January 15, 2019].
- Stark, Luke. 2018. Algorithmic Psychometrics and the Scalable Subject. *Social Studies of Science* 48 (2): 204-231.
- Tufekci, Zeynep. 2014. Engineering the Public: Big Data, Surveillance and Computational Politics. *First Monday* 19 (7): n.p. <https://firstmonday.org/article/view/4901/4097>
- Wang, Y., and M. Kosinski. 2018. Deep Neural Networks Are More Accurate Than Humans at Detecting Sexual Orientation from Facial Images. *Journal of Personality and Social Psychology* 114 (2): 246-257.
- Warnock, Jason. 2015. How Facial-Recognition Software Will Shape the Future of Email Marketing. *Marketing Land*, Dec. 8. <http://marketingland.com/facial-recognition-software-will-shape-future-email-marketing-154713> [accessed March 2, 2016].
- Wu, Xiaolin, and Xi Zhang. 2017. Automated Inference on Criminality Using Face images. *Computing Science Repository*. abs/1611.04135. <https://arxiv.org>
- Zuboff, Shoshana. 2019. *The Age of Surveillance Capitalism. The Fight for a Human Future at the New Frontier of Power*. New York: Public Affairs.