Facial recognition and the right to appear

Infrastructural challenges in anti-surveillance resistance

Chapter 12

FACIAL RECOGNITION AND THE RIGHT TO APPEAR: INFRASTRUCTURAL CHALLENGES IN ANTI-SURVEILLANCE RESISTANCE

Benedetta Catanzariti

> Now that they are preparing the way for the automation of perception, for the innovation of artificial vision, delegating the analysis of objective reality to a machine, it might be appropriate to have another look at the nature of the virtual image. This is the formation of optical imagery with no apparent base, no permanency beyond that of mental or instrumental visual memory. Today it is impossible to talk about the development of the audio-visual without also talking about the development of virtual imagery and its influence on human behaviour, or without pointing to the new industrialisation of vision, to the growth of a veritable market in synthetic perception and all the ethical questions this entails. (Virilio 1994: 59)

The protests sparked by the murder of George Floyd in the summer of 2020 might represent the largest social movement in the history of the United States (Buchanan, Bui and Patel 2020). Data suggests that, between May and June, a little over 25 million people assembled across the country and marched together, speaking out against police brutality and structural racism, asking to divest, defund and abolish the police. Discussions about the demilitarisation of law enforcement have included the divestment of a wider set of surveillance practices: predictive policing, social network analysis and facial recognition. However, as Black Lives Matter (BLM) protests continued, US law enforcement deployed a wide array of surveillance techniques to monitor, target and arrest demonstrators (Kanno-Youngs 2020), undermining their very right to protest. This chapter considers the challenges that the police use of facial recognition poses to public assembly politics. I outline

these challenges by discussing notions of visibility and recognition in the context of public assemblies, which relates to the rights of urban citizens to assemble in public space. I bring together Judith Butler's reflections on public assembly politics and science and technology studies (STS) on infrastructures to illuminate how the classification choices underpinning the development of facial recognition, and particularly face datasets, affect the social and material conditions necessary to appear in public and protest – the infrastructural conditions of public assembly politics and public visibility.

FACIAL RECOGNITION IN 2020: NOTES ON ABOLITION

At the beginning of 2020, the *New York Times* published a story about Clearview AI, a small American tech company that had developed a facial recognition app matching faces to a database of three billion images scraped from the Internet. By uploading a face to the system, one could find any publicly available information matching that particular face. Hoan Ton-That, Clearview's founder, had envisioned his app to serve as a ground-breaking tool for law enforcement in the US and Canada. A month later, a BuzzFeed report uncovered that Clearview had provided its technology to more than 2,000 government agencies, law enforcement, and private companies across twenty-seven countries (Mac et al. 2020). Among these, the company had signed contracts with US Immigration and Customs Enforcement (ICE), the Attorney's Office for the Southern District of New York, the FBI, US Customs and Border Protection (CBP), Interpol, local police departments and the US department store Macy's. In the summer of 2020, the Miami police department and the New York police department (NYPD) used Clearview to identify BLM protesters (Lopatto 2020).

In June, 2,435 researchers in technical, legal and social science signed an open letter asking the Springer Editorial Committee to rescind a paper previously accepted for publication. The authors of the paper, among which was a NYPD veteran, claimed to have developed a machine learning system capable of predicting whether someone is a criminal based on a picture of their face. As the open letter notes, this type of claim is characteristic of a larger trend in machine learning research grounded on the assumption that biometric data can serve as an objective marker for criminal behaviour. This assumption is based in turn on physiognomy, the practice of interpreting someone's facial

features to infer their moral value, which in the nineteenth century became part of a larger project of social classification that harnessed empirical evidence to justify racism and sexism. Despite being discredited as pseudoscience, this practice continues to re-emerge under the guise of objective statistical models such as machine learning. In recent years, law enforcement and government agencies have adopted machine learning systems claiming that they improve objectivity and transparency – however, these tools perpetuate and reinforce discrimination and racial profiling, hidden behind the rhetoric of scientific and mathematical objectivity (Stop LAPD Spying 2018; Brayne 2020).

In reaction to the increasing awareness in civil society of facial recognition's limitations, errors and biases, and the potential discrimination towards marginalised communities, several US cities have recently banned the use of facial recognition by law enforcement and other municipal authorities. At the beginning of 2020, the European Commission considered a five-year ban on facial recognition in public spaces. Later in the year, Amazon, Microsoft and IBM announced a moratorium on the sale of facial recognition to police departments. These outcomes are the result of years of advocacy by scholars and community activists committed to critical and anti-racist technology (Buolamwini and Gebru 2018; Benjamin 2019; Nkonde 2019; Raji and Buolamwini 2019). However, there are several reasons why these solutions might not be sufficient to mitigate the harms of surveillance technology. First, big tech corporations such as Amazon, Microsoft and IBM are not the biggest players in the facial recognition market. As the Clearview story shows, a constellation of smaller and less well known actors provide policing tools to law enforcement, government agencies, insurance companies and banks with little to no public oversight. Second, while these corporations might have temporarily suspended their sales of facial recognition software, they haven't halted their engagement with law enforcement. For instance, in recent years police departments across the US have partnered with Ring, Amazon's home surveillance camera company, to get access to home owners' security footage without a warrant (Matzakis 2020). Further, Ring's users can upload their footage on Neighbors, a crime-reporting app owned by Amazon that has been known to promote peer-to-peer surveillance and reinforce racism (Haskins 2019). Finally, moratoria often centre the discourse on facial recognition on the concept of accuracy – the use of facial recognition is suspended until it can be proven to

be more accurate and 'fair'. Misclassification and misidentification are clearly harmful – most commercial facial recognition systems heavily misclassify darker-skinned individuals, and especially Black women, while the error rate for white men is close to 0 per cent (Buolamwini and Gebru 2018).

However, 'accurate' surveillance tools are equally problematic. Scholars have pointed out how surveillance technologies continue to be employed disproportionately on vulnerable and marginalised communities (Browne 2015; Eubanks 2018). Virginia Eubanks has shown how low-income communities of colour, especially women of colour and immigrants, are subject to intrusive surveillance by US government agencies, who perceived them as inherently dishonest and untrustworthy. As Eubanks notes, 'marginalized people are subject to some of the most technologically sophisticated and comprehensive forms of scrutiny and observation in law enforcement, the welfare system, and the low-wage workplace' (2014: n.p.). In this respect, the Carceral Tech Resistance Network[1] calls for the abolition of the 'police industrial complex': the entanglement of police, academia and private corporations responsible for the creation of police and surveillance tools as solutions to social, political and economic issues.

As millions of people have gathered in the streets across the globe, speaking out against police brutality and structural racism, asking to divest, defund and abolish the police, the question remains, then, how can a person demand an end to police brutality when the very right to gather in a city's public spaces to make those demands may be compromised by ubiquitous surveillance techniques?

NORMS OF RECOGNITION AND THE 'RIGHT TO APPEAR'

In *Notes Toward a Performative Theory of Assembly*, Judith Butler (2015) argues that the material and social conditions of appearance are prerequisites for political participation. While embodied forms of public assembly have a performative dimension that goes beyond that of political speech, not everyone, Butler points out, 'can appear in a bodily form' (2015: 8). Building on her work on gender performativity (1990), Butler argues that social norms shape people's 'conditions of appearance' (2015: 19), excluding from public recognition those who don't align with such norms, such as the gender binary. However, as she points out, in the context of public assembly, exclusion from public

appearance extends beyond gender nonconformity to all marginalised and vulnerable groups. The conditions of appearance, she notes, include 'infrastructural conditions of staging as well as technological means of capturing and conveying a gathering, a coming together, in the visual and acoustic fields' (p. 19). Here, Butler's emphasis is on the role that both mainstream and social media have in framing social movements' identity. The media do not simply amplify or silence the voices of those who demand justice; they participate in the very definition of people's identity. Media constitute 'the site of the hegemonic struggle over who "we" are' (p. 20). Butler draws on Hannah Arendt's definition of the public sphere as 'the space of appearance' where political action takes place (pp. 72–3). According to Arendt, this space can be 'anywhere and anytime', as political action happens 'between people' and does not require any specific location (p. 73). However, Butler argues that Arendt's definition of the public sphere is grounded on the classic notion of the Greek *polis*, where the domain of public politics was restricted to male citizens – while women, children, foreigners, slaves and lower-status subjects were deemed 'prepolitical' or 'extrapolitical' (p. 78).

Moving away from Arendt, Butler points out that 'we cannot presume the enclosed and well-fed space of the polis, where all the material needs are somehow being taken care of elsewhere by beings whose gender, race or status render them ineligible for public recognition' (p. 96). Moreover, this view ignores the importance of the body for political participation. Political action requires material supports – it is 'invariably bodily, even . . . in its virtual forms' (p. 73). These bodily supports – shelter, food, social and medical care, employment – are what make the action possible and are, quite often, part of the political struggle, or even its very object. If we take into account the material, bodily conditions underpinning Arendt's space of appearance, we ought to understand how these conditions are created, who can or cannot enter such space and who controls it. 'Significantly,' Butler argues, 'it is precisely this operation of power – the foreclosure and differential allocation of whether and how the body may appear – that is excluded from Arendt's explicit account of the political' (p. 88). An example of such power is the French ban on face coverings that, since 2011, prohibits women from wearing the niqab,[2] the Islamic face veil. Measures such as this are grounded on a supposed principle of universalism, as opposed to ethnic or religious separatism, that, nonetheless,

fails to recognise the right of those who do not conform with secular norms of appearance.

As I illustrate in the next section and beyond, the right to appear in public spaces is a constitutive property of the right to the city. Often outside the purview of democratic citizen oversight, data infrastructures can condition public visibility by excluding, discriminating against and harming those who don't conform with dominant norms of appearance. From this perspective, a critical discussion on the data infrastructures that regulate access to the public sphere contributes to the formulation of a notion of citizenship that hinges on the right to appear, rather than on legal status.

RECOGNITION AND EXPOSURE: VISIBILITY IN PUBLIC ASSEMBLIES

Public appearance, or visibility, entails recognition. Protesters gathered, and still gather, in the streets to make visible the reality of violence and marginalisation against Black and people of colour (POC) communities. This recognition is premised on an ideal notion of 'public privacy' (Slobogin 2002; Goold 2009; Aston 2017): the right to peacefully gather together without fear of prolonged and pervasive state surveillance, identification and interference. According to this notion, visibility within public assemblies is conditioned on assumed characteristics of public anonymity and homogeneity. In reality, visibility can lead to harmful exposure. This 'paradoxical double bind' of visibility (Brighenti 2007) means that we demand recognition by exposing ourselves to potential harm. This 'heightened bodily exposure' can happen, as Butler points out, when protesters and assembled crowds confront police in the streets, or in situations of territorial occupation, such that encountering a checkpoint, or even walking in public, can make a person vulnerable to violence, incarceration, even death.

Such unwanted and precarious bodily exposure is both the prerequisite and the aim of public assembly politics. Bodily exposure can be a form of political resistance. However, if the aim of politics is to establish better and equal life conditions such as shelter, health care, food and work, then Butler argues that the question of recognition cannot be treated as separate from the built environment, which shapes political action (2015: 127). How we are able to access and use the space of politics – the space of visibility – is critical for recognition. Writes Brighenti:

'It is not simply true that if I am disempowered or a society's outsider, then I am invisible. Rather, what happens is that I access visibility places in ways that are largely or completely out of my control' (2007: 333). When our faces are captured on camera and assigned an identity – a set of demographic characteristics, emotions and personality traits, in ways that are outside our purview – these inferences can have a detrimental impact not only on our individual rights, but also on collective practices of participatory democracy, as scholars and activists have warned about the chilling effect of facial and emotion recognition technologies on protesters and its adverse impact on freedom of assembly and the right to protest (Chowdhury 2020; ARTICLE 19 2021).

Much discussion has taken place about the asymmetrical relationship of visibility underpinning surveillance technologies, and how this asymmetry is informed by practices of social classification and forms of systemic oppression – racism, sexism, ableism and capitalism (Browne 2015; Dubrofsky and Magnet 2015; Zuboff 2019). Facial recognition is and has been deeply entwined with these practices since its very inception. In the next sections, I offer a historical and material account of the practices that inform the work of classification underpinning facial recognition technologies and, in particular, of face datasets. Building on STS work on infrastructures, Stevens and Keyes have argued that facial recognition technology 'has no unifying essence', but it is rather a 'shifting web of programmers, algorithms, datasets, testing standards, formatting requirements, law enforcement agents and other operators and users' (2021: 2). It is only by looking at specific facets of facial recognition technologies that we can recognise what values and politics they perpetuate.

BUILDING INFRASTRUCTURES: FACIAL RECOGNITION IN HISTORY

Technologies that use computer vision to process digital images of the face differ in use and scope. Facial identification – the task of matching a face to a unique identity – is not the same as facial analysis – the task of classifying facial features by race, gender, age, emotions or personality types. However, while surveillance generally involves a process of individual identification, its ultimate goal is social classification, which is, as Brighenti argues, 'a type of classification that is essentially grounded in the *summa divisio* between safe and dangerous subjects'

(2007: 333). In this respect, facial identification and facial analysis are both functions of a project of social classification that is infused with specific cultural and social assumptions.

Early artificial intelligence (AI) research on facial recognition developed in the context of suspect identification, started and funded by US intelligence agencies (Raji and Fried 2021; Stevens and Keyes 2021). Despite the effort of governments and corporations to frame current applications of facial recognition as benign or even humanitarian, this military history 'has shaped everything from the nature of the data collected for benchmarks to the nature of evaluation metrics, and certainly the definition of tasks' (Raji and Fried 2021: 8). AI systems that can 'see' – detect, recognise and classify a face – are trained on large-scale datasets with thousands of facial images categorised according to some form of taxonomy. From the early 1960s until the late 2000s, when social media made large quantities of face images suddenly available, the US government supported facial recognition research by providing mugshot databases as training data. The MEDS dataset, for instance, contained a collection of mugshots of previously arrested and deceased subjects. Write Stevens and Keyes (2021), these subjects, 'whether convicted or not, under laws that may or may not have been valid in the mid-2000s – had their mugshots taken and reused, in some cases up to 40 years after their arrest, for the purpose of further refining law enforcement tools of surveillance and control' (p. 10). Moreover, in the process of data curation – the organisation and standardisation of inconsistent images and metadata – researchers assigned values of gender and race according to their own judgement. Images of injured and bruised subjects, 'screaming and looking away from the camera', were treated as inconsistencies posing a problem to the tagging process, a decision that goes unremarked upon, 'and evidently poses no issue for the purposes of the dataset's developers and users – blood and bruises are not part of their remit unless it interferes with the algorithmic gaze' (p. 12).

In 2019, in response to growing critiques around the disproportionate impact of facial recognition on Black and POC communities, IBM created Diversity in Faces (DiF), a dataset of one million face images sourced from Flickr 'for advancing the study of facial diversity' (Merler et al. 2019: 1). DiF sets out to achieve statistical fairness in face representation by including annotations of craniofacial features such as craniofacial distances, ratios and areas and facial symmetry. As mentioned

earlier, the idea that facial features can serve as an objective marker of race and gender identity echoes nineteenth-century efforts to harness empirical evidence – in this case, anthropometric measurements – to justify the classification and discrimination of social groups according by race, gender and class (Gould 1981). The invention of photography and the rise of social statistics further grounded anthropometry as a method of social regulation. Francis Galton's composite portrait applied statistical methods to group classification: by overprinting hundreds of photographs of individuals he believed to be of the same type – 'the criminal', 'the Jew' – Galton hoped to generate portraits of ideal characters and materialise the visual evidence of hereditarian laws (Sekula 1986). The shift from individual identification to group classification relies on statistical inference and pattern recognition that can also be observed in early facial recognition research (Lee-Morrison 2018) and underpins the logics of today's machine learning and Big Data.

MAINTAINING INFRASTRUCTURES: VISIBLE BODIES, INVISIBLE LABOUR

In machine learning, algorithms 'learn' to recognise or classify faces from labelled data. This approach is called supervised learning, and the training data consists of a set of images that have been manually labelled – categorised according to some form of taxonomy, such as 'male' or 'female'. Bowker and Star have famously argued that the technical choices we make about classifications have significant yet often invisible ethical and political implications (Bowker and Star 2000). Similarly, the classification choices embedded in datasets can reinforce and normalise appearances and behaviours that are aligned with cultural and social norms (Uliasz 2020). Datasets govern the way AI systems see the world, and as these are increasingly embedded into our social life, the project of interpreting and labelling images is political, rather than merely technical (Crawford and Paglen 2019). For instance, work in human–computer interaction (HCI) has shown how limitations in gender representation – the exclusion of transgender and gender non-conforming identities from datasets – can cause harm and reinforce stereotypical assumptions about gender appearance (Hamidi, Scheuerman and Branham 2018; Keyes 2018). Furthermore, even if the inclusion of every possible gender expression was technically feasible, this line of research questions the idea that it is possible (and

meaningful) to infer gender from facial features. Klaus Scheuerman and colleagues analysed over ninety facial image datasets to investigate the categories and assumptions at work in the definition of gender and race underpinning computer vision; they found that only a few of these datasets contain underlying information on how gender and race identity is classified and that, when such information is available, gender and race are presented as unquestionable and apolitical and assigned on the basis of visible and physical appearance (Scheuerman et al. 2020). However, as the authors note, identity is sociohistorical, rather than merely physical. Ironically, this dissonance is reflected in the variety of interchangeable concepts ('skin type', 'race', 'ethnicity') and categories ('Black', 'African-American', 'Caucasian', 'White', 'Hispanic', 'Indian', 'Middle Eastern', 'Other') that are employed to define race in most datasets (2020: 16–17).

In addition to the social and cultural assumptions that inform how race and gender are framed, the power structures that shape the AI design pipeline can also affect the labelling process. Often, researchers and practitioners rely on out-sourced annotation workers to label face images. These workers can be either employed by annotation companies with traditional managerial structures, or work for on-demand work platforms, such as Amazon's Mechanical Turk. In both cases, workers are often recruited from poor and vulnerable populations to provide data annotation services at competitive prices (Irani 2013; Gray and Suri 2019; Miceli, Schuessler and Yang 2020). After conducting ethnographic observations within three different annotation companies, Miceli and colleagues (2020) concluded that multiple factors contribute to the interpretation and annotation of data. First, standards imposed by commissioning clients (their expectations and requirements) shape the workers' interpretation of the images they are labelling. Second, the distribution of workload among clients, team leaders, reviewers and annotators, internal or outsourced, creates multiple layers of meaning that can infuse the annotation process. Third, standards and layers reinforce the idea that labels are self-evident and discourage annotators from questioning them, normalising the classification choices. Fourth, these annotation companies and their clients are mostly concerned with speed and cost-efficiency, rather than the ethical and social implications of their work. Embedded in the work of classification are technical assumptions that require the annotators to assign mutually exclusive labels to categories like race and gender.

As the authors point out, 'whether such categorisation captures the realities of data subjects or coincides with the values and beliefs of data workers is not negotiated' (Miceli et al. 2020: 5). Finally, while the authors chose companies with traditional management structures, they note how crowdsourced annotation companies, such as Mechanical Turk, make hierarchies of power more opaque and further obfuscate accountability. Here, the material conditions of 'ghost work' (Gray and Suri 2019) – the vulnerable status of crowdsourced workers – can further infuse the annotation process.

HAPPY OR HOSTILE? WHAT FACES (CAN'T) TELL

In recent years, emotion recognition technologies have found applications in sectors ranging from monitoring emotional states of patients for improved health care delivery, to illuminating consumer behaviour in the retail sector and monitoring drivers' attention to enhance road safety. Recent research in computer science suggests that emotion recognition could be used to monitor emotional states of people in airports, public spaces and borders for security reasons and by law enforcement to monitor crowds and identify violent group behaviour (Holder and Tapamo 2017).

What makes possible the idea that AI systems could recognise human emotions, moods or personality types is a popular notion of contemporary psychology: the theory that there is a set of discrete, fixed emotions universally conveyed through the same patterns of facial expressions. Psychologist Paul Ekman developed this theory in the 1960s; he identified six universal emotions (fear, disgust, anger, sadness, surprise and happiness) through the measurement and quantification of facial micro-expressions. In 1978, together with his colleague Wallace Friesen, Ekman published the Facial Action Coding System (FACS). FACS breaks down the movements of facial muscles into twenty-eight action units (AU), along with additional codes for head and eye movements, responsible for the spectrum of facial expressions. Ekman was influenced by Charles Darwin's 1872 *The Expression of Emotions in Man and Animals*, as well as a Princeton Professor, Silvan Tomkins, who insisted on the existence of discrete, biologically determined, affective expressions. Tomkins's fascination with the face manifested in his legendary ability to interpret human expressions. According to a *New Yorker* piece on facial recognition, Tomkins was known to be 'the best

face reader there ever was' and that he could 'walk into a post office, go over to the "Wanted" posters, and just by looking at mug shots, tell you what crimes the various fugitives had committed. Tomkins felt that emotion was the code to life and that with enough attention to particulars the code could be cracked' (Gladwell 2002).

Psychologists, cognitive scientists and anthropologists have all challenged Ekman's theory. Recently, psychologist Lisa Feldman Barrett and colleagues have argued that the available scientific evidence fails to support Ekman's view, concluding that the ways in which people communicate emotions vary significantly across cultures and even across the same individual. More importantly, Ekman's theory assumes that it is possible to infer emotional states from facial movements and that the relationship between these is measurable, universal and consistent. Yet research has shown that the face alone is not sufficient to understand someone's emotional state (Barrett et al. 2019). Visual context changes our perception of affective behaviour in unique ways, and contextual information helps us recognise other people's facial expressions. Other factors influence the reliability of Ekman's theory: most experiments ask participants to use a set of predetermined labels (joy, anger, sadness, etc.) to recognise emotions, acting as a constraint on their choices. Moreover, research on emotion perception often employs images of actors performing emotions or uses computer-generated facial expressions, which does not correspond to some person's real, existing emotional state (Barrett et al. 2019: 29).

Despite the ongoing scientific debate, Ekman's theory is the tenet of contemporary emotion recognition technologies. Most of these systems build on Ekman's Emotion FACS (EMFACS), a system designed to match emotions with the facial expressions coded within the FACS. Ekman and his colleagues built a database for the EMFACS data called Facial Action Coding System Affect Interpretation Dictionary (FACSAID), which contains images of coded facial expressions and their assigned emotional meaning (Gates 2011). This approach performs a reverse inference or, as Kelly Gates has noted, a sort of reverse engineering of the emotion it aims to classify. As Gates pointed out, the accuracy of EMFACS depends on the authority of the experts who interpreted facial behaviours and attached meanings to them. Similarly, emotion recognition systems use machine learning algorithms to track the shape and movement of facial features, such as the corners of the mouth, the corners of the eyebrows or the outline of the nose, and

compare them to images of facial expressions stored in a training dataset and labelled as specific emotions. As in Ekman's system, the face is treated as information that can be extracted from its context and separated from the subject's intentions.

BRIDGES AND BARRIERS: LESSONS FROM INFRASTRUCTURE STUDIES

A critical analysis of the classification practices and choices that underpin the creation of datasets is useful to understand the limitations of facial recognition technologies. From this perspective, debates about bias or technical solutions, such as improving the information on the phenotypic composition of training datasets or deleting offensive and problematic labels, fail to address the political role of datasets. As Crawford and Paglen (2019) argue, the whole project of 'collecting images, categorizing them, and labeling them is itself a form of politics, filled with questions about who gets to decide what images mean and what kinds of social and political work those representations perform' (p. 33).

Indeed, the machine does not impose the conditions of appearance; they are instead the result of multiple layers of technical assumptions and choices that reinforce dominant social and cultural norms. However, if the right to appear is the prerequisite for political participation, as Butler notes, how do we reconcile recognition with the technical constraints imposed on the bodies entering the visual field? Butler (2015) points out how the 'highly regulated field of appearance does not admit everyone, requiring zones where many are expected not to appear or are legally proscribed from doing so' (p. 35). The norms embedded in facial recognition datasets structure the visual field so that certain identities are excluded from it. Paraphrasing Butler, these identities are *technically proscribed* from appearing. Importantly, Butler notes how

> the compulsory demand to appear in one way rather than another functions as a precondition of appearing at all. And this means that embodying the norm or norms by which one gains recognizable status is a way of ratifying and reproducing certain norms of recognition over others, and so constraining the field of the recognizable. (Ibid.)

Building on Butler's theory of gender performativity, Tobias Matzner argues that data-driven surveillance practices participate in the production of suspect subjects, by containing 'every bit of data under the performative of suspicion', so that a person loses all resources or recourse to contest the suspicions about them (Matzner 2016: 209). Similarly, facial recognition datasets exert performative power over the production of racialised, gendered and criminal identities. In this context, when those who are excluded by the norm that 'they are expected to embody' (Butler 2015: 37) gather in the streets to protest, they are first and foremost reclaiming the field of appearance: the ability to appear in public unconditioned by dominant norms of visibility.

If we look at facial recognition datasets in terms of infrastructure, we can recognise how those norms and conditions of appearance are produced. Building on Bowker and Star's (2000) definition of infrastructure (p. 35), datasets are similarly embedded into larger sociotechnical systems by being integrated into multiple and varied applications; they blur boundaries between public and private use; and their scope goes beyond a single application or use-case. Moreover, as research on data annotation has shown, a dataset, like infrastructure, is part of the norms of a 'community of practice' (p. 35). Power dynamics and classification standards within annotation companies affect the interpretation and labelling of data. Importantly, infrastructure is most visible when it breaks (p. 35). Similar to a power outage manifesting the reach and scope of the electric grid, we realise the extent of facial recognition when cameras cannot see dark skin or when they prompt a wrongful arrest.

Building on the classic example of Robert Moses, the New York city planner who designed the bridges over the Grand Central Parkway low enough so that public transport – hence, low-income families – could not reach the wealthy suburbs of Long Island (Winner 1980), Susan Star (1999) has famously argued that moral values are inscribed in the information environment through everyday design choices and standards. As Star notes, some of these choices are flexible and could be changed with time and knowledge. Others 'present barriers to users that may only be changed by a full-scale social movement' (p. 389). In the light of these considerations, what are the implications for anti-facial recognition resistance? Looking at facial recognition through the lens of infrastructure can inform activist interventions and critiques of these systems in the form of what Bowker (1994) has called

'infrastructural inversion', which is an effort to attend to the often hidden choices, standards and constraints underpinning the development of facial recognition systems.

Star's investigation of infrastructure offers some useful strategies: infrastructures can be inspected through looking for the invisible labour that maintains them and for finding their 'master narratives' and their 'others' (pp. 384–6). Critical literature on data work has shown the often invisible and precarious labour behind much of AI-fuelled automation (Irani 2015; Gray and Suri 2019). Critiques of surveillance technologies are not fully separable from accounts of the capitalist ecosystem that supports the development of such technologies. Identifying and deconstructing master narratives means recognising how we lose sight of the circumstances that gave rise to a given accepted, scientific fact, and how, despite the equivocations that gave rise to it, it achieves the status of truth (Star 1999: 385). Infrastructural reading is a way to challenge and contest those narratives premised on claims that facial recognition systems can actually make a range of inferences from a person's face. In this respect, exposing the culturally and politically situated classification work underpinning facial recognition might have practical consequences for intervention: what would resistance look like if it rejected claims of gender and race recognition, mood or personality detection? Also, this approach might prove useful for academics investigating the implications of facial recognition systems and contesting the imaginaries at work in the use and development of AI surveillance techniques.

CONCLUSION

By looking at facial recognition technologies through the lens of infrastructure, I have sought to illuminate the ways in which these technologies condition public visibility. Public assemblies sit at the uncomfortable intersection of recognition and anonymity. In truth, far from an ideal notion of public privacy, assemblies take place in a visual field within which gender and race, along with other assumptions about identity, are unproblematically assigned to individuals and in which technical choices reinforce the unequal distribution of power. Questions about who can or cannot enter such spaces and who makes these decisions are of foremost importance. If we look at the classification choices embedded in the AI pipeline, we can observe how cultural and political contexts, organisational settings, power structures,

standards and discretionary decisions all shape the performative power of these systems. Famously, Bowker and Star (2000) have shown that resistance can take the form of exposing how standards and classifications themselves do political work in the world (p. 49). And since the right to public assembly and private publicity is a component of the broader right to the city, the data infrastructuress constituting these spaces are also a key part of these struggles. From this perspective, illuminating the visual epistemologies that inform the use and development of facial recognition can open up new possibilities for resistance movements to challenge them.

NOTES

1. https://www.carceral.tech/
2. At the time of this writing, and under the ongoing COVID-19 pandemic, France has enforced face mask rules, while simultaneously implementing measures to ban women under the age of 18 from wearing the hijab in public spaces.

REFERENCES

ARTICLE 19. (2021) *Emotion Recognition Technology Report*. Available from https://www.article19.org/emotion-recognition-technology-report/ (last accessed 13 January 2022).

Aston, V. (2017) State surveillance of protest and the rights to privacy and freedom of assembly: a comparison of judicial and protester perspectives. *European Journal of Law and Technology*, 8(1).

Barrett, L. F., Adolphs, R., Marsella, S., Martinez, A. M. and Pollak, S. D. (2019) Emotional expressions reconsidered: challenges to inferring emotion from human facial movements. *Psychological Science in the Public Interest*, 20(1), 1–68.

Benjamin, R. (2019) *Race after Technology: Abolitionist Tools for the New Jim Code*. Cambridge: Polity Press.

Bowker, G. C. (1994) *Science on the Run: Information Management and Industrial Geophysics at Schlumberger, 1920–1940*. Cambridge, MA: MIT Press.

Bowker, G. C. and Star, S. L. (2000) *Sorting Things Out: Classification and Its Consequences*. Cambridge, MA: MIT Press.

Brayne, S. (2020) *Predict and Surveil: Data, Discretion, and the Future of Policing*. New York: Oxford University Press.

Brighenti, A. (2007) Visibility: a category for the social sciences. *Current Sociology*, 55(3), 323–42.
Browne, S. (2015) *Dark Matters: On the Surveillance of Blackness*. Durham, NC: Duke University Press.
Buchanan, L., Bui, Q. and Patel, J. K. (2020) Black Lives Matter may be the largest movement in U.S. history. *New York Times*. Available at https://www.nytimes.com/interactive/2020/07/03/us/george-floyd-protests-crowd-size.html (last accessed 13 January 2022).
Buolamwini, J. and Gebru, T. (2018) Gender shades: intersectional accuracy disparities in commercial gender classification. In *Conference on Fairness, Accountability and Transparency* (pp. 77–91). PMLR.
Butler, J. (1990) *Gender Trouble, Feminism and the Subversion of Identity*. London: Routledge.
Butler, J. (2015) *Notes Toward a Performative Theory of Assembly*. Cambridge, MA: Harvard University Press.
Chowdhury, A. (2020) Unmasking facial recognition: an exploration of the racial bias implication. Available at https://webrootsdemocracy.files.wordpress.com/2020/08/unmasking-facial-recognition-webroots-democracy.pdf (last accessed 13 January 2022).
Crawford, K. and Paglen, T. (2019*) Excavating AI*. Available at https://excavating.ai (last accessed 13 January 2022).
Dubrofsky, R. E. and Magnet, S. A. (eds) (2015) *Feminist Surveillance Studies*. Durham, NC: Duke University Press.
Eubanks, V. (2014) Want to predict the future of surveillance? Ask poor communities. *The American Prospect*. Available at https://prospect.org/api/content/36656b9e-c446-5205-9257-0120f64aabdb/ (last accessed 13 January 2022).
Eubanks, V. (2018) *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. New York: St. Martin's Press.
Gates, K. A. (2011) *Our Biometric Future: Facial Recognition Technology and the Culture of Surveillance*, vol. 2. New York: New York University Press.
Gladwell, M. (2002) The naked face. *The New Yorker*. Available at https://www.newyorker.com/magazine/2002/08/05/the-naked-face (last accessed 13 January 2022).
Goold, B. (2009) Surveillance and the political value of privacy. *Amsterdam Law Forum*, 1(4).
Gould, S. J. (1981) *Mismeasure of Man*. New York: W. W. Norton.
Gray, M. L. and Suri, S. (2019) *Ghost Work: How to stop Silicon Valley from Building a New Global Underclass*. Eamon Dolan Books.
Hamidi, F., Scheuerman, M. K. and Branham, S. M. (2018) Gender recogni-

tion or gender reductionism? The social implications of embedded gender recognition systems. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* (pp. 1–13). New York: ACM.

Haskins, C. (2019) Amazon's home security company is turning everyone into cops. *Vice*. Available at https://www.vice.com/en/article/qvyvzd/amazons-home-security-company-is-turning-everyone-into-cops (last accessed 13 January 2022).

Holder, R. P. and Tapamo, J. R. (2017) using facial expression recognition for crowd monitoring. In *Pacific-Rim Symposium on Image and Video Technology* (pp. 463–76). Cham: Springer.

Irani, L. (2013) The cultural work of microwork. *New Media & Society*, 17(5), 720–39.

Irani, L. (2015) Justice for 'data janitors'. Public Books. Available at https://www.publicbooks.org/justice-for-data-janitors/ (last accessed 13 January 2022).

Kanno-Youngs, Z. (2020) U.S. watched George Floyd protests in 15 cities using aerial surveillance. *New York Times*. Available at https://www.nytimes.com/2020/06/19/us/politics/george-floyd-protests-surveillance.html (last accessed 13 January 2022).

Keyes, O. (2018) The misgendering machines: trans/HCI implications of automatic gender recognition. In *Proceedings of the ACM on Human–Computer Interaction*, vol. 2 (CSCW) (pp. 1–22).

Lee-Morrison, L. (2018) A portrait of facial recognition: tracing a history of a statistical way of seeing. *Philosophy of Photography*, 9(2), 107–30.

Lopatto, E. (2020) Clearview AI CEO says 'over 2,400 police agencies' are using its facial recognition software. *The Verge*. Available at https://www.theverge.com/2020/8/26/21402978/clearview-ai-ceo-interview-2400-police-agencies-facial-recognition (last accessed 13 January 2022).

Mac, R., Haskins, C., Sacks, B. and McDonald, M. (2020) How a facial recognition tool found its way into hundreds of US police departments, schools, and taxpayer-funded organizations. BuzzFeed News. Available at https://www.buzzfeednews.com/article/ryanmac/clearview-ai-local-police-facial-recognition (last accessed 13 January 2022).

Matzakis, L. (2020) Amazon doubles down on ring partnerships with law enforcement. *WIRED*. Available at https://www.wired.com/story/ces-2020-amazon-defends-ring-police-partnerships/ (last accessed 13 January 2022).

Matzner, T. (2016) Beyond data as representation: the performativity of Big Data in surveillance. *Surveillance & Society* 14, 197–210.

Merler, M., Ratha, N., Feris, R. S. and Smith, J. R. (2019) Diversity in faces. *arXiv preprint* arXiv:1901.10436.

Miceli, M., Schuessler, M. and Yang, T. (2020) between subjectivity and imposition: power dynamics in data annotation for computer vision. *arXiv:2007.14886 [cs]*.

Nkonde, M. (2019) Automated anti-Blackness: facial recognition in Brooklyn, New York. *Harvard Journal of African American Public Policy*, 20, 30–6.

Raji, I. D, and Buolamwini, J. (2019) Actionable auditing: investigating the impact of publicly naming biased performance results of commercial ai products. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society* (pp. 429–35).

Raji, I. D. and Fried, G. (2021) About face: a survey of facial recognition evaluation. *arXiv:2102.00813 [cs]*.

Scheuerman, M. K., Wade, K., Lustig, C. and Brubaker, J. R. (2020) How we've taught algorithms to see identity: constructing race and gender in image databases for facial analysis. In *Proceedings of the ACM on Human–Computer Interaction*, 4 (CSCW1), pp. 1–35.

Sekula, A. (1986) The body and the archive. *October*, 39, 3–64.

Slobogin, C. (2002) Public privacy: camera surveillance of public places and the right to anonymity. Miss. 72 *Mississippi Law Journal*, 213

Star, S. L. (1999) The ethnography of infrastructure. *American Behavioral Scientist*, 43(3), 377–91.

Stevens, N. and Keyes, O. (2021) Seeing infrastructure: race, facial recognition and the politics of data. *Cultural Studies*, 35(4–5), 833–53.

Stop LAPD Spying Coalition (2018). Before the bullet hits the body. Available at https://stoplapdspying.org/before-the-bullet-hits-the-body-dismantling-predictive-policing-in-los-angeles/ (last accessed 13 January 2022).

Uliasz, R. (2020) Seeing like an algorithm: operative images and emergent subjects. *AI & Society*, 36(6), 1–9.

Virilio, P. (1994) *The Vision Machine*. Bloomington: Indiana University Press.

Winner, L. (1980) Do artifacts have politics? *Daedalus*, 121–36.

Zuboff, S. (2019) *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. London: Profile Books.